

**LIFE CYCLE INVENTORY UNCERTAINTY IN
RESOURCE-BASED INDUSTRIES -
A FOCUS ON COAL-BASED POWER GENERATION**

Philippa Josephine Notten

**Thesis Presented for the Degree of
DOCTOR OF PHILOSOPHY
in the Department of Chemical Engineering
UNIVERSITY OF CAPE TOWN**

October 2001

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

ACKNOWLEDGEMENTS

The financial support of Eskom and the National Research Foundation is gratefully acknowledged. Opinions expressed in this work, or conclusions arrived at, are those of the author alone, and should not necessarily be attributed to either sponsor.

My thanks to Jim Petrie, for the many opportunities this project has afforded me, and for his guidance and supervision, through the difficult circumstances of rarely being on the same continent at the same time. This work would not have been possible without his continued vision and encouragement.

The work of Yvonne Hansen in providing the “footprint” analyses used in the case studies is gratefully acknowledged. Thanks and appreciation are especially due to her for always managing to co-ordinate her work schedules with mine, and getting these to me in such good time and with such good humour.

The time and assistance of the many Eskom and affiliated personnel, too numerous to mention by name, who fitted me into their busy schedules and put up with endless requests for information, is gratefully acknowledged. In particular, thanks are due to Brian Gore, for his co-ordinating role.

A great deal of thanks are owed to the many people that have offered me help and advice, and supported and encouraged me over the years of this work. To mention but a few: Mary Stewart, for many helpful and stimulating conversations; Lauren Basson, for many a shared “Eskom experience”; the many “Greenhouse” members I have crossed over with, for friendship and always a pleasant working environment, especially Shehnaaz Moosa, for always being there to talk to, and Sue Buerger for keeping it all running so smoothly. Special thanks are also due to Sue for help in handing in this thesis long-distance.

My thanks to my family for putting up with me through difficult times, and for keeping me going by reminding me what it is really all about. I owe too much to my parents to ever be able to adequately express my gratitude, without their letting me believe that I could do this, and providing me the means to do so, none of this would have been possible, or even worth contemplating.

Finally, to Petrus, where thanks just seem so terribly inadequate, but it would take another thesis to cover all that he has done for me, and all that his love and support has meant to me. Suffice to say, I could not have done it without him.

ABSTRACT

The aim of this thesis is to develop an approach to support prospective environmental decision-making in resource-based industries. The specific focus is on coal-based power generation. The objectives of the approach are that it be able to adequately reflect the environmental burdens arising from primary industries, and to make explicit the trade-offs often encountered in environmental decisions. In addition, it needs to take into account that the context in which the assessment takes place affects data availability and quality significantly, and consequently the certainty with which systems can be evaluated.

Resource-based processes typically involve large-scale disruption of the local and regional environments, with imprecise processes and diffuse emissions. The modelling of the environmental performance of such processes therefore raises significant challenges, where many disparate sources of data, available at different levels of aggregation, and over various time intervals, have to be brought together into a coherent assessment. An “uncertain” definition of the system is therefore much more meaningful, in which variables are defined over ranges of values to cover inconsistencies and imbalances in the system. The inherently high variability of mining and minerals processes further supports their modelling as ranges of potential performance rather than “typical” operations, where the relevant process of interest must be identified and the variability within the particular process incorporated into the assessment.

Life cycle assessment (LCA) has received increasing attention for its role in environmental decision-making processes, where it supports the process of defining the contribution of human activities to (at least the environmental dimension of) sustainable development. It is therefore the structured approach to environmental decision-making investigated in this thesis to organise the large data sets of varying quality and completeness available around resource-based industries into useful information, able to provide the environmental objective in a decision-making process. LCA is an inherently uncertain procedure in that it combines data sources of varying accuracy and representativeness, and employs subjective judgement in applying this data to future operating systems. Subjective judgements are also present in the definition of the systems, and in the modelling choices determining the accuracy and complexity of the inventory and impact models used. Nonetheless, LCA results are most often presented as single values, which in a comparative analysis, gives the often incorrect impression that one system is always better or worse than another system.

A framework has been developed in this thesis to include all relevant sources of uncertainty encountered in LCA models explicitly, where empirical parameter uncertainty, model parameter uncertainty, and uncertainty in model form are investigated in a looped fashion. The innermost loop assesses empirical uncertainty in an iterative probabilistic analysis, using Latin Hypercube sampling of the uncertain input distributions to propagate the data uncertainty to the output, and rank-order correlation analyses to determine the relative uncertainty importance of the parameters input into the model. Model parameter uncertainty is assessed next, by a parametric analysis, or by a combination of sensitivity analyses and a parametric analysis, if a large number of model parameters require consideration. The top-most layer is an assessment of model form, in which alternative model forms are investigated in a sensitivity analysis.

The LCIA impact categories in commonly used assessment methodologies were found to be deficient for a full assessment of primary industries, partly because of the level of inventory detail available, and partly because of LCIA's inherent limitations with respect to predicting site-specific impacts, and those with a time dependency. This is of particular consequence to resource-based industries. Water-related impacts, particularly those arising from the extremely large dedicated waste deposits of resource-based processes, are identified as particularly poorly assessed by LCIA. The "impacted land footprint" is proposed as a "proxy" indicator to provide an assessment of the high volume, low-grade waste deposits of resource-based processes, whilst avoiding the problems associated with predicting the individual metal species leaching from complex waste materials, and the fact that the toxicological factors developed for metal species are of questionable accuracy.

Case studies in the context of coal-based power generation present the key features of the methods developed, and cover the three broad decision contexts identified for resource-based industries. The first case study presents the situation where the decision context is in parallel to production system, i.e. the decision-maker merely imports information from another system (e.g. background LCI data). In this case study an inventory representing the average grid mix of South African coal-based power is developed. The other two case studies explore the use of LCA to support prospective decision making, ranging from strategic to operational levels. In the former, the production system falls within the decision context (i.e. the decision-system contains elements external to the system generating the information), whilst in the latter the decision context is within the production system (i.e. the decision-system falls wholly within the system generating the information). Since the different decision models are based on the flow of information to characterise the system, their choice has significant

implications on the quantity and quality of data required of the study, and consequently on the uncertainty present in the system. The decision context has particular implications with regards to the definition of the problem, and thus the relative importance of the various sources of uncertainty.

This work contributes to the development of an approach to support prospective environmental decision-making. A procedure to guide the assessment of loosely defined option sets is developed, along with an ability to incorporate the often high empirical data uncertainty of future systems. The outcome of the analysis is an explicit consideration of the full decision-space in which the future systems could act. The powerful data analysis and presentation features of Principal Component Analysis (PCA) are used to provide an enhanced interpretation of the large information sets resulting from the incorporation of an uncertainty analysis. PCA is able to provide information with respect to the strength and independence of the criteria chosen to evaluate the systems, as well as enabling a powerful “graphical summary” of the results. This allows an exploration of the full operating space, clearly identifying the trade-offs between the various options. A stochastic modelling approach considerably increases the relevance of the LCA results, by allowing a degree of confidence to be attached to the perceived differences between the options. The iterative uncertainty analysis procedure developed facilitates the incorporation of empirical data uncertainty by using rank order correlation analyses to focus the analysis on those parameters with high uncertainty importance. The combined model parameter and empirical parameter uncertainty analysis therefore considerably enhances the decision-support capabilities of the LCA process.

This work also has significance for the environmental management of coal-based power generation. The need for a comprehensive assessment, incorporating all aspects of power generation and fuel procurement, has been demonstrated. The case studies show that firm conclusions can not be drawn around technology choice within either the boiler or water plants without explicit consideration of the impacts of solid waste disposal. The quantitative evaluation of the leachate potential of the solid waste dumps, provided by the “impacted land footprint”, is therefore invaluable. The enhanced decision-support possible with the uncertainty analysis procedure developed is demonstrated in the case studies, which show that incorrect conclusions could be drawn without a consideration of the high uncertainty and variability in these processes. The uncertainty analysis allows the operating parameters to be determined for which the system has a high probability of meeting the desired objective. In particular, the analysis allowed the combination of decision variables to be selected for which a definitive statement could be made that fluidised combustion of coal discards is environmentally preferable to

pulverised combustion of run-of-mine coal for power plant re-powering. In the context of South Africa's development needs, this is an important observation, unlocking as it does, a secondary energy source, whilst addressing the significant adverse environmental impacts of current discard management practices.

TABLE OF CONTENTS

Abstract.....	i
List of Abbreviations.....	ix
Terminology.....	xi

CHAPTER 1: INTRODUCTION

1.1. LIFE CYCLE ASSESSMENT AS A DECISION-MAKING PROCESS.....	1-2
1.1.1. Steps of the Decision-Making Process.....	1-5
1.1.2. Overview of Life Cycle Assessment.....	1-6
1.2. THE ASSESSMENT OF PRIMARY INDUSTRIES.....	1-10
1.2.1. Features of Primary Industries.....	1-10
1.2.2. Life Cycle Assessment of Primary Industries.....	1-12
1.2.3. A Focus on Coal Based Power Generation.....	1-13
1.3. THEMES AND THESIS STRUCTURE.....	1-15
1.3.1. Inventory Modelling for Decision Support in Primary Industries.....	1-15
1.3.2. Modelling Uncertain Systems.....	1-16
1.3.3. Indicators for Resource-Based Industries.....	1-17
1.3.4. Case Studies.....	1-18

SECTION I: METHODOLOGY AND MODEL DEVELOPMENT

CHAPTER 2: CONTEXT SPECIFIC INVENTORY MODEL DEVELOPMENT

2.1. INVENTORY MODELLING.....	2-2
2.1.1. Systems Modelling in LCA.....	2-2
2.1.2. System Boundary Definition.....	2-6
2.1.3. System Complexity.....	2-8
2.1.4. Inventory Model Structure.....	2-11
2.2. APPLICATION DEPENDENCY OF LCA.....	2-13
2.3. DECISION CONTEXTS FOR PRIMARY INDUSTRIES.....	2-17
2.3.1. Type I (Background data / Product declarations).....	2-20
2.3.2. Type II (Tactical / Strategic Decisions).....	2-22
2.3.3. Type III (Operational Decisions).....	2-24
2.3.4. Case Study Decision Contexts.....	2-25

CHAPTER 3: REVIEW OF ANALYSIS AND MANAGEMENT OF UNCERTAINTY IN LCI

3.1. SOURCES OF UNCERTAINTY.....	3-2
3.1.1. Sources of Uncertainty in Models.....	3-2
3.1.2. Sources of Uncertainty Relevant to LCA Models.....	3-8
3.2. METHODS FOR THE ASSESSMENT OF UNCERTAINTY.....	3-19
3.2.1. Propagation and Analysis of Uncertainty.....	3-20
3.2.2. Reduction of Uncertainty.....	3-22
3.2.3. Reporting and Visualising Uncertain Results.....	3-27
3.3. UNCERTAINTY ASSESSMENT METHODS DEMONSTRATED IN LCA.....	3-31
3.3.1. Qualitative and Semi-Quantitative Data Quality Assessment Methods.....	3-31
3.3.2. Quantitative Measures to Address Data Uncertainty.....	3-35
3.3.3. Model Parameter Uncertainty and Uncertainty about Model Form.....	3-42
3.3.4. The Reduction and Analysis of Uncertainty.....	3-43

CHAPTER 4: DEVELOPMENT OF INVENTORY MODELLING TO INCLUDE UNCERTAINTY

4.1.	A FRAMEWORK FOR THE QUANTITATIVE ASSESSMENT OF UNCERTAINTY.....	4-1
4.1.1.	A Framework for Empirical Parameter Uncertainty.....	4-1
4.1.2.	Extending the Framework.....	4-6
4.1.3.	Important Considerations in the Assessment of Uncertainty.....	4-8
4.2.	INVENTORY MODELLING TO INCLUDE UNCERTAINTY.....	4-13
4.2.1.	Modelling with Variable and Correlated Inputs.....	4-13
4.2.2.	Characterising Uncertain Parameters.....	4-14
4.2.3.	Analysis of Uncertainty.....	4-24
4.3.	INTERFACE WITH IMPACT ASSESSMENT METHODS.....	4-26
4.4.	SUMMARY OF THE UNCERTAINTY ANALYSIS PROCEDURE.....	4-28

CHAPTER 5: IMPACT ASSESSMENT OF RESOURCE-BASED INDUSTRIES

5.1.	THE ASSESSMENT OF RESOURCE-BASED INDUSTRIES.....	5-1
5.1.1.	Resources.....	5-2
5.1.2.	Aerial Pollutants.....	5-4
5.1.3.	Solid and Liquid Wastes.....	5-5
5.2.	CONSIDERATION OF SOLID WASTE IN LCA.....	5-7
5.2.1.	Problems with Solid Waste Characterisation.....	5-7
5.2.2.	LCI Studies Addressing Solid Waste Disposal.....	5-9
5.3.	CONSIDERATION OF REGIONAL AND SITE-SPECIFIC IMPACTS IN LCA.....	5-10
5.4.	POTENTIAL SOLUTIONS FOR ADDRESSING SOLID WASTE IMPACTS.....	5-16
5.4.1.	Qualitative Indicators.....	5-16
5.4.2.	Impacted Land Footprint.....	5-17

SECTION II: CASE STUDIES

CHAPTER 6: LCI OF SOUTH AFRICAN COAL BASED POWER GENERATION

6.1.	THE POWER INDUSTRY IN SOUTH AFRICA.....	6-1
6.1.1.	Coal Mining and its Impacts.....	6-5
6.1.2.	South Africa's Power Generation Technologies and Plant Mix.....	6-7
6.2.	THE COAL-ELECTRICITY LIFE CYCLE.....	6-10
6.2.1.	A Life Cycle Inventory Model of South African Coal-fired Power Generation.....	6-10
6.2.2.	Foreground and Background Systems.....	6-13
6.2.3.	Flowsheet Construction.....	6-14
6.2.4.	Some Methodological Considerations.....	6-15
6.2.5.	Data Sources.....	6-17
6.2.6.	Limitations of the Assessment of Waste Dumps and Coal Stockpiles.....	6-20
6.3.	INVENTORY OF SOUTH AFRICAN ELECTRICITY.....	6-23
6.3.1.	South African Generating Mix.....	6-23
6.3.2.	A Consideration of Uncertainty.....	6-28
6.4.	CASE STUDIES.....	6-36
6.4.1.	Comparison with European LCA data.....	6-36
6.4.2.	Electricity Input into a Zinc Refinery.....	6-39

CHAPTER 7: LCA MODELLING FOR DECISION MAKING IN PRIMARY INDUSTRIES: TECHNOLOGY CHOICE CASE STUDY

7.1.	DECISION CONTEXT.....	7-2
7.2.	PRELIMINARY TECHNOLOGY SCANNING STUDY.....	7-3
7.2.1.	System Definition.....	7-3

7.2.2.	Selected Results and Conclusions.....	7-6
7.3.	AN INVESTIGATION INTO THE COMBUSTION OF DISCARD COAL IN FLUIDISED BED BOILERS AS AN OPTION FOR RECONDITIONING OLD POWER STATIONS.....	7-10
7.4.	METHODOLOGICAL FRAMEWORK.....	7-11
7.4.1.	Scenario Definition.....	7-11
7.4.2.	System Definition.....	7-12
7.4.3.	System Evaluation.....	7-16
7.4.4.	Definition of Parameter Uncertainty.....	7-18
7.5.	RESULTS AND INTERPRETATIONS.....	7-24
7.5.1.	System Definition.....	7-27
7.5.2.	System Evaluation.....	7-29
7.5.3.	Empirical Parameter Uncertainty.....	7-32
7.5.4.	Model Parameter Uncertainty.....	7-41
7.5.5.	Further Analysis: Choice of Desired FBC Operating Region.....	7-43
7.6.	CONCLUSIONS.....	7-46

CHAPTER 8: LCA MODELLING TO INFORM OPERATIONAL DECISIONS: CASE STUDY

8.1.	DECISION CONTEXT.....	8-1
8.1.1.	Problem Statement.....	8-2
8.2.	METHODOLOGICAL FRAMEWORK.....	8-5
8.2.1.	Scenario Definition.....	8-5
8.2.2.	System Definition.....	8-6
8.2.3.	System Evaluation.....	8-9
8.2.4.	Definition of Parameter Uncertainty.....	8-10
8.2.5.	Model Uncertainties.....	8-12
8.3.	RESULTS AND DISCUSSION.....	8-15
8.3.1.	Identification of Operating Range / Selection of Options.....	8-15
8.3.2.	Selection of Criteria for Comparison.....	8-21
8.3.3.	Model and Empirical Parameter Uncertainty.....	8-26
8.4.	CONCLUSIONS.....	8-34

CHAPTER 9: CONCLUSIONS

9.1.	SIGNIFICANT FINDINGS.....	9-1
9.1.1.	Modelling the Environmental Performance of Resource-Based Industries.....	9-1
9.1.2.	Use of LCA for Environmental Decision Support.....	9-2
9.1.3.	Decision Contexts of Resource-Based Industries.....	9-4
9.1.4.	Uncertainty Analysis.....	9-6
9.1.5.	Key Features of the Case Studies.....	9-14
9.2.	THE WAY FORWARD.....	9-16

REFERENCES.....	R-1
-----------------	-----

APPENDICES*

APPENDIX A: STATISTICAL METHODS

A.1.	Propagation and Analysis of Uncertainty.....	A-1
A.2.	Probability Distributions.....	A-8
A.3.	Principal Component Theory.....	A-15
*A.4.	Glossary of Statistical Terms.....	A-22

APPENDIX B: MEASURES TO ADDRESS DATA UNCERTAINTY IN LCA

B.1.	Qualitative and Semi-Quantitative Data Quality Assessment Methods.....	B-1
B.2.	Quantitative Methods To Address Data Uncertainty Demonstrated in LCA.....	B-6

APPENDIX C: PERTINENT ASPECTS OF LCIA

*C.1.	LCA Studies Addressing Impact Assessment Uncertainty.....	C-1
*C.2.	The Eco-Indicator 99 Impact Assessment Method.....	C-3
*C.3.	Impact Assessment Methods for Resource Use.....	C-5
*C.4.	Methods Developed to Include Fate and Exposure in LCIA.....	C-8

APPENDIX D: INFORMATION SOURCES

*D.1.	Documents and Databases Sourced.....	D-1
*D.2.	Personnel Interviewed.....	D-5

APPENDIX E: INVENTORY MODELS

E.1.	General Model Structure.....	E-1
*E.2.	Foreground Processes.....	E-7
*E.3.	Background Processes and Transportation.....	E-27

APPENDIX F: INVENTORY MODELS AND THEIR DATA INPUTS

*F.1.	SA Inventory (chapter 6).....	F-2
*F.2.	PF vs FBC Case Study (chapter 7).....	F-3
*F.3.	Mine Water Case Study (chapter 8).....	F-4

APPENDIX G: SUPPORTING INFORMATION FOR SA COAL-FIRED ELECTRICITY LCI

*G.1.	Description of South African Coal Mining and Power Generation Technologies.....	G-1
G.2.	Analysis of the SA Electricity Inventories.....	G-9
G.3.	Uncertainty Considerations.....	G-14

APPENDIX H: SUPPORTING INFORMATION FOR TECHNOLOGY CHOICE CASE STUDY

*H.1.	Technology Scanning Study: LCA of Process Technologies for Sulphur Removal in Coal-Fired Power Generation in South Africa.....	H-1
H.2.	Additional Results and Explanations.....	H-21
H.3.	Inventory Versus Impact Uncertainty.....	H-31
H.4.	Analysis of Variance: Example of Carcinogenic Effects.....	H-38
H.5.	FBC Operating Space Parameter Analysis.....	H-43

APPENDIX I: SUPPORTING INFORMATION FOR OPERATIONAL CASE STUDY

I.1.	Analysis of the Principal Component Tables.....	I-1
I.2.	Probability Plots for Mine Water Options.....	I-3
I.3.	Box And Whisker Plots of Mine Water Extreme Options.....	I-6
I.4.	Model Parameter Analysis.....	I-7

*APPENDIX J: OVERVIEW OF LCA METHODOLOGICAL FRAMEWORK.....	J-1
--	-----

Note: The full appendices are available in Adobe Acrobat format on the CD in the back cover. Those appendix sections marked with an asterisk () are only available electronically (an installation file for Acrobat Reader is also available on the CD)

LIST OF ABBREVIATIONS

AMD	acid mine drainage (also sometimes called ARD – acid rock drainage)
CHAINET	European Network on Chain Analysis for Environmental Decision Support
CAPCO	chief air pollution control officer (air pollution regulatory body)
CCW	circulating cooling water
CDF	cumulative density function
CFBC	circulating fluidised bed combustor
C.V.	calorific value
CV	coefficient of variance
CPP	condensate polishing plant
DALYs	disability adjusted life years
DQG	data quality goal
DQI	data quality indicator
DMEA	Department of Mineral and Energy Affairs
EDR	electrodialysis reversal (membrane water treatment plant)
EI 99	Eco-Indicator 99 LCIA method
EIA	environmental impact assessment
EMPR	Environmental Management Programmes Report
ERA	ecological risk assessment
ESP	electrostatic precipitator (particulate removal device)
FBC	fluidised bed combustion
FF	fabric filter (particulate removal device)
FGD	flue gas desulphurisation
GDP	gross domestic product
GWP	global warming potential
ISO	International Standards Organisation
ISP	Imperial Smelting Process
LCA	life cycle assessment
LCI	life cycle inventory
LCIA	life cycle impact assessment
LHV	lower heating value
MWhSO	Megawatt hours sent out (net electricity produced)
NOEC	no observable effects concentration
ODP	ozone depleting potential
PC	principal components
PCA	principal component analysis
PDF	probability density function; or in the context of impact assessment, potentially disappeared fraction of vascular plant species (units of ecosystem quality)
PF	pulverised fuel
PM10s	airborne particulate matter less than 10 microns in diameter
POCP	photochemical oxidant formation potential
RA	risk assessment
ROM	run-of-mine (untreated coal)
SAPP	Southern African Power Pool
SDA	spray dryer absorption (flue gas desulphurisation process)
SETAC	The Society of Environmental Toxicology and Chemistry
TRO	tubular reverse osmosis (membrane water treatment plant)
USES	Uniform System for the Evaluation of Substances (model including fate in the characterisation of toxic releases)
WIA-2	SETAC-Europe second working group on LCIA
ZLED	Zero Liquid Effluent Discharge (policy under which all effluent and stormwater run-off must be used or retained on site)

TERMINOLOGY

- The *parameter types* (uncertain quantities) are defined in section 3.1.1.
- Frequently used *statistical terms* are defined in appendix A.21.
- The following terminology applies to *Life Cycle Assessment* (see overview in appendix J):
 - *Life Cycle Inventory* (LCI) refers to the first phase of LCA, the compilation of a quantitative list of extractions from and emissions to the environment, and other variables at the boundary of the system and the environment (e.g. types of land use). The term *environmental interventions* is used to collectively represent the inventory elements (the environmental inputs and outputs).
 - the *functional unit* is a consistent measure of the function that the system delivers, chosen as a relevant unit of service or quantity of product, so that different systems can be compared on an equivalent function.
 - *Life Cycle Impact Assessment* (LCIA) refers to the second phase of LCA, in which the environmental interventions are aggregated into *impact categories*, classes covering all relevant areas of environmental concern (also called *damage categories*).
 - The *impact indicator* is a quantifiable representation of the impact category, and can be defined at various levels of the environmental mechanism connecting the environmental interventions to the impact categories. *Category endpoints* are the variables of direct societal concern (e.g. human lives lost) (also called the *damage level*), whilst the *category midpoints* are the intermediate variables in the environmental mechanism i.e. they fall between the interventions and the endpoints (e.g. concentration of toxic substances).
 - The terms *characterisation factors* or *equivalency factors* (or *potentials*) are used interchangeably in this thesis to represent the factors derived from an impact model which convert the environmental interventions to the common unit of the impact indicator, and allow them to be aggregated into a single impact category (these are called *damage factors* in the EI 99 method).
 - *Selection criteria* is used to encompass the subset of impact categories and environmental interventions chosen to represent the pertinent environmental impacts of the particular system under evaluation.

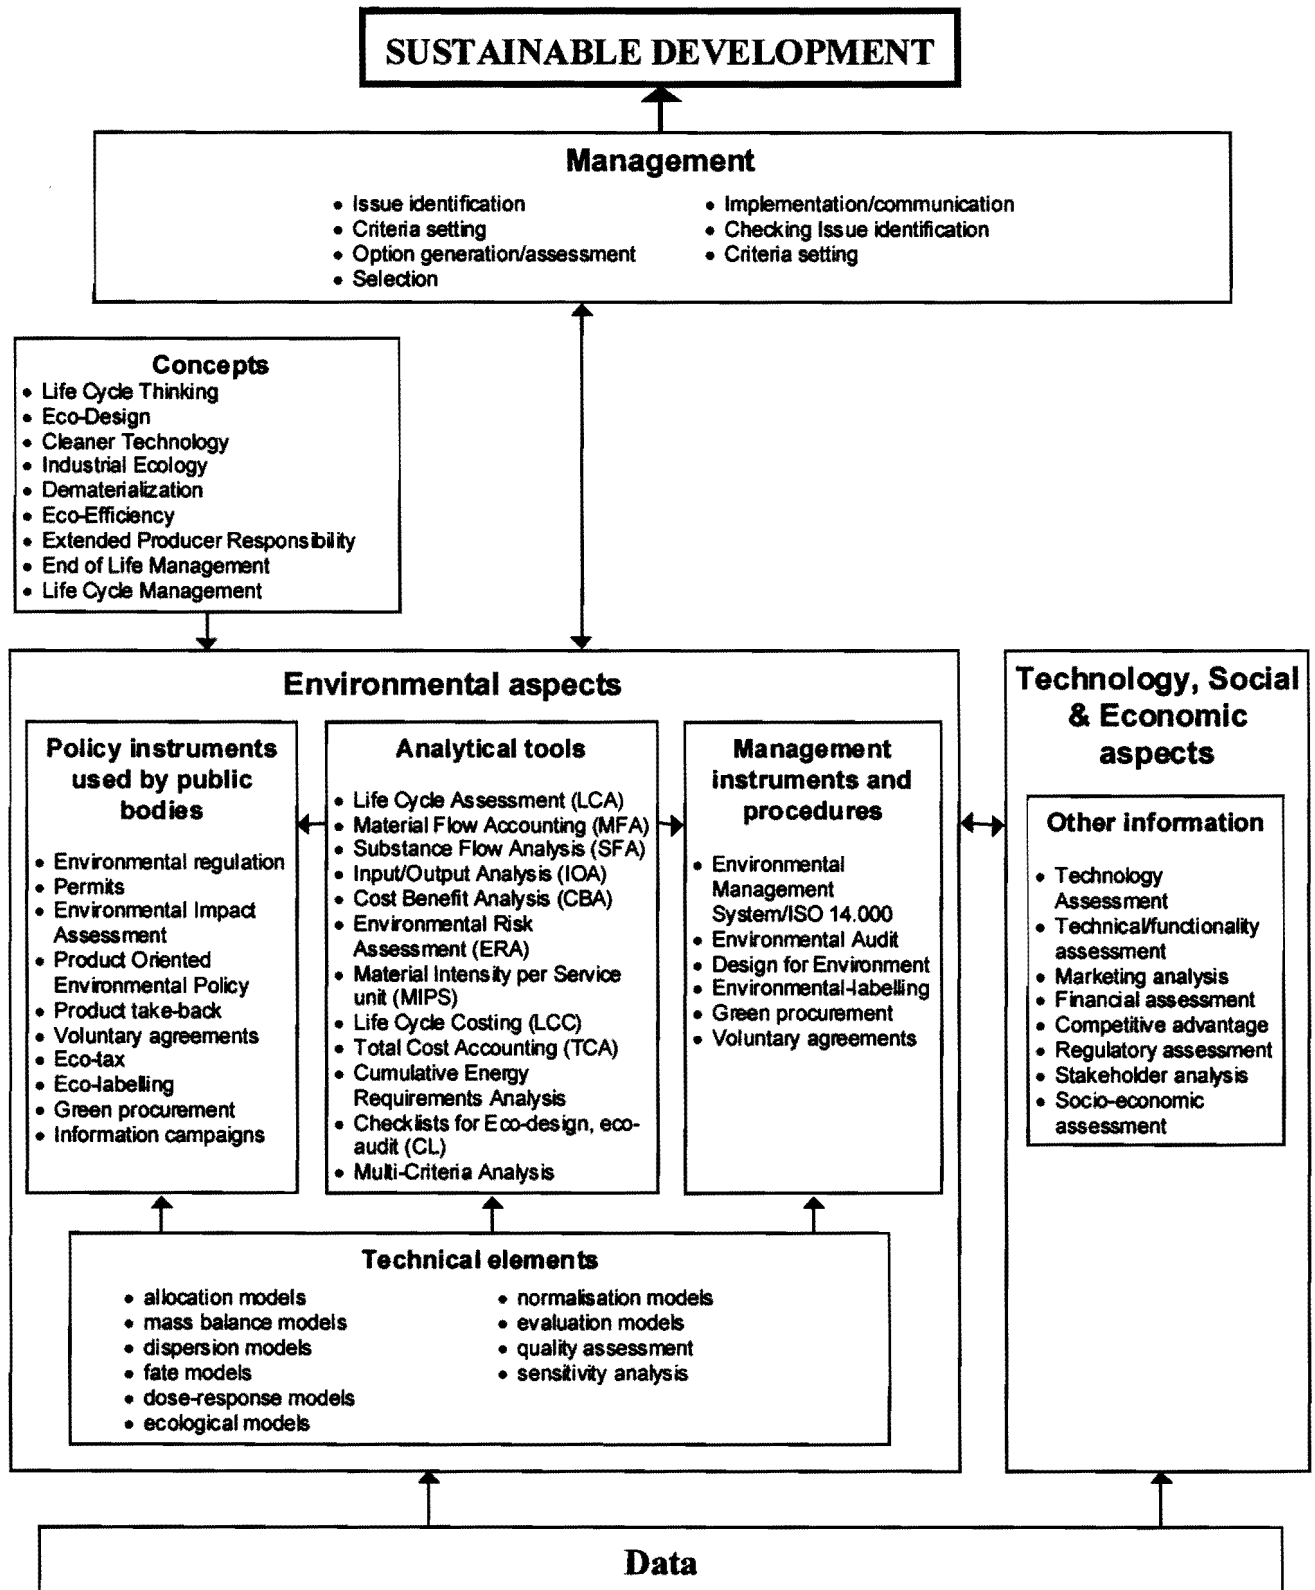


Figure 1-1 The relationship between concepts, analytical tools, procedural tools, technical elements, data and other elements in the decision-making process (CHAINET, 1999).

Brief summaries of the major analytical tools listed in Figure 1-1 can be found in Wrisberg and Gameson (1998). Generally, no one tool can provide all the information required for making decisions, so combinations of tools are used. The choice of these tools, and the integration of the results, constitute a process. One or more stakeholders will be involved in this process in order to define which tools are chosen, and whether other factors should be taken into consideration. A process can therefore be described as a way of using and integrating different tools with stakeholder expectations and other decision parameters to meet the requirements for a decision (Cowell et al., 1997). The decision context determines which combination of tools is most appropriate, and how the output of these tools is used. Decision-making processes can therefore not be standardised, as tools can. However, it is possible to develop guidelines on processes to be used in different contexts (Cowell et al., 1997).

This distinction between tools and processes is useful because it helps to unravel some of the confusion that has arisen over use of approaches such as EIA, ERA and LCA (where “approaches” is used as a generic term for tools and processes). In fact, some of these approaches can be viewed as either tools or processes depending upon how they are incorporated into decision-making. As tools, they would be regarded as procedures undertaken by experts and producing relatively objective results that are subsequently considered in a decision-making process. As processes, the tool-like properties of the approaches are subsumed in a decision-making process that addresses the concerns of stakeholders and takes into account the context of the decision (Cowell et al., 1997).

The use of LCA to support decision-making in resource-based industries is investigated in this thesis. It is necessary to distinguish between life cycle thinking (as a concept), and LCA as a decision-making process or analytical tool. Life cycle thinking proves the philosophical basis to the decision-making process in that it supports the process of defining the contribution of human activities to the environmental aspects of sustainable development (Cowell, 2001). This it does by identifying and assessing the environmental impacts associated with services delivered to societies, regardless of their geographical and temporal location (where identifying the relative magnitude of these impact is a first step towards minimising them). It is thus able to incorporate the issues of intra-generational equity (impacts occurring on different geographical scales) and inter-generational equity (impacts occurring on different time scales), integral to defining the contribution of human activities towards sustainable development (Cowell, 2001).

Life cycle thinking is the philosophy that shapes the quantitative LCA tool, i.e. LCA provides the quantitative information around environmental performance from a life cycle perspective. The distinction of Cowell et al. (1997) between tools and processes is important, as it helps to clarify the different ways in which the LCA method is being developed. On one hand, LCA is being developed to assess potential environmental impacts using generic weighting factors, and without regard to site-specific conditions, e.g. Heijungs et al. (1992). This is analogous to using LCA as a tool. On the other hand, some practitioners argue that the more site-specific and subjective, evaluative components of LCA are crucial to the results, and should be recognised in the methodology. This is analogous to using LCA as a process (Cowell et al., 1997). As applied in this thesis, LCA falls more towards its definition as a process, which is demanded by the nature of the primary industries assessed (see section 1.2). The case studies investigated show how the LCA method requires adaptation and consideration of other tools to achieve a full assessment of the impacts from primary industries. One may argue that this is strictly no longer LCA, and borrowing from technology assessment and risk assessment blurs the distinction between these tools. Nonetheless, it can be argued that drawing from the strengths of related environmental tools leads to a more powerful and versatile LCA process.

1.1.1. Steps of the Decision-Making Process

Decision-making is an iterative process, with a number of feedback loops between the steps. Wrisberg and Gameson (1998) identify the following six steps:

- Issue identification: defines the problem at hand,
- Criteria setting: sets the criteria to be used for judging decision options,
- Option generation/assessment: identifies options and assesses these against the previously determined criteria,
- Selection: a decision is made on the best decision option,
- Implementation of the decision
- Checking / reviewing the result.

The iterative decision-making process first needs to be placed in the context in which the decision is to be made, as this defines the type of information required to support the decision. Wrisberg and Gameson (1998) relate the demand for information to the stakeholders involved and their motives, and the type of decision to be made (where stakeholders are defined as someone with a legitimate interest in the decision, e.g. regulators, suppliers, investors, trade unions etc.). The type of decision to be made

depends on the decision subject (e.g. product, process, technology etc.), whether the decision involves the comparison of existing systems or the design of a new system, and the desired levels of improvement (ranging from incremental improvement, through redesign of existing concepts, to a complete change of functionality) (Wrisberg and Gameson, 1998). Additional characteristics shaping the context of the decision involve the level of spatial specification (site-specific or not), the spatial scale (local, regional or global), and the temporal scale (short or long time frame) (CHAINET, 1999).

1.1.2. Overview of Life Cycle Assessment

Life cycle assessment entails evaluating the effects of a product or process, or more specifically, the function a product or process is designed to perform, over the entire course of its material life cycle. It thus aims to quantify the environmental consequences of performing that function, and is perhaps best described as a systems analysis tool able to provide a picture of the interactions of an activity with the environment. LCA's integrative approach avoids substituting one set of environmental problems for a different set of problems, and guards against three common types of problem shifting: shifting from one stage of the life cycle to another, e.g. substituting a hazardous raw material for a less hazardous one, but one which involves more intensive (and waste producing) pre-processing; shifting from one problem to another, e.g. gas scrubbers substituting a solid waste and liquid effluent problem for a gaseous emission problem; and shifting from one location to another, e.g. a plant switching to electric power from coal combustion, shifts the emissions away from the plant to the power station site (UNEP-IE, 1996).

Although the principles of LCA have been in use for around 20 years, it has only been since the early 1990s that it gained attention as a promising environmental management tool. The Society of Environmental Toxicology and Chemistry (SETAC) took the lead in trying to consolidate and standardise the emerging tool, which culminated in the formulation of a Code of Practice for LCA (Consoli et al., 1993; Fava et al., 1991). More recently, the International Standards Organisation (ISO) has taken over this role, and has established a uniform framework and terminology. LCA is formally defined by SETAC as "a process to evaluate the environmental burdens associated with a product, process or activity by identifying and quantifying energy and materials used and wastes released to the environment; to assess the impact of those energy and material uses and releases to the environment; and to identify and evaluate opportunities to effect environmental improvements. The assessment includes the entire life cycle of a product, process or activity, encompassing extracting and processing

raw materials; manufacturing, transportation and distribution; use, re-use, maintenance; recycling, and disposal” (Consoli et al., 1993). This definition is consistent with that put forward in the ISO environmental management standard on LCA (ISO, 1997).

The methodological framework, defined by the ISO standards listed below, is broken down into four main phases:

1. ISO-14040 (Principles and Framework)
2. ISO-14041 (Goal and Scope Definition, Inventory Analysis)
3. ISO-14042 (Life Cycle Impact Assessment)
4. ISO-14043 (Life Cycle Interpretation)

Although usually presented as a step-wise process, in practice LCA is an iterative process with considerable feedback between the phases. The four phases are therefore better represented as in Figure 1-2. A brief overview of what is embodied in each of these steps is given in appendix J, and a detailed description is available in a very large number of references, including Consoli et al. (1993), Guinée et al. (1993a-b), Heijungs et al. (1992), ISO (1997, 1998, 2000a-b), Jensen et al. (1997), and UNEP-IE (1996).

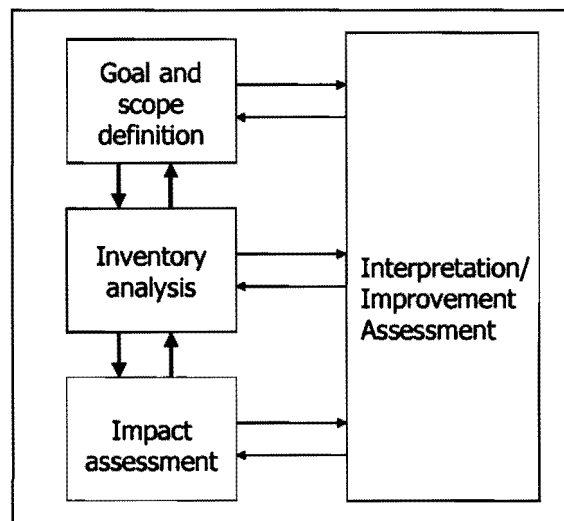


Figure 1-2 Components of a life cycle assessment (ISO 14040).

Goal and scope definition sets the overall frame conditions for the study (the level of complexity, certainty and completeness of the inventory and impact models). Inventory analysis (LCI) identifies and quantifies the environmental interventions, where the term “environmental interventions” is used to encompass all variables at the system boundary (including extractions from and emissions to the environment, as well as other interactions, such as land use). Impact assessment (LCIA) is the phase aimed at evaluating the magnitude and significance of the potential environmental impacts of the system, typically by using equivalency factors to quantify the relative severity of each environmental intervention’s contribution to a number of impact categories. In life cycle interpretation, an analysis of the results identifies the most promising possibilities for minimising the impacts of the system.

LCA was originally conceived as a quantitative tool, designed to provide objective, scientific information to enhance the credibility of the environmental argument in decision-making processes, and to allow for reproducible, defensible answers. However, more recently the neutrality and objectivity of LCA has been questioned increasingly (Finnveden, 1997; Tukker, 1998). Even allowing for the more overtly subjective phases of LCA, such as valuation, increasingly there is the realisation that LCA is rooted in a particular way of viewing, and making sense of the world. Of particular note is the underlying premise of LCA that trade-offs can be made between different environmental impacts, and the fact that LCA only addresses impacts that are known and quantifiable (i.e. although LCA results appear certain and complete, they are limited to those processes for which data are available, and to impacts that are known to occur). The fact that LCA presents a single particular framing of the decision situation may bias the discussion because the results appear comprehensive, unambiguous and concise compared with other types of information presented during the decision-making process (Cowell, 2001).

Historically LCA has been predominantly product-focussed, with internal product improvement cited by companies as the most common reason for performing an LCA, with marketing and product design as the next most common uses (Baumann, 1996). More recently, increasing application has been found in using LCA as a tool for process selection, design and optimisation (Azapagic, 1999; Azapagic and Clift, 1999a; Diwekar, 1999; Kniel et al., 1996; Stewart and Petrie, 1996). There has also been an increasing distinction between “retrospective” and “prospective” LCA studies (Frischknecht, 1997; Weidema et al., 1999; Wenzel, 1999). A consequence of a decision focus on LCA is that by necessity LCA will always study the future, as decisions affect the future and not the past. It can therefore be argued that “retrospective” or “historical” LCA studies are largely without meaning (Wenzel, 1999).

However, retrospective studies can perhaps also be seen as decision orientated, where they find application as background data in a prospective LCA study (see section 2.3).

A key requirement of an LCA is that it reflects the environmental change caused by the decision. The decision to be supported therefore forms the basis of the LCA method, and the consequences of the decision should be well foreseen and well reflected by the system models within the LCA (Wenzel, 1999). This suggests the need for an adaptable LCA method, able to provide the different levels of information required by different decision contexts. This is reflected by the development of various levels of LCA, ranging from matrix LCA using qualitative or semi-quantitative data (Graedel, 1998; Graedel et al., 1995), to screening LCA using readily available quantitative data (e.g. the many publicly available LCI databases, such as the ESU database), to full LCA, where the collection of new inventory data is required (Wenzel, 1998).

In addition, there is the growing awareness that LCA is just one of a number of environmental tools and processes, and it is agreed that stand-alone, comprehensive LCAs are unlikely to be sufficient (or entirely appropriate) as a sole basis for environmental decision-making (Cowell et al., 1997). This again raises the question of using LCA as a process to support decision-making, rather than a strict analytical tool. Both approaches have their strengths and weaknesses. Using LCA as a tool allows for an easier and more transparent analysis, because a standardised set of impact categories and equivalency factors can be supplied. The systematic, prescriptive framework provided by the LCA tool is identified as a significant strength of LCA, in that it provides a problem structuring approach for assessing the environmental impacts of human activities (Hertwich and Hammitt, 2001). Given the complexity of many environmental problems, this benefit should not be underestimated (Cowell, 2001). However, it may lead to misleading results, because the predicted potential impacts may be very different from the actual impacts.

Using LCA as a process may produce results more relevant to the decision under consideration, and consequently greater acceptance by stakeholders, because their concerns have been incorporated into the process. However, it may allow stakeholders to exclude impacts from the analysis that fall outside their immediate interests, thereby compromising LCA's comprehensiveness and its ability to identify trade-offs in environmental impacts (Cowell et al., 1997). Defining LCA as a process also facilitates LCA's use for learning and awareness raising, rather than purely decision-taking, by placing a greater emphasis on problem structuring and the understandability of the LCA model and its results (Cowell,

2001). Ultimately it needs to be recognised that there are different ways of framing decision situations, and there needs to be greater emphasis on placing the LCA process within the particular decision context at hand.

1.2. THE ASSESSMENT OF PRIMARY INDUSTRIES

Whilst good environmental management of primary industries is important because of their considerable impact on the environment, an assessment of their environmental performance is especially important from an LCA perspective. This is because they inform the inventories of all products or processes further along the supply chain.

1.2.1. Features of Primary Industries

Primary industries are inherently resource intensive, in that they supply material and energy to industries further down the supply chain. This involves the extraction and concentration of resource stocks, processes typically associated with the consumption of considerable energy, and the generation of large waste volumes. The continuing depletion of high-grade resources and increasing exploitation of lower grade ores will inevitably lead to a steady increase in these impacts. The particular focus of this thesis is on the provision of primary energy (coal-fired power generation), although many of the issues are generic to all abiotic resource-based industries. Primary industries based on biotic resources (forestry, fishing etc.) are not dealt with here.

1.2.1.a Consumption of Resources

The extractive industries (mining and minerals processing) are reliant on the consumption of abiotic resources. As such, they can be seen as inherently non-sustainable processes. For example, The Natural Step's first principle decrees that "substances extracted from the lithosphere must not be systematically accumulated in the ecosphere". This requires reduced mining to establish sustainable material exchange between society and the ecosphere (Holmburg et al., 1996). Similarly, the World Commission on Environment and Development (1987) states that sustainability ultimately requires "not using non-renewable resources faster than renewable substitutes can be found for them". This apparent paradox has led to a more broad definition of sustainability for the minerals sector. Natural Resource Canada (1997) define it as (amongst other factors) "finding, extracting, producing, adding value to, transporting, using, reusing, recycling, and disposing of mineral and metal products....in the most efficient, competitive, and environmentally responsible manner possible". Thereby acknowledging that

given the apparently indispensable nature of these industries from both a social and economical perspective, the best that can be done for now is to minimise the resource consumption (and associated impacts).

Extractive industries are energy intensive processes, with the heavy machinery used during mining consuming considerable diesel fuel or electricity. Primary and secondary processing are also very energy intensive processes, with furnaces and smelters extremely high consumers of electricity, coal or fuel oil. Water is used in high quantities in most primary industries, primarily as a slurring medium and a coolant. Once in contact with sulphidic ores, the resulting high salinity means that the water quality is degraded to the extent that it is unfit for most industrial uses. In addition, water impacts extend beyond merely the process water requirements, as rain water coming into contact with exposed ores and waste dumps becomes contaminated. Mining and primary processes are therefore often characterised by large volumes of dilute effluent. Historically this water has been discharged to natural water courses, but increasingly this is no longer acceptable, and the water requires treatment (often an energy and resource intensive process in itself) or containment and evaporation. Mining processes are also associated with extensive surface disruption. Opencast mining operations can significantly change the topographical features and land use patterns, whilst tailings dumps and dams can extend for many kilometres, and often result in land sterilisation. Even where the dumps and disturbed surface areas are rehabilitated (as is called for by mining legislation), the land areas are rarely returned to their original state (Russell, 1991).

1.2.1.b Air Pollution

Many air pollutants are emitted during mining and primary processing. Mining, particularly opencast methods, cause significant particulate emissions from blasting and the use of heavy machinery. Other emissions to air from mining include methane, and vehicle emissions from diesel fuelled trucks and machinery. Minerals processing also releases considerable air pollutants. The burning of coal and roasting of sulphidic ores release acid gases and toxic trace elements present in the ores. Dust blown up from waste dumps can also cause considerable particulate emissions. Whilst stack emissions are usually able to be quantified by monitoring or a knowledge of the composition of the feed materials, non-stack emissions (e.g. dust from blasting, or haul roads) are generally only able to be roughly estimated.

1.2.1.c Solid Waste and Liquid Effluents

Primary processes are responsible for huge volumes of relatively inert waste. Generally the proportion of hazardous waste to total waste is small. Most mining and associated primary processing operations are characterised by very large “mine dumps” or tailings dams (large dedicated waste repositories), often stretching for several kilometres. In some cases, mining waste is used as mining backfill or as road-fill, but the very large volumes usually mean that the bulk of the waste ends up in massive dumps. In addition, the waste rock and/or tailings are often slightly radioactive or contain leachable elements, which make them unsuitable for most uses. Typically only the mass or volume of the waste, or the area of the waste dump is quantified. Emissions from the waste dumps are generally disregarded because of the difficulties involved in their prediction. Section 5.2. explores these issues in greater detail.

The quantity of pollutants emitted to water is generally quantifiable if emitted from a definite point source, i.e. where the effluent is leaving the system in a discharge pipe, and the flow and quality of the water is monitored. Even when not monitored, from mass balance principles and a knowledge of the process, the quantity of pollutants to water can usually be estimated. However, diffuse pollution sources are much more difficult to quantify and are consequently most often over-looked e.g. water becoming contaminated through contact with tailings dams, waste rock dumps, and exposed mining seams. Pollutants are released into the environment either as run-off, or as leachate seeping through the dumps, which subsequently migrate into the underlying aquifers or nearby surface-water bodies, contaminating these environments. The practice of slurring tailings with water for hydraulic transport to dumps, where the excess water is decanted off, notably increases this potential for groundwater contamination.

1.2.2. Life Cycle Assessment of Primary Industries

Much of LCA’s methodology has been developed around the assessment of secondary products, e.g. the classic LCA studies addressed such issues as plastic vs. paper packaging, disposable vs. cloth nappies et cetera. LCA’s potential for process assessment and design is increasingly being exploited and this has led to some different development issues (Azapagic, 1999; Kniel et al., 1996; Pessoa, 1993; Stewart and Petrie, 1996). Early studies usually involved difficult to define functional units and system boundaries, and consequently much of the emphasis on inventory methodology development was on system definition and allocation rules. The products assessed were usually fairly far down the

production chain, which meant that upstream supplies came from widely disparate locations. The global nature of LCA meant that it was able to assess these on an equal basis, and was therefore an important strength of the tool.

In contrast, primary industries are at the start of the supply chain. The emphasis of the inventory assessment is thus less on system definition, which remains important (e.g. allocating burdens to multiple products), but on characterising variable and imprecise processes. Significant impacts associated with these industries are a consequence of the extensive volumes of liquid and solid waste generated, and given that these impacts are site-specific, the current globally orientated LCIA framework is not able to sufficiently address them. A particular emphasis of this study is therefore on quantifying the potential impacts around solid waste deposits and the associated diffuse water pollution. A discussion on the ability of the current LCIA methodology to address resource-based industries, in light of their features mentioned above, is given in section 5.1.

1.2.3. A Focus on Coal Based Power Generation

Resource-based industries play an extremely significant role in South Africa's economy. Coal is the mainstay of South Africa's energy industry, supplying 71% of the country's energy needs in 1996 (Grobbelaar et al., 1997). A considerable portion of this energy is consumed in primary industries. Of the coal sold for domestic consumption, some 60% is burnt to generate electricity (Prevost, 1998). Other major coal users are the production of synthetic fuels and chemicals (33%), the production of iron and steel (3.1%) and mining (0.8%) (Prevost, 1998). Of the electricity produced by South Africa's principal supplier, 19% is sold direct to mining clients and 31% to major industries (Eskom, 1999a). In addition to being a major consumer of energy, the mining sector is the largest generator and accumulator of solid wastes in South Africa, accounting for 74% of the total waste stream in 1990 (CSIR, 1991).

Coal forms the basis of many countries' energy mix, with almost 40% of the world's electricity generated from coal (IEA, 1993), whilst over 90% of South Africa's electricity production is coal-based (Eskom, 1999a). Extensive coal reserves mean that coal is likely to remain the dominant power source for the foreseeable future, especially in developing countries where the affordability of power is an important consideration. The inherently polluting nature of coal-combustion processes has been recognised and many technologies developed to mitigate these effects (Tavoulareas and Charpentier,

1995). This presents the dilemma of choosing between disparate technologies that offer relative degrees of improvement in environmental performance, and often result in a trade-off between environmental impacts. A reliable framework for technology selection for power production is therefore necessary.

The provision of energy is especially significant from an LCA perspective, as it is used in some form or another in practically all product and process systems. A complete LCA of practically any system thus requires the impacts associated with energy provision to be included in the assessment. The energy data used can have a substantial effect on the overall conclusions of an LCA, and the use of inconsistent energy data has been cited as a major reason for the seemingly contradictory conclusions of different studies on the same product (Weidema, 1993). Applicable energy data is therefore essential for meaningful LCA results. Many LCA database tools are available, but the majority of these are of European or American origin (Menke et al., 1996). An important component of this study is therefore to provide an inventory of coal-generated power reflective of South African conditions. This inventory is presented in chapter 6.

1.3. THEMES AND THESIS STRUCTURE

Figure 1-3 gives an overview of the themes addressed in this thesis, how they inform each other, and where in the thesis they are addressed. The thesis is divided into two parts. The first part deals with methodology development, and in the second part case studies are presented which support the methods developed. An introduction to the main topics covered is given below.

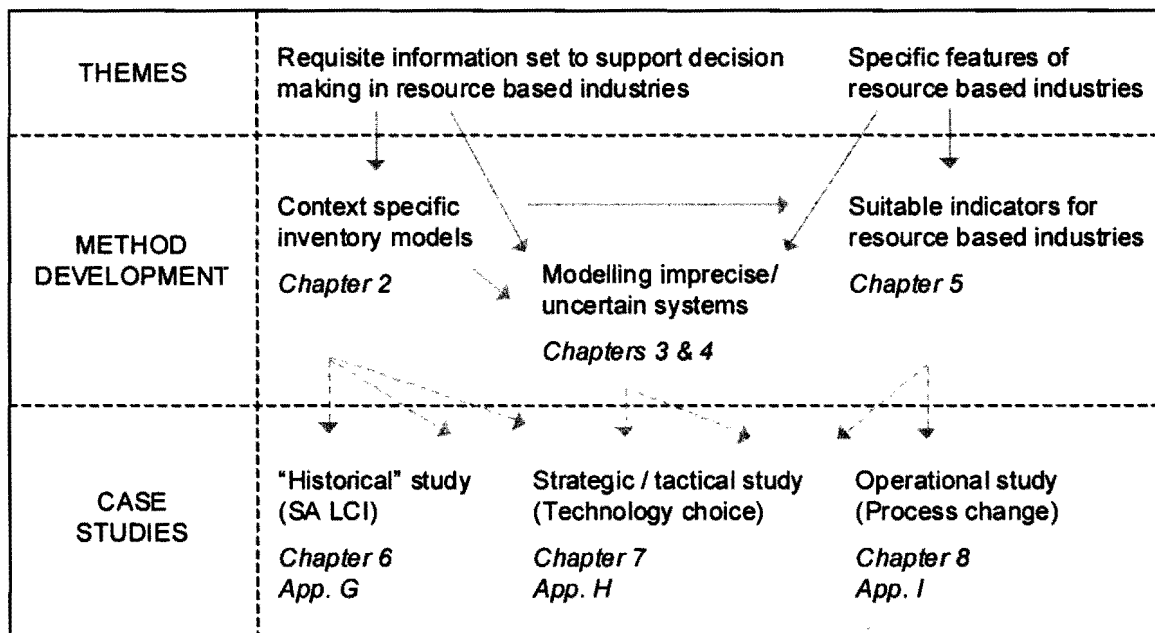


Figure 1-3 Specific themes of the thesis, and where they are addressed.

1.3.1. Inventory Modelling for Decision Support in Primary Industries

The LCA approach is based on a rigorous analysis of all material and energy flows occurring within the system boundary, and consequently has much in common with the conventional mass and energy balance approach to process technology assessment. The environmental interventions are linked to the impacts to which they could potentially give rise, according to a pre-defined set of environmental impact factors. Thus the LCI provides a causal link between process design and performance, and potential environmental impact. This link is an important tool in the analysis and design of process systems for better overall environmental performance.

A common criticism of LCA relates to the large data set required to carry out a comprehensive LCA (Nash and Stoughton, 1994; Udo de Haes, 1993). To some degree this has been mitigated by the development of LCI databases, but generally considerable data collection and some sort of process model is required to construct the inventory. Furthermore, for meaningful decision support, the inventory must necessarily reflect future systems, and is thus inherently uncertain. The degree of uncertainty will depend on the decision context, i.e. how far into the future the inventory needs to reflect. The emphasis in this study is on generating LCI information to support decisions. A static inflexible inventory model is therefore of limited use, and the inventory must be able to be predictive and respond to changing input variables, at a level of detail consistent with the decision context.

Chapter 2 presents some of the main considerations with respect to drawing up an inventory model for primary industries, e.g. system boundary definition, flowsheet detail, spatial and temporal considerations et cetera. The importance of tailoring the inventory method to the application of the LCI (i.e. to the decision to be supported) is recognised, and after a review of the application typologies in the literature, a differentiation based on information flow, relevant to the assessment of primary industries, is defined.

1.3.2. Modelling Uncertain Systems

A number of different sources of uncertainty are present in the system models underlying the inventory and impact assessments of LCA, and different types of uncertainty require different methods for their assessment. Inventory models generate and collate data from a number of different sources to provide a consistent listing of environmental interventions for the system, whilst impact models approximate the causal mechanism linking the environmental interventions to a relevant indicator of environmental impact. Broadly, three types of uncertainty can be discerned: uncertainties in the form of the models used, uncertainties in the definition of the system, and uncertainties in the data (empirical parameters) used to characterise the system.

Empirical parameter uncertainty can be assessed using stochastic modelling techniques. In these methods uncertain input variables are entered as probability distributions, the shape and range of which are governed by the amount of information known about the variable. Uncertainties can then be efficiently propagated through the model by simulation using probability sampling techniques (Morgan and Henrion, 1990). Uncertainties arising from the definition of the system can be assessed using

sensitivity analyses, or more formally, in a parametric analysis. Modelling uncertainties are the most difficult to internalise, as an alternative model form is required for an assessment using a sensitivity analysis. Uncertainty due to LCA choices (e.g. allocation method) can be assessed using such an analysis.

The use of quantitative uncertainty analyses allows results to be presented as ranges and confidence intervals, thereby enhancing decision-making by focussing on impact categories or interventions for which there are real differences between the alternatives. Challenges to the inclusion of uncertainty analyses in inventory models include definition of the uncertain quantities, and suitable methods of presentation and analysis of the output samples. Chapter 3 presents a review of the classification and characterisation of inventory uncertainty from a statistical and LCI standpoint, whilst chapter 4 pulls these together into the development of a framework for including uncertainties in LCI models.

1.3.3. Indicators for Resource Based Industries

No consensus has yet been reached for one single default list of impact categories, although the categories strive to cover the three generally accepted relevant areas of protection: human health, the natural environment and natural resources. Chapter 5 reviews the ability of the current LCIA methods to characterise the impacts of resource-based industries, particularly the site-specific impacts arising from liquid effluents and solid wastes. For operability purposes, impact categories have predominantly been based on midpoint indicators and not on the endpoints or actual damage functions. They are thus limited inherently to the prediction of potential impacts. Research on making the equivalency potentials more environmentally relevant (moving the indicators closer to the endpoints) is on-going, especially around toxicity where large uncertainties still exist (Udo de Haes, 1999a; Guinée et al., 1996b). However this requires dramatically more complex models, as it requires the inclusion of site-specific fate and transport models, and introduces the global vs. site-specific argument prevalent in LCA methodology development (Potting et al., 1999). These issues are covered in detail in section 5.3.

Some of the current limitations of impact assessment methods have important consequences for the assessment of resource-based industries, specifically the limitations around the prediction of site-specific impacts associated with solid waste management. Toxicity from metal species is poorly assessed in the current LCA impact categories, because the fate and exposure models used are not strictly applicable to the assessment of metal species (Guinée et al., 1996a). In addition, the toxicity

categories only address the metals emitted, and not possible synergistic effects, e.g. the potential of sulphate emissions to mobilise metal species. The fact that questionable toxicity data for only a few metal species exist, as well as the lack of trace monitoring data, and the difficulty to quantify the potential mobilisation of metal species from waste dumps, have led to this study using a proxy-indicator based on salinity for indicating impacts to water. This is presented in section 5.4.

1.3.4. Case Studies

Case studies showing the key features of the methods developed are presented in chapters 6, 7 and 8. The case studies cover the broad decision contexts identified for resource-based industries in section 2.3. Chapter 6 presents a “historical” LCI of coal-based power generation in South Africa. The study provides an inventory of the average grid mix, taking into account the various technologies employed at the power stations and their associated collieries. The study aims to be more relevant for use in future decision systems by estimating the inventory for a near-future generating mix, and estimating the uncertainty in the inventory this implies.

Chapters 7 and 8 explore the use of LCA to support prospective systems. Chapter 7 investigates a decision with a long time-frame and significant implications (investigates technology options to re-power an in-storage power station), whilst chapter 8 investigates a decision with a considerably shorter time-frame and smaller decision consequences (investigates incorporating a mining effluent into the power station water plant). Both studies place a considerable emphasis on uncertainty assessment (both data uncertainty and uncertainty in the definition of the systems), and the studies are presented more as an exploration of the decision space, and the trade-offs between environmental impacts within this, than unequivocally finding the optimal system, i.e. the emphasis is on learning, rather than on decision-taking.

The case studies draw together the various elements of this thesis. They demonstrate the challenges of inventory modelling, i.e. synthesising data sources of various levels of certainty and completeness into a coherent inventory, and the need to adapt the LCA method to include a representation of all relevant impacts of resource-based industries. The resultant analysis, incorporating an explicit investigation into the definition of the system and the empirical uncertainty, aims to provide insights unable to be achieved by conventional “single-point” LCA, and remove the arbitrariness and sense of false accuracy for which LCA studies are often criticised.

CHAPTER 2

CONTEXT SPECIFIC INVENTORY MODEL DEVELOPMENT

Any environmental assessment process necessarily rests on some sort of a model generating or collating the environmental data. The core of an LCA is thus the inventory model, and decisions and assumptions made in its construction will underpin the integrity of the entire assessment. The underlying concepts of inventory modelling are introduced in this chapter. Definition of the system(s) to be modelled falls under the goal and scope phase of LCA, which formalises the problem structuring phase of the decision-making process (identifying the issues, setting the criteria and identifying the options). The scope of the study sets what is to be included in the assessment, and at what level of detail. It should be defined in sufficient detail to ensure that the breadth and depth of the study are compatible with and sufficient to address the stated goal (Consoli et al., 1993; ISO, 1997). These decisions cannot be taken in isolation and are very much a function of the context of the study. It is thus important to gain a clear understanding of the decision to be supported at an early stage of the LCA process.

Goal definition and scoping is not only essential for defining the inventory phase of the LCA, but also for decisions taken to shape the impact assessment and improvement assessment phases of LCA. For example, the temporal and spatial placement of the system, the impact assessment method followed, and the impact categories considered, are all vital decisions influencing the outcome of the impact assessment. Also, the scope of the study will determine what level of improvements in the system can be identified. The central nature of goal definition and scoping has led to the suggestion of goal-dependent methodology (Frischknecht, 1998; Wenzel, 1998). In this way, the broad scope of decisions to be taken during goal definition and scoping can be tailored to the type of decision to be supported, thereby simplifying the LCA process. The various application typologies that have been identified in the literature, and how they affect the necessary decisions made during goal definition and scoping, are investigated in this chapter, and a conceptual model covering the decision contexts encountered in primary industries is presented. This chapter thus presents the basis for the distinction between the three case studies presented in chapters 6, 7 and 8, and why particular methodological decisions are taken in each.

2.1. INVENTORY MODELLING

The emphasis in this study is on generating LCI information to support decisions. A static inflexible inventory model is therefore of limited use, and the inventory must be able to be predictive and respond to changing input variables, at a level of detail consistent with the decision context. It must be sufficiently detailed to track emissions to their source, and to the operating variables responsible for their control. This represents a departure somewhat from traditional LCI methods, and a merging of technology assessment and LCI methodologies. The key considerations in constructing inventory models are discussed below.

2.1.1. Systems Modelling in LCA

Grisel et al. (1997) define system modelling as building relationships between processes, based on information about how parameters influence the system and information about the relationships so that the system might act dynamically. A better reflection of the interplay of these parameters in the inventory model gives better possibilities for analysis and improve the reliability of the results. However, the influence of a more sophisticated model on the results has to be balanced by gathering and treatment of much additional information on the physical characteristics of the links within the system (Grisel et al., 1997). Although LCA aims to be comprehensive, it is recognised that it is not always practical to study all the relationships between all the processes in a product system, or all the relationships between a product system and the system environment (ISO, 1998). Frischknecht (1997) states that the inventory analysis needs to address two central questions:

1. How to adequately model reality in the inventory phase?
2. How to relate the different models developed and used to the questions asked, i.e. to the goals of the study?

This is in agreement with Wenzel (1999), who recognises that the aim of the system models is to reflect the environmental consequences of the decision as well as possible, whilst matching the model to the resulting consequences of the decision (Wenzel, 1999). Frischknecht (1997) warns that the growing volume of information may reach the limits of calculation power or personal receptivity of information, and tradeoffs between simplification and robustness of results should be identified. More sophisticated system modelling is also expected to enhance the sharing of data between LCA and other environmental management tools (Grisel et al., 1997), e.g. including a site-specific dimension will

enable a closer co-operation with the tools of ERA. A key aim of inventory modelling is thus to develop system models that are able to answer the questions asked with a sufficient degree of confidence, whilst expending minimum time and expense. On the other hand, the degree of sophistication able to be obtained is limited by the specific characteristics of the system, particularly the amount of information available to model the system. This “fitness of purpose” of the inventory model, together with the inherent limitations to the system models, is discussed in the context of primary industries in section 2.3.

2.1.1.a Scenario Development

Scenarios are an integral part of an LCA study, although a systematic guide to their development is not included in any of the formal LCA methodological frameworks (Pesonen et al., 2000). They are required to organise disparate economic, technical, political and social information into a framework where it is able to be modelled and judged. A scenario is usually used to describe a snapshot in time, i.e. the condition of important variables at some particular time in the future, or it is used to describe a future history, i.e. the evolution from present conditions to one of several futures (Pesonen et al., 2000). “Scenario” is used in this thesis to refer to the setting of the frame conditions of the system to be modelled. The definition of the scenario thus includes the specific assumptions underlying the scenario, which provide the basis for modelling each alternative system.

Pesonen et al. (2000) differentiate between three types of scenario applications in LCA: technology scenarios, environment scenarios and valuation scenarios. Technology scenarios are applicable to the inventory modelling phase, whilst environment and valuation scenarios fall into the phases of impact assessment and interpretation. In this thesis a distinction is made between scenario analysis and sensitivity analysis. Environment and valuation scenarios are thought to be better described as sensitivity analyses, where the effect of the time horizon and geographical boundary applied during the modelling of environmental effects, the environmental model itself, and the weighting factors used can be analysed. In inventory modelling, both sensitivity and scenario analysis play a role. Sensitivity analyses involve analysing the effect of varying key parameters, e.g. major process variables, allocation rules etc. (more on using sensitivity and parametric analyses for an assessment of model uncertainties follows in chapter 3). Scenario analysis, on the other hand, involves the specification of the various technological systems to be modelled.

2.1.1.b Process modelling in LCA

Process modelling characterises the relationship between process inputs and outputs. A number of independent parameters are introduced that provide information on the relationship between the inputs and outputs, allowing the effect of various influences on the process to be explored. Grisel et al. (1997) identify three types of process modelling employed in LCA studies:

1. Black box models of processes
2. Models of processes with linear functional relations
3. Models of processes with a mix of non-linear and linear functional relations

Black box models are the most commonly applied, as they require the least information, and all publicly available LCA databases are an example of this approach. All environmental interventions are linked with homogenous linear functions to a single variable (the functional unit), i.e. doubling the functional unit leads to a doubling of all resource inputs and emissions to the environment (Grisel et al., 1997). In reality, the process function will be dependent on several variables, with the function either linear homogenous, linear non-homogenous or non-linear (Frischknecht, 1997). Doubling the functional unit then leads to different linear and/or non-linear relations between the input and output streams. The increase in effort and complexity has meant that few linear and non-linear process models have been performed for LCAs, especially so for the non-linear case (Grisel et al., 1997).

Homogenous linear functions are inappropriate for primary industries (or indeed for many processes), as doubling the functional unit does not result in a two-fold increase in the environmental interventions. In addition, Fleischer and Schmidt (1995) found that the true relationship between inputs and outputs is essential for adequately performing a sensitivity analysis. The inventory models developed in this thesis are therefore predominantly based on linear mass-balance factors specifying the relationship between the inputs and outputs. Although also required for a meaningful uncertainty analysis (see following chapter), this approach was somewhat forced on the assessment through the type of data available. Emissions monitoring is not required by legislation in South Africa, other than for particulates. Process outputs and emissions to the environment are therefore calculated from data monitored for process operation and control purposes, and from the input streams (whose quantities are mostly known, as they are associated with a monetary value and their consumption therefore recorded). This supports the assertion of the Grisel et al. (1997) that an expected benefit of linear and non-linear process modelling is to enhance data collection. This is in agreement with Keoleian (1993), who states

that to be cost-effective and facilitate the collection of data, LCA should be integrated into the existing environmental management and information systems within a company .

2.1.1.c Temporal and Spatial Restrictions

LCA models currently integrate results over time and space. They are therefore inherently limited to an assessment of impact potentials, and not actual impacts, as can be determined using approaches such as ERA. This topic is covered in greater detail in section 5.3, and is only discussed here as it pertains to inventory modelling.

The need for an inventory disaggregated in time and space is a result of the requirements of impact assessment, where, increasingly there is a call for the determination of actual site-specific impacts, requiring the use of site specific, time-dependent inventory data (Finnveden and Lindfors, 1997) (the determination of actual impacts requires a knowledge of the background load and sensitivity of the receiving environment, and the duration of the emission). Disaggregation of the inventory according to time and location has been demonstrated by several studies (Frischknecht, 1997). However, these were shown to require significantly more detailed inventory models, and resulted in difficult to interpret solutions, as spatially and/or temporally differentiated results increase the dimensions of the already extensive results requiring presentation in an LCA study.

Spatially differentiated damage factors have been developed for some impact categories, e.g. receiving environment-dependent acidification factors (Potting et al., 1998), and a greater breakdown of damage factors according to release medium is exhibited in the more recent updates of impact assessment methods (Goedkoop and Spriensma, 1999). However, applying such spatially differentiated damage factors requires a similar level of spatial differentiation in the inventory, which requires significantly increased modelling detail and additional data to generate. In many cases, such spatial detail is simply not known, particularly for background processes (secondary materials and processes).

Time-dependent information is even more difficult to internalise in an LCA. The current practice of integrating over time in an LCA arises because LCA is supposed to support decisions whose effects will only be felt in the future, although for an actual system, a substantial part of the processes will have taken place in the past (Frischknecht, 1997). Emissions varying over the life time of a project are usually summed or averaged when incorporated into the inventory. In addition to changing over the years, emissions usually exhibit short-term variability, i.e. daily or monthly fluctuations, which are

averaged out in the inventory. Inventory data can therefore not be matched with the real time emissions that are required for an accurate assessment of environmental impact (Huijbregts, 1998a). As with spatial differentiation, presenting a time-dependent inventory will dramatically increase the modelling detail and data required. Calculating an inventory incorporating the daily or hourly variability required by many damage models is unlikely to be feasible, although an inventory with annual temporal differentiation may add a valuable dimension.

The need and/or possibility to include site and/or time dependent information depends on the decision context. In certain studies the decision to be supported affects processes and products whose location is fixed and known, whilst in others the location of the processes are not fixed and remain unknown. The latter case often involves decisions that affect future systems (long-term forecasting), where it may be impossible or even undesirable to include site-specific information. In this way, time and space considerations are interrelated, in that the time horizon to some degree affects the availability and the need for site specific information, with the need for site dependent information decreasing as the time horizon increases (Wenzel, 1999).

2.1.2. System Boundary Definition

A necessary first step in an LCA is the drawing up of the system flowsheet. This identifies all material and energy flows involved in the provision of the product or service, tracing them back up or down the flowsheet to flows either directly from or to the environment. In most studies, there is not sufficient time, data or resources to include all processes identified in the system. Decisions to omit life cycle stages, processes or inputs and outputs are dictated by the application of the study, i.e. each simplification made to the system inevitably comes with an associated increase in uncertainty, which must not exceed that acceptable to the goal of the study. The system boundaries define the exact placement of the system in the overall system environment. Lindfors et al. (1995) identify three different boundaries requiring definition:

- Geographical boundaries
- Life cycle boundaries
- Boundaries between the technosphere and biosphere

Heijungs et al. (1992) add a fourth: The boundary between the product system under consideration and other product systems, which involves the application of allocation procedures.

The geographical boundary specifies the area for which the study is representative. The boundary can be political, e.g. national boundaries, or natural, e.g. ecosystems, or climatically or geologically defined regions. The geographical boundary has implications for scenario development, choice of impact categories and data collection/uncertainty. Certain impact categories may not be relevant on a local scale, and similarly, certain site-specific impacts, are not able to be meaningfully assessed on a global scale. Data collected from an area outside the specified geographical area will introduce an element of uncertainty, depending on the similarity of the areas, and thus the applicability of the data to the system of interest. The spatial positioning also has a large effect on the quantity of data required, particularly around impact assessment, with the level of information required increasing as study moves from a global to a local scale.

To be in accordance with LCA principles, the life cycle boundary should include all stages of the life cycle from primary resource extraction through to final waste disposal, including use, re-use and recycling. However, the products from primary industries spread out to many diverse uses, and a study could not feasibly follow all of these. The system boundary for these systems is therefore “cradle-to-gate” rather than “cradle-to-grave”, as it ends where the material spreads out to its many uses. The amended life cycle analysis tracks the flow of material from its extraction and preparation through to primary processing, and includes all relevant associated processes, including the provision of ancillary materials and services.

The following processes, identified by Lindfors et. al (1995) as usually not quantified as part of an LCA study, are excluded from the life cycle boundary considered in the case studies in this thesis:

- capital equipment and infrastructure,
- accidental spills,
- impacts caused by personnel, and
- human resources.

For many systems the share of releases due to capital goods and infrastructure are less than 5-10% (some notable exceptions are in the transportation sector, renewable energy and agricultural systems) (Frischknecht, 1997). Similarly, personnel related impacts, e.g. canteens or sanitary facilities, are usually not considered because the differences between the scenarios are likely to be small (Lindfors et al., 1995). Accidental spills are not included where process modelling is based on normal or design operating conditions, and even where they are incorporated, e.g. in total annual or monthly emissions, they are usually not considered explicitly. The use of risk assessment has been suggested to address the

possibility of severe accidents (Frischknecht, 1997). Additional scope for system simplification arises in comparative assessments, where processes that remain essentially unchanged across the scenarios can be excluded, although it is essential that the systems be comparatively complete, or the more complete systems will be negatively biased.

The boundary between the technosphere and biosphere delineates the boundary between the technical processes and the environment. This is usually clear, other than in biomass production/agricultural processes and in landfill processes (Frischknecht, 1997). The latter are of significant importance in primary industries, and these seemingly “environmental” processes require placement within the system boundary or significant impacts can be missed. The life cycle inventory records all material and energy flows crossing the system boundary, so it is important that the containment of mining and processing residues are explicitly included within the system boundary. If merely the mass of solid waste produced is recorded, the waste effectively leaves the system, and the potential emissions associated with landfilling the waste escape assessment. Waste deposits therefore need to be included as “pseudo unit-processes” within the system, allowing dust and leachate to be included in the LCI. Landfill processes can be seen to be eventually “returned” to the natural system, although there are conceptual problems in defining the time-scale over which this occurs. Frischknecht (1997) suggests that information on the dissipative use of resources can be used, i.e. the difference in concentration between a resource in an ore when extracting it and the concentration of that resource when finally dumped in a landfill. More common, and the method followed in the case studies, is the definition of a time-span for which emissions from the landfill are included in the inventory (Finnveden et al., 1995).

2.1.3. System Complexity

The system needs to be specified at a level of detail consistent with the type of decision to be supported. The level of disaggregation in the flowsheet is most often driven by information availability, with higher levels of aggregation inevitable where information is deficient. A high degree of breakdown in the flowsheet (usually to a unit operation level) allows a detailed inventory to be constructed, and the construction of an audit trail, which links the impacts back to the unit operations responsible for generating them (Stewart, 1999). A high degree of flowsheet breakdown is therefore essential if the LCA is to be used for the analysis and improvement of existing process systems, or in a design study to explore the effect of the various operating variables. To maintain the audit trail, Stewart

and Petrie (1996) suggest the following heuristics to define the appropriate level of flowsheet breakdown (in the context of minerals processing):

- Common function:* If the function of units is common they are integrated into one unit.
- Mass flow rate:* Units which have a high through-put, high reagent addition or high waste generation must be kept separate so that their effect on the total waste stream is not diluted.
- Hazardous waste:* If a unit gives rise to a hazardous emission it must be kept separate so that its point of emission can be pinpointed.
- Energy Intensity:* Units with a high energy consumption must be identified.
- Waste generation:* A unit which generates a specific waste is not always the point at which that waste crosses the system boundary, so it is important to couple waste generation with the process unit where it arises, and not only where it crosses the system boundary.

Another reason for a disaggregated inventory is to facilitate the integration of LCA with tools capable of predicting site-specific impacts. In its current format, LCIA is limited to predicting potential environmental impacts, most often of global significance (see section 5.3). Site-specific time-dependant impacts fall into the gambit of ecological risk assessment (ERA), whose methodology integrates the results of exposure and effects assessments to obtain a probabilistic expression of the effects of environmental change (van Leeuwen and Hermens, 1993). The fate and transport models of ERA require information on the point source of pollution (the type and quantity of emission) as well as spatial information (background information). Information on the point source of pollution is consistent with the type of information generated by a disaggregated LCI study, so to facilitate this integration of LCI with ERA (or other site-specific environmental assessments), the inventory must be structured so that the information is available at a disaggregated level.

However, modelling at a detailed level may not be possible or even desirable in all circumstances. It has been suggested that only the portion of the system of direct interest, i.e. the group of processes over which influence can be exerted by the decision makers, should be modelled in detail (Clift et al., 1998; Trinius and Le Téno, 1999). Clift et al. (1998) distinguish between a foreground and a background system, where the background system is defined as the set of processes whose operation is not directly affected by decisions based on the study, other than the quantity of material (or magnitude of the function) input into the foreground system. Ancillary or secondary material inputs usually fall into this

definition, e.g. fuels, treatment chemicals etc. An important distinction between the foreground and background sub-systems is that data at a high degree of aggregation can be used to describe the background processes (Clift et al., 1998), i.e. secondary or generic data, as in published LCI databases.

Similarly, Trinius and Le Téo (1999) suggest splitting the system into different sub-systems depending on the sphere of influence or focus of the decision maker. They define three different focal zones, where a focal zone is defined as “a system model or part of the physical reality based on part of a given actor’s decision scope” (Trinius and Le Téo, 1999). The system model comprises processes that are identified as being part of a primary or secondary focal zone, where the primary focal zone consists of the processes identified by the decision-maker as the primary field of concern, and the secondary focal zone comprises those processes indirectly influenced by the product or process in main focus. Aspects outside the decision scope of the actors, i.e. those process ruled out as not influenced at all by the decision-maker, are relegated to the non-focal zone and are not included in the system model. Positioning the unit processes in the primary or secondary focal zone affects the level of flowsheet breakdown and the quality of the applied data (Trinius and Le Téo, 1999). Thus, as in the foreground/background approach, aggregated generic data can be used for processes in the secondary focal zone.

In both approaches, the use of detailed process modelling and high quality data for the processes over which direct influence can be exerted makes good sense, as it is likely that higher quality information will be available to the decision makers around the processes that are of direct concern to them. The use of a combination of detailed modelling with high quality data and low resolution with secondary data is supported by the findings of Hunt et al. (1999). Although they base the differentiation on a sensitivity analysis, i.e. by finding the parts of the system that contribute little to the totals, they found that such a combined approach deviated least from the results of a fully detailed LCA, compared with a number of different streamlining methods (Hunt et al., 1999).

A similar conclusion can be reached by looking at the system from a marginal technology perspective. The proponents of this approach state that for comparative life cycle assessments the actual environmental impacts are most correctly modelled by using data on the marginal production facilities, since marginal data reflect better the actual consequences of a decision (Weidema et al., 1999). This approach calls for site specific, detailed modelling around those processes whose production volume is

directly affected by the studied change, i.e. the so-called marginal technologies. This is somewhat analogous to placing certain processes in the foreground system (Weidema et al., 1999).

2.1.4. Inventory Model Structure

The underlying concept of an inventory model is the conservation of mass and energy. The inventory table records all material flows crossing the system boundary. Ideally, mass flows entering should equal mass flows leaving, but usually this is not possible, as there will be some degree of simplification. Also, the vast differences in the magnitudes of certain environmental interventions (often several orders of magnitude) mean that an exact balance is infeasible. However, the mass balance is an indispensable check to see if all the major flows have been accounted for. Similarly, energy entering the system should equal the energy leaving. However, the energy balance is a whole lot more complex to specify, as it requires data on the energy contents and chemical energies of the various inputs and outputs (Heijungs et al., 1992). The lack of standardisation on the definition of energy content, as well as on reference temperatures and pressures, also hampers the drawing up of an energy balance. It is therefore common practice to limit the assessment to a global/approximate energy balance, in which only the major energy flows (usually only electricity, fuels and steam) are incorporated (Heijungs et al., 1992; ISO, 1998).

For practical reasons, the inventory is usually built up from a number of smaller material balances around the individual processes or unit operations included in the system, which are then summed to yield the overall inventory table. The level of breakdown of these sub-balances is essential, as this is the lowest level at which inventory information is available (see section 2.1.3). The sub-balances are calculated on some relevant basis, e.g. 100kg product, which needs to be related to the actual amount required by the system. This is done through the specified reference flow, termed the functional unit, which is the flow required to fulfil the function of the system. Two calculation procedures have been identified to scale the individual balances to the overall inventory; the commonly used sequential method and the mathematically elegant matrix method (Heijungs et al., 1992).

In the sequential method, references are made to earlier processes in a certain ratio, based on the product, material and energy needs of the process which supplies the functional unit, and the quantity of waste processing services required, based on the waste produced by the functional unit. In turn, the requirements of these processes initiate the next round of calculations, following up the entire process

chain until there is no further demand or supply of product, materials, etc. A problem of the sequential method is that it requires an iterative calculation to deal with processes which refer to themselves, commonly encountered in energy production processes, e.g. coal mining requires electricity, which in turn requires coal to be produced. In the matrix method, the occurrence of the processes is determined simultaneously, making it possible to deal with feedback immediately. The matrix method provides algebraic expressions for the quantification and occurrence of each process. Details of the linear algebra manipulations involved can be found in the relevant literature (Frischknecht, 1998; Heijungs, 1994).

Although less concise and accurate in its calculation method, the simpler and easily followed sequential method is used in the models developed in this thesis. A key requirement of the models is that they be auditable and allow for an explicit investigation of uncertainty in the model parameters, i.e. that the effect of key process parameters be traceable through the model. This was judged to be more efficiently and transparently achieved with the sequential method.

2.2. APPLICATION DEPENDENCY OF LCA

The above discussion has consistently stressed the dependence of the various methodological choices governing inventory modelling on the decision to be addressed by the LCA. In fact, all the items requiring consideration under the scope definition phase of LCA are to some degree a function of the goal or application of the study. The complexity of the choice of functional unit is dependent on the decision to be supported, and can range from very global to very specific, limiting the alternatives for consideration as it moves towards the specific (Heijungs et al., 1992). The choice of impact categories and the methodology of impact assessment requires justification in relation to the goal of the study. Whether global, regional or local impacts are to be considered, and the time horizon over which they will be assessed, are dictated by the spatial and temporal context of the decisions to be supported. Data collection is directed by the application of the study, and will dictate the age and the degree of geographical and technical specificity deemed acceptable. The data availability will prescribe certain assumptions that must be taken, which, amongst other factors, will limit the degree of confidence that can be attained in the results. The need for a critical review process and the requirements for the report are also wholly dependent on the application and intended audience.

Consequently, from early on in the development of LCA there has been an emphasis on defining the purpose or application of the study, and what this implies for the rest of the study. Early methodological work differentiated between applications internal to a company and external applications. The distinction between the two broad application areas related primarily to the quality assurance of the results (e.g. the need for critical review, data quality assessment etc.) and the reporting structure required (Consoli et al., 1993; Fava et al., 1991). More recent work has been to try and classify the application in relation to the relevant system model type. This progression in classification is developed below.

Consistent with the early SETAC classification, Udo de Haes (1993) distinguishes between internal and external applications, but also makes a distinction between the main users (governments, companies and NGOs), and the application level (operational or strategic, where the former deals with concrete objects, such as products and materials, and the latter with general objects, such as policies and strategies). Each of these combinations of applications were further subdivided according to their level of sophistication, where he distinguishes between LCA as a concept, qualitative LCA and quantitative

LCA, and their level of completeness, i.e. whether all stages in the life cycle are considered, whether all phases of LCA are considered and, if LCIA is included, whether all relevant impacts are considered (Udo de Haes, 1993). The dependence of the application area (i.e. main user and application level) on the detail and depth of the study is clearly shown in his review, with NGOs found to be performing only qualitative LCAs; governments performing mostly qualitative LCA, some simplified quantitative LCA but few detailed quantitative studies; and companies performing all three, but notably more detailed and simplified quantitative LCAs, especially around operational issues (Udo de Haes, 1993).

Heijungs et al. (1992) identified three application areas relating to products; information, innovation and regulation, as well as a fourth area; policy strategies, which they define as pertaining to comparative assessments of scenarios rather than products. Lindfors et al. (1995) distinguish between applications in the private and public sector, which is similar to the internal/external split of SETAC. The ISO standard makes a distinction between hot-spot identification, decision-making in industry, government or NGOs, the identification of relevant indicators of environmental performance, and marketing purposes. Fleischer et al. (1995) differentiate between applications according to development, selection and optimisation, further distinguishing between new or existing products, and one or several products.

Weidema (1998a) distinguishes between enterprise specific versus generic applications, between applications for information versus decisions supporting change, and on the level of application, i.e. operational, tactical or strategic applications. Here operational decisions are characterised by being non-comparative and the results used on the product itself, tactical decisions by improvements being evaluated by comparisons between products and the results being used to influence the surroundings of the product, and strategic by improvements being evaluated in relation to an environmental target and the result being used to place the product in a larger context (Weidema, 1998a). He thus develops a matrix of application areas, which encompasses the application areas listed in the ISO standard and Nordic guidelines (ISO, 1997; Lindfors et al., 1995).

Wenzel (1998) identifies three key methodological parameters whose choice significantly depend on the application of the study. The first two are a function of the decision consequence, the environmental and socio-economic consequence, and the third a function of the decision context. The environmental consequence of the decision, i.e. the nature and extent of the environmental change caused by the decision, relates to the form of system model required to model the changes. This in turn relates to the

LCA, and their level of completeness, i.e. whether all stages in the life cycle are considered, whether all phases of LCA are considered and, if LCIA is included, whether all relevant impacts are considered (Udo de Haes, 1993). The dependence of the application area (i.e. main user and application level) on the detail and depth of the study is clearly shown in his review, with NGOs found to be performing only qualitative LCAs; governments performing mostly qualitative LCA, some simplified quantitative LCA but few detailed quantitative studies; and companies performing all three, but notably more detailed and simplified quantitative LCAs, especially around operational issues (Udo de Haes, 1993).

Heijungs et al. (1992) identified three application areas relating to products; information, innovation and regulation, as well as a fourth area; policy strategies, which they define as pertaining to comparative assessments of scenarios rather than products. Lindfors et al. (1995) distinguish between applications in the private and public sector, which is similar to the internal/external split of SETAC. The ISO standard makes a distinction between hot-spot identification, decision-making in industry, government or NGOs, the identification of relevant indicators of environmental performance, and marketing purposes. Fleischer et al. (1995) differentiate between applications according to development, selection and optimisation, further distinguishing between new or existing products, and one or several products.

Weidema (1998a) distinguishes between enterprise specific versus generic applications, between applications for information versus decisions supporting change, and on the level of application, i.e. operational, tactical or strategic applications. Here operational decisions are characterised by being non-comparative and the results used on the product itself, tactical decisions by improvements being evaluated by comparisons between products and the results being used to influence the surroundings of the product, and strategic by improvements being evaluated in relation to an environmental target and the result being used to place the product in a larger context (Weidema, 1998a). He thus develops a matrix of application areas, which encompasses the application areas listed in the ISO standard and Nordic guidelines (ISO, 1997; Lindfors et al., 1995).

Wenzel (1998) identifies three key methodological parameters whose choice significantly depend on the application of the study. The first two are a function of the decision consequence, the environmental and socio-economic consequence, and the third a function of the decision context. The environmental consequence of the decision, i.e. the nature and extent of the environmental change caused by the decision, relates to the form of system model required to model the changes. This in turn relates to the

level of LCA modelling required and the positioning of the analysis in time and space, i.e. the need for trend analysis and projection, and the need for site specific information. The social and economic consequence relates to the burden of proof within the study, and therefore influences the level of certainty, transparency and documentation required. The socio-economic consequence has nothing to do with the prescribed level of LCA required but merely with the requirements for documentation and the validity of the conclusions. The decision context relates to the choice of impact assessment and weighting criteria, i.e. those elements that implicitly derive from the problem to be studied, and influence the definition of impact assessment criteria, including considerations for site dependency and values in the weighting (Wenzel, 1998).

Recent literature shows an increasing distinction between “retrospective and prospective” (Weidema et al., 1999) or “informative and effect-orientated” applications (Frischknecht, 1997), where the former deal with hot spot identification and product declarations, and the latter with comparative studies. Clift et al. (1998) suggest the time-horizon as the key parameter for shaping the methodology to the application, and distinguish between four main application areas (Clift et al., 1998):

- Historical (not relating to any choice), e.g. environmental reports
- Short term (time span short compared to capital replacement cycles), e.g. system optimisation
- Mid-term (time span in the order of capital replacement cycles), e.g. hot spot identification and elimination, product optimisation and development
- Long-term (time span in the order of technology-mix replacement cycles), e.g. strategic planning

An important distinction is that average data are only relevant for the historical case, and that identification of the short-term and mid-term marginal technologies is required for the short-term and mid-term cases, whilst data on the anticipated future changes of technologies and technology mixes are required for the long-term category (Clift et al., 1998). Pesonen et al. (2000) similarly distinguish applications dependent on the time horizon of the study, and define the two opposite extremes, what-if scenarios and cornerstone scenarios. The what-if scenario is used to compare two or more options in a well-defined situation with a relatively short time-horizon, whilst the cornerstone scenario applies to the opposite extreme, and is relevant to long-term planning, with long time horizons and ill-defined systems. Most systems fall somewhere between these two extremes (Pesonen et al., 2000). A definition based on time horizon is consistent with the “levels of improvement” suggested in Wrisberg and Gameson (1998), after Stevels (1997). “Level 1” requires operational information to describe small

changes within small scale systems with short time horizons, whilst “Level 4” requires strategic information to describe large and possibly qualitative changes of large scale systems with long time horizons. All other decision-types fall somewhere between these extremes.

Weidema (1998a) updates his typology with respect to the parameters identified in the literature, unifying the various approaches. His classification and the relevant determining parameters are summarised in Figure 2-1. The information vs. change distinction of Weidema (1998a) does not fit into this representation. This distinction is largely artificial, since it can be argued that the purpose of any LCA is ultimately to affect a choice, or the study is largely without meaning (Wenzel, 1999). Information studies should perhaps rather be seen also as change or effect orientated, but where the decision maker has only a passive influence and not an active role (Weidema, 1998a). The decision maker exerts his/her influence on the study, as different decision-makers will have different resources at their disposal, and are able to afford different levels of data collection and inventory sophistication (Wenzel, 1998). The decision-maker should therefore be seen as an independent parameter in Figure 2-1 (Weidema, 1998a).

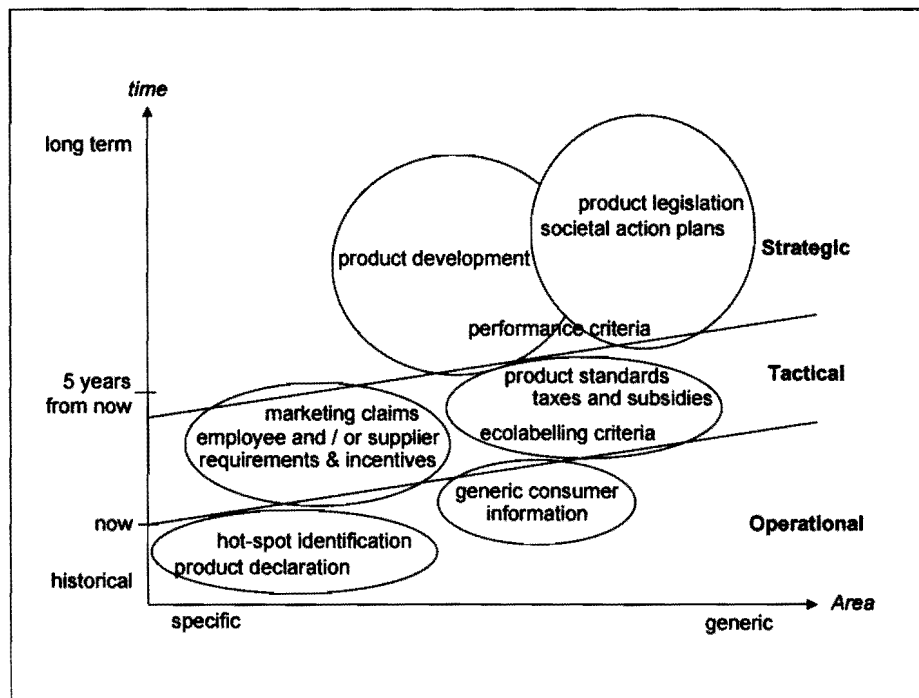


Figure 2-1 Application areas of LCA studies in relation to their determining parameters. The vertical axis represents time, while the horizontal axis may be interpreted as both geographical space and as the more abstract “area of affected products and interest groups” (Weidema, 1998a).

2.3. DECISION CONTEXTS FOR PRIMARY INDUSTRIES

To enable more relevant use of LCA in decision making, it has been suggested that the methodology be shaped to fit the decision in hand. The application typologies discussed above are a step towards achieving this. In this section, three decision contexts are identified, which broadly encompass the types of decision systems arising in primary industries. The various typologies presented above have a number of common threads running through them, and in many cases it is more a matter of definition or semantics than a true difference. The models proposed attempt to distil the main points from the various typologies, and are discussed with reference to the main methodological features affected by the decision type.

The decision models proposed are based on the placement of the decision context in relation to the system generating the information to support the decision. The definitions are thus a convenient way of seeing information flows between the system generating the information and where it is applied, and how this affects the quantity and quality of the information. The following three information mappings are identified and represented schematically in Figure 2-2:

- Type I: Decision context in parallel to production system
- Type II: Production system within decision context
- Type III: Decision context within production system

The production system is analogous to the foreground system or primary focal zone defined by Clift et al. (1998) and Trinius (1999) respectively. The relationship between the decision context and the production system is thus from the perspective of the decision maker. Type I encompasses decisions where the decision maker has no input into the generation of the information, other than the choice to use information from that production system. In this case the production system falls into the background system or secondary sphere of interest with respect to the decision being taken, i.e. the information is used as secondary or background information in a study separate to the one in which it was generated.

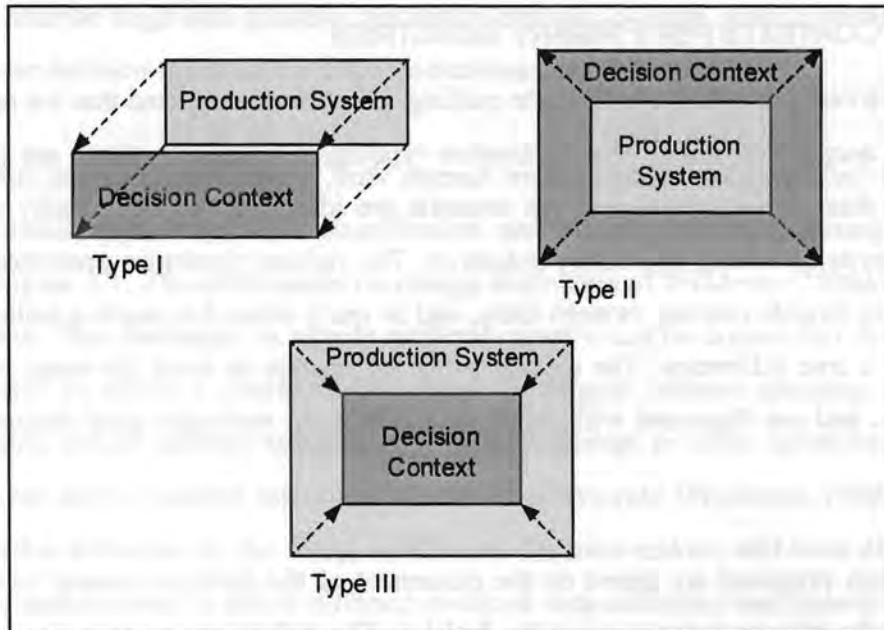


Figure 2-2 Schematic of decision model systems.

In the Type II and III decision models the decision maker is actively involved in the generation of information. In Type III systems, the decision falls wholly within the decision maker's sphere of influence, whilst in Type II systems the decision encompasses some elements wider to the decision maker's particular sphere of influence. Type II systems therefore involve decisions of a more tactical or strategic nature, whilst Type III involve more operational-type decisions. The Type II and III models represent the extremes of the decision types, and most decisions will fall somewhere between the two extremes, e.g. in process design, a decision may start out as purely tactical (a technology assessment, or product scanning assessment), and become more operational as the design process advances.

The system model and its implications on data issues (information flows) are thus the central focus of the decision models. As the "internal vs. external" distinction made by a number of authors does not affect these aspects of the methodology, it is not considered relevant here. This is in agreement with Weidema (1998a), who states that there is no apparent link between the intention to publish and the methodology of the study. The background and education of the intended audience and the time available for decision making will no doubt depend on the application area, but this will primarily affect the reporting structure, which is independent of the methodological issues under consideration.

Consoli et al. (1993) also relate the required quality assurance of the study to whether it is applied internally or externally. The quality of the study results, and therefore the certainty with which the decision can be implemented, is a function of a number of key parameters, including the data certainty and completeness, the detail of the modelling, data uncertainty and sensitivity analyses performed etc. The quality of the results are therefore very much a function of the methodological choices taken. However, quality assurance is more meaningfully related to the decision consequences than to the intended users of the study (Wenzel, 1998). The implications of the environmental and socio-economic consequences of the decision on the methodology are incorporated into the decision models by recognising the possibility of different sub-contexts within the Type II and III decision models, relating to the level of risk able to be tolerated, i.e. the required degree of certainty in the results.

The differentiation between the Type I model and the other two types is equivalent to the “information vs. change”, the “prospective vs. retrospective” and the “informative vs. effect-orientated” application types of Weidema (1998a), Weidema et al. (1999), and Frischknecht (1997) respectively. The progression from the Type II to Type III models has parallels with the “operational, tactical or strategic” application type of Weidema (1998a), the “operational vs. strategic level” of Udo de Haes (1993) and the time horizon-dependent application types of Clift et al. (1998).

However, Weidema (1998a) and Udo de Haes (1993) offer different definitions of “operational” and “strategic” to the sense in which they are applied to the Type III and II models, respectively. Weidema’s definition is product-focused, with operational applications being defined as non-comparative, tactical as comparisons between products, and strategic as comparisons in relation to an environmental target. Udo de Haes makes the distinction between concrete objects (products and materials) in operational applications and general objects (policies and strategies) in strategic studies. However, the definitions are applied to the decision models more from a systems than a product perspective, in accordance with the definition by CHAINET (1998), i.e. operational information describes small changes of small scale systems with a short time horizon, whilst strategic refers to large and possibly qualitative changes of large scale systems with long time horizons (Wrisberg and Gameson, 1998). As used here the time horizon applies to the decision cycle, i.e. the time from which the study is commenced to the time when the chosen solution is implemented. It therefore does not relate to either the duration of the study, the period over which data is collected, or the time over which impacts are seen to act. These temporal considerations relate more to the consequences of the decision, i.e. the level of certainty required in the results.

This definition of a decision continuum from an operational to a strategic context correlates with the time horizon-dependent definitions of Clift et al. (1998), i.e. a short, mid and long-term decision perspective. The Type III model is equivalent to the short time perspective, whilst the mid and long-term are combined in the Type II model. The definition of an “operational” application as used by Weidema (1998a) (i.e. non-comparative), is actually closer to the “historical” category of Clift et al. (1998) or the Type I model definition, than the short-term decision or Type III systems of the other definitions. The number of stakeholders can also be related to the decision context, with strategic decisions expected to have a far more stakeholders and interested and affected parties than an operational decision. This roughly parallels Weidema’s distinction between “enterprise specific vs. generic” applications (Weidema, 1998a).

2.3.1. Type I (Background data / Product declarations)

The Type I decision model represents systems where the decision context lies outside the sphere of influence of the person generating the data, i.e. system models which generate information without knowing the exact destination of this information (e.g. LCI information to be placed in a database, or product declaration information). Although decisions may well be taken as a result of the use of the data, the compiler of the inventory does not exert any direct influence on the decision which the data is used to support, and the nature of the decision in which it will be used is not known at the time of generating the information.

The exact origin of secondary or background inputs is usually not known, i.e. the geographical location of the plant and the technology it employs is usually detail beyond the scope of the study. LCA databases typically reflect sector-wide averages, sometimes in national or technological sub-divisions, e.g. European electricity production, electricity from coal etc. It has been acknowledged in the literature that if a decision is involved (i.e. a choice between systems), average data should not be used, and that only a marginal analysis is valid, i.e. data on the actual technological systems affected must be collected (Frischknecht, 1997; Weidema et al., 1999; Wenzel, 1999). It has also been argued that all inventories are ultimately generated to be used in a decision, as information for its own sake is meaningless. It can therefore be argued that the use of average data is never valid. However, this becomes a circular argument, as the reason for placing products/processes in the background system in the first place is that the detail around their production is beyond the scope of the study. The reasoning behind the foreground/background system (or primary/secondary focal zone) split thus becomes

apparent, where generic, lower quality data is judged acceptable for systems exerting a lesser influence on the main system. The Type I system model, using average data, is therefore still of relevance.

A similar argument ensues with defining the system in the temporal sense. Historical data are typically used in generating life-cycle inventories. As with the marginal technology argument, where the inventory is to be used to inform a decision, it has to reflect a future scenario (Wenzel, 1999). Thus, the use of historical data is not valid for inventories used in a decision system. However, product declaration-type inventories are not developed with the time-frame of the decision they may possibly be used to support in mind. Indeed this is not feasible, as an inventory drawn from a database could be used in the background systems of a number of studies, ranging over short to long time spans. The generation of an historical Type I system model is therefore still of relevance, although it should be augmented with data on future projections where possible. A similar approach to the regional and technological subdivisions of average data could be used, i.e. short, mid and long term projections could be developed for certain aspects of the inventory, e.g. fuel mixes etc.

High uncertainty is inherently associated with average data, as it is collected from a number of different sources over different time periods. The Type I system inventories are therefore associated with high uncertainty. Methods to assess the uncertainty are available, and are covered in detail in section 3. Although quantitative uncertainty assessments have been performed in LCI studies (see section 3.3), it is by no means commonplace yet. It is therefore unlikely that probabilistic output ranges would be utilised in a subsequent study at the present time, and the significant time and effort expended in their calculation would be wasted. A full quantification of the uncertainty may therefore be relevant for Type I models in the future, but is not considered necessary at the present time. However, since a key feature of these inventories is that they require interpretation by outside users, they should include some sort of representation or discussion on uncertainty (e.g. a qualitative or semi-quantitative uncertainty analysis).

System resolution (flowsheet detail) in Type I models should be sufficient to generate a high quality inventory. As with the consideration of uncertainty, that fact that the inventories require interpretation by outside users, requires that they be detailed enough to allow for transparency, i.e. it should be clear what has been included in the analysis. However, where aggregated data is available at no loss of quality, this should be used, as there is no additional benefit per se of using an increased level of detail (since the inventory is to be used at an aggregated level). The available level of data will therefore determine the system resolution of the Type I model. The data of highest quality should be used,

regardless of its resolution, e.g. a metered overall electricity consumption figure may be known, but not the individual machinery requirements. The measured overall value should be used rather than the sum of estimated requirements of pumps, fans etc., as it is of higher quality even though it gives no insight into which unit operations are consuming the most power.

The use of average, historical data precludes any site specific considerations in Type I models, other than may be achieved using broad subdivisions, e.g. national, regional or technology-wide averages. A lack of detailed referencing information in databases (Grisel et al., 1997), or a lack of applicable data, results in inventories from different regions from the one under consideration often used in LCA studies (e.g. an inventory will be chosen that reflects the same technology, even if implemented in another country). Thus detailed site-specific information in Type I systems would be superfluous. In addition, large LCA studies typically aggregate background data from a wide array of regions, countries and even continents. This is an inherent limitation of the LCA methodology to providing real, site-specific impacts (see section 5.3). Although it is theoretically possible that the inventory data of each background element be collected with sufficient regional or local detail to make a disaggregated inventory possible, as with the marginal analysis problem, this is infeasible within the scope of most LCA studies. The time and effort spent on data collection will increase dramatically, and the reporting structures of LCI databases would have to be revolutionised. The impacts considered by Type I models are thus limited to those of a global or regional nature. Any other impacts which may be considered can only be seen as impact potentials, and an actual consideration of site-specific impacts is not within the scope of this model type.

2.3.2. Type II (Tactical / Strategic Decisions)

Decision systems of a tactical or strategic nature are encompassed by this model type. Whether more “strategic” or “tactical” will depend on the time horizon of the decision, with decisions of a longer time horizon tending towards “strategic”. However, the same decision elements are applicable to both, so they are incorporated into a single decision model. Tactical and strategic decisions involve elements beyond the direct control of the decision maker, so although primarily involved with decisions in their primary sphere of interest, outside systems will have a significant input into the decision, e.g. long term expansion plans.

The Type II system model necessarily reflects future scenarios. Average data is thus not applicable, and data will have to be predicted for the relevant mid or long term marginal technology. This will obviously be associated with a fair degree of uncertainty, which increases as the time horizon of the study increases. Tactical studies may involve the planning and implementation of already demonstrated technologies or those demonstrated on a pilot scale, so will be less uncertain than strategic studies, which may involve the modelling of laboratory scale or merely conceptually-developed technologies. Uncertainty in the scenarios considered also increases as the decision moves from more tactical to strategic in nature, as the number and effect of outside influences become increasingly difficult to predict, e.g. future consumption patterns. The option set therefore tends to become more loosely defined as the time horizon increases, and consequently a larger number of scenarios require consideration. Strategic and tactical decisions also tend to have a more loosely defined option set, and consequently more scenarios under consideration, than operational decisions.

Uncertainties in the Type II decision model are dominated by those related to the definition of the system, i.e. the decision variables (as classified in Figure 3-1), with uncertainty associated with the empirical parameters or choice variables expected to play a smaller role. Identifying the relevant scenarios for comparison and exploring the uncertainties around the definition of these scenarios is therefore key to an assessment of uncertainty in tactical/strategic decision systems.

Data availability is often the over-riding factor in determining the system resolution, especially in strategic studies, where future systems may be only sketchily defined (the heuristic “rules” of Stewart and Petrie (1996) are applicable to the foreground system where the data availability allows). The foreground/background convention of Clift et al. (1998) is followed, so aggregated LCI data is used for the background processes. A significant consideration of tactical/strategic decision systems is to ensure all the scenarios are constructed at the same level of detail. There is a danger that data deficient systems be modelled in less detail, which prejudices the more complete systems.

Impact assessment in tactical and strategic decision systems is constrained by the same general arguments as in the Type I systems, i.e. they are made up of spatially and temporally aggregated data, limiting LCIA to a prediction of global, potential impacts (see section 5.3). However, an additional limitation in these systems is that their geographical location is most often not yet known, particularly as the decision moves towards that of a more strategic nature. This system type is thus inherently limited to a consideration of global and possibly regional effects, and only potential local effects, i.e. a

consideration of local effects can only be based on characterisation models employing generic information, as site specific information is not able to be collected.

2.3.3. Type III (Operational Decisions)

Operational decisions systems are represented by the Type III decision model, in which the change to be effected by the decision lies wholly within the sphere of influence of the decision maker. The decisions are mostly of a relatively short time scale, and typically involve the day to day operation of the system, e.g. process optimisation type decisions. Different decision contexts arising from different decision consequences can be identified, which are related primarily to the burden of proof of the results, e.g. a study involving an investment in capital equipment will require greater certainty than one which merely involves adjusting a temperature or flowrate.

Operational decisions apply to well defined systems (e.g. a definite change to an existing system). There is thus less of an emphasis on scenario analysis and the exploration of possible operating states, as in tactical/strategic studies. The uncertainty in decision variables can therefore be limited to the few model parameters directly affected by the proposed change. This allows for a more systematic parametric analysis to investigate uncertainty in model parameters, than is possible in tactical/strategic studies, where the large number of significant parameters make this infeasible.

Operational decisions are usually supported by high quality information, because they involve existing systems, i.e. measured process data from the specific process can be used to model the system. There is thus not the additional, and often substantial, uncertainty from extrapolating data from other sources, as in tactical/strategic studies. In addition, operational decisions are defined to have short decision time cycles (typically less than a year). The time from identifying the problem to implementing the solution is therefore shorter than in strategic/tactical studies, and the uncertainty in extrapolating data to represent future conditions is small or negligible. The empirical parameter uncertainty exhibited by operational systems is therefore primarily due to the nature of the data itself (its completeness and reliability), with less emphasis on its applicability for use in the study. Operational decisions relate to a well defined change within an existing system, so model parameter uncertainty is expected to be slight, and uncertainties associated with empirical variables are expected to dominate the overall uncertainty in the results. In addition, process data is typically available over long time-periods, which allows for a

good characterisation of data variability. A probabilistic assessment of empirical uncertainty is therefore well suited to these systems.

The Type III decision model requires system models at high resolution as they usually involve a decision at a high level of flowsheet detail, i.e. involve an aspect of the technology or unit operation rather than the whole technological system in place. As in the other decision systems, low resolution is applicable for background processes. A well considered definition of the foreground/background system boundary split is therefore essential to avoid unnecessarily detailed modelling of processes not directly affected by the decision, i.e. those processes that are only affected by an increase or decrease in their product volume can be incorporated at a high level of aggregation, and high resolution modelling of these processes would be expending unnecessary time and effort.

The Type III decision model applies to existing production systems, thus their specific geographical location is known, and no constraints, other than the expense of additional data collection, restrict the collection of site-specific information. In fact, a consideration of site-specific impacts is often essential in these systems, as the effects of operational decisions are often only felt on a local scale. The work on interfacing LCIA methods with those of ERA is therefore especially relevant for these systems. However, it is important to note that the interface is only possible for the foreground processes, as the same problem of spatially and temporally aggregated background data still applies. It is likely that these systems may have to be characterised by a dual set of local impact categories, one reflecting the actual impacts of the foreground system, and one the potential impacts of the overall system. Only one set of global impacts is required, and one or two of regional impacts, depending on the geographical variability across the background processes.

2.3.4. Case Study Decision Contexts

The decision models as presented above represent the extremes of the decision types. In reality there will be a fair degree of overlap between the decision systems. The case studies presented in this thesis demonstrate the main features of the decision models, although the delineation between them is not always exact. This is particularly true of the Type II and III decision models, which really form a continuum from a strategic down to an operational context, with the decision system showing features more towards that of a Type II or III system depending on where in the continuum it falls.

A case study representing the Type I decision model is developed in chapter 6. This study involves the compilation of an inventory of South African coal-fired power generation. The goal of the study is to provide a comprehensive inventory that can be used in LCAs of South African products. The inventory is built up from a consideration of the average technology mix of South African power plants.

A case study typical of the Type II decision model is presented in chapter 7. The study has two parts and demonstrates the progression from a more strategic level down to a tactical level (and ultimately down to an operational level). The first part is a broad assessment of a range of technology options for decreasing sulphur dioxide emissions from coal fired power plants. The study is typical of that on a strategic level, as it considers technologies at different stages of development, and with very different levels of data availability and certainty. The second part of the study looks in more detail at the preferred sulphur-abatement technology, and at the operating conditions required for it to perform better than a suitably selected base case scenario.

A case study typical of systems represented by the Type III decision model is presented in chapter 8. This study looks at the effect of incorporating a mine effluent into a power station's water system, where only operational changes can be made in the water plant to incorporate the poor quality water, i.e. no new equipment is to be considered.

CHAPTER 3

REVIEW OF ANALYSIS AND MANAGEMENT OF UNCERTAINTY IN LCI MODELS

It is well recognised that an analysis of significance is an essential component of the improvement assessment/interpretation phase of LCA, as it is not possible to rate certain environmental interventions or impact categories as more or less important if the magnitude and certainty of the differences occurring between them are not taken into account (Barnthouse et al., 1998). An analysis of significance refers to assessing the relative magnitude of the numerical differences between alternative systems, and the importance of these differences in making decisions between the systems. An analysis of uncertainty forms the core of an analysis of significance, because it determines whether the calculated difference represents actual empirical differences or whether it falls within the range of uncertainty of the results, i.e. it enables the degree of confidence to be stated for an observed difference between options. The other aspect of an analysis of significance, i.e. the importance attached to the differences, requires information external to the LCA process for its assessment, e.g. the actual occurrence and magnitude of environmental impacts, national or regional priorities etc. (Barnthouse et al., 1998). The degree of difference large enough to be considered significant, and the level of confidence considered reliable, are highly subjective judgements, and will depend very much on the type of study and its application. This chapter concentrates on the uncertainty component of significance, although it should be remembered that these uncertainties will need to be placed in the wider picture of how significant they are in the context of the study.

Methods to quantify the uncertainty and variability in LCA results are essential if LCA is to be used with any confidence as a decision support tool. Methods to date have tended to focus on data quality issues, and have progressed from early qualitative indicator approaches to the more recent quantitative approaches. The methods and approaches demonstrated in LCA studies are reviewed in this chapter, after a general introduction to potential sources of uncertainty in LCA models and the tools to analyse them. The strengths and experiences from these studies are drawn upon to develop a procedure applicable to the inventory modelling of primary industries, which is presented in chapter 4.

3.1. SOURCES OF UNCERTAINTY

The term uncertainty is confusing as it is used to encompass a number of concepts. Uncertainty may result from incomplete or imprecise information, from disagreement between information sources, or from inherent variability in the system. Uncertainty can be about a quantity or about the structure of a model, and arises from the simplifications and approximations inevitably made during model construction. In addition to model and data uncertainty, subjective elements, such as values and preferences, may also be uncertain. An understanding of the different types and sources of uncertainty is necessary if a meaningful assessment of an uncertain system is to be undertaken, especially since different types of uncertainty require different methods for their analysis and reduction.

3.1.1. Sources of Uncertainty in Models

Two broad schools of thought can be discerned in the management of uncertainty in policy models, with the emphasis either on a quantitative probabilistic treatment of uncertainty, or on its qualitative elucidation. The leading proponents of the former, Morgan and Henrion (1990), distinguish between uncertainties in model form and uncertainties arising from quantities input into models. The uncertainty of each type of parameter they define can be quantified and its effect assessed using statistical procedures. This quantitative statistical treatment of uncertain quantities is dealt with in the following sections.

The second approach to uncertainty is essentially complementary to the former, but its primary area of concern is where the uncertainties are too great to be effectively managed by techniques deriving from statistical contexts (Funtowicz and Ravetz, 1990). Whilst incorporating statistical representations of data quality wherever appropriate, the approach is aimed at capturing those (often extremely significant) aspects of uncertainty that usually miss quantitative assessment. Funtowicz and Ravetz (1990) differentiate between three “sorts” of uncertainty; inexactness, unreliability and “border with ignorance”, which correspond to a distinction between the technical, methodological and epistemological levels of uncertainty, respectively. Inexactness is the simplest sort of uncertainty, and is expressed by significant digits and error bars, whilst unreliability relates to the level of confidence that can be placed in a quantitative statement, and is usually represented by a confidence level. Thus knowledge of the behaviour of the data gives the spread (inexactness), and knowledge of the process gives the assessment (unreliability), both which can be captured by standard statistical techniques.

However, Funtowicz and Ravetz (1990) notice that there is still “something more”, which they define as the “realm of ignorance”. This includes all sorts of gaps in our knowledge not encompassed in the previous sorts of uncertainty. Thus they develop a quantitative/qualitative notation (see appendix B.1.4), aimed at enhancing the communication of uncertainties in quantitative information, which goes “beyond what statistics has provided in its mathematical approach to the management of uncertainty”.

3.1.1.a Quantities Input into Models

Not all quantities input into models are uncertain and not all can be treated by the same uncertainty analyses. Morgan and Henrion (1990) distinguish between seven main types of quantity in policy models, listed in Table 3-1, and defined below. For some variables, the category into which it falls depends on the context in which it is used, and on the intent and perspective of the decision maker. Thus a variable can not be irrevocably assigned to a category.

Empirical or chance quantities represent properties or states of the world, and are measurable, at least in principle, either now or at some time in the past or future. The majority of data input into LCA models fall into this category, and sources of uncertainty in empirical quantities are dealt with in detail in the following section. In principle, all empirical quantities are uncertain. There is never absolute certainty about the exact value of any continuous empirical quantity, because no matter how great its precision, no experiment can measure a real-valued quantity with zero error. However, in many instances the uncertainty is negligible for practical purposes, and the quantity can be treated as certain.

Defined constants are by definition not associated with uncertainty. This is true of some constants, e.g. the number of degrees in a circle, but not the fundamental physical constants, e.g. the universal gas constant. These are in fact measurable empirical quantities, and are therefore inherently uncertain, although mostly to such a small degree that to all intents and purposes they can be treated as certain.

Decision variables are quantities over which the decision maker exercises direct control, e.g. the specification of a particular unit operation. By definition, a decision variable has no true value, as it is up to the decision maker to select its value. However, it is very likely there will be uncertainty about the “best” value for the variable.

Value parameters represent aspects of the preferences of the decision makers or people they represent, e.g. the degree of risk aversion when comparing uncertain outcomes.

Table 3-1 Summary of types of quantity in policy models (Morgan and Henrion, 1990).

Type of Quantity	Treatment of Uncertainty	Example
Empirical parameter or chance variable	Probabilistic ¹ , parametric ² or switchover ³	Thermal efficiency
Defined constant	Certain by definition	Atomic weights
Decision variable	Parametric or switchover	Plant capacity
Value parameter	Parametric or switchover	Risk tolerance
Index variable	Certain by definition	Time period
Model domain parameter	Parametric or switchover	Geographic region
Outcome criterion	Determined by treatment of its inputs	Utility

1. The input quantity can be varied according to a fitted probability distribution

2. The input quantity can be varied deterministically

3. The model is examined analytically or numerically to discover the value(s) of a specified uncertain input at which the optimal decision changes.

Index variables, often also referred to as *independent variables*, are used to identify a location or cell in the spatial or temporal domain of a model, e.g. a particular year in a multi-year model. They are used to specify a member from a possible set of elements, and can logically not be seen as uncertain.

Model domain parameters specify the spatial or temporal domain of the system being modelled, generally by specifying the range and increments for index variables, and the base line properties of the system. These parameters are often ignored during uncertainty analysis, despite having very considerable potential impact. They essentially control both the precision of the representation and the computational complexity of the model. A balanced trade-off between these two is thus required, and, although they have no “true” value, it is usual to be uncertain about the appropriate value to choose. Although not often done, a parametric analysis can be used to examine how the model domain parameters affect the results of the analysis, and can allow a more informed choice of parameters in subsequent modelling.

Outcome criteria are the variables used to rank or measure the desirability of possible outcomes, e.g. expected utility. These quantities will be deterministic or probabilistic according to whether the input quantities on which they depend are deterministic or probabilistic.

Inextricably linked to the definition of the types of uncertain quantities, is the relevant method for their treatment, which is also shown in Table 3-1. Probability is the most often used measure of uncertain belief, where probability is defined according to the Bayesian or subjectivist view (i.e. the probability

of an event occurring is the degree of belief that a person holds that it will occur, given all the relevant information currently known to that person), as distinct from the classical or frequentist view (i.e. the probability of an event occurring in a particular trial is the frequency with which it occurs in a long sequence of similar trials). It is not meaningful to define probability according to the frequentist view in a modelling context because the primary source of uncertainty is lack of appropriate knowledge of the system of interest. Classical probability estimates require relevant populations of trials of similar events to be known, and so are not able to consider a lack of knowledge (Morgan and Henrion, 1990).

Expressing uncertainty as subjective probability distributions is a simple and commonly used method to capture uncertainty. However, empirical quantities are the only type of quantity whose uncertainty may appropriately be represented in probabilistic terms, as they are the only type of quantity that is both uncertain and can be said to have a true, as opposed to an appropriate or good, value (Morgan and Henrion, 1990). It is generally inappropriate to represent uncertainty about decision variables and model domain parameters by probability distributions, as these have appropriate rather than true values. For these quantities it is more meaningful to conduct parametric sensitivity analyses, i.e. to examine the effect on the output of deterministic changes to the uncertain quantity (Morgan and Henrion, 1990). Value parameters are often treated probabilistically, although Morgan and Henrion (1990) warn against this, as value parameters tend to be among those quantities people are most unsure about. A probabilistic treatment of the uncertainty may hide the impact of this uncertainty, and a parametric treatment of uncertainty is expected to yield more insight into the situation.

3.1.1.b Sources of Uncertainty in Empirical Quantities

Empirical quantities generally constitute the majority of quantities in models, and attempts to categorise different kinds of uncertainty have generally concentrated on empirical uncertainties. Table 3-2 lists the different sources of empirical uncertainty identified by Morgan and Henrion (1990). The reasons for their differentiation follow below.

Statistical variation resulting from random error gives rise to the most easily handled type of uncertainty. No measuring instrument is perfect, and variations in the quantity measured will inevitably be present, with the resultant uncertainty dependant on the size of the variations and the number of measurements taken. This kind of uncertainty has been well studied, and a number of well established statistical techniques have been developed to quantify it, e.g. standard deviation, confidence intervals etc.

Table 3-2 Sources of uncertainty in empirical quantities (based on Morgan and Henrion, 1990).

Source of Uncertainty	Example	Uncertainty Input
Statistical variation	Process monitoring data	Standard or custom fit probability distributions
Subjective judgement	Measuring proxy quantity	Sensitivity analysis, uniform distributions
Linguistic imprecision	"Small" increase	N/A
Variability	Raw water quality	Standard or custom fit probability distributions
Inherent randomness	Rainfall	Standard or custom fit probability distributions
Disagreement	Data from various sources	Sensitivity analysis, custom fit distributions
Approximation	% rain to surface run-off	Uniform or triangular probability distributions

The measurement of any quantity involves not only random error but also systematic error. Systematic error is defined as the difference between the true value of the quantity of interest and the value to which the mean of the measurements converges as the sample size increases. Systematic errors arise from biases in the measuring apparatus and experimental procedure, and because the sources of these biases are unknown (assuming all known sources have been adjusted for), estimating their magnitude necessarily involves *subjective judgement*. Unlike random errors, which can simply be reduced by taking additional measurements, systematic errors are not reduced by increasing sample sizes, and so tend to dominate the overall error. The measured quantity and the quantity of interest are virtually never the same (e.g. toxic effect of large doses in mice to denote the effect of small doses in humans), and the systematic errors this gives rise to are frequently ignored or consistently underestimated.

Uncertainties arise when quantities are specified using imprecise language. *Linguistic imprecision* can often be removed by simply providing a clear specification of all quantities, although it is not always possible to obtain well-specified quantities. Considerable research has been undertaken on translating verbal phrases to numerical values (e.g. "quite likely", "highly probable" etc.), although the considerable variation in interpretation of such phrases, and their high context-dependency, suggest that simple mappings between words and numbers are unlikely to be adequate.

Many quantities exhibit *variability* over time and space, and, if sufficiently well specified, can usually be described by a frequency distribution (as distinct from a probability distribution). The uncertainty around a quantity drawn from a population has two components, the uncertainty from sampling the frequency distribution as well as the uncertainty about the frequency distribution. Distinction between the two sources is necessary if the uncertainty is to be meaningfully analysed, e.g. an analysis of the

uncertainty due to sampling the frequency distribution can point to the desired level of disaggregation, whilst an analysis of the uncertainty about the frequency distribution can point to the need for more precise measurements or increased sample sizes.

Uncertainty is said to be due to *inherent randomness* where the uncertainty is irreducible. A quantity is seen as random if no pattern or model is known to account for its variation. In this definition, randomness is a function of the knowledge available to the assessor, and a quantity may be random to one assessor, but deterministic to another who knows its underlying generating process. In the context of modelling, inherent randomness is perhaps better termed “practical unpredictability”.

Disagreement between informed experts about the value of an uncertain quantity arises frequently and for a number of different reasons, e.g. as a result of different interpretations of scientific data arising from the scientists’ different backgrounds and perspectives. Disagreements are commonly resolved by combining the opinions of the experts using weights, with the weights derived either by simple rating techniques or more complex Bayesian methods.

Approximations arise because a model is only ever a simplified version of a real-world system. In addition to the model approximations introduced by the model domain parameters, which specify the model’s spatial and temporal resolution, approximations are present in the data used to populate the model, e.g. the use of discrete probability distributions to represent continuous distributions. In many instances “best guess” values are used in the absence of measured data.

Empirical uncertainties can be characterised by probability distributions or by sensitivity analyses. What sort of probability distribution is appropriate, or whether the quantity should rather be varied parametrically, is a function of how much is known about the quantity. If sufficient data is available, and the quantity is seen to fit a particular standard probability distribution (normal, beta etc.), the defining parameters can be derived (e.g. mean, standard deviation etc.) using well known statistical techniques. Alternatively, a custom distribution can be defined by specifying points on the probability density curve. If little information other than a range and possibly the mode are known, a uniform or a triangular distribution can be assigned. When very little is known about the uncertainty, a sensitivity analysis can be used to explore the effect of extreme values. The uncertainty inputs corresponding to the different sources of empirical uncertainty given in Table 3-2 are therefore indications of the amount

of information typically associated with the various types of empirical quantities, and should not be seen as hard and fast rules.

3.1.1.c Model Uncertainty

The discussion so far has focussed on the uncertainty about the parameters of a model. Equally important, but more difficult to comprehend, is uncertainty about the form or structure of the model itself. Although there has been relatively little research into situations in which there is uncertainty about what form of model to use, uncertainty about structure is usually more important, and more likely to have a substantial effect on the results of the analysis (Morgan and Henrion, 1990). A simple model may have few input quantities and therefore little parameter uncertainty, but it is likely to have significant model uncertainty. Conversely, a large and complex model may have less uncertainty about the model form, but will have many more input quantities and therefore significant uncertainty about the model parameters.

An assessment of uncertainty about model form is complicated by the rather grey area separating it from uncertainty about parameter values. In some cases it is possible to internalise the problem by incorporating two or more model structures into a single “metamodel”, which contains the models as special cases according to a parameter value, e.g. an exponential dose-response function, which reduces to a linear function when the exponent is one. In such cases the assessment of model form will then depend on what form the internalised parameter takes (e.g. probabilistic if an empirical parameter, parametric if a value parameter etc.). Any model is unavoidably a simplification of reality, so even if a model is a good approximation of a particular real world system, it can never be exact. A competing model may be said to be better, in that it may give better predictions, but it is not meaningful to call it a more probable model. A sensitivity analysis is therefore most appropriate to examine the effect of model uncertainty (Morgan and Henrion, 1990).

3.1.2. Sources of Uncertainty Relevant to LCA Models

Now that an understanding of the general types of quantities input into models has been obtained, these can be applied to the more specific case of LCA models. An understanding of the types of quantities contributing to the uncertainty in LCA results is required to devise strategies to better address the uncertainty, since the source of uncertainty determines the appropriate method for its analysis.

3.2.1.a Sources of Uncertainty Identified in the LCA Literature

The sources and quality of data used in LCA studies have been under consideration since the early attempts at defining the LCA method (Consoli et al., 1993; Fava et al., 1994; Fava et al., 1991; Heijungs et al., 1992; Lindfors et al., 1995). The “Code of Practice” (Consoli et al., 1993) focuses on data sources rather than sources of uncertainty, and distinguishes between primary data and secondary data, where primary data is that taken directly from individual companies, whilst secondary data are obtained from published sources. Consoli et al. (1993) identify a number of factors as reducing the quality of primary data, broadly covering the same sources of empirical uncertainty identified above, except for “allocation problems”, which implicitly recognises model uncertainties. They recognise that LCI data are unlikely to be obtained from every possible facility, and thus represent a sample from the population, which is generally not random, and is rarely truly representative. Some data variation will always be present (from differences in technology, locations etc.), which can often not be characterised by standard statistical measures due to the small sample sizes typically used in LCA studies. Other factors identified include systematic and random errors in measurements; the accuracy and format of records, which require manipulation and thus possible quality degradation during conversion to useful LCI units; and the sharing of common facilities, leading to allocation problems. Secondary data are affected by all of the data quality issues of primary data, as well as issues associated with the averaging and homogenisation of data from different sources (Consoli et al., 1993).

The earliest measure proposed to combat poor quality data was the use of data quality assessment frameworks, which use data quality indicators (DQIs) to identify the possible contributing factors to poor quality data. Although not particularly designed to define the sources of uncertainty, the DQIs are intended to address all possible sources of poor quality results, so amount to much the same thing. A comparative list of some of the DQIs proposed for LCA data is given in Table 3-3. The DQIs have been broadly partitioned into three categories to facilitate this discussion, and generally encompass:

- data variability (in most cases split into geographical, temporal and technological indicators or termed representativeness);
- measurement errors, identified in the statistical DQIs, such as precision, accuracy, bias etc., or summarised in the DQI completeness; and,
- subjective judgements, embodied in such DQIs as reliability of the source, nature of the data, applicability, suitability, et cetera.

Table 3-3 Data quality indicators/descriptors identified for use in LCA.

	SETAC	USEPA	ISO 14041	Weidema
1	<i>Completeness^a</i> <i>Precision</i> Representativeness Stability <i>Homogeneity</i> <i>Data distribution</i> Applicability / Suitability / Compatibility <i>Bias</i> <i>Accuracy</i> Identification of Anomalies	Precision Representative Data collection method Bias	<i>Completeness</i> <i>Precision</i> Representativeness Time-related coverage Geographical coverage Technology coverage Nature of the data	Completeness Temporal correlation Geographical correlation Technological correlation Reliability of source
2	Consistency Derived models <i>Uncertainty</i>		Consistency	
3	Comparability Accessibility Transparency Peer review	Comparability Referenced Acceptability	Reproducibility	

1. DQIs applicable to empirical quantities

2. DQIs to assess model uncertainties

3. DQIs to assess practical aspects of the study as a whole.

a. Quantitative DQIs are given in italics. All DQIs are treated semi-quantitatively in the pedigree matrix approach (Weidema, 1998b) qualitatively in the USEPA worksheet approach (USEPA, 1995a).

Whilst the DQI methods predominantly address data quality and not model uncertainties, the ISO 14041 and SETAC list of DQIs include consistency, which refers more to model uncertainty, since it relates to how uniformly the methodology is applied across the various components of the study across comparative studies. The additional DQIs listed in category 3 in Table 3-3 relate more to reporting structures of the study, and are thus not relevant to this discussion.

In a distinction similar to that made by Morgan and Henrion (1990) between empirical variables and decision variables, Weidema and Wesnæs (1996) distinguish between environmental data, relating to data on the investigated process, system data, relating to the flow of materials, energy and products through the investigated processes, and performance data, related to the definition of the functional unit. The factors determining data quality are given by Weidema and Wesnæs (1996) as:

- its uncertainty (spread and pattern of distribution),
- its reliability (measurement error, calculations, assumptions and quality control),
- its completeness (sample size and time period, sample representativeness),
- its age,
- the geographical area for which the data is representative, and
- the process technology or technological level for which the data is representative.

A distinction is made between ‘basic’ uncertainty, which is a property of all sampled data as it relates to measurement errors and inherent fluctuations, and ‘additional’ uncertainty, which is related to the data not being of the optimal quality (Weidema and Wesnæs, 1996). The ‘additional’ uncertainty can be broadly understood as how well the particular data available fits the specific application, and encompasses such attributes of the data as variability, subjective judgements, and approximations.

In a similar distinction to Weidema and Wesnæs’ (1996) distinction between “basic” and “additional” uncertainty, Steen (1997) identifies “at least two types of uncertainty” causing uncertainty in LCA studies. The first he identifies as the “normal uncertainty associated with the determination of a parameter in a given system”, and the second as “associated with the choice of such a parameter value to represent a value in another similar system”. Steen (1997) also implicitly identifies model uncertainties by recognising the introduction of further uncertainty by using linear relations instead of non-linear ones (Steen, 1997). Similarly, Coulon et al. (1997) recognise the distinction between data variability, data quality aspects related to the source and application of the data, and model uncertainty. They identify obstacles to the reliability and comparability of LCA results as “the variability between the measurements within industrial plants, the discrepancy between bibliographical and actual site data, and the sensitivity of the results to core methodological choices” (Coulon et al., 1997).

Huijbregts (1998a) has looked more systematically at sources of uncertainty in LCA models, and published a general framework classifying the uncertainties present, shown in Table 3-4. Huijbregts’ classification is basically a regrouping of the more general classification of uncertainty outlined in section 3.1.1, which is to be expected, since LCA is merely a subset of general policy models, and would be expected to share their characteristics. Huijbregts places an increased emphasis on variability, which he breaks down into spatial variability, temporal variability and variability between objects and sources. His classification of variability differs somewhat from that defined in section 3.1.1, which only incorporates variability in empirical quantities, whilst that defined by Huijbregts includes both empirical variability, and, in the case of spatial and temporal variability, variability causing model uncertainty. In addition, Huijbregts (1998a) divides non data-related uncertainties between model uncertainty and uncertainty due to choices. The uncertainty due to choices includes well-known LCA choices, e.g. allocation method, weighting method, as well as the choice of certain model parameters, e.g. functional unit.

Table 3-4 Types of uncertainty identified by Huijbregts (1998a), and their relationship to the various phases of LCA.

Phase Source	Goal and Scope	Inventory	Choice of Impact Categories	Classification	Characterisation	Weighting
Parameter Uncertainty		Inaccurate emission measurements			Uncertainty in life time of substances	Inaccurate normalisation data
Model Uncertainty		Linear instead of non-linear modelling	Impact categories are not known	Contribution to impact category is not known	Characterisation factors are not known	Weighting criteria are not operational
Uncertainty Due to Choices	Functional unit	Use of several allocation methods	Leaving out known impact categories		Using several characterisation methods within one category	Using several weighting methods
Temporal Variability		Difference in yearly emission inventories			Change of temperature over time	Change of social preferences over time
Spatial Variability		Regional differences in emission inventories			Regional differences in environmental sensitivity	Regional differences in distance to (political) targets
Variability Between Objects and Sources		Differences in emissions between factories which produce the same product			Differences in human characteristics	Differences in individual preferences, when using panel method

Meier (1997) also undertakes a detailed assessment of the different types of uncertainty at all levels of LCA studies. His characterisation is given in Table 3-5, in which he broadly distinguishes between “model”, “data” and “valuation” uncertainties. “Systematic errors” give rise to “model” uncertainties (e.g. errors arising from inadequate definition of the system), which are either not quantifiable or require estimation by default estimations. “Data” uncertainties arise from variability, as well as systematic and stochastic errors, where stochastic errors are defined as uncertainties resulting from errors in measurements or determination of the data value, and variability as those uncertainties linked to averaging data. “Valuation” uncertainties arise from stochastic uncertainty. Meier’s distinction between “data” and “valuation” uncertainties appears to relate primarily to the placement of the uncertain quantity, i.e. whether it falls into the inventory or impact assessment phase of LCA. Another distinction appears to be that “valuation” uncertainties arise purely from data gaps and measurement errors, whilst “data” uncertainties arise predominantly from variability and data averaging, in addition to measurement errors. Furthermore, the name “valuation” uncertainties appears to suggest uncertainties due to value judgements in selecting the applicable data for impact assessment.

Table 3-5 Identification and characterisation of uncertainties in LCIA methods according to Meier (1997).

Origin of Uncertainty	Uncertainty Due To	Uncertainty Related To
<i>Goal and Scope Definition</i>		
General assumption (e.g. lifetime of system)	Systematic errors	Model
Non-consideration of all processes in the inventory	Systematic errors	Model
Uncertainty due to inadequate definition of the functional unit	Systematic errors	Model
<i>Foreground Inventory Data</i>		
Uncertainty due to emission allocation	Systematic errors	Model
Uncertainty in data measuring	Stochastic uncertainty	Data
Uncertainty due to age of data / temporal variation	Stochastic uncertainty	Data
<i>Background Inventory Data</i>		
Uncertainty due to data module choice	Variability	Data
Uncertainty in determination / calculation of emissions	Stochastic uncertainty	Data
Uncertainty due to averaging data (spatial, technology, time)	Variability	Data
Non-consideration of emitted substances in process modules	Systematic errors	Data
<i>Impact Assessment</i>		
Uncertainty due to wrong valuation model / assumptions	Systematic errors	Model
Uncertainty due to non-consideration of impact categories	Systematic errors	Model
Uncertainty in determination of classification factors	Stochastic uncertainty	Valuation
Uncertainty in determination of the actual flow (normalisation factor)	Stochastic uncertainty	Valuation
Uncertainty in determination of the reduction factor	Stochastic uncertainty	Valuation

3.2.1.b A Classification of Sources of Uncertainties in LCA Related to their Assessment

A classification of uncertainties based on the appropriate method of analysing the uncertainty was found to be valuable for structuring the uncertainty assessment (see section 4.1). The classification is shown in Figure 3-1, and broadly distinguishes between empirical parameters, for which a probabilistic assessment is appropriate, model parameters, for which a parametric assessment is appropriate, and uncertainties in model form, which can be assessed using sensitivity analyses. Some general features of the classification are presented, after which the sources of uncertainty are discussed with reference to their occurrence in LCA models.

The classification is primarily distilled from that of Morgan and Henrion (1990), but also draws on those from the LCA literature. The emphasis on variability by both Huijbregts (1998a) and Weidema and Wesnæs (1996) is incorporated into the classification, by breaking down empirical variability into geographical, temporal and technological variability after Weidema and Wesnæs (1996). Huijbregts' (1998a) definition of variability is more broad, and includes both empirical uncertainty arising from sampling variable populations, and model uncertainties arising from LCA's inability to incorporate spatial and temporal variability into its model form. As these two sources of variability require very different measures for their assessment, they are split between empirical parameter variability and uncertainty in model form in Figure 3-1.

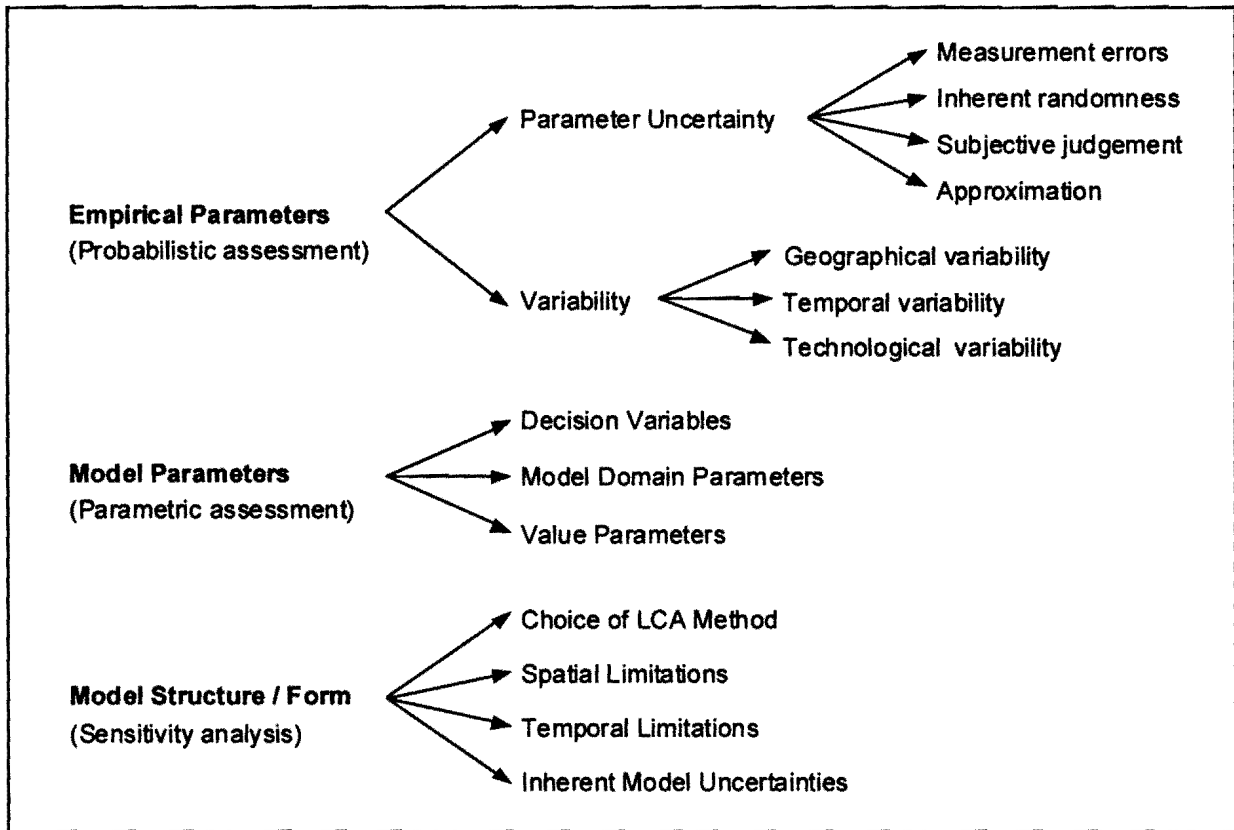


Figure 3-1 Classification of uncertainties present in LCA models (parameter types defined in section 3.1.1).

Huijbregts (1998a) differentiates between “model uncertainty” and “uncertainty due to choices”. The distinction between these two is not completely clear, in that many of the “choices” in LCA have a profound effect on model uncertainty, e.g. choice of characterisation method, whilst model uncertainty may also involve a choice, e.g. whether to use linear or non-linear modelling. The “uncertainty due to choices”, as defined by Huijbregts, appears to include well-known LCA choices, e.g. allocation method, weighting method etc. The category appears to be a valid one for LCA, where many such choices exist. It is therefore incorporated in Figure 3-1, but is slightly refined in that it is defined as “choice of LCA method”, and falls under uncertainties in model form. This is to keep it distinct from choices of model parameters. Whilst both involve choices, the choice of method is best suited to a sensitivity analysis, where two or three discrete options are considered, whilst model parameters are best suited to a parametric analysis, where the parameter range can be assessed in combination with the other model parameters. The distinction between choices of model form and choices of model parameters is not always clear. The choice of model form can perhaps be best seen as introducing a set of model parameters, which subsequently require definition.

Meier (1997) distinguishes between model uncertainties that can not be quantified and those that can be estimated by a scenario analysis. This is roughly analogous to Huijbregts' distinction between model uncertainties and uncertainty due to choices. A similar distinction is made in Figure 3-1, with the breakdown of uncertainty in model form into "choice of LCA method" and "inherent model uncertainties". Spatial and temporal limitations are kept separate from inherent model uncertainties to emphasise their importance in LCA models.

No distinction is made in Figure 3-1 between the sources of uncertainty and where in the LCA process they occur, i.e. the same basic sources of uncertainty are equally applicable to an analysis of uncertainty in inventory or impact assessment models. However, the relative importance of the various sources, and thus the emphasis of the uncertainty assessment, depends very much on the different phases of the LCA process. For example, the complex models of impact assessment are associated with high model uncertainties, so these are likely to dominate the uncertainty analysis, although empirical parameter uncertainty is still significant. Inventory models, on the other hand, are fairly simple, with consequently lower uncertainty in model form, and are likely to be dominated by uncertainty in empirical or model parameters. Figure 3-2 attempts to portray this, with the sources of uncertainty identified in Figure 3-1 ordered according to their likely importance in that particular phase of LCA. The emphasis of the uncertainty analysis is also likely to vary according to the decision type, since different levels of information are available for different decision contexts. The relative importance of the sources of uncertainty is therefore given according to the three broad decision types defined in section 2.3. The order is given for illustrative purposes only, and although based on typical observations, cannot be taken as a hard and fast rule. Goal and scope definition is not included in Figure 3-2 because the representation is of where the uncertainty manifests itself, i.e. although the choices governing the degree of uncertainty in the LCA results are made in the scope definition phase of LCA, the effect of these choices are felt in one of the three phases depicted.

Empirical parameters

These form the bulk of data inputs into both inventory and impact assessment models. They can be broken down broadly into those parameters that are uncertain due to variability and the incorrect application of the data, and those that are uncertain due to the manner in which the data is obtained and our imperfect knowledge of the system.

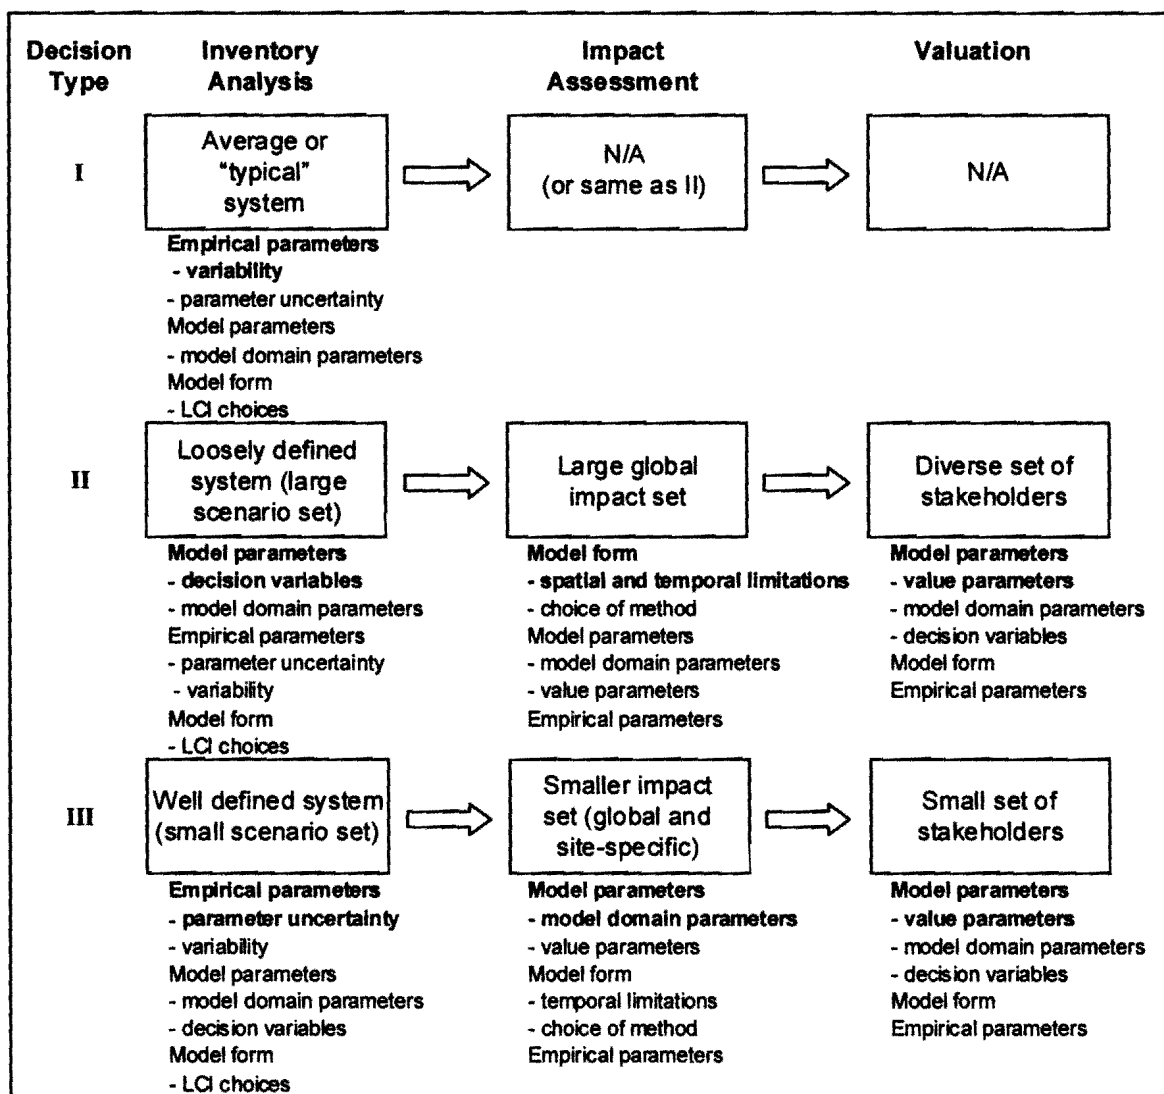


Figure 3-2 Relative importance of the sources of uncertainty according to the phase of LCA and the decision type (as defined in section 2.3). The source of uncertainty likely to dominate the analysis is given in bold print.

Many quantities exhibit *variability* in time and space, which in most cases can be quantified by fairly precise frequency distributions. Process emissions and material consumption patterns vary over the life of the plant, and usually also with variations in load and external conditions, all of which give rise to *temporal variability*. Similarly, variations occur across different geographical regions and countries. *Geographical variability* is especially significant in characterisation models, as different receiving environments have different sensitivities and background concentrations. *Technological variability* refers to the fact that two seemingly similar systems will still exhibit variability. Processes producing the same product but using different technologies are likely to exhibit large variability, whilst those operating under the same technology are unlikely to employ identical operating conditions, and will

still exhibit some, albeit less, variability. It is important that technological variability be kept distinct from spatial and temporal variability, so that dual counting does not occur in the subsequent quantification of uncertainty, i.e. for the variance attributed to the various sources of uncertainty to be additive, the sources must be independent (Weidema and Wesnæs, 1996).

Measurement errors are a familiar and easily handled source of uncertainty, provided a sufficient data set is available for their characterisation. However, this is seldom the case in LCA studies, and more often there are too few values to statistically quantify the uncertainty. In addition, often only measurements of a proxy quantity are available, e.g. applying the data from an existing process to a future process. The uncertainty associated with extrapolating the measured quantity to the actual item of interest is often significant, but is difficult to quantify because of the largely *subjective judgement* involved. In many LCA studies, measured data are lacking altogether, and recourse has to be made to highly uncertain *approximations* or best guesses. As the parameter moves from highly qualified (large data sample of relevant quantity) to an approximation, the representation of uncertainty moves from a well defined standard or custom fit probability distribution, through to a rough uniform distribution, where only an equally probable range of values can be proposed. Certain data are termed inherently random when no underlying pattern for the variation in values can be determined, or it is beyond the scope of the model to determine the causal relationships. *Inherent randomness* or variability is assumed to incorporate all variability that can not be reduced through a better understanding or definition of the system, i.e. that not due to the spatial, temporal and technological placement of the process. Inherently variable data can be precisely represented by a standard or custom fit probability distribution, or roughly approximated by a uniform distribution, depending on the number of data samples available to characterise the frequency distribution.

Model Parameters

These quantities define the system model, so their choice is informed by the particular goals and scope of the study. The *model domain parameters* define the resolution of the system, and so determine the level of accuracy, detail and cost of the model. Model domain parameters include the familiar choices made during the scoping phase of LCA, e.g. the system boundary, the functional unit, the number of impact categories considered, the level of spatial differentiation etc. *Decision variables* are those variables under the control of the decision makers, and define the operating state of the system, e.g. the load at which the plant is operated. *Value parameters* represent the values and preferences held by the decision makers. These define the weights assigned to the impact categories in a simple valuation step,

or determine the relative importance of the criteria in a more complex multi-criteria decision making procedure. Value parameters are also implicit in certain characterisation methods, e.g. setting the toxic dose (e.g. NOEC, LD₅₀ etc.) in calculating toxicity factors. It is not always possible to make a clear distinction between decision variables and value parameters, as the choice of many decision variables implicitly depend on the values of the decision makers.

Model Form

Uncertainties in *model structure or form* arise as all models are a simplification of reality. The higher the degree of simplification, the higher the associated model uncertainties. Various degrees of sophistication have been developed for certain aspects of the LCA method (e.g. allocation method, characterisation models etc.), and which of these is applied impacts considerably on the uncertainty of the study. For example, whether a simple mass-based allocation procedure is followed or whether a complex marginal analysis is conducted, can potentially have a large effect on model uncertainty, and possibly the overall outcome of the study. Characterisation models are associated with large model uncertainties, as the very complex environmental systems they represent mostly have to be severely simplified to make them operational. More complex models are being developed (e.g. to include the fate and transport of the pollutants and the sensitivity of the receiving environment), although the associated parameter uncertainties of these models are vastly increased, and to some degree counter-act the reduced model uncertainties. Inherent model uncertainties are always present, in that it is always possible to further increase the complexity of the model. Inherent model uncertainties therefore represent such improvements in model form as using non-linear relationships instead of linear relationships, including additional parameters (increasing the resolution) or including feedback loops.

Spatial and temporal limitations are intrinsic to the LCA method, and the feasibility of dealing with them within the current LCA framework is limited (see section 5.3.1). Spatial limitations arise because the LCA method allows emissions from various locations to be aggregated into a single inventory, although the derivation of accurate damage factors requires that specific regional information be incorporated. Similarly, temporal limitations arise because emissions occurring over various time-scales are aggregated in the inventory, and thus no longer represent real time emissions matched with real time variation in a characterisation model. LCA's aggregation of inventory data over spatial and temporal dimensions is one of its key features, i.e. its ability to characterise global systems. However, this sits fundamentally in opposition to the determination of real site-specific impacts. The traditional

global approach of an aggregated inventory is thus associated with an unavoidable degree of model uncertainty.

3.2. METHODS FOR THE ASSESSMENT OF UNCERTAINTY

The preceding discussion identifies the various sources of uncertainty in models. Fundamental to an uncertainty analysis is to determine the effect of the uncertainty in input quantities on the predicted outcome of the model. Methods to propagate uncertainties through the model are thus of key interest. The ultimate aim of an uncertainty analysis is surely to reduce the uncertainty in the output, i.e. to reduce the uncertainty in those quantities that are found to be hampering decisions able to be taken based on the model output. In most cases, merely obtaining the range of uncertainty in the output is thus not sufficient, and methods for comparing the relative importance of the uncertain input parameters, and strategies for reducing the uncertainties once the analysis has highlighted where in the model effort should be focussed, are therefore essential. A three stage process to addressing uncertainty is therefore required:

1. Identifying and characterising uncertainties
2. Uncertainty propagation and analysis
3. Reducing the highlighted uncertainties

A full uncertainty analysis is therefore a lengthy process that dramatically increases the time and effort spent on a study. The returns are usually well worth the effort, i.e. obtaining meaningful results that can be applied with a fair degree of confidence. However, a full quantitative uncertainty analysis is not feasible in all studies, and in some cases qualitative measures are more applicable. Although better than including no representation of uncertainty, qualitative methods are not able to give any information on the magnitude of the effect the uncertain quantity has on the model outcome, and consequently the degree to which the uncertainty needs to be addressed. This discussion covers only quantitative statistical measures to address uncertainties, although qualitative and semi-quantitative methods developed for use in LCA are discussed in section 3.3.1.

The preceding section dealt with the first stage of an uncertainty analysis, that of identifying and characterising the uncertainty, whilst the following section deals with the subsequent propagation and analysis of the uncertain parameters. As discussed above, only empirical parameter uncertainty can be meaningfully assessed by a probabilistic assessment, so certain of the methods discussed below are

only applicable to empirical parameters. The methods all assume that the uncertainty in the parameters is able to be described, i.e. that the range and/or probability distribution covered by the parameter is known. This is not a trivial component of the study, and is perhaps the most difficult part of the assessment. The characterisation of uncertain parameters in LCA models is postponed to section 4.2.2.

3.2.1. Propagation and Analysis of Uncertainty

A variety of methods to propagate and analyse uncertain inputs has been developed, differing in their approach, complexity and accuracy of results. An overview of these approaches is given in appendix A.1. The methods are discussed fairly briefly, as all are well established statistical methods and details can be found in the extensive statistical literature. Methods that have been employed in LCA studies are reviewed in section 3.3.2.

The analysis of uncertainty involves measuring the degree to which each input contributes to the uncertainty in the output. Methods to quantify this range from analytical methods (simple sensitivity analyses to more complex Gaussian analyses), to semi-graphical parametric methods employing scenario- or probability- trees. Criteria to be used in selecting a method include the nature of the model, the requirements of the analysis, the resources available and the relative importance of uncertainty about model form versus parameter uncertainty, i.e. if there are large uncertainties about the model structure, an extensive evaluation of parameter uncertainty may be pointless. If the model is simple and has small uncertainties, Gaussian approximation is typically best suited to the analysis. However, this only provides the mean and variance of the output distribution, so is not suitable when the tails of the distribution are of interest. Monte Carlo and other sampling methods have the distinct advantage that they are easy to apply to existing models (i.e. as an “add-on” analysis), and they avoid the analytic complexity of obtaining partial derivatives. Moreover, the development of more efficient sampling strategies has decreased the number of sample runs required to the extent that the computation advantages of Gaussian approximation are often not significant. These advantages result in sampling methods being the technique of choice for all but the simplest near-linear models with small uncertainties.

The sampling methods also show distinct advantages over the probability tree approaches. In the sampling methods, the computational effort is essentially linear in the number of uncertain quantities, whilst it is exponential for the tree schemes. Therefore, tree schemes are only appropriate for models

with a small number of uncertain inputs. Sampling methods are also able to employ more powerful measures of uncertainty importance than tree schemes, as well as being able to estimate the imprecision in the propagated distribution using standard statistical techniques. Both continuous and discrete distributions are able to be used in the sampling methods, whilst tree-based methods require that continuous distributions be discretised, which introduces an associated approximation error. It is generally easier to express probabilistic dependencies among variables as discrete conditional probability distributions, so tree-based methods are favoured in systems with complex probabilistic dependencies. Similarly, it is generally easier to apply Bayes' rule to discrete variables, thus favouring tree schemes. However, the clear advantages of the sampling methods for models with a large number of uncertain inputs, as well as their ability to be relatively simply applied to existing models, suggests they are the method of choice for the assessment of empirical uncertainties in LCA models (details on the sampling methods are given in appendix A.1.1).

The sampling methods enable the uncertainty in the output data to be calculated, and thus enhance the interpretation of the results (e.g. it is possible to determine whether observed differences between two options are statistically significant). However, to obtain the full benefit of an uncertainty analysis, those data elements that significantly contribute to the variance of the results must be identified. This allows the model developers to know where to focus their efforts to reduce data uncertainty, and where possible, to reduce the variance of the model results so that it is possible to discriminate between alternatives with greater confidence. Three measures of uncertainty importance are presented in appendix A.1.2. A major limitation of using sensitivity analyses is that any correlation effects between the uncertainties are ignored, i.e. combinations of uncertain inputs that could possibly be significant are missed by a simple sensitivity analysis. The accuracy obtained using Gaussian approximation to identify the key elements is slightly compromised by it being a local method (i.e. not suited to large uncertainties) and the requirement of linearity (i.e. the use of absolute error). Although, rank order correlations are also inaccurate where the relationships are highly non-linear, they are deemed equally good at finding key data elements as the analytical method, and are favoured because they are compatible with the easy to implement simulation-type uncertainty analyses.

3.2.2. Reduction of Uncertainty

Implicit in any uncertainty analysis is that once the key uncertainties have been highlighted by the analysis, these must somehow be managed or reduced. In this section possible approaches to reducing uncertainties are discussed, according to the same broad categorisation of uncertainties as before, and then with particular reference to LCA models.

3.2.2.a Uncertainty Reduction According to Source

The way in which the uncertainty can be reduced is generally dependent on the source of the uncertainty. Possible approaches to addressing the uncertainty are therefore discussed with reference to the sources identified in Figure 3-1.

Empirical Parameters

Uncertainty reduction measures tend to focus on empirical parameters, partly because empirical parameters typically receive the most attention in an uncertainty analysis (they constitute the majority of quantities in models), and partly because they are easier to effect, i.e. the uncertainty can be reduced within the current model form (Morgan and Henrion, 1990). Broadly speaking, uncertainties in model inputs arise because simplifications are made and shortcuts are taken to expedite the modelling process. All input parameters therefore benefit from a more detailed and considered analysis of the quantity or quantities from which the parameters are derived, and the necessity of additional modelling effort is common to reducing the uncertainty of all empirical parameters. Hence the importance of methodologies to identify and prioritise the uncertain input quantities.

Measurement errors are usually the easiest of all to reduce, as all that is typically required is that more measurements be taken. Additional measurements allow the error due to the measuring process to be better defined using well known statistical techniques, and therefore a better determination of the true value to be made. If this still does not allow for a sufficiently precise determination of the parameter, recourse would have to be made to a more accurate measuring procedure, should one exist.

Inherent randomness also benefits from taking additional measurements, as this enables a more accurate characterisation of the random quantity. If the parameter is truly random, this is the most that can be done to reduce the uncertainty. However, in many cases, a more detailed examination of the

quantity reveals that there is some underlying cause to the variability. If the quantity can be seen to vary spatially or temporally, then the uncertainty can be reduced by the measures discussed for variable parameters below. For “pseudo-random” quantities, the apparently random variation can be removed by modelling the quantity in greater detail. However, in most cases, the quantities are initially characterised as random because their causal mechanisms are extremely complex (e.g. weather patterns), so modelling these processes in detail is likely to be beyond the scope of the study.

The uncertainty due to *subjective judgement* is less straightforward to reduce. A better understanding of the mechanisms involved may allow for the estimate of the error to be reduced. However, it could equally lead to a realisation that the error is in fact underestimated, in which case it may be necessary to refine the measurement process to better relate the measured quantity to the quantity of interest, or to find a more appropriate quantity on which to base the measurement.

Uncertainty due to *approximations* can be removed by replacing the approximation with the underlying mechanisms and measured data. However, approximations are present in the first place because the mechanisms are too complex to model, or measured data are not available. Uncertainty due to approximations is thus not able to be removed without increasing the complexity of the model.

Two sources of uncertainty arise in *variable quantities*. The uncertainty about the frequency distribution itself (i.e. the uncertainty in the sampled values of the variable quantity), is a result of measurement errors and/or subjective judgement, and can be reduced by taking more or better measurements, as outlined above. The other component of uncertainty is due to the sampling of the frequency distribution. This can be reduced by a better definition of the temporal, spatial and technological placing of the quantity of interest. “Zoning in” on the particular region, time-span or technology of interest results in a disaggregation of the frequency distribution, leaving a narrower span of variation over the actual zone of interest, and consequently less sampling uncertainty.

Model Parameters

Reducing the uncertainty resulting from the choice of model domain parameters is much more difficult to achieve than reducing that resulting from empirical parameters. Usually it requires a significant redefinition of the model structure, with a host of knock-on effects, e.g. additional model inputs, and, at worst, complete model redevelopment. Because of this, uncertainty due to model domain parameters is often not addressed, or addressed only qualitatively. Reducing model parameter uncertainty is also

hampered by the fact that there is no “true” value for these quantities. A carefully considered choice can ensure the best value is selected for the given circumstances, but as this inherently involves some degree of subjectivity, it will always contain some degree of uncertainty. Nevertheless, an uncertainty analysis will certainly facilitate the choice of the best value, as well as providing information as to the significance of the uncertainty. In many cases, particularly with decision variables, no clear “best” value exists, in which case overall model uncertainty is reduced by presenting the results for various scenarios, thereby ensuring the possible spread of results is not hidden.

Model Form

Perhaps most difficult of all is reducing the uncertainty due to the form of the model. Significant improvements can usually only be achieved by a completely different modelling approach, which essentially requires building a new model. However, some, albeit less, reduction in uncertainty may be achieved by refining the existing model. For example, modelling selected components in greater detail, or increasing the complexity of certain key relationships, e.g. replacing linear approximations with more representative non-linear relationships.

3.2.2.b Reduction of Uncertainty in LCA Models

Huijbregts (1998a) identifies certain “tools” which are able to address uncertainty and variability in LCA. The list of tools presented in Huijbregts (1998a) does not distinguish between those to characterise the uncertainties and those to reduce them, so it is reordered and grouped in Table 3-6 so as to be comparable with the preceding discussion (i.e. broken up into analysis, prioritisation and reduction of uncertainties). In agreement with Morgan and Henrion (1990), Huijbregts suggests the use of probabilistic simulation only with empirical parameters (called parameter uncertainty and variability in objects/sources in his framework). He excludes spatial and temporal variability from such a probabilistic assessment, suggesting multi-media modelling and scenario modelling respectively for their assessment. This is true for the model uncertainty associated with spatial and temporal variability, but does not address the parameter uncertainty due to variability in time and space. By breaking down spatial and temporal variability into two components, i.e. into model and parameter uncertainty, as is done in Figure 3-1, they can also be analysed in a probabilistic simulation along with technological variability (called variability in objects and sources by Huijbregts).

Table 3-6 Modified layout of tools suggested by Huijbregts (1998a) to address uncertainty and variability in LCAs.

Tools	Types	Parameter Uncertainty	Model Uncertainty	Uncertainty Due to Choices	Spatial Variability	Temporal Variability	Variability in Objects / Sources
1	Probabilistic simulation	x					x
	Scenario modelling			x		x	
2	Correlation and regression analysis	x					x
	Additional measurements	x					x
3	Standardisation			x			
	Expert judgement / peer review	x		x			x
	Non-linear modelling		x				
	Multi-media modelling		x		x		

1. Tools to analyse uncertainty

2. Tools to prioritise between uncertain inputs

3. Tools to reduce uncertainty

Even more importantly, uncertainty arising from variability can be reduced by disaggregation of the data set. This is consistent with the marginal technology approach to LCA, which is the most methodologically defensible approach to inventory analysis (Frischknecht, 1997; Wenzel, 1999). The marginal technology approach avoids the use of average data by only using data on the technology whose production volume is affected directly by the studied change (Weidema et al., 1999). It therefore leads to less data uncertainty, since average data is inherently uncertain (it combines data from a large number of different technologies, often from different sources).

Marginal data always relate to a well-defined technology, and their use is therefore analogous to “zoning in” on the actual technology of interest, and its actual placement in time and space. The classification of the sources of uncertainty in LCA in Figure 3-1, specifically groups variability as a sub-class of empirical uncertainty, as uncertainty arising from variability can uniquely be reduced by the use of marginal production technologies (i.e. by moving away from highly variable average systems to the actual system of interest). Where the system has been divided into foreground and background sub-systems, the use of the marginal production technology is somewhat analogous to placing the process in the foreground system (Weidema et al., 1999). Background processes are, by definition, treated at a high level of data aggregation, and thus incorporate high variability. The only way in which high variability in background processes can be reduced is thus to move the process from the background into the foreground system, where its marginal production technology can be evaluated.

It is apparent from the reordering of the “tools” suggested by Huijbregts (1998a) to analyse and reduce uncertainties in LCA (see Table 3-6), that not all the sources of uncertainty are assigned tools to both analyse and reduce the uncertainties. This is important, as a reduction in uncertainty is largely meaningless if no mechanism exists to analyse the benefits and significance of the increased modelling effort. The strength of the categorisation of the relevant sources of uncertainty in Figure 3-1 is that the applicable uncertainty analysis method is clearly highlighted, in that it is the primary classifying factor, i.e. the sources of uncertainty are divided according to their method of analysis (probabilistic, parametric or sensitivity analysis). A broad consensus does exist between Huijbregts classification and that in Figure 3-1, although different terminology obscures this somewhat. The term “scenario analysis” is used in this thesis to denote a particular operating state of a system, i.e. a particular set of decision variables (see section 2.1.1.a), whilst “scenario modelling”, as used by Huijbregts (1998a), broadly covers both the terms “sensitivity analysis” and “parametric analysis”, as applied in the preceding discussion.

The choice of model parameters is complicated by the fact that they only have an appropriate and not a true value. Huijbregts (1998a) suggests standardisation and expert judgement/peer review as tools to address uncertainty due to choices in LCA, and the latter to address parameter uncertainty and technological variability as well. These two tools are concerned primarily with reducing the uncertainty across studies, thereby improving the reproducibility of LCA results, and allowing better comparisons (i.e. two studies conducted using the same allocation and impact assessment method will obviously be more comparable than if different methods are used). Meier (1997) also suggests the use of rules and guidelines to minimise what he terms “qualitative uncertainties”, i.e. those uncertainties not able to be characterised by probability distribution functions or in a sensitivity analysis.

The ISO standards set the broad outline and method requirements of an LCA study (ISO, 1997; 1998; 2000a-b). Although representing a notable step towards standardisation, these supply only guidelines, and choices still have to be made when selecting the exact method to follow. This is especially pertinent for the impact assessment and improvement assessment phases of LCA, although choices influencing inventory analysis also have to be made (e.g. allocation method). Addressing model uncertainty in LCA is therefore somewhat unusual, in that for certain aspects of the method, a set of pre-defined sub-models exist from which to choose. It thus lends itself to an investigation of model uncertainties via sensitivity analyses. Impact assessment methods range from the extremely simple (e.g. aggregation according to release medium) to the extremely complex (e.g. damage models including

fate, effect and background level information). Consequently, very different levels of uncertainty are implicit in the various methods, and the method implemented can have significant consequences on the outcome of the study. Because of the level of effort involved, relatively few studies comparing the different methods have been undertaken, although those that have done so found large discrepancies in the results using different methods (Baumann and Rydberg, 1994; Huppel et al., 1999; Tukker, 1998).

Certain methods are more applicable to certain decision types, since the applicability of the method depends largely on the level of information available and the requirements of the results (see discussion in section 2.3). Thus it may be possible to decide before hand whether a more complex method is warranted by the particular decision to be supported, thereby reducing the options requiring assessment in a sensitivity analysis. In addition, the particular requirements and limitations of the different decision contexts suggest that a standardisation across all types of LCA models to ensure uniformity would be more damaging than useful. Different sets of standards therefore have to be developed, tailored to the particular types of decisions to be supported. Also potentially able to be standardised according to the decision in hand is the acceptable level of uncertainty in the results, and thus the degree of uncertainty reduction required.

3.2.3. Reporting and Visualising Uncertain Results

An uncertainty analysis is pointless if the insights the analysis yields are not communicated clearly to those who will be using the results. The overall degree of uncertainty about the conclusions, and an understanding of those sources of uncertainty most critical to the conclusions, needs to be clearly presented. Although common statistical measures can be used to define the uncertainty (e.g. variances, confidence intervals etc.), graphical methods are typically the most efficient in effectively communicating the necessary insights (Morgan and Henrion, 1990). Presenting and communicating the large amounts of data resulting from an LCA is essential to its successful use as a decision support tool. This is made even more so by the inclusion of an uncertainty assessment, as this dramatically increases the amount of data that needs to be communicated.

The three basic ways of presenting probabilistic output are:

- the probability density function (PDF),
- the integral of the PDF, the cumulative density function (CDF), and
- displaying selected fractiles, as in box and whisker plots.

The PDFs and CDFs of some common probability distributions are shown in appendix A.2, and Figure 6-3 is an example of a box and whisker representation of uncertain data. Each emphasise different aspects of the probability distribution. Box plots emphasise confidence intervals and means, whilst PDFs clearly show the relative probabilities of different values, and the shapes of the distributions. The ability of the PDF to indicate small variations in the probability density may in fact be a disadvantage for stochastic modelling, where the small variations arise simply from random sampling noise, and the corresponding CDF will look a lot less noisy with the equivalent sample size. In addition, like the box and whisker plot, the CDF is able to display the probability that the quantity lies in specific intervals, and is thus also superior to the PDF if specified fractiles or confidence intervals are of interest. The simpler box and whisker plots are a useful way to represent the results when the range and not the distribution shape is primarily of interest. They are also an easy way to visualise where overlaps between two systems are occurring, although they do not quantify this overlap (as is possible with CDFs). The CDF is best for determining stochastic dominance between two distributions, as this can be clearly seen by whether or not their CDFs intersect (Morgan and Henrion, 1990).

Comparative studies often share a number of sub-processes for which identical data may have been used. Thus significant correlation between the two options may exist, which could cause the differences between the two options to be overemphasised. Coulon et al. (1997) suggest that to eliminate such a correlation, the probability distribution of their normalised difference should be studied, i.e. for the comparison of two options, A and B, the following ratio can be used to indicate how option A compares to option B:

$$\frac{A - B}{A}$$

Similarly, Meier (1997) finds that the comparison of absolute values of performance indicators does not consider the correlation of uncertainties among the options, and recommends that the significance of the relative difference in the performance indicators be used to distinguish between the options (Meier, 1997). For similar reasons, Huijbregts (1998b) recommends plotting a ratio of the two options, which he terms the “comparison indicator”, i.e.:

$$\frac{B}{A}$$

and Steen (1997) recommends plotting the difference between the options.

The normalised (or un-normalised) difference or ratio is used to indicate the significance of the difference between the two product systems, in that the y-intercept on a plot of the CDF of the difference or ratio shows at what probability the one system is always preferred to the other, i.e. when the quotient is greater than 1 (or the normalised difference is negative), option A contributes less to the category than option B, and the reverse is true when the quotient is greater than 1 (or the difference is positive). Huijbregts (1998b) defines a certain result as “significant” when 95% of the simulation iterations lie above or below 1, although the acceptable degree of confidence in the results will vary according to the “burden of proof” required of the study. The effective normalisation between comparative options carried out by Coulon et al. (1997), Meier (1997) and Huijbregts (1998b) is not the same as the normalisation that typically takes place as part of impact assessment/valuation, i.e. a ratio against background levels for that category to determine the relative significance of the impact categories. The normalisation as carried out here is merely to enhance the presentation, and to remove correlations in the uncertainty between the two systems. Whilst effective for comparisons of two systems and single indicator LCIA methods, many versions and combinations of ratios or differences have to be plotted for studies comparing many options across many categories.

Thus, a significant problem of probability plots is that they are limited in the number of dimensions they can display (i.e. typically two-dimensional plots). Although multi-dimensional displays are possible through adaptations and combinations of various plots, these can be complex to interpret, and the loss in transparency can sometimes outweigh the benefits of using a single plot (Morgan and Henrion, 1990). Conclusions thus have to be drawn across a large number of single plots if many options are to be compared, or if the options are compared across many impact categories or environmental interventions. This shortcoming is apparent in studies presented in the literature, which are limited to the presentation of a few major impact or inventory categories (usually two) (Coulon et al., 1997; Huijbregts, 1998b; Kennedy et al., 1996; Maurice et al., 2000). It is thus difficult to get an overview of the results. Box and whisker plots are slightly less limited, in that they can be used as a sort of “uncertainty enhanced” bar chart, and can be used to display more or less the same level of information as conventional bar charts.

A potential solution to viewing large multi-dimensional data sets is to use principal component analysis (PCA), which is able to reduce the dimensionality of the data set by producing a planar view of the data. In many cases, the first two principal components retain sufficient information so that the principal component plane can be plotted in two dimensions, and the solution set is able to be

visualised on a single plot (Jackson, 1991). Potential benefits of using PCA include reducing the dimensionality, selecting the most useful variables, visualising multidimensional data and identifying the underlying variables (Murtagh and Heck, 1987). The theory behind PCA is given in appendix A.3.

The use of principal component analysis (PCA) has been demonstrated as a powerful tool for visualising the inner structure of LCA data sets, and can provide significant insights into the interpretation of either deterministic or stochastic LCA results (i.e. inventories calculated from single point data inputs, or in those using simulation techniques to generate probabilistic results) (Heule and van den Berg, 1998; Le Téo, 1999). In deterministic LCA the options plot as single points on the PCA plane, whilst for stochastic LCA they plot as “clouds” of points, which can be interpreted as “zones of confidence”. The overlap between the “clouds” visually identifies the significance of the rankings between the options (Le Téo, 1999). In addition to enabling the data to be presented on a single plot, PCA provides significant insights into the data structure. It is able to discern highly correlated (and therefore redundant) impact categories, and to identify the most influential impact categories. This allows negligible categories to be identified and discarded (Le Téo, 1999). The relative strengths of the various representations of uncertain output are demonstrated and discussed in the case studies in chapters 7 and 8.

3.3. UNCERTAINTY ASSESSMENT METHODS DEMONSTRATED IN LCA

Guidelines for data quality assessment are included in most LCA reference sources (Heijungs et al., 1992; ISO, 1998; Lindfors et al., 1995), but their implementation is by no means standard and surprisingly few published LCA studies include formal data quality assessment. This can partly be explained by a lack of available methods and partly by a lack of reliable data for conducting such analyses (Maurice et al., 2000). Methods proposed to incorporate data quality assessment range from the early “post-it note” qualitative type approaches, to the more recent quantitative summaries of judgement-based DQI characterisations (Barnthouse et al., 1998). Appendix B summarises the approaches taken to incorporating an assessment of uncertainty into LCI studies. This section discusses the relative strengths and weaknesses of these studies, and draws on their experiences so as to inform the framework presented in the following chapter.

3.3.1. Qualitative and Semi-Quantitative Data Quality Assessment Methods

Qualitative data quality frameworks have tended to focus on setting data quality goals (DQGs) and guiding data collection, rather than assessing the quality of the data in hand. The DQGs are first defined according to the goals and scope of the study, after which the data is assessed according to a set of data quality indicators (DQIs). Table 3-3 lists some DQIs suggested to qualify LCA data. The data quality “score” is either qualitatively or semi-quantitatively represented, and displayed with the data element from which it arises. Weidema and Wesnæs (1996) recommend placing the “scores” in brackets behind the numeric value of the data element, in accordance with the “data pedigree” approach of Funtowicz and Ravetz (1990), whilst the SETAC data quality framework suggests a matrix representation or the use of value trees.

A matrix representation implies some level of aggregation, and thus avoids having to report the data quality for each data element. It allows the uncertainties to be borne in mind while drawing conclusions from the results, without getting lost in the detail of individual data scores. The emphasis of this thesis is on the quantitative assessment of uncertainty, although it is recognised that a quantitative assessment is not always feasible. This is the case in the technology scanning study presented in appendix H-1, where the pictorial matrix method developed by Graedel and Allenby (1995) for use in qualitative LCA applications is adapted to provide a representation of data quality. A matrix is drawn up with one axis comprising the life cycle stages of the product or process, and the other broad inventory or impact

categories, depending on the level of analysis. An oval symbol, able to represent two DQIs, is used to fill each matrix cell. The data quality “score” is represented by the degree of shading and the degree of fill of the oval, with the degree of “blackness” representing the degree of concern attached to that element (see Table H-6).

An overview of the major frameworks and guidelines developed for use in LCA is given in appendix B, and briefly reviewed below.

The qualitative frameworks give only broad guidelines with no specific criteria for choosing relevant DQIs or assigning data quality scores. The SETAC framework aims to be all encompassing and is thus very general, with the disadvantage that the exact steps required are not particularly clear. Similarly, the ISO 14041 standards only provide general guidelines. The USEPA data quality worksheets provide a definite structure to follow, and are therefore more useful than the SETAC and ISO guidelines, especially when it comes to data quality documentation. However, a clear procedure for data quality assessment is lacking in all three guidelines. The pedigree matrix approach of Weidema and Wesnæs (1996) is a departure from this, and gives a practical semi-quantitative procedure for assessing LCI data quality. The method provides a clear framework against which the data can be assessed, and unlike the previous frameworks, provides guidelines as to rating the DQIs, with clear descriptions of broad default categories (see Table 3-7).

Data Quality Indicators

The set of generic DQIs listed in the various guidelines are fairly extensive and cover broad definition areas (see Table 3.3). Some are overlapping (e.g. comparability and consistency), whilst others are too broad to be clearly applied (e.g. uncertainty, derived models etc.). Although the SETAC guidelines are clear that not all the listed DQIs should be applied, there is no clear definition of a necessary and sufficient set for most LCAs, or a clear mechanism for the selection of a relevant set, other than that their choice will depend on the context of the decision. A more manageable set of DQIs are presented in the USEPA and ISO 14041 guidelines, whilst Weidema and Wesnæs (1996) define a set of DQIs which are “necessary and sufficient to describe those aspects of data quality which influence the reliability of the result”. Their set of DQIs is driven particularly by the requirement to find an independent set of indicators, i.e. where each indicator reflects a unique aspect of the data quality. This suggests that a large set of generic indicators is not necessary and merely adds confusion to the analysis, especially where the indicators are not independent.

Table 3-7 Pedigree matrix with guidelines for scoring data quality indicators (Weidema, 1998b).

Indicator Score	1	2	3	4	5
Indicators which are independent of the study in which the data are applied:					
Reliability of source	Verified data based on measurements	Verified data partly based on assumptions or non-verified data based on measurements	Non-verified data partly based on assumptions	Qualified estimate (e.g. by industrial expert)	Non-qualified estimate or unknown origin
Completeness	Representative data from a sufficient sample of sites over an adequate period to even out normal fluctuations	Representative data from a smaller number of sites but for adequate periods	Representative data from an adequate number of sites but from shorter periods	Representative data but from a smaller number of sites and shorter periods or incomplete data from an adequate number of sites and periods	Representativeness unknown or incomplete data from a smaller number of sites and/or from shorter periods
Indicators relating to the technological and natural production conditions under which the data are valid, and therefore dependent on the data quality goals for the study in which the data are applied:					
Temporal correlation	Less than 3 years of difference to year of study	Less than 6 years of difference	Less than 10 years of difference	Less than 15 years of difference	Age of data unknown or more than 15 years of difference
Geographical correlation	Data from area under study	Average data from larger area in which the area under study is included	Data from area with similar production conditions	Data from area with slightly similar production conditions	Data from unknown area or area with very different production conditions
Further technological correlation	Data from enterprises, processes and materials under study	Data from processes and material under study but from different enterprises	Data from processes and materials under study but from different technology	Data on related processes or materials but from same technology	Unknown technology or data on related processes or materials, but from different technology

Assessment Framework

Even with the contained set of DQIs of Weidema and Wesnæs (1996), assigning DQIs to every data element in an LCA is extremely time consuming, and could become impractical for very large LCA studies. However, an even more significant criticism of the DQI approach than their extensive time requirements, is their inability to provide a summary statement of the data quality of the final results. Although the individual data quality items may be precisely “pedigreed”, no condensed index can be provided for the overall results without significant value judgements. At best, highly uncertain data inputs can be “flagged”, and the “flag” carried through to the results calculated from this data. If these results are seen to be influential in the final decision, the data input can be revisited (the so-called “post-it-note” approach promoted under ISO). Thus, the qualitative approaches enable each process step to be accurately described, but no aggregate information can reliably be produced (Coulon et al., 1997).

Combining the individual data scores or pedigrees is not very informative, since the combined pedigree can never be better than the pedigree of the data contributing to the total, and will therefore reflect the lowest score for each DQI. The combined pedigree can therefore only indicate that a data quality problem has occurred somewhere in the study, but not where or how serious the problem is. Thus, when used alone, DQIs do not give adequate information on the quality of the results. They are only useful when seen in combination with information on the uncertainty of the data (Weidema and Wesnæs, 1996). Calculating a weighted average of indicator scores to obtain an “output” score, as is done with the value tree approach, is not recommended, because of the significant value judgement averaging across disparate systems introduces (Coulon et al., 1997), and the false sense of accuracy this implies (Fava et al., 1994).

The general concept of the SETAC data quality framework is that input data be screened through the application of DQIs, and are adjusted to satisfy specified criteria contained in the data quality goals (DQGs). The up-front definition of preliminary DQGs and their associated DQIs allows consistency in internal levels of detail and application of decision rules, as well as facilitating peer review and validation of the study results (Fava et al., 1994). A disadvantage of an up-front definition of DQGs is that it could result in unnecessarily detailed data collection for some areas of the life cycle that may turn out to have little or no effect on the overall results of the LCA. Attaining consistency in internal levels of detail may thus be a waste of effort, and time could better be spent on going into greater detail in the more influential areas. In addition, unrealistic DQGs may be set, and effort will be wasted trying to collect better quality data where it simply is not feasible to do so. The framework does allow for the refining of DQGs as the data collection progresses, which addresses the second criticism, although redefining the DQGs for certain areas of the life cycle to some degree erodes a stated strength of DQGs, i.e. consistency in internal levels of detail and application of decision rules.

In theory, pre-screening data for use in LCA is a valuable practice, particularly where no prior information is available, but in practice the question tends to be “with what certainty can I support decisions with the data currently available to me?”. The influence of data quality on the results of the study is therefore the most important aspect of an LCA data quality framework, although it receives relatively little attention in the SETAC data quality conceptual framework. Starting at the bottom of the framework with the data quality assessment (see Figure B-1), and the study first carried out with currently or easily available data, would enforce the iterative nature of LCA. The ability of the results to address the decisions requiring support by the LCA can then be evaluated in light of the data quality

assessment, and the areas of the LCA highlighted where more certainty is required. Explicitly stated DQGs, as defined by the associated DQIs, are required for this evaluation, and can be derived in an identical process to that defined by the SETAC conceptual framework. A “bottom-up” framework is therefore suggested as being more useful in the context of many LCA studies.

3.3.2. Quantitative Measures to Address Data Uncertainty

Inventory data have traditionally been compiled and reported as point estimates, which reflect no information about the range, uncertainty or other characteristics underlying the raw data and inventory calculations. This leads to deceptive results, where one inventory category always appears definitively better or worse in a comparative study (Barnthouse et al., 1998). This is despite the fact that the SETAC “Code of Practice” emphasises the importance of including an estimate of the variability associated with the data, although they acknowledge the lack of accepted methodology for producing quantitative estimates of data quality (Consoli et al., 1993). Similarly the ISO standard recommends characterising uncertainty in results using ranges and/or probability distributions whenever feasible, but acknowledges that “uncertainty analysis as applied to LCI is a technique in its infancy” (ISO, 1998).

Chevalier and Le Téo (1996) argue from a different perspective for the inclusion of a quantitative uncertainty analysis. They argue that most LCA data cannot be realistically known with accuracy because they are taken from a wide number of inherently differing sources, and the “realism” of the data is in fact decreased by increasing the accuracy, i.e. there is no “true” single value. LCA data is therefore fundamentally fuzzy, where fuzzy data are defined as that whose realism decreases as its accuracy increases. Using mean values in LCA models forces an assumed data accuracy, and consequently yields unrealistic results. Realistic results and not accuracy should thus be the aim of the study. Realistic results can be obtained either by restoring the realism afterwards by computing error bounds on the results, or by retaining as much realism as possible from the start and modelling with ranges of data using simulation or analytical techniques (Chevalier and Le Téo, 1996).

Sensitivity analyses have been the primary mechanism for quantitatively assessing uncertain data in LCA studies (Coulon et al., 1997). They have predominantly been applied to assess the influence of key methodological choices, e.g. allocation rules, system boundaries etc., but they also have been used to determine the significance of individual data inputs on the results. Whilst useful in highlighting influential data elements, sensitivity analyses are not able to examine the combined effect of uncertain

data inputs, and are often based on applying arbitrary variance to the available single point estimate (e.g. $\pm 10\%$) (Kennedy et al., 1996). Also, although able to indicate the extreme values that a system can take, they do not provide any information on the distribution within that range (Coulon et al., 1997). Appendix B.2 summaries the relatively few studies that have quantitatively estimated the uncertainty in the results by examining the combined effect of input data uncertainties. The majority of these use simulation and random sampling of the input distributions, although a few analytical methods have been demonstrated.

The use of Monte Carlo simulation and other related sampling methods for the propagation and analysis of uncertainty is explained in appendix A.1. This method employs a relatively straightforward procedure of sampling the distributions of the input parameters a large number of times, and computing the model output with each of these samples, thereby generating a distribution of output values. A number of widely available software programs and spreadsheet routines are able to run Monte Carlo and related simulation analyses. The challenge of using simulation methods for the assessment of uncertainty in LCA therefore lies not with the method itself, which is well established, but in defining relevant input distributions that realistically characterise the uncertainty in the input data

The emphasis of the earlier studies is on the demonstration of the method, and the input probability distributions are partly based on arbitrary distributions and ranges (Huijbregts, 1998b; Kennedy et al., 1996). Whilst this was a necessary first step towards demonstrating the efficacy of the method, the definition of relevant input distributions is probably the most difficult aspect of the method, and therefore requires serious consideration if the use of stochastic modelling in LCA is to be made operational. Of issue is not only the shape of the probability distribution, but also the extent of empirical uncertainty that has been incorporated. In certain of the studies only actual data uncertainty is included, i.e. the statistically measurable variance of the data sample (Coulon et al., 1997; Maurice et al., 2000), whilst in others an estimate of the qualitative aspects of data quality is incorporated (e.g. its adequacy for use in the particular study), although, with the exception of Meier (2000), a transparent procedure for incorporating the subjective judgement involved is not provided.

Although identified as a shortcoming by some of the studies (Kennedy et al., 1996; Maurice et al., 2000), none of the studies explicitly address the implications that stochastic modelling has for the construction of the inventory model, i.e. the meaningless results that can be obtained by sampling correlated input data, and the inadequacy of “black box” LCA models for stochastic modelling (see

section 4.2.1). Also, with the exception of Huijbregts (1998b) and Meier (2000), the studies are primarily concerned with generating probabilistic output, with little or no emphasis on interpreting and reducing the uncertainties, i.e. the emphasis is primarily on distinguishing between uncertain options. Brief summaries of these studies can be found in appendix B.2.1, whilst a discussion on their relative strengths and weaknesses follows below.

The stated intention of Kennedy et al. (1996) for using Monte Carlo simulation is to generate inventory model results that can be analysed using well established and accepted statistical methods to differentiate between the study alternatives. Similarly, Coulon et al. (1997) conclude that “the use of stochastic models and the presentation of ranges and confidence intervals enhances decision making by helping to focus on categories for which there are real differences between alternatives”. However, they also caution that since the use of stochastic models is still a research field, its “application should be reserved to selected case studies”. Reporting a range of output is therefore undeniably an enhancement of the LCA method, although the reliability of the uncertainty assessment is still questionable and depends predominantly on how reliably the uncertainty in the input data is characterised. The statistical method used to propagate the uncertain input distributions through the model has a lesser contribution to the reliability of the analysis, and their limitations are well documented in the statistical literature (see appendix A.1).

Uncertainty Propagation Method

The choice of method will probably often depend on the resources available to the analyst (computer programs etc.). However, the following are of consideration. Propagation of uncertainties using simulation is intuitive and easy to apply to practically any form of existing model. It is able to be incorporated into most modelling platforms and can be reliably applied without a detailed understanding of the mathematics involved. The analytical methods, however, require that the inventory be set up in a matrix form and fairly demanding mathematical manipulations are required for their implementation. In addition, whilst an “exact” solution is obtained by analytical methods, limitations of the methods (e.g. the fact that first-order approximation is a local approach) raise questions as to their accuracy, whilst the accuracy of the simulated solution can be increased by simply increasing the number of samples taken. Finally, the manner in which input uncertainty is able to be incorporated is very flexible in the simulation approaches, and can range from simple uniform intervals of values to precisely defined probability distributions, whilst because of the mathematical complexity involved, most analytical methods can only incorporate uncertainty intervals or standard deviations.

Characterising Input Uncertainty

Coulon et al.'s (1997) assertion that actual data provides the best source of information for characterising the distribution of the input is questionable, as in many instances the applicability of the data will introduce greater uncertainty than any measured statistical uncertainty. Coulon et al. recognise that stochastic modelling conducted using actual data to characterise the inputs falls short of addressing the adequacy of the data with regard to the goal and scope of the project (Coulon et al., 1997). They recommend such aspects be addressed by an accompanying qualitative assessment. However, methods to quantitatively include the adequacy of the data have been developed (Kennedy et al., 1996; Meier, 1997; Weidema and Wesnæs, 1996).

Whilst relatively easy to apply, the accuracy of using a single overall DQI as suggested by Kennedy et al. (1996) is questionable, especially since there is such a wide range of data quality issues which must be intuitively considered by the "expert". A manageable and well-defined set of DQIs, as proposed by Weidema and Wesnæs (1996) or Meier (1997), is therefore preferred. In addition, the DQI set developed by Weidema and Wesnæs can be determined independently of each other, thereby allowing the influence of each indicator to be considered additively, important if the overall uncertainty is to be determined. Weidema and Wesnæs' method especially bridges the gap between the methods using actual data and those using "expert" judgement, in that it explicitly incorporates both measured uncertainty and uncertainty due to the application of the data. However, Coulon et al. (1997) caution against the significant value judgement such mixed approaches can entail.

Kennedy et al. defend the inherent subjectivity of their method by stating that an explicit intention of the method is to quantitatively incorporate expert judgement into the models, and that the information contained in the input data is enhanced by incorporating this qualitative assessment of data quality. They acknowledge the reliance of the method on subjective judgements, and caution against the inappropriate use of the model results, suggesting that sensitivity analyses be conducted to determine the sensitivity of the model results to the distribution parameters. Another significant criticism of Kennedy et al. (1996) is the apparently arbitrary allocation of DQI values to the variance and probability distribution shape applied to that data element. They acknowledge that no standard was established for the amount of variance to apply to varying levels of input data uncertainty, and that the variance of the results presented should be viewed conservatively. Although the structured approach and well-delineated guidelines of Weidema and Wesnæs' method is expected to improve the

reproducibility of applying subjective judgements, their method is subject to the same general cautions as that of Kennedy et al. (1996).

A range of probability distribution shapes have been recommended for characterising LCI input data. Uniform and triangular distributions are typically used where little information is known about the parameter (Coulon et al., 1997; Huijbregts, 1998b; Maurice et al., 2000), with triangular distributions recommended when a central tendency to the data is expected. Normal, lognormal and beta distributions are recommended for more accurate representations of the sample distribution (Coulon et al., 1997; Kennedy et al., 1996; Maurice et al., 2000; Meier, 1997). Meier recommends normal and lognormal distributions as both are well-known and easily definable statistical functions, whilst Maurice et al. choose a version of the beta distribution, as they find these easier to define than the “less intuitive” standard deviation of normal distributions. Reasons for choosing the particular distributions are therefore contradictory, and the choice appears primarily to be related to whatever distribution the analyst is familiar with and feels comfortable using.

Level of Inventory Modelling

The need for the nature and influence of parameter interdependency to be better understood has been highlighted by stochastic LCA studies (Coulon et al., 1997; Kennedy et al., 1996; Steen, 1997). Steen (1997) concludes that a dependency amongst input data will result in an overestimate of uncertainty in the final results, and recommends that, as far as possible, LCA calculations should be made “from the cradle”, i.e. that a primary value is entered only once in the model. Kennedy et al. (1996) recommend that any correlations between the input data elements be specified, to enable control of the physical and chemical properties in the model, although they acknowledge the increase in modelling complexity incorporating correlation data would entail. Another advancement of the method suggested by Kennedy et al. is that the assessment of the data quality should occur at a lower level of aggregation. This would allow for a more thorough qualitative assessment, and in some cases, where sample sizes are sufficient, the actual probability distributions could be determined.

Streamlining

A criticism common to all studies incorporating stochastic uncertainty analysis is that it increases the already arduous data collection phase of LCA (Coulon et al., 1997). Huijbregts (1998b) observes that it will probably not be feasible in LCA case studies to determine the uncertainty ranges of all the input data used in the inventory, and to perform an extensive parameter uncertainty analysis in the

characterisation phase. Ways to streamline the process are therefore required. Huijbregts recommends focussing on the key parameters, which he claims will increase the feasibility of the analysis whilst decreasing its validity only to a limited extent. A formal procedure for such an approach has been developed by Maurice et al. (2000). The basic principle of the method is to select a limited number of key inputs. Relevant probability distributions are assigned to these flows and the influence of their uncertainty on the cumulative results calculated using Monte Carlo simulation. The basic premise of the method is that the selected elements will account for the bulk of the uncertainty in the results, and that the non-selected data will not modify the conclusions to any significant degree. This is checked by assigning rough upper and lower limits to the non-selected data, and the simulation repeated. If the influence of the non-selected data on the overall uncertainty is found to be appreciable, an uncertainty importance analysis is conducted to identify those input elements contributing substantially to the uncertainty, and better approximations are made of their distributions.

Maurice et al. (2000) base their selection of the “main data” on a qualitative assessment of their uncertainty and a quantitative assessment of their contribution to the overall results, according to Heijungs’ (1996) definition of “key” data (see Figure 3-3). Groups of elementary flows from unit processes contributing between certain percentage limits to the cumulative elementary flow (or category score) of the functional unit are selected as having a significant contribution. To place the data on the abscissa of Figure 3-3, a qualitative data quality “score” of 1 to 5 on an ordinal scale is assigned to each data element. Although there are differences in terminology, the DQIs selected by Maurice et al. are essentially identical to those defined by Weidema and Wesnæs (1996) as the smallest independent set of indicators possible. However, unlike Weidema and Wesnæs, who warn strongly against aggregating the indicators (because assigning a score on an ordinal scale does not relate to any additive “amount” of data quality), Maurice et al. weight the DQIs and sum them to yield a single indicator. They defend the use of an aggregated DQI by pointing out that the score is not directly transformed into a probability distribution, but merely used as an intermediate indicator for the identification of the “main data”.

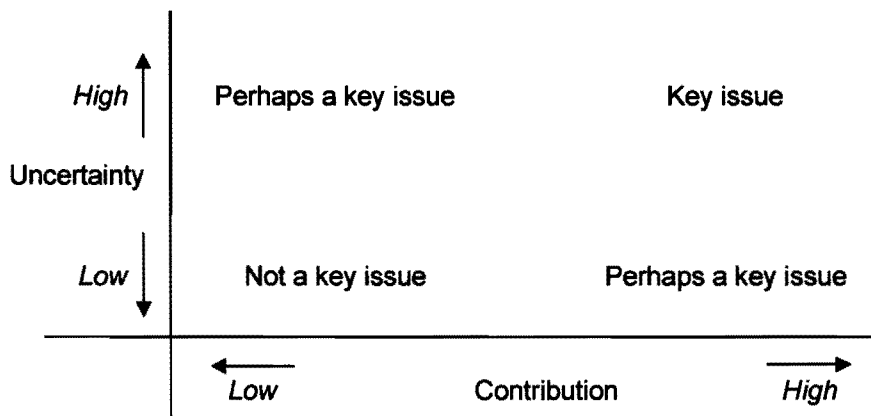


Figure 3-3 Graphical representation of 'key' data with respect to its uncertainty and the magnitude of its contribution to the overall results (Heijungs, 1996)

To avoid the setting of arbitrary cut-off rules for selecting key parameters, a broad sensitivity analysis, using standard uncertainty estimates, can first be performed to determine which parameters may contribute substantially to the overall uncertainty (Heijungs, 1996). However, the lack of standard uncertainty estimates in the literature leads to the use of arbitrary uncertainty ranges, often 10%, which can result in parameters with large unknown uncertainty ranges being incorrectly thrown out. The use of standard sensitivity ranges for different types of environmental interventions could help avoid this (Huijbregts, 1998b). Similarly, Weidema and Wesnæs (1996) suggest using default uncertainties generated from similar sets of data, although they suggest these for semi-quantitative DQI scores. A matrix of default uncertainties for each DQI score could be generated for different types of studies (e.g. for different types of data or different industries) (Weidema and Wesnæs, 1996). This approach is demonstrated by Meier (1997) in a study of gas purification systems.

Huijbregts (1998b) also suggests that applying the stochastic modelling at a higher level of aggregation, as is done by Kennedy et al. (1996), could help streamline the uncertainty analysis process. However, this leads to a fairly arbitrary allocation of uncertainty (as discussed above), and a conclusion of Kennedy et al. (1996) is to assign probability distributions at a more disaggregated data level. Nevertheless, assigning uncertainty ranges to accumulated inventories has useful application in the analysis of the potential importance of background data (Huijbregts, 1998b).

3.3.3. Model Parameter Uncertainty and Uncertainty about Model Form

The majority of uncertainty assessments demonstrated in LCA studies have been around empirical uncertainty. However, the use of sensitivity analyses to determine the effect of the many choices in LCA studies is fairly well established, and there is a general consensus on the use of parametric or sensitivity analyses for the assessment of model uncertainties (Huijbregts, 1998b; Maurice et al., 2000; Meier, 1997). Studies assessing uncertainties arising from model choices have predominantly been around the choice of allocation method (Azapagic and Clift, 1999b; Azapagic and Clift, 2000; Huijbregts, 1998b), major model parameters (e.g. system lifetime) (Huijbregts, 1998b; Maurice et al., 2000; Meier, 1997), and the choice of impact assessment method (Baumann and Rydberg, 1994; Meier, 1997; Tolle et al., 2001; Tukker, 1998). These studies demonstrate the extremely significant consequences the choice of LCA method can have on the results.

For the assessment of model parameters, Steen (1997) suggests that a probability number could be assigned to each alternative choice, although Morgan and Henrion (1990) argue strongly against this (see section 3.1.1. for reasons). Huijbregts (1998b) identifies that sensitivity and parametric analyses are cumbersome to perform and report, and that it is not feasible to analyse the effect of all possible combinations of choices, and to compute the model output for all ranges of model parameters. Huijbregts (1998b) therefore suggests finding the two extreme options for every choice, and then constructing two extreme combinations of options, and computing the effect of these on the LCA results.

Meier (1997) recognises the need for scenario modelling of discrete options where uncertainties are caused by preferences within a decision-making process. For example, in an assessment of gas purification systems, Meier (1997) repeats the probabilistic assessment with a different background electricity production model. For one metric of ecological indicators, this causes the rankings of the different technologies to be nearly reversed, thereby demonstrating the potential significance of the choice of model parameters. Meier (1997) further distinguishes model uncertainties whose importance can only be qualitatively addressed, i.e. no clear scenarios can be defined to test the model choices by sensitivity analysis. These are model uncertainties based on inherent model assumptions and definitions (Meier, 1997), e.g. spatial and temporal limitations. This is in agreement with Huijbregts (1998b), who recognises that certain model uncertainties (such as the inherent inability of the current LCA method to

deal with spatial and temporal uncertainties) are not able to be addressed, although their effects may be serious, and cautions that it is important to avoid the illusion that the uncertainty in an LCA study is totally quantified when stochastic modelling and sensitivity analyses have been applied (Huijbregts, 1998b).

3.3.4. The Reduction and Analysis of Uncertainty

Key issues or “hot spots” are defined by Heijungs (1996) as the areas in a study where a small change has a large consequence. By addressing the key issues first, modelling effort is directed to where it is really needed, thereby streamlining the LCA process and returning an overall LCA of higher quality. Heijungs (1996) distinguishes between key issues as areas to concentrate on in a more detailed LCI, and key issues as areas that make up a large contribution to the total burden. It is the former which can be highlighted by an uncertainty analysis, and are of interest when the aim of the analysis is to improve the quality of the LCA results. The latter are of importance in an improvement analysis. This distinction is evident in the graphical representation of key issues shown in Figure 3-3 (Heijungs, 1996). The graphic highlights that highly uncertain data may have a negligible influence on the results, and that it is therefore necessary to identify the data that is both uncertain and for which the final results are sensitive.

Heijungs (1996) demonstrates the use of Gaussian approximation to determine the key data elements, whilst Huijbregts (1998b) demonstrates the use of correlation coefficients in a study of two different gutter system. Rank-order correlation coefficients between every uncertain parameter and every model outcome are computed, and a high correlation indicates that the parameter has a relatively large impact on the uncertainty of the model outcome (Huijbregts, 1998b). Different parameters are likely to be highlighted by the uncertainty importance analysis depending on whether the analysis is performed for each system separately, or whether the analysis is performed for a comparison of systems. Where the systems being compared have a number of processes in common, the parameters causing uncertainty in these common processes are divided out in the comparison, whereas they may be significant in an analysis of the individual product system (Huijbregts, 1998b). Maurice et al. (2000) also demonstrate the use of correlation coefficients, although they do so to direct their iterative streamlining method, rather than to analyse the relative uncertainty importance of the input parameters.

Steen (1997) formalises a methodology for using sensitivity to determine the most significant inputs, by defining the “relative sensitivity”:

$$\sigma_x/\Delta x$$

This is ratio of the uncertainty of the data element, x , expressed as a standard deviation, to the magnitude of the change in x required to change the priority obtained in a comparative LCA. The ratio gives an indication of how important the precision in the estimation of x is, for the “best” option to remain the best (Steen, 1997). Meier (2000) also demonstrates the use of sensitivity analyses for the identification of the key data elements contributing to the uncertainty in his assessment of waste gas purification systems.

Many definitions of life cycle assessment stress its iterative nature, although criteria on which to base the necessity for performing a subsequent iteration are generally lacking, i.e. guidelines for deciding whether a sufficient level of detail has been reached, and if not, which elements require a more detailed assessment (Heijungs, 1996). An exception is Steen (1997), who defines the term “net improvement efficiency” to describe the consequences of uncertain input data on improvement of the environment, which is a ratio of the average environmental improvement to the highest possible improvement obtained if all decisions are correct. However, this measure is only applicable if two systems are being compared, and if a single indicator method is used. The following chapter develops an iterative framework for incorporating an analysis of uncertainty into LCIs, with a specific emphasis on directing the analysis back into the inventory to determine which input parameters are responsible for the most variance. This allows the “key” data to be determined, and directs the analysis to those parameters that should be addressed first for the greatest reduction in output uncertainty to be achieved.

CHAPTER 4

DEVELOPMENT OF INVENTORY MODELLING TO INCLUDE UNCERTAINTY

LCA studies that have included rigorous quantitative uncertainty analyses have shown enhanced decision making capabilities and significantly increased the credibility of the results (Kennedy et al., 1996; Maurice et al., 2000; Meier, 1997). Indeed, one author has gone as far as saying that resorting to presenting single value LCA results is a “dangerous convenience ruining LCA’s predictive credibility” (Le Teno, 1999). However, uncertainty analyses are by no means standard, and their inclusion into the LCA methodological framework certainly not seamless. This chapter provides a framework to include a quantitative assessment of data quality in LCI, and looks at how inventory modelling needs to be adapted to include uncertainty.

4.1. A FRAMEWORK FOR THE QUANTITATIVE ASSESSMENT OF UNCERTAINTY IN LCI

A successful framework for incorporating uncertainty into LCI needs to incorporate the strengths of both qualitative and quantitative approaches. In addition, it needs to be sensitive to the fact that different types of uncertainty require different methods for their assessment. The framework presented below combines the strengths of the DQI approaches (the pedigree matrix) with quantitative simulation techniques, whilst explicitly differentiating between uncertainties in model form, model parameter uncertainty and empirical parameter uncertainty.

4.1.1. A Framework for Empirical Parameter Uncertainty

Maurice et al. (2000) identify three problems common to all quantitative uncertainty assessments:

1. Different types of uncertainty exist in LCA and it is not possible to analyse the importance of each of them in detail.
2. A full quantitative uncertainty assessment is much too time consuming to be applied to every LCI as a default.
3. The determination of a probability distribution for each data element input into an LCA model is problematic.

The framework developed by Maurice et al. (2000) aims to address these problems. The second problem is addressed by employing an iterative procedure that selects the data most likely to contribute to the uncertainty, determines accurate probability distributions for this data only, and then evaluates the effect of this simplification, adding more data to be precisely quantified, if necessary. The third problem is addressed by tailoring the type of probability distribution to the amount of information known about the data element under consideration. The framework consists of five steps, with the possibility of iterative feedback loops (Maurice et al., 2000):

1. Compilation of a first inventory.
2. Selection of main data based on a qualitative evaluation of their uncertainties and on their quantitative contribution to the cumulative results.
3. For the data selected in step 2, selection of an appropriate probability distribution depending on the information available.
4. Calculation of the uncertainty spread and average value of the cumulative results using stochastic simulation, firstly with just the main data characterised as input distributions, and then also with the non-selected data characterised with roughly specified maximum and minimum values. The effect of the non-selected data is evaluated, and a return to step 2 or 3 directed if its effect is significant.
5. Analysis and discussion of the results, with reference to the goal and scope of the study.

A weakness of the above framework is that the compilation of a first inventory is required as the starting point of the assessment. The inclusion of margins or error or variation should not necessarily be viewed as an add on to a completed inventory, and as Chevalier and Le Téo (1996) point out, the very nature of most LCI data is such that only if the modelling is conducted with a range of values can realistic results be obtained. Ideally, distributions or ranges of values should be modelled from the start, as this removes the often arbitrary declaration of mean or most likely values. The assertion that an uncertainty assessment need not start with a completed mid-point inventory is a major point of departure between the framework developed here and that postulated by Maurice et al. (2000).

A single point inventory is the necessary starting point for Maurice et al.'s method, as it is required to calculate each data element's contribution to the cumulative inventory results. This, together with a qualitative assessment of their uncertainty, selects the data for which detailed probability distributions are defined. This step is identified as the weakest part of the framework, in that not only does it force

the uncertainty assessment into being an “add-on” analysis to a completed mid-point inventory, but it requires arbitrary “cut-off” rules to be set, to determine at which point data can be judged to be unimportant. The qualitative uncertainty assessment has a similar problem in setting the “cut-off” DQI score, in addition to the reservations regarding weighting and aggregating DQI scores, i.e. the fact that ordinal scores can not meaningfully be aggregated, and that weighting assigns the same importance to the DQIs for all types of elementary flows, although their relative importance is likely to change according to the type of flow.

In addition to setting the “cut-off” criteria, subjective judgement is required in setting the criteria at which the “non-selected” data is judged to have no appreciable affect on the results. Maurice et al. identify different criteria according to the type of LCI study. For a relative comparative assessment of products, the aim is to identify the best option rather than to have the best estimation of the confidence interval of the results. The two simulation cases, that including the roughly characterised non-selected data, and that not including the non-selected data, are compared, and if the change in the mean and standard deviation is less than a pre-defined maximum tolerable difference (e.g. 10%), then the influence of the non-selected data can be considered negligible. In an evaluation of cumulative results for a product the aim is to determine the most probable confidence interval of the cumulative results. In this case, it is checked that the simulation including the roughly characterised non-selected data does not change the standard deviation by more than 10% or cause the mean to increase by more than 5%. The percentage change tolerated will depend on the scope of the study, and the values of 10% and 5% are the result of personal experience of Maurice et al. (2000).

The framework developed in this thesis dispenses with either calculating an initial mid-point inventory or a qualitative ranking of uncertainty. It starts from the premise of using broad uncertainty estimates to identify those parameters contributing substantially to the overall uncertainty. An uncertainty importance analysis identifies the influential parameters, so that their definitions can be refined, and an improved estimate of the output uncertainty obtained. This process is shown by the steps in the central shaded area in Figure 4-1. Data ranges are incorporated as they arise during data collection, thereby removing the need to isolate an often unrealistically defined “most likely value”. The model output is computed by simulation, and rank-order correlation coefficients are calculated from the simulated inputs and outputs. The rank-order correlation coefficients provide a relative measure of each input parameter’s contribution to the output uncertainty, which, in turn, is a function both of the magnitude of

its numerical contribution and its uncertainty contribution. The parameters displaying the highest correlation coefficients (i.e. those contributing most to the variance in the output) are identified, and their input definitions revisited (i.e. their variance reduced). The procedure is repeated until the output uncertainty has been reduced to an acceptable level.

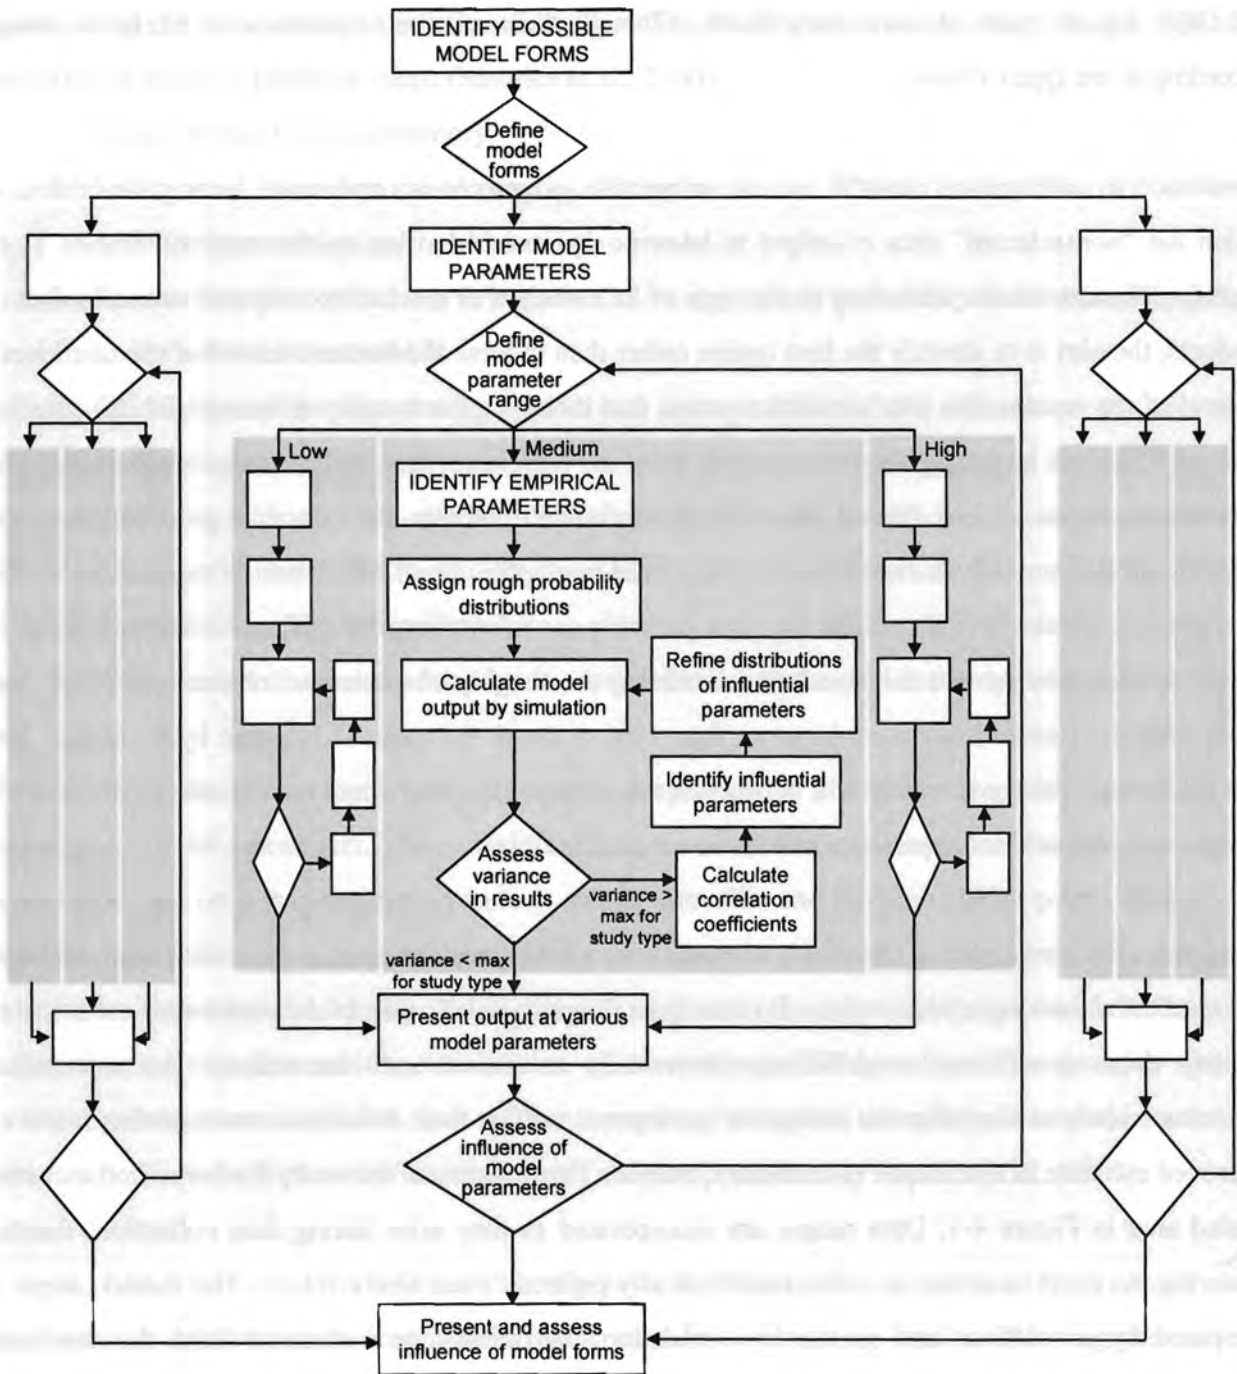


Figure 4-1 Procedure for an iterative uncertainty assessment of empirical parameters, model parameters and model form.

The emphasis of the correlation analysis is different to that employed in the framework of Maurice et al (2000). If the uncertainty analysis is incorporated from the start of a study, it is used to shape data collection, e.g. generic estimates can be updated with specific process data if the parameter is returned as important. If the process is applied to a study in which data collection is completed, the emphasis is on identifying those parameters whose distributions require better definition (as in Maurice et al.), and on identifying those parameters limiting the certainty in the results, i.e. those parameters that have high uncertainty importance even once they have been defined as accurately as the current level of data allows.

As in Maurice et al's framework, "stopping" criteria are required, i.e. to judge when a sufficient number of iterations have been completed. This point is thought to be more meaningfully related to the degree of confidence in the results, than a percentage change between the iterations. The parameters can be successively refined until an acceptable variance in the results is obtained, where the acceptable variance will depend on the goal and scope of the study. The acceptable variance is most meaningfully defined in a comparative study, where it can be related to a required degree of confidence in the differences between the options, e.g. iterative model refinements can be made until the confidence intervals show 90% probability of being non-overlapping.

The acceptable confidence limit will depend on the decision context (i.e. the level of risk able to be tolerated by the decision makers), and on the degree of separation between the options. The level of acceptable variance is likely to be revisited as the study progresses, and may differ between selection criteria (i.e. according to the particular impact category or environmental intervention under consideration). Data availability is often the limiting factor in determining the acceptable variance in the results. Additional data collection may be outside the scope of the study, or better data may simply not be available. In this case, the acceptable or limit variance is reached when successive iterations yield little or no reduction in the variance of the results. In this case, a percentage change between the iterations can be taken as the "stopping" criterion, e.g. the 10% suggested by Maurice et al. (2000).

For the first iteration, the uncertainty of each parameter is conservatively estimated (a simple distribution with an overestimate of the range). The rough definitions of the parameters returned with high uncertainty importance are then amended by combining their measurable uncertainty (i.e. the statistically measurable variance of the data sample) with an estimate of the increase in variance due to

sources of uncertainty not captured by the data sample (e.g. its applicability to the study). The quantification of the more qualitative and subjective components of uncertainty is formalised by the use of the pedigree matrix developed by Weidema and Wesnæs (1996) (see section 4.3.2 for detail on the characterisation of uncertain parameters). Subsequent iterations may find that a more precisely defined distribution is not sufficient to reduce a parameter's contribution to the output uncertainty (i.e. the parameter is still returned with high uncertainty importance), in which case, a larger data sample or more applicable data requires collection if the uncertainty is to be reduced further.

The procedure presented here is not expected to be any more arduous than that of Maurice et al. (2000), in that they have to check for the effect of neglecting the “non-selected” data. This somewhat negates the streamlining benefits of removing this data in the first place, and results in their computing the model output twice, once without the uncertainty of the non-selected data, and again with it roughly included. The streamlining procedure itself is computationally intensive and thus also time intensive, especially if there are a large number of unit processes and cumulative elementary flows whose contributions must be evaluated against each other. Thus if a large number of iterations need to be done (i.e. if it is found that the non-selected data is contributing significantly to the uncertainty) the benefits of the streamlining may not be apparent. How much benefit is derived from combining a qualitative and quantitative approach will probably depend on the case study in hand. For the case study conducted by Maurice et al. (2000) they found they required two iterations to get a stable estimate of the parameter uncertainty of the cumulative results. A potential problem of incorporating all empirical parameters as data ranges from the start is that extensive computing power is required to calculate the rank-order correlation coefficients. However, higher-end desktop systems have been found to be sufficient for the case studies conducted in this thesis.

4.1.2. Extending the Framework

The procedure outlined above is relevant only to an assessment of empirical uncertainties. The majority of studies addressing uncertainty have tended to look only at empirical uncertainties, and those that have included an assessment of model uncertainties have not followed a systematic approach. The framework developed here aims to address this and explicitly places the different sources of uncertainty in relation to each other. The types of uncertain quantities relevant to LCA are identified in Figure 3-1. The distinction between uncertainty in empirical parameters, model domain parameters and model form is important as it relates to how these uncertainties can be assessed.

An assessment of the three types of uncertainty can be visualised as consisting of an onion-like structure, as shown in Figure 4-2. Initially a single model form and a single set of “most likely” model parameters are selected. The innermost level of Figure 4-2 consists of the probabilistic assessment of empirical parameter uncertainty described above. The analysis then moves up to the second level of assessment, that of a parametric assessment of model parameters. The model is computed for a few discrete values of model parameters, or sets of model parameters, covering the feasible operating range of each model parameter. Varying the model parameters may affect which empirical parameters significantly contribute to the uncertainty in the overall results (i.e. different key parameters may emerge), and the innermost iteration may have to be repeated until the empirical uncertainty is returned to the acceptance level. If the results are found to be sensitive to the value of the model parameter, the discrete versions of the results are retained and discussed as part of the final results.

The top-most layer of the analysis is the assessment of model uncertainty. If known, a different form or structure of the model can be implemented, which in most cases will introduce new model parameters and empirical parameters. This will consequently require the assessments at the lower levels to be repeated. Once the uncertainty due to the empirical parameters has been reduced to an acceptable level and that due to the model parameters explicitly investigated, the effect of the different model form on the results can be evaluated. The analysis of model forms may result in the “best” version being chosen, if this can be supported qualitatively, otherwise the range of results resulting from the different model forms are retained for presentation and discussion in the final results.

A representation of the framework is given in Figure 4-1, which presents the individual steps of the assessment and highlights the iterative nature of the procedure. The assessment is shown for only three model forms and three model parameter values. If the model is found to be sensitive to a number of the model parameters, so that a single “best” value can not be meaningfully defined, and if a number of different model forms are available, the number of cases requiring comparison in the final results quickly explodes to an unmanageable number. The level of sensitivity to model parameters that can be tolerated is therefore an important consideration. As with setting the acceptable variance limit, whether or not a parameter is considered sensitive or not will depend on the particular decision to be supported, and the relative magnitude of the empirical and model parameter uncertainty (see following section).

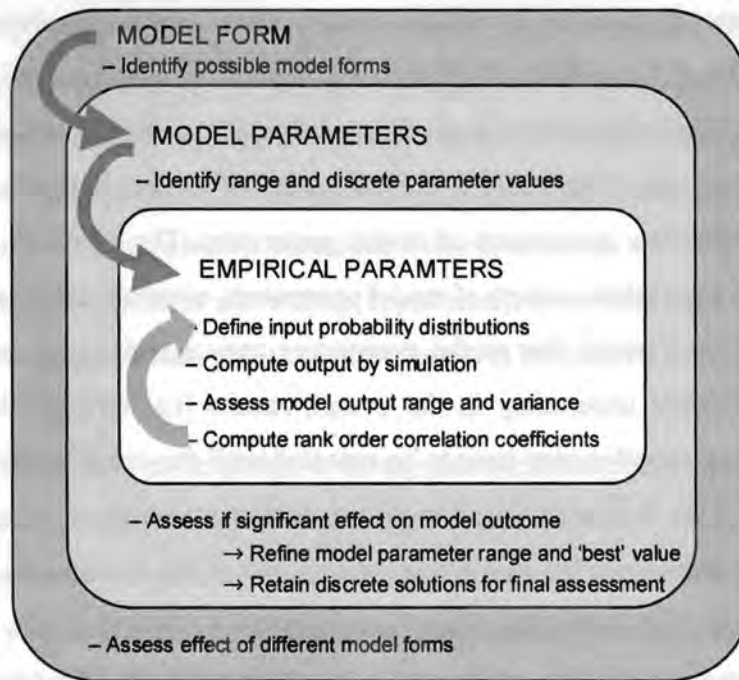


Figure 4-2 Schematic of the nested relationship of uncertainty assessments for the three major types of uncertainty present in LCI models.

4.1.3. Important Considerations in the Assessment of Uncertainty

The required accuracy of the uncertainty assessment, as well as the degree of variance and sensitivity considered acceptable by the decision-makers, are governed by the decision context of the study. Hard and fast rules can therefore not be given, as the actual values chosen will depend very much on the particular study. However, the following considerations can be taken into account.

4.1.3.a Accuracy of the Uncertainty Analysis Results

The type of decision to be supported will dictate the appropriate level of precision in the results of the uncertainty analysis. A system with highly uncertain data and crudely estimated input probability distributions will hardly benefit from a meticulous uncertainty assessment with tens of thousands of sampling runs. A few hundred will probably suffice, as at best, the range of output will be all that can be meaningfully discerned. Whilst a well characterised system with representative input distributions will require a sample size sufficiently large to accurately calculate the output distribution parameters. Thus, the desired precision in the results needs to be selected with reference to the application of the uncertainty analysis, e.g. whether a rough estimate of the robustness of the results is required, or whether an accurate estimate is required to differentiate between capital intensive projects.

Two important parameters governing the accuracy of the probabilistic simulation are the choice of sampling method, and the number of samples to be taken. An overview on the various sampling routines are given in appendix A.1.1. The advantages of a stratified sampling routine outweigh their slight disadvantage (that of not being able to simply calculate the precision of the output), so Median Latin Hypercube sampling is the recommended sampling method. The use of stratified sampling means that the equations for estimating the required sample size given in appendix A.1.1. are no longer strictly applicable (as these are based on Monte Carlo sampling). However, since stratified sampling always returns a precision equal to, or better than, Monte Carlo sampling, these can still be used to calculate the maximum number of samples required for a desired precision, i.e. for stratified sampling they will give an over-estimate of the required sample size (Morgan and Henrion, 1990). For Monte Carlo sampling, 1000 samples give a 95% confidence interval that the 50th percentile (as the least accurate percentile) will be plus or minus 3.5 estimated percentiles. Given that Median Latin Hypercube sampling should be more accurate and the high model parameter uncertainty, sample sizes in excess of this would appear to be unnecessarily precise for most LCI studies.

Whilst the precision of the simulated output can be simply estimated from the sample size, this does not convey the overall accuracy of the probabilistic analysis, which contains aspects that can not be quantitatively determined, e.g. the applicability or comprehensiveness of the assigned input distributions. The accuracy of the model parameter and model form analyses is even less easy to evaluate. In order to streamline the parametric analysis, “best” and “worst” combinations of variables are grouped together, which could result in possible combinatorial effects being overlooked. A trade-off therefore needs to be found between a manageable number of runs and the potential loss of accuracy through the grouping of variables. The accuracy of the model parameter analysis will also depend on the definition of meaningful parameter ranges and intervals at which the parameters are evaluated. The accuracy of an assessment of different model forms can not usually be determined because in most cases the full range of potential model forms is not known, and are continually changing as advancements are made.

The level of accuracy aimed at for each level of the assessment (i.e. each layer in Figure 4-2) will depend on the accuracy of the other levels. For example, there is little point in conducting a highly accurate empirical parameter assessment, with a large sample size and precisely defined probability distributions, in a study where model parameter uncertainty is high, and completely overshadows the

empirical parameter uncertainty. Similarly, if results from different model forms are shown to be widely divergent, a parametric analysis of every possible combination of model parameters would be redundant. The relative importance of the various levels depends on the decision context (see Figure 3-2).

4.1.3.b Limiting Variance / Sensitivity

In the framework outlined above an acceptable variance in the probabilistic output needs to be defined in order to determine whether further iterations are required. The value chosen will be closely related to the decision context of the study, e.g. a decision resulting in large capital expenditure will require far greater certainty than one resulting in a simple process modification. A confidence level in excess of 95% is likely to be required for the former, whilst 60 or 70% may be sufficient for the latter. For comparative systems, the variance does not need to be less than that showing significant non-overlap of the options, i.e. if the options differ widely, a high variance in the output of each option can be tolerated.

The chosen variance limit is likely to be revisited during the course of the study, e.g. if targets have been set that are infeasible within the scope or data availability of the study. In such cases the minimum variance achievable with the data currently available to the study will be reached, and either the acceptable variance must be relaxed or the scope of the study must be amended (e.g. significantly increased modelling detail or data collection undertaken). In some cases, the inherently uncertain nature of the system means that the output variance will never be able to be reduced, even given limitless resources, e.g. strategic modelling of systems for implementation well into the future. A revision of the acceptable variance must be done in accordance with the goals of the study, as increasing the level of uncertainty able to be tolerated in the results may mean the goals of the study can no longer be attained, and will have to be simultaneously revised.

The sensitivity at which the model parameters are considered significant depends on similar considerations, i.e. how robust the results need to be. In a comparative study, it will depend on the degree of overlap between the options, with the choice of a parameter obviously not significant if it does not cause the options to overlap. If the options do overlap, the combined "uncertainty range" from the variance arising from the empirical parameters and the range arising from the model parameters can be set according to the required degree of certainty (e.g. less than 5% overlap between the options). As

empirical parameter uncertainty. Similarly, if results from different model forms are shown to be widely divergent, a parametric analysis of every possible combination of model parameters would be redundant. The relative importance of the various levels depends on the decision context (see Figure 3-2).

4.1.3.b Limiting Variance / Sensitivity

In the framework outlined above an acceptable variance in the probabilistic output needs to be defined in order to determine whether further iterations are required. The value chosen will be closely related to the decision context of the study, e.g. a decision resulting in large capital expenditure will require far greater certainty than one resulting in a simple process modification. A confidence level in excess of 95% is likely to be required for the former, whilst 60 or 70% may be sufficient for the latter. For comparative systems, the variance does not need to be less than that showing significant non-overlap of the options, i.e. if the options differ widely, a high variance in the output of each option can be tolerated.

The chosen variance limit is likely to be revisited during the course of the study, e.g. if targets have been set that are infeasible within the scope or data availability of the study. In such cases the minimum variance achievable with the data currently available to the study will be reached, and either the acceptable variance must be relaxed or the scope of the study must be amended (e.g. significantly increased modelling detail or data collection undertaken). In some cases, the inherently uncertain nature of the system means that the output variance will never be able to be reduced, even given limitless resources, e.g. strategic modelling of systems for implementation well into the future. A revision of the acceptable variance must be done in accordance with the goals of the study, as increasing the level of uncertainty able to be tolerated in the results may mean the goals of the study can no longer be attained, and will have to be simultaneously revised.

The sensitivity at which the model parameters are considered significant depends on similar considerations, i.e. how robust the results need to be. In a comparative study, it will depend on the degree of overlap between the options, with the choice of a parameter obviously not significant if it does not cause the options to overlap. If the options do overlap, the combined "uncertainty range" from the variance arising from the empirical parameters and the range arising from the model parameters can be set according to the required degree of certainty (e.g. less than 5% overlap between the options). As

with the acceptable variance, the required degree of certainty is defined according to the goals and scope of the study, and is likely to require revision as the study progresses. High model parameter uncertainty is often due to a loosely defined problem, and can be reduced by revisiting the problem definition phase of the study. On the other hand, the very nature of some problems is a loosely defined model parameter set (e.g. long-term scenario planning studies), in which case the model parameter uncertainty can not be reduced, but can be managed by grouping the model parameters into likely scenarios, i.e. the results presented at a number of separate model parameter states.

As with the relative accuracy required at each level of the analysis, the variance and sensitivity considered tolerable will depend to some degree on each other, and on the model uncertainty. Low empirical uncertainty, evaluated at an exhaustive breakdown of model parameters, will not be of much value if the model forms are shown to give contrasting results. Similarly, there is little point in reducing empirical uncertainty to very low level, if model parameter uncertainty is high. The number of iterations in the core of Figure 4-2 therefore depends on the contributions to the overall uncertainty of the two outer layers. The acceptable variance, which essentially controls the number of iterations required, thus needs to be defined bearing in mind which layer is expected to dominate the uncertainty assessment, which in turn is dependent on the decision context.

4.1.3.c Influence of the Decision Context

The influence of the decision context on model parameter uncertainty is due to its influence on the choice of decision variables, with that due to the choice of model domain parameters arising from the nature of LCA. Thus, these are common to all three decision types, and the discussion below only relates to model parameter uncertainty arising from decision variables. Similarly, uncertainty in model form largely arises from the nature of LCA, although the degree of sophistication of the LCA model chosen is constrained by the decision context, particularly around impact assessment (see section 2.3).

Three broad decision-types are identified in section 2.3. In Type I systems there is no uncertainty related to the definition of the system, as these are purely descriptive systems. Model parameter uncertainty is consequently low. Empirical uncertainty encompasses two distinct sources, that of the data used to describe the processes, i.e. an indication of what compromises have been taken in data collection, and the variability across the processes, i.e. how broad a definition the Type I system covers (e.g. technology-wide, sector-wide or regional production etc.). For the former source of uncertainty,

the acceptable variance is free to be set according to the scope of the study, and since Type I systems do not explicitly involve a comparative analysis (i.e. the inventory is not generated with a specific comparative system in mind, although they may ultimately be used as part of a comparative analysis), the choice of acceptable variance is purely subjective. As LCI uncertainty assessments become more commonplace, a consensus on the acceptable uncertainty for various LCI applications is likely to be established. The latter source of uncertainty is intrinsic to an averaged system, and the variability of the particular system will set the limit variance. In this case, a more meaningful breakdown of the system boundary (e.g. into sub-regions) would be the only way in which the variability could be reduced. The acceptable or limit variance for the system will depend on which source of uncertainty dominates the assessment.

The uncertainty in strategic or tactical decision systems (Type II) is expected to be dominated by the choice of decision variables. The focus of the analysis is therefore on model parameter uncertainty, with a consequently large variance acceptable in the empirical uncertainty. Strategic decision systems necessarily reflect future systems, so are inherently associated with high data uncertainty. This also supports setting a high acceptable variance limit, since the quality of the available data is likely to limit the degree to which the variance can be reduced. For particularly poor data quality, it may be shown that no statistically significant differences can be shown between the options, unless the options are redefined (i.e. their performance made to move further apart).

In operational decision systems (Type III) an assessment of model parameter uncertainty can usually be limited to the few key parameters controlling the decision, and the emphasis of the uncertainty analysis thus falls on empirical uncertainty. As with Type I systems, because operational studies are applied to existing systems, the data exists to be collected and the acceptable uncertainty in the results is able to be based on the scope of the study (i.e. on how much non-specific data is able to be tolerated), whilst the system's inherent variability determines the lowest possible variance. However, unlike Type I systems, operational studies involve comparisons. The acceptable variance in the output is thus based on the degree of confidence required in the differences between the options, which is dependent on the particular consequences of the decision, e.g. those with high environmental or financial risks will require confidence levels in excess of 90%.

4.2. INVENTORY MODELLING TO INCLUDE UNCERTAINTY

This thesis aims to promote uncertainty analysis as an integral part of the LCI model, and not as an add-on to a completed inventory study. The implications that this has on inventory modelling are covered in this section.

4.2.1. Modelling with Variable and Correlated Inputs

In a probabilistic uncertainty analysis a parameter is assigned a range of values, instead of a single value. The parameter is then assigned a value from this range at random, in accordance with the probability that this value is likely to occur. A different value of each uncertain parameter is taken each sample, and the output calculated for each set of values. This random sampling can lead to nonsensical results for correlated variables, i.e. it results in combinations of parameter that are not feasible. The problem of correlations between the input data elements has been identified by Kennedy et al. (1996) and Maurice et al. (2000). Both sets of authors recognise that independently varying certain data elements will result in a violation of the mass balance. Kennedy et al. (1996) recommend that any correlations between the input data elements be specified and incorporated in the model, although they recognise that this would considerably increase the modelling complexity. Similarly, Maurice et al. (2000) recommend that advanced models should be able to represent the interaction between the parameters. They suggest that correlation analyses be used to highlight dependencies between the data and bring to light possible modelling errors.

A probabilistic uncertainty analysis is therefore not applicable to a “black box” modelling approach. In such models nothing is known about the interactions between the variables. Hidden correlations abound, and nonsense is likely to result by varying all variables independently of each other. Deconstructing the model is the best way to avoid correlations. At sufficiently detailed levels, apparent correlations can be broken down into the individual variables and the relationships between them. However, a balance needs to be found between removing the correlations and increasing the number of variables and complexity of the model. From this study it has been found that, at a minimum, emissions need to be related directly to the flow responsible for their generation. Deconstructing certain correlations may be beyond the scope or understanding of the modeller. In such cases conditional probability distributions can be used to specify probabilistic dependencies between variables. These

specify the variable, given the value of the correlated variable. Continuous conditional probability distributions are extremely complex to specify, so discrete conditional probability distributions are recommended, although these result in a loss in accuracy due to discretising the distributions.

Another problem with randomly varying input data is that unrealistic extreme values can be generated, especially since certain distribution shapes tend to exaggerate the extremes of the distribution. In practice, it is not feasible to characterise all correlations, and unlikely combinations of variables will inevitably occur. It is therefore important to check that anomalous results are not being generated at the extremes. Internal constraints (e.g. minimum flow rates) can be specified to prevent these sample runs from being included in the final output sample. The importance that can be attached to the extreme values is thus dependent on how well inconsistent runs have been excluded from the final output sample.

In addition to avoiding infeasible combinations of variables, inventory modelling at a high level of resolution is also required for meaningful data quality characterisation. This enables more relevant probability distributions to be defined, as the probability distribution can be determined from the actual data sample and not arbitrarily assigned to a higher level of aggregated data. However, modelling at a high level of detail is time and energy intensive. The system is therefore broken down into a foreground and background system, with only the foreground processes modelled in the detail required for meaningful probabilistic assessment. The breakdown into foreground/background systems is in line with current thinking on inventory modelling (Clift et al., 1998; Trinius and Le Téo, 1999; Weidema et al., 1999) (see section 2.1.3). The background processes are modelled at an high level of aggregation using generic LCI data, with uncertainty factors applied to the single-point inventories according to the type of environmental intervention (see following section). Although this involves the problems associated with “black-box” modelling stated above, this method is necessary to enable the use of generic LCI data, and because of the necessary simplification in LCI modelling this allows.

4.2.2. Characterising Uncertain Parameters

The crux of an uncertainty assessment is characterising the uncertainty of the input data, as the assessment can only be as relevant as the information from which it is comprised. This requires the definition of a representative range or probability distribution of possible values the parameter could have. The first step in the characterisation of parameter uncertainty is to classify the parameters

according to whether they are empirical parameters (chance variables) or model parameters (decision variables). For some parameters, the delineation depends on the decision-maker's perspective. Decision variables are those parameters whose value is directly under the control of the decision maker, whilst the empirical parameter values are independent of the decision maker, e.g. in a study involving commissioning a new power plant, the size of the generating units will fall under the control of the decision maker, whilst in a study looking at regional average power production, the unit size is an empirical parameter.

4.2.2.a Empirical Parameters

Empirical parameters or chance variables make up the majority of quantities input into models, and are responsible for the familiar "data uncertainty" aspect of uncertainty assessments. The manner in which the empirical parameters are assessed depends on whether they are applied in the background or foreground system. Little is known about the source or quality of the data used in the background system, so only standard uncertainty ranges can be applied. However, processes falling into the foreground system are modelled in detail using data from a variety of sources. This allows a more detailed investigation of the types of uncertainty giving rise to a parameter's overall uncertainty, and each data element requires assessment against a set of data quality criteria. The procedures used to quantify both foreground and background empirical parameter uncertainty follow below.

Foreground Data

As with most LCI models, the foreground inventory models used in the case studies in this thesis are built up using a variety of data sources. Routinely monitored process data are used wherever possible, although this is not available or relevant for all aspects of the system, and recourse has to be made to other data sources. High quality data, such as data from on-site measurements and detailed information from company-specific records, are frequently not available, and it is necessary to rely on previous studies, databases, literature sources and handbooks to some degree. This data is often old, imprecise and poorly documented, and qualified guessing often has to be used to fill in gaps and to select and adjust data from different sources. Even where highly applicable process data are available these are subject to a certain degree of natural variation and measurement error. Very different degrees of uncertainty are thus present in the different sources of data, and different representations are applicable for their uncertainty characterisation.

An understanding of the various sources of empirical uncertainty is required to ensure all aspects of a parameter's uncertainty and variability are taken into account (i.e. to choose a necessary and sufficient set of DQIs against which the data can be assessed). Section 3.1.1.b describes these sources, which are summarised in Figure 3-1 and the accompanying discussion. Whilst some aspects of data uncertainty are independent of the particular study in hand (i.e. those associated with its acquisition), others depend strongly on the goals of the study (i.e. the applicability of the data). In Figure 3-1 the study-independent sources of uncertainty are related to the parameter uncertainty, i.e. measurement errors, approximations, inherent randomness and subjective judgement, whilst the study-dependent uncertainty is related to the variability in the parameter, since this gives rise to uncertainty when the data is applied in a different context. Variability in a parameter not related to its application (i.e. variability arising even when the parameter is applied in the context in which it is measured) is incorporated into "inherent randomness". "Subjective judgement" refers to the choice of quantity measured to reflect a particular parameter, and not to the judgement applied in using this data in a particular study.

The uncertainty associated with the acquisition of an empirical parameter (i.e. the study-independent uncertainty) is predominantly a function of its source, with the uncertainty increasing as the source moves from the specific to the general (see the "reliability of source" DQI in Table 3.7). In addition to the source of the data, the completeness of the data sample is an important contributor to the uncertainty (i.e. the size of the sample and the duration over which it is collected). The source and completeness of the data sample are sufficient DQIs to fully characterise the study-independent component of data uncertainty (Weidema and Wesnæs, 1996). These together encompass the sources of uncertainty identified as giving rise to parameter uncertainty in Figure 3-1, where inherent randomness is incorporated under the "completeness" indicator, whilst measurement errors, subjective judgement and approximations are incorporated under the "source" indicator.

The second component of data uncertainty, that of its adequacy with respect to the context of the study, tends to be overlooked, as its characterisation is inherently subjective. This is of particular concern to LCI models, since their high data requirements often require that readily available data be adapted for use in the LCI to streamline data collection. Often over-looked in an assessment of uncertainty, is that well characterised data (i.e. many measurements) does not necessarily mean low contribution to uncertainty, since well characterised data can introduce significant uncertainty when applied in a different context to that for which the data sample is representative. The use of actual data uncertainty

(i.e. statistically measurable uncertainty estimates) is only relevant where the data sample takes into account all underlying sources of variability. However, in most cases, the sample is not representative, and the missing sources of variability need to be incorporated into the uncertainty estimate via subjective judgement. Although this uncertainty “add on” is inherently subjective, its incorporation is judged far more meaningful than leaving off the potentially limiting source of uncertainty. A good understanding of the context of the study is required to determine the degree of uncertainty that should be assigned to this data (e.g. its geographical features, so that the degree of uncertainty introduced by applying data from other geographical locations can be estimated).

An assessment against the DQI set developed by Weidema and Wesnæs (1996), as laid out in their pedigree matrix (see Table 3.7), is considered the most workable method. The five DQIs defined by them incorporate all sources of empirical parameter uncertainty listed in Figure 3-1, whilst the structured approach and guidelines minimise the subjectivity of the method. The subjectivity, can however never be eliminated, and the “score” should not be seen as objective, but rather as a representation of a subjective judgement of the data quality (Weidema, 1998b). An important feature of the pedigree matrix DQI set is that the indicators can be determined independently of each other, thereby allowing the influence of each indicator to be additive, and an estimate of the overall uncertainty to be determined. A quantitative estimate of the uncertainty, introduced by a poor score on an indicator, is made by assigning a coefficient of variance (CV) to each indicator score, where the CV is the ratio of the standard deviation to the mean. An overall CV is then calculated for the parameter (the square root of the sum of the squares of the individual coefficients). The pedigree matrix is thus used to structure the quantitative estimate of otherwise qualitative uncertainty elements, enabling the incorporation of uncertainty or variability that is not exhibited in the data sample itself.

The data quality scores do not represent a “quantity” of uncertainty, and it is only their influence and not the scores themselves that are additive. This means that estimates of variance must be assigned to each indicator, and not their sum. This also means that it is not meaningful to pre-assign CVs to each pedigree score. The score assigned will depend on the nature of the specific parameter under consideration. For example, two parameters may each score badly on technological variability, i.e. data on each are obtained from a different process to the one under study. However the one is known to exhibit high variability across technologies and is assigned a CV of 60%, whilst for the other, variability across technologies is known to be small, and it is only assigned a variability of 10%, even

though both had a DQI score of 4. The CV can therefore not be assigned in isolation, and requires “expert” judgement and knowledge of the application system to be meaningful. The estimate of CV is informed by the variance in the data sample, and relevant literature data and samples of similar systems. A conservatively large CV is chosen where data is not available to support the estimate.

The slightly different method to that proposed by Weidema and Wesnæs (1996) is used to extrapolate the pedigree score to an estimate of total variance. Their distinction between “basic” uncertainty, which they define as that which is statistically measurable in the data sample, and “additional” uncertainty, which they infer from the pedigree matrix scores, is not retained. It was felt that this distinction is not always clear, and can lead to double-counting if the data sample is not well defined. A slightly different approach is therefore taken here, in which the DQIs are seen to represent the full uncertainty of the quantity, i.e. to incorporate all possible sources of uncertainty (both “basic” and “additional”). The “basic” or measurable component of uncertainty, as defined by Weidema and Wesnæs (1996), is seen to be incorporated primarily under the “completeness” indicator (incorporating natural fluctuations and variations over time, evident as daily, hourly or, monthly variation in the data sample) or the “reliability of source” indicator (incorporating measurement errors). The variance of the data sample (i.e. the “basic” uncertainty) is thus not evaluated separately and combined with the variance estimated from the pedigree matrix (i.e. the “additional” uncertainty), but rather used to inform the estimate of variance due to each indicator, where relevant. The estimates of the variance introduced by each DQI are then summed to get an estimate of the total variance. The sum of the variance due to the set of indicators can not be less than that measured for the data sample, and in most cases is much larger (i.e. the data sample very rarely incorporates all sources of uncertainty).

Once the variance of the parameter has been specified, an appropriate probability distribution has to be assigned. A range of distribution shapes is required to characterise the various types of data encountered in an LCI study. A brief description of appropriate probability distributions, and the parameters required for their definition, are given in appendix A.2. The shape of the distribution is chosen according to the amount of information known about that parameter. Where a sufficient number of data samples exist, the shape of the distribution is inferred from a frequency plot of the data, i.e. a frequency table of the data is compiled, and the resultant histogram plotted. An appropriate distribution shape is estimated by eye to fit the frequency plot (e.g. whether the data is normally distributed, positively or negatively skewed etc.). The relevant distribution parameters are calculated from the data

sample (e.g. median, mean, minimum and maximum values etc), and the estimate of variance for the parameter.

The uncertain parameters requiring characterisation in LCI studies are practically all constrained to be positive. The normal distribution is therefore generally found to be inappropriate as it can produce negative samples, and the Gamma distribution is used to approximate normally distributed quantities instead. The Gamma distribution is found to be a useful distribution type, as it generates only positive samples, and has great flexibility of shape, ranging from an exponential distribution to approximating a normal distribution. It is also straightforward to specify, as its parameters are easily related to the mean and variance (see appendix A.2). The lognormal distribution also generates non-negative samples, and is thus a useful distribution shape for positively-skewed quantities. However, the Gamma distribution is generally preferred for positively-skewed quantities as well, as it prescribes a lower probability to the extremes of the distribution than the lognormal distribution (it consequently lessens the generation of unlikely extreme values). A disadvantage of the Gamma distribution is that it is more computationally intensive to generate samples from, and can thus potentially increase the computation time of the model. The beta distribution is found to be a useful distribution where the values need to be contained to a specific range, and is thus especially suited to parameters expressed as percentages. However, the parameters of this distribution are less easily related to the mean and variance, and thus less straightforward to specify. It is therefore not used unless a contained range is specifically required.

These distributions and the method described above is suitable for parameters defined by process monitoring data, i.e. for which data is available over a number of months or years. For parameters described by few literature values or approximate factors, a uniform or triangular distribution is more appropriate. A uniform distribution is used if only the data range is known or estimated, whilst if a most probable value is also known, or a strong central tendency is expected, a triangular distribution is used. In rare cases, a uniform or triangular distribution may be known to be inaccurate (i.e. there may be some knowledge of the likely distribution shape, although only a few data values are known). In such cases, the more likely distribution shape can be applied instead, albeit unsupported by actual data, e.g. the uncertainty associated with the percentage of discard dumps burning, where there is a high probability of zero, but a small probability of a high value. A uniform distribution would be meaningless, and a highly negatively-skewed beta distribution is used instead.

according to whether they are empirical parameters (chance variables) or model parameters (decision variables). For some parameters, the delineation depends on the decision-maker's perspective. Decision variables are those parameters whose value is directly under the control of the decision maker, whilst the empirical parameter values are independent of the decision maker, e.g. in a study involving commissioning a new power plant, the size of the generating units will fall under the control of the decision maker, whilst in a study looking at regional average power production, the unit size is an empirical parameter.

4.2.2.a Empirical Parameters

Empirical parameters or chance variables make up the majority of quantities input into models, and are responsible for the familiar "data uncertainty" aspect of uncertainty assessments. The manner in which the empirical parameters are assessed depends on whether they are applied in the background or foreground system. Little is known about the source or quality of the data used in the background system, so only standard uncertainty ranges can be applied. However, processes falling into the foreground system are modelled in detail using data from a variety of sources. This allows a more detailed investigation of the types of uncertainty giving rise to a parameter's overall uncertainty, and each data element requires assessment against a set of data quality criteria. The procedures used to quantify both foreground and background empirical parameter uncertainty follow below.

Foreground Data

As with most LCI models, the foreground inventory models used in the case studies in this thesis are built up using a variety of data sources. Routinely monitored process data are used wherever possible, although this is not available or relevant for all aspects of the system, and recourse has to be made to other data sources. High quality data, such as data from on-site measurements and detailed information from company-specific records, are frequently not available, and it is necessary to rely on previous studies, databases, literature sources and handbooks to some degree. This data is often old, imprecise and poorly documented, and qualified guessing often has to be used to fill in gaps and to select and adjust data from different sources. Even where highly applicable process data are available these are subject to a certain degree of natural variation and measurement error. Very different degrees of uncertainty are thus present in the different sources of data, and different representations are applicable for their uncertainty characterisation.

The simple method of estimating the general distribution shape from a frequency plot of the data, and using the calculated mean or median and estimated variance to define the distribution, is judged sufficient for the level of accuracy of the study, i.e. a more accurate definition would be inappropriate given the rough definitions of the less well specified parameters. A more rigorous method could be applied to get a more accurate determination of the distribution shapes, should this be justified by the scope of the study. For example, custom fit distributions can be fitted to the data samples, although these are cumbersome, and more computationally intensive. However, their use is justified in rare cases where none of the standard distribution shapes are applicable (e.g. a bimodal distribution). Also, a more statistically justifiable method of updating the sample distribution shape with the information from the pedigree matrix could perhaps be warranted. For example, Bayes' rule could be used to combine the sample probability distribution (measured) with that from the pedigree matrix (estimated). However, this was deemed beyond the scope of this study, and the distribution is merely adjusted to reflect the increased variance, whilst the mean and probability distribution shape is kept the same.

Estimating a relevant probability distribution shape and range for each parameter is a time consuming process. The model is therefore initially run with the parameters roughly characterised by a conservative uniform or triangular probability distribution, regardless of whether that parameter is well specified or not. The results of the uncertainty importance analysis are then analysed and those parameters contributing significantly to the overall uncertainty are identified. The method outlined above is then applied to these parameters, i.e. a CV inferred from scoring the pedigree matrix and an appropriate relevant distribution shape assigned. Subsequent model runs are used to refine the data characterisation further, or to highlight where better data needs to be sourced. This process is shown schematically in the context of the case study in Figure 7-2.

Background Data

Generic LCI data from published LCI databases are used for the background processes. The data quality of cumulative inventory results is only able to be assessed if provided by the respective database, which is typically not the case for LCI databases. The published source of the databases (PEMS LCA software) used in this study did not contain such uncertainty information, so literature estimates of inventory data uncertainty had to be used (De Smet and Stalmans, 1996; Finnveden and Lindfors, 1998; Hanssen and Asbjørnsen, 1996). These papers broadly agree with each other, in that all recognise very high variance in literature data. Approximate ranges for three broad classes of inventory elements;

material resources, fossil fuels and emissions resulting directly from their use, and all other emissions, are adapted from Finnveden and Lindfors' (1998) "rules of thumb" for expected variations in LCI data. An uncertainty factor of 10, 2 and 100 is used for each class respectively. Each inventory item is then specified as a uniform distribution, with the minimum value taken as the mean divided by the uncertainty factor, and the maximum value as the mean multiplied by the uncertainty factor. The specification of such a rough distribution and increase in mean value, is aimed at inflating the uncertainty importance of the background data, as there is considerable uncertainty in applying European data to South African processes and conditions.

Even with these large uncertainty ranges, the background data is found not to dominate the uncertainty in the case studies. The limited information available means that the background empirical parameters are not able to be iteratively refined in subsequent model runs, other than replacing the uniform distribution with a lognormal distribution where this is seen to be warranted. The background processes therefore introduce a certain amount of irreducible uncertainty into the assessment. Apart from obtaining more detailed information on the database, the only way in which the uncertainty could be reduced further would be to bring the process into the foreground and model it in detail.

4.2.2.b Model Parameters

The assessment of model parameters is separated from that of empirical parameters, as they can not meaningfully be analysed probabilistically. Model parameters are broken down into decision variables, model domain parameters and value parameters. Value parameters and decision variables are considered together here, as the same principles apply for their assessment. The treatment of uncertainty in both model domain parameters and decision variables is essentially the same, as both are analysed in a parametric assessment. Their differentiation is thus primarily a function of their different roles within the model.

Decision Variables

Decision variables are those parameters, specifying the scenario, which fall directly under the control of the decision maker. Their most appropriate choice is often not known, and a range of possible values therefore requires investigation. The feasible operating range of the parameter is first identified. The number of samples to be taken over this range then needs to be decided. A very influential parameter would require the range to be sampled at frequent intervals (e.g. 10% intervals), whilst minimum,

“most likely” and maximum values may be all that is required for less influential parameters. Not all decision variables can be defined as feasible ranges, and instead require a discrete choice that involves the choice of a number of related parameters, e.g. wet or dry ash disposal. In these cases, the choice of one parameter determines the choice of others, and a series of operating “states” (groups of dependent parameters) are defined. If a system is governed by a relatively small number of model parameters, the effect of their operating range can be assessed in a systematic parametric analysis (i.e. one parameter “stepped through” its range (or states), with the others held constant, until the model output has been computed at all possible combinations of the parameters).

However, in many cases, an assessment of all possible combinations of decision variables quickly becomes unmanageable. A full parametric analysis is therefore not attempted, and the parameters are first “screened” with a sensitivity analysis. The model is run with each parameter at its minimum and maximum values and the spread in results caused by the change assessed. If the parameter varied to its extremes causes little increase over the empirical variance in the results, i.e. if the increase in variance is small compared to the empirical variance, the effect of the parameter is considered negligible, and all future runs are conducted with it set to its “most likely” value. The sensitivity considered significant will thus depend on the particular study (see section 4.1.3.c), e.g. in the case study in chapter 7, a parameter is considered significant if the median of the output sample, calculated with the parameter at either extreme (i.e. the minimum and maximum values), lies outside the 50% confidence interval of the sample calculated with the most likely value; and less significant (but still retained in the analysis) if the median value of the extreme output sample changes that of the most likely sample by more than 10%, but there is a large overlap between the samples (the median of the extreme sample lies within the 50% confidence interval of the most likely sample).

Variables with sensitivities lower than these criteria are excluded from the parametric analysis, unless it is suspected that they could have a significant effect when coupled with certain other parameters at their extremes. In this case their cumulative effect can be checked by calculating the outcome at a few combinations of parameters that, through a qualitative understanding of the system, are known to potentially lead to extreme situations. The less significant model parameters are either set at their most likely value for all remaining analyses, or if the nature of the parameter allows, are incorporated into the probabilistic assessment (whether a variable is defined as a decision variable or an empirical parameter depends purely on the perspective of the decision maker, and thus, if a parameter can be considered to

no longer be defined as a “controlled” variable, it can be treated probabilistically). Those parameters shown to be important are retained for the parametric assessment. As before, a number of operating states spanning the operating range are defined (low, mid and high values are chosen as the initial set), and the model outcome computed at all possible combinations of these variables. A scenario tree can be used to facilitate the process of combining the parameter operating states into the scenarios at which the system is evaluated.

In studies with high model parameter uncertainty (a large number of significant decision variables), even after the less-significant parameters have been discarded, a strict scenario tree analysis quickly becomes unmanageable as an explosive number of scenarios arise. To keep the number of runs to a manageable number, the parameters can be grouped to predict “worst case” and “best case” scenarios, which are defined using either a qualitative knowledge of likely system operating states, or strictly from the sensitivity analysis. These represent the extremes of the system, with all other combinations falling between them, and thus represent the full range of possible outcomes. A problem with this streamlining approach is that if a single indicator method is not being used, a change in a parameter value will not necessarily cause the scenario to be “worse” or “better” in all selection criteria (i.e. in each impact category or environmental intervention considered). Either “best” and “worst” parameter states for each criterion can be defined (although for a large number of criteria this may result in an equally unmanageably number of scenarios), or a priority order of the criteria has to be established. In this case, if a change in the parameter value is found to cause some criteria to perform worse, and others better, it is assigned to the “best” or “worst” scenario according to how it performs against the highest priority criterion. If there is no change against this criterion, the criterion with the next highest priority is used, and so on.

Model Domain Parameters

Model domain parameters specify the particular model form. Their distinction from choices of model form or structure is not always clear, as it is the choice of model form that results in a set of model domain parameters requiring definition. Their distinction comes in their assessment, as model domain parameters are varied parametrically (i.e. varied together with other model parameters), whilst the choices of model form can only be treated in a sensitivity analysis (i.e. a top-level assessment). The various possible “states” of the model domain parameters are first identified (e.g. daily, monthly or annual variability). In many cases, the possible values the parameter can have will be limited by the

scope of the study (i.e. the level of detail applicable). Where the parameter can have many values (e.g. time horizon considered), the range is broken down into applicable parameter “states” (e.g. short, medium and long).

The assessment of the uncertainty in model domain parameters is less intuitive than that of decision variables, and is thus often neglected. In LCA studies, guidelines and “common practice” dictate the choice of the most appropriate model parameters, and the effect of this choice is generally not considered. Where the guidelines are rigorously researched and applicable to the study context, a set of “best choice” parameters could remove the need for an assessment of model domain parameters.

4.2.3. Analysis of Uncertainty

The aim of the uncertainty analysis is twofold: firstly to calculate the variance in the output (i.e. the certainty with which decisions can be supported), and secondly, to provide a mechanism by which the variance in the output can be reduced. The latter is achieved by highlighting the key parameters contributing to the variance in the output. A powerful measure of uncertainty importance, compatible with the simulation methods of uncertainty propagation, is the correlation of the sample output values with the corresponding sample of rank-ordered input values for each variable (see section A.1.2). The parameters displaying the highest correlation coefficients are those contributing most to the variance in the output.

The correlation coefficients can therefore be interpreted as a measure of the variables “uncertainty importance”, and allows a prioritised list of input parameters to be drawn up, with the parameters at the top of the list responsible for a high proportion of the overall uncertainty. It therefore makes sense to address the definitions of these parameters first, as decreasing their variance will have the largest effect on decreasing the variance in the results. This prioritisation of input parameters makes the iterative analysis of empirical parameter uncertainty workable. The initially rough distributions of the parameters returned with high uncertainty importance are refined until the best possible distributions within the constraints of data availability are defined, whilst no additional effort is expended on parameters that have low uncertainty importance (i.e. the rough initial distributions are retained for the parameters falling low down on the ordered list of rank-order correlation coefficients).

Calculating rank-order correlations is extremely computationally intensive, so to decrease the run time it was found better to do duplicate runs with smaller sample sizes, than a single run with a high sample size (a sample size of 700 is found to be sufficient in the case studies). As highlighting the top few variables is all that is required by the analysis, and not the value of the correlation coefficients as such, a smaller sample size does not introduce a loss in accuracy. However, the duplicate run is required because a problem with the analysis, particularly at lower sample sizes, is that co-incidental variance between the parameters is returned as important, i.e. purely random and not deterministic co-variance is picked up between a parameter and the output. By repeating the analysis with a different set of random samples (i.e. re-initialising the random seed of the random number generator between runs), the deterministically co-varying variables can be identified as those returned by both analyses.

The acceptable variance is chosen through a consideration of the decision context and the consequent goals and scope of the study (i.e. the time and resources that can be spent on model development and data collection, and the certainty required in the results). In practice this is likely to be decided upon by first reaching the minimum variance achievable with the data in hand (and thus the confidence with which the decision can be supported), and the decision then taken whether additional data collection is warranted. The uncertainty importance analysis is able to direct this additional data collection effort by ranking the data inputs according to their contribution to the uncertainty, and thus prioritising the data collection.

Ultimately the irreducible variance is reached, where the variance cannot be reduced further without significant changes to the model form, e.g. if an approximate factor or “black box” element is returned with high uncertainty importance, the variance in the output can not be reduced without modelling that element in greater detail. For the background processes, this is analogous to moving the identified parameter into the foreground system, and a disaggregated inventory modelled for the process. Breaking down approximate factors into the variables and relationships underlying them introduces additional uncertain parameters and increases the size and complexity of the inventory models. Increased modelling detail may therefore not always achieve the desired reduction in output variance, or be feasible within the scope of a study. In such cases, the goals of the study will have to be revised instead, to allow for a higher level of variance to be tolerated in the results.

In addition to highlighting where greater modelling depth is required, the uncertainty analysis may allow certain elements of the study to be cut, by identifying the unimportant parameters (i.e. those falling low down on the ordered list of correlation coefficients). This may be valuable where simulation time is important. Cuts may be applicable to both the environmental intervention list and sub-processes. Where the study is being assessed on the results of an impact assessment procedure, a fairly trivial observation is that all environmental interventions not directly related to an equivalency factor may be omitted. If no data inputs for certain unit operations or sub-processes are identified as significant, the entire sub-process may be omitted from the uncertainty analysis (i.e. modelled as point estimates), although it must be borne in mind that elements appearing insignificant in a comparative assessment (due to them being common to the two systems being compared) may be of consequence when the system is taken alone.

4.3. INTERFACE WITH IMPACT ASSESSMENT METHODS

Uncertainty assessment in LCA has tended to focus on the inventory phase. However, the more recently developed LCIA methods have a far greater emphasis on uncertainty, both in terms of qualitatively discussing shortcomings and providing quantitative estimates of data uncertainty (Goedkoop and Spriensma, 1999). Extending the inventory assessment to impacts increases the relevance of the assessment (basing the assessment on environmental interventions has higher valuation uncertainties than basing the assessment on potential impacts). However, as the relevance of the indicator increases (i.e. as it moves from environmental interventions to a prediction of actual site-specific impacts), so the uncertainty in the results increase, from the high data and model uncertainties associated with calculating the equivalency factors.

The more accurate and representative the characterisation models (e.g. those including site-specific fate, exposure and effect factors), the greater the number of data inputs required, and generally the greater the associated data uncertainty. Thus decreasing model uncertainty is to some degree offset by increasing empirical uncertainty. Whilst this empirical uncertainty is relatively straightforward to estimate (the same methods are applicable as the assessment of empirical uncertainty in inventory models), the increased model complexity and greater number of uncertain inputs in characterisation models has meant that their application has been less frequently attempted. To relevantly determine uncertainty ranges for the equivalency factors, incorporating both model and data uncertainties, requires

an intimate knowledge of the ecological and toxicological models on which they are based. Since the emphasis of this thesis is on capturing inventory uncertainty, this is beyond the scope of this study. However, it is recognised that it would be meaningless to rigorously include inventory data uncertainty whilst completely disregarding impact assessment uncertainty. Uncertainty estimates for the equivalency factors are therefore obtained from the literature, and a brief review of studies addressing impact assessment uncertainty is presented in appendix C.

The Eco-indicator 99 method is used in the case studies in this thesis, as it is judged to be representative of the current state-of-the-art in impact assessment development, and since it has a particular focus on addressing the uncertainties inherent to LCIA methods (Goedkoop and Spriensma, 1999). Details on the method can be found in appendix C. Quantitative estimates for the equivalency factors are not yet comprehensive, so estimates and extrapolations had to be made for some categories, whilst for mineral and fossil fuel resource depletion, no uncertainty could be included (see appendix C). The uncertainty arising from impact assessment is simply combined with the assessment of inventory uncertainty by modelling the equivalency factors with a lognormal probability distribution.

CHAPTER 1

INTRODUCTION

The aim of this thesis is to develop an approach to support prospective environmental decision-making in resource-based industries. The particular objectives of the approach are that it be able to adequately reflect the environmental burdens arising from primary industries, and to make explicit the trade-offs often encountered in environmental decisions. In addition, it needs to be able to structure large, disparate data sets of varying quality and completeness into useful information able to provide the environmental objective in a decision-making process. Life cycle assessment (LCA) has received increasing attention for its role in environmental decision-making processes, where it supports the process of defining the contribution of human activities to sustainable development (Cowell, 2001). The efficacy of LCA to provide a basis for environmental decision-making in resource-based industries is explored in this thesis through case studies in the context of coal-based power generation.

Primary industries involve vast consumption of resources, generally have a large spatial footprint, and generate significant volumes of solid waste and liquid effluent. A comprehensive management strategy is therefore essential to minimise the impacts of these industries on the environment. In particular, a consistent framework for technology selection is necessary, recognising that the context in which the assessment takes place has important consequences on data availability and quality, and consequently the certainty with which technological systems can be evaluated. The particular focus of this thesis is thus on technology selection and design in primary industries. The need to include environmental considerations into process design has increasingly been recognised, and tools and methods have been developed to incorporate environmental objectives into the usual economic optimisation of chemical processes (Azapagic and Clift, 1999a; Diwekar, 1999; Pistikopoulos et al, 1994; Rossiter, 1994; Shonnard, 1999; Sorin and Paris, 1997; Spengler et al., 1998; Stefanis et al., 1995; Stewart and Petrie, 1996).

The chapter starts with a brief overview of environmental decision making and LCA. The features of resource-based industries are then discussed, as these give rise to the particular considerations of the thesis. This chapter concludes with an introduction to the topics to be covered, and guidelines as to where in the subsequent chapters they will be discussed.

4.4. SUMMARY OF THE UNCERTAINTY ANALYSIS PROCEDURE

Figure 4-3 summarises the main steps of the procedure developed in this chapter to allow a systematic assessment of all sources of uncertainty in LCI models. The process is briefly described below to bring together the different elements of this chapter, and provides a summary of the procedure followed in the case studies presented in chapters 7 and 8.

The first step is to identify all quantities input into the LCI model, and to characterise them according to their variable type. Variables are classified broadly into empirical parameters and model parameters. The empirical parameters are assigned probability distributions, according to the data available to characterise them, whilst the feasible operating range and most likely value is defined for each model parameter.

The model is first run for a particular model form, with the model parameters set to their “most likely” values and the empirical parameters defined by simple probability distributions. Latin Hypercube sampling, with a sufficient number of samples to meet the specified accuracy, is used to calculate the output. The mean and variance of the results are analysed, and correlation coefficients computed to determine the uncertainty importance of each empirical parameter. The most influential parameters are identified, and the definition of their probability ranges revisited. Where possible, the variance of each influential variable is decreased through a better definition of their uncertainty (more meaningful distribution shape and range assigned). However, if the variable is already assigned a well-fitting distribution accurately reflecting the source of the data, increased data collection and investigation is required. The effect of the adjustments on the variance of the results is checked, and the “tuning” process continued until the acceptable variance is reached, or until the variance can no longer be reduced within the scope of the study (i.e. significantly increased data collection or modelling detail is required to further reduce the variance in the output). Once an estimate of the empirical variance in each selection criterion has been obtained, an idea of the magnitude of the difference between the options required to distinguish a significant difference between the options can be obtained, i.e. the difference between the options must, at least, be significantly greater than that due to the background variance of the system for it to be stated with any degree of confidence that the one system is preferred over the other.

Figure 4-3 Step by step plan of uncertainty analysis of model and empirical parameters.

<ol style="list-style-type: none"> 1. List all parameters input into model. 2. Sort parameters according to parameter type. 			
<u>Model parameters</u>		<u>Empirical parameters</u>	
3. Split into model domain parameters and decision variables.		3. Split into parameters applied in the foreground and the background systems.	
<u>Model Domain Parameters</u>	<u>Decision Variables</u>	<u>Foreground parameters</u>	<u>Background parameters</u>
4a. Define relevant model "states".	4a. Define feasible parameter range.	4a. Obtain set of qualitative data quality scores according to Table 3.7.	4a. Define uncertainty factor according to type of background parameter.
	4b. Define states (e.g. minimum, most likely and maximum) or intervals at which to sample the range.	4b. Estimate variance due to each data quality score by assigning a CV to each indicator.	4b. Calculate minimum and maximum values for uniform distribution, or "order of magnitude" change for lognormal distribution.
		4c. Calculate overall variance, checking that this agrees with what is known from the data sample.	
		4d. Decide on relevant probability distribution shape.	
		4e. Define distribution parameters to adequately specify the distribution.	
5. Run sensitivity analysis to determine influential model parameters.		5. Run probabilistic simulation to propagate uncertainties to output.	
6. Combine influential parameters into likely scenarios, or combine parameters using scenario tree to construct all possible parameter combinations.		6. Analyse results (correlation analysis) to identify influential parameters, and repeat above analysis with better distribution definitions and/or better data for these parameters (if possible).	
7. Compute model outcome at all possible scenarios, and evaluate model parameter uncertainty in light of empirical parameter uncertainty, refining the scenarios, if necessary.			
8. Run sensitivity analysis on possible model forms, repeating the empirical and model parameter assessments, as required.			

The next step is to evaluate the possible range in model output arising from an unsure definition of the system, i.e. uncertainty about the appropriate definition of the decision variables and model domain parameters. The feasible operating range of the parameters first requires definition, and possible operating states across this range are selected. These are either assessed systematically by computing all possible combinations of the selected states of each model parameter, or in a streamlined assessment, where a sensitivity analysis is first used to discard the less significant parameters, and the operating states of the remaining parameters are grouped into possible operating scenarios. The empirical variance of the model output is checked with each model parameter “combination set”, and if found to be necessary (i.e. if the variance changes significantly from the “most likely” case), the correlation coefficients are computed, the influential parameters identified, and the definition of their probability distributions investigated.

The outermost layer of the uncertainty assessment evaluates the sensitivity of the model outcome to a change in model form or structure. If the change is seen to have a significant effect (e.g. if it changes the preferred ordering of the options), both sets of results must be presented and discussed. For example, if a change in the allocation method is seen to change the outcome of a scenario, results from both allocation methods must be presented, and the discussion motivate why one may be considered more valid than the other. However, probably the most significant model uncertainties are those inherent to the LCA method, e.g. the spatial and temporal limitations, which can only be assessed qualitatively because no clear scenarios can be defined to assess them in a sensitivity analysis (Meier, 1997). It must therefore be remembered not to be over confident when interpreting LCA results, even those covered by full uncertainty analyses, as the probabilistic output, or “certainty” ranges in the results, can not represent all sources of uncertainty.

CHAPTER 5

IMPACT ASSESSMENT OF RESOURCE-BASED INDUSTRIES

The impact assessment component of LCA aims to increase the relevance of the inventory results by extending the mass loadings to a consideration of their potential environmental effects. A brief overview of life cycle impact assessment (LCIA) can be found in section 1.1.2. This phase of LCA is not without controversy, and although greater consensus on methods and approaches has been reached in latter years (Barnthouse et al., 1998; ISO, 2000a; Udo de Haes et al., 1999a), there is still considerable debate over the relevance of impact assessment and the extent to which it should be carried out. The assessment of primary industries raises some particular limitations of LCIA, notably its shortcomings with respect to the characterisation of impacts arising from waste deposits.

The focus of this thesis is on inventory development for resource-based industries. However, the current inability of impact assessment methods to characterise certain key features of these systems, has meant that a “default” impact assessment method could not simply be applied to extrapolate inventory data to impact indicators, and has forced the inclusion of these topics in the thesis. This chapter only briefly raises the issues (as their full consideration is beyond the scope of this thesis), and the reader is referred to the many references given in the discussion below for a more detailed consideration of these topics. In addition, the method proposed in section 5.4.2 borrows heavily from the work of Yvonne Hansen, who has been an invaluable co-researcher in a wider project from which the intellectual property of this thesis has been distilled. Readers are referred to her thesis (Hansen, 2001), and the papers referenced in section 5.4 for comprehensive coverage of the method developed.

5.1. THE ASSESSMENT OF RESOURCE-BASED INDUSTRIES

The aim of impact assessment is to capture all relevant effects of the system. However, this is typically not achieved with the “standard” set of LCIA impact categories, as they tend not to adequately assess local impacts. The particular requirements of the impact assessment stem partly from the features of the system being evaluated, and partly from the decision context in which they are applied. A discussion on relevant decision contexts identified for primary industries, and the implications their information flows have on impact assessment, is given in section 2.3. Primary industries are inherently resource intensive, in that they are the primary material and energy suppliers to industries further along the

supply chain. They typically involve the extraction and concentration of resource stocks, processes associated with the generation of significant solid waste volumes and surface disruption, in addition to extensive emissions to air and water (a brief discussion on the features of these industries is given in section 1.2.1). The ability of the current LCIA methods to address these main features of resource-based industries is discussed here.

5.1.1. Resources

The input-related categories in LCIA are required to incorporate all aspects of resource consumption by the system. However, an internationally recognised standard or unambiguously favoured method has not yet been developed (Goedkoop and Spriensma, 1999; Heijungs et al., 1997). This is largely due to the fact that characterisation methods for resource extraction and land use inherently contain more subjective judgements and are less “scientific” than most output-related categories. It is necessary to design methods on the basis of logic and theoretical reasoning, as there is no empirically correct or experimentally verifiable method for aggregating extractions of resources to an overall resource depletion score (Guinée and Heijungs, 1995).

Even the categorisation of resources is problematic. Heijungs et al. (1992) distinguish between resources of an abiotic (e.g. mineral ores) and a biotic (e.g. trees) origin, whilst Fava et al. (1993) define flow resources and stock resources, corresponding to renewable and non-renewable resources respectively. In a similar distinction, the EI 99 method distinguishes between dissipative and non-dissipative resources, after Müller-Wenk (1998) (Goedkoop and Spriensma, 1999). Finnveden’s (1996) distinction between deposits, funds and flows is the categorisation taken up by SETAC-Europe’s second working group on LCIA (WIA-2). These refer to the speed of recovery in relation to the speed of extraction. Deposits are basically depleted as the renewal rate is extremely low. Funds can be depleted but can recover, whilst flows cannot be depleted and can only give rise to competition (Udo de Haes et al., 1999b). This categorisation is preferred as it is able to explicitly include water use (surface water use under flows, and groundwater use under funds). Deposits, funds and flows are applicable to abiotic resources, with biotic resources and land use proposed to form separate categories, in accordance with the distinction between taking something out of the environment vs. different types of land use (Heijungs et al., 1997; Udo de Haes et al., 1999b).

Of the three input related categories, namely extraction of abiotic resources, biotic resources and land use, abiotic resources have been the most intensely studied in an LCA context. This can be attributed to LCA's historic routes in energy analysis and the resultant desire to aggregate different energy carriers (and later resources) into a single score (Heijungs et al., 1997). A characterisation method for abiotic resources must only consider the reduced availability of the resource, as the other impacts (notably land disruption) are covered under separate indicators, i.e. the categories must be kept as exclusive as possible (Udo de Haes et al., 1999b). An assessment of resource use needs to consider such topics as the depletion of the reserve, the loss of use options for future generations, and the increase in future mining impacts because the easily accessible ores will be depleted first. A number of methods of varying degrees of sophistication have been developed to assess resource depletion. These are briefly reviewed in appendix C.3.1, including the approach taken in the EI 99 method, which is the method followed in the case studies. This uses the concept of "surplus energy" to relate the impact of using resources to the decreasing quality of the reserves, and is defined as the difference between the energy needed to extract the resource now and at some point in the future (Goedkoop and Spriensma, 1999).

Notably lacking from the methods available in the literature (see appendix C.3.1) is a consideration of flow resources. Water use is an essential consideration in the regional context of the case studies undertaken in this thesis, and this is often likely to be the case for resource-based industries (given their water-intensive nature and the regions in which many take place). Udo de Haes et al. (1999b) suggest the use of an indicator relating the consumption of flow resources to the size of the flow. However, no operable method could be found, so the flow per functional unit is used as an indicator in the case studies in chapters 6-8. However, this flow is defined to include not only water actually consumed by the process, but also (for the foreground processes) the change in catchment for the area. It therefore reflects the degree to which water catchment is disrupted by the operations (e.g. by collection in opencast mining pits, run-off from contoured waste dumps etc.) (see section 6.2.6.c).

The impacts associated with land use are often not addressed, due to limited data availability, and characterisation methods that are incomplete and quite diverse (Heijungs et al., 1997). Surface disruption and land degradation are unavoidable side-effects of practically all abiotic resource extractions, and even if the land is rehabilitated after the operation has ceased, the land is most often not returned to its original state. Thus the ability to characterise land use and degradation is essential in an assessment of resource-based industries. Appendix C.3.2 briefly summarises the main considerations in the characterisation of land use impacts. An important distinction to be made is that

between land as a resource, and the loss of biodiversity as a consequence of the land use (Finnveden, 1996). Early methods, such as that of Heijungs et al. (1992), combine these two aspects of land use *impacts by defining a number of land use categories (e.g. natural systems, modified systems, cultivated systems, built systems and degraded systems)*. The change in land use (i.e. the transition from one category to another) is measured by the duration of the occupation, and area of the land occupied (e.g. $\text{m}^2 \times \text{yr}$). More recent proposals recommend keeping the distinction explicit by defining two (or more) separate indicators (Lindeijer et al., 2000; Udo de Haes et al., 1999b).

A severely simplified version of the EI 99 land use method is used in the case studies of Chapters 7 and 8 to characterise land use impacts (Goedkoop and Spriensma, 1999). This method only addresses the second aspect of land use, i.e. loss of biodiversity, and expresses the damage as the potentially disappeared fraction of vascular plant species ($\text{PDF} \times \text{m}^2 \times \text{yr}$). A highly approximated form of the method is required because of the need to interface with the land transformation categories defined by Heijungs et al. (1992), which are used by the background LCI databases. For compatibility, foreground land use is also allocated to these categories, with some adaptations made as to the allocation of land occupied by waste dumps (see section 6.2.6.d).

5.1.2. Aerial Pollutants

The LCIA methodology is probably best developed to reflect the effects of gaseous emissions. Where the pollutants have a global effect, for example greenhouse gases or ozone depleting substances, the LCIA methodology is especially well suited for their assessment. Although arguably less applicable, methodologies for the assessment of regional pollutants (acid gases, nutrifying substances etc.) also exist. Less well assessed are pollutants that act on a local level (e.g. human and eco-toxicity), although methods are in an advanced state of development (Udo de Haes et al., 1999b). There has, however, been considerable debate over whether it is reasonable (or even desirable) for LCIA to predict local effects (see section 5.3).

Dust emissions are particularly ill-fitting to the LCA methodology, although they tend to be a particular feature of resource-based processes. Non-stack emissions are typically not easily quantified in an LCA framework, since dust arising as a result of blasting, or that blown from a tailings dam, is a function of the prevalent weather conditions and the features of the particular site. The potential health effects of particulate emissions are also very much a function of local and regional conditions. Acid gas

emissions are also typical of resource-based processes, as sulphur compounds are present in many mineral ores and fossil fuels. These also act on a regional and local scale, as do toxic emissions (e.g. heavy metals released by smelting and fossil fuel burning processes). Thus the effects of aerial pollutants of resource-based processes are largely felt on a local and regional level, so the debate on the prediction of actual vs. global effects is especially relevant for these processes. It is therefore devoted a separate section (section 5.3).

There is a high degree of consensus on the relevant categories to characterise the effects of air emissions, even if the exact indicator to be used is still somewhat under discussion (particularly in relation to moving the indicator closer to the actual category endpoint) (Udo de Haes et al., 1999b). The EI 99 method represents the state of the art for these categories, as it models effects at their endpoints (damage to human and ecosystem health) (Goedkoop and Spriensma, 1999). However, the considerable model uncertainty this implies for certain categories (e.g. estimating the damage to human health from climate change) renders this of questionable benefit in some cases.

5.1.3. Solid and Liquid Wastes

The emission and containment of solid and liquid wastes will have a local and sometimes regional effect. LCA's global approach to impact assessment has meant that these effects have tended to be under emphasised in LCA studies, or missed completely. This is a very significant oversight in primary industries, where these effects are often extensive.

In addition to the usual pipe discharge type effluent releases, resource-based industries are usually associated with large volumes of dilute effluent (e.g. from rain water coming into contact with exposed mining seams and waste dumps). These effluents are not easily quantifiable, and thus often not included in the inventory. A consideration of diffuse water pollution sources highlights the inextricable linking of liquid and solid waste management in resource-based industries (e.g. the stormwater management procedures and the prevention of water coming into contact with exposed mining seams or collecting on dumps will significantly impact on water contamination and leachate formation). This is further intensified by the common practice of using liquid effluents for dust suppression in mining operations (e.g. on the haul roads and waste dumps), and for slurring tailings in mineral processing operations. These practices lead to a greater potential for leachate generation from the dumps. Such interactions are totally ignored by a system model that only considers the total mass of solid waste, or

merely the volume or area required for landfill, as is typically done in LCA studies. However, there have been some exceptions, where models have been developed to predict emissions from solid waste systems. These, and the reasons for their limitations, are discussed in section 5.2.

Emissions to water are generally less well characterised by LCIA methods than emissions to air, as a result of LCA's global approach (emissions to water inherently act on a local or a regional scale). Equivalency factors are available for emissions of toxic elements to water, although generally for fewer components than for emissions to air (Goedkoop and Spriensma, 1999; Guinée et al., 1996a). However, for other categories (e.g. acidification and nitrification), equivalency factors for emissions to water are generally lacking (Goedkoop and Spriensma, 1999). In addition, toxicity factors are generally not sufficient to capture the impacts of all waterborne emissions. Toxicity factors are not present for major elements (e.g. sulphates, chlorides etc.), although increased salinity has been shown to have a detrimental effect on aquatic ecosystems (Carlson and Adriano, 1993).

It is common in water quality monitoring programs to monitor an overall indicator (e.g. conductivity, BOD) rather than expensively monitoring a range of individual components. However, this is extremely problematic for LCA, as these broad indicators are generally not able to be addressed by impact assessment, i.e. the problem lies with the substances that they indicate, and not with the indicator itself. For example, conductivity can be related to TDS (total dissolved solids). Thus a high conductivity implies a high dissolved salt content, which may include toxic metal species. However, unless the individual species are listed, the stream will not be picked up as potentially problematic by the LCIA toxicity methods. This problem is encountered in the case studies in chapters 7 and 8, where data is predominantly only available for the major ions. An additional problem encountered is that the same set of elements is not monitored in all streams, and across all the mines and power stations surveyed. Sulphates are the most routinely monitored, since they are typically the greatest contributor to the overall salinity of mining-related effluents. The environmental intervention "waterborne sulphates" is therefore chosen as a "proxy" indicator to represent potential water quality impacts in the case studies in chapters 7 and 8. Although it is acknowledged that an assessment based on indicators at such different levels of sophistication (ranging from environmental interventions to endpoints) introduces significant uncertainty in interpretation (Owens, 1997a; Udo de Haes et al., 1999a), this was deemed preferable to ignoring a significant component of the overall impact from the systems under review.

5.2. CONSIDERATION OF SOLID WASTE IMPACTS IN LCA

The effects of solid waste disposal have been called the forgotten part of LCA, as traditionally they have been given limited attention in LCA studies, and often only the mass of solid waste produced has been noted. Possible emissions from landfills, to air or by leachate, are generally completely disregarded (with a few exceptions), largely because of the difficulty involved in predicting both the quantity and rate of these emissions (Finnveden, 1992). The solid waste from resource-based processes are the particular focus of this thesis (e.g. minerals processing tailings, furnace ashes, mining waste rock etc.), whilst municipal or post-consumer waste has received the greatest attention from LCA studies. This section briefly looks at the difficulties associated with predicting emissions from solid waste containments, and the methods that have been developed to do so.

5.2.1. Problems with Solid Waste Characterisation

In the assessment of landfills, emissions can not be measured because of the long time-spans over which they act. Indeed, leachate emissions from landfills may only become cause for concern long after the processes generating them have ceased to operate (i.e. once the natural buffering capacity of the soil has been exceeded and pollution “break through” occurs). The best that can be done is to predict potential emissions. However, these are complex systems with many interacting parameters, so predicting emissions from solid waste deposits is a considerable and data intensive task. In addition, the inclusion of the time-dependent emissions presents difficulties within the LCA methodological framework.

Wastes are often insufficiently characterised and the mechanisms controlling the rate and quality of leachate generation poorly understood. Significant effort must therefore be expended to model the hydrodynamic behaviour, and the complex chemical reactions (e.g. dissolution, precipitation, sorption and desorption etc.) of the multiple species present (Hansen, 2001). Whilst such models for solid waste deposits have been developed (see next section), comprehensive leachate generation models are beyond the scope of most LCA studies. Models developed to predict leachate emissions from industrial waste deposits, e.g. Petersen (1998) and Hansen (2001), offer a potential solution, as these can be adapted to provide LCI information. However, the interface between LCA and the output of such models needs to be developed (see section 5.4.3). The ideal is to be able to treat solid waste deposits as additional unit operations within the process, with well-defined inputs and measurable variables, and thus predictable

1.1. LIFE CYCLE ASSESSMENT AS A DECISION-MAKING PROCESS

The overall goal of environmental decision-making is sustainable development (WCED, 1987), where in addition to striving for inter- and intra-generational equity, we seek to deliver economic, environmental, and social outcomes from our collected anthropogenic activity which systematically improve quality of life without compromising the planet's carrying capacity. Three broad areas are generally considered relevant in decision-making as practical constraints on human activities (Cowell et al., 1997):

- Natural and physical sciences, including ecology and thermodynamics (the physical laws and relationships that shape ecosystems).
- Micro-economics and technology (the economic relationships, structures and products that shape business systems).
- Social issues and macro-economics (the social structures and issues that shape society)

Decisions which promote sustainable development occur when all three broad sets of constraints are satisfied (i.e. decisions which satisfy the so-called “triple bottom line”). Thus for a decision-making process to contribute positively to the overall goal of sustainable development, environmental information has to be integrated with economic, social and technological information. Figure 1-1 shows how various concepts, tools and technical elements, providing and structuring environmental, operational, economic and social information, build up the decision-making process.

The goals (sustainable development in Figure 1-1) set the overall direction for decision-making, i.e. all environmental management decisions should result in progress towards the goals (Cowell et al., 1997). The concepts provide the decision maker with ideas on how to approach environmental aspects in decision-making, e.g. Industrial Ecology, Waste Minimisation etc., whilst the tools are a means of combining information in a form which can be used in the decision-making processes. Tools are thus the operational methods supporting the concepts, and can be further broken down into analytical and procedural tools. In the former the focus is on computational algorithms or checklists aimed at finding a better decision, whilst the latter focus on guiding the best way to reach a decision. The technical elements are the methods of obtaining data, processing and presenting information, e.g. dispersion modelling, material balances etc. (Wrisberg and Gameson, 1998).

outputs. In this way the effect of upstream operational or technological changes on the emissions from the solid waste deposits can be assessed.

Methodological problems of including emissions from solid waste deposits include the time-frame over which they are generated and the variability in emission quality over time. Whilst most emissions incorporated into an LCI are steady-state and act over the life time of the operation, emissions from landfills will be spread over hundreds or even thousands of years. They must therefore be reconciled with the far shorter time period of most other emissions, which usually correspond to the operating life of the process, and are subsequently normalised to the chosen functional unit. This is usually done by integrating the potential emissions from landfills over a chosen time period. However, this requires the choice of a relevant time-horizon. Whilst the definition of LCA does not explicitly mention time-frames, the general consensus is that impacts on future generations should not be disregarded. Heijungs et al. (1992) therefore suggest integrating the emissions from the point where the waste is landfilled to infinity. Other researchers have used a time period of 100 years (Bez et al., 1998; Nielsen and Hauschild, 1998). Finnveden and Nielsen (1999) warn against neglecting the time period after 100 years and refute the assumption that an almost inert residue remains after 100 years, as suggested by Bez et al. (1998). Studies have shown that the emission of landfilled metals may be underestimated by a factor of more than a thousand if only emissions during the first century are considered (Finnveden and Nielsen, 1999).

A further complication is that the emission rate from landfills will vary as a function of time, but can not be assumed to decline exponentially with time (Finnveden et al., 1995). Finnveden (1995) suggests looking at two time perspectives: a short-term or surveyable time period, and a long-term or hypothetical, infinite time period. He suggests that the short-term period be defined as the time it takes to reach a pseudo-steady state. In this period, the chemistry is expected to be changing only slowly in response to external changes and the concentrations may be controlled by kinetic as well as equilibrium reactions. The surveyable time period is expected to correspond to roughly one century. The hypothetical, infinite time period is defined by the time it takes for complete degradation and spreading of all landfilled material (Finnveden et al., 1995).

outputs. In this way the effect of upstream operational or technological changes on the emissions from the solid waste deposits can be assessed.

Methodological problems of including emissions from solid waste deposits include the time-frame over which they are generated and the variability in emission quality over time. Whilst most emissions incorporated into an LCI are steady-state and act over the life time of the operation, emissions from landfills will be spread over hundreds or even thousands of years. They must therefore be reconciled with the far shorter time period of most other emissions, which usually correspond to the operating life of the process, and are subsequently normalised to the chosen functional unit. This is usually done by integrating the potential emissions from landfills over a chosen time period. However, this requires the choice of a relevant time-horizon. Whilst the definition of LCA does not explicitly mention time-frames, the general consensus is that impacts on future generations should not be disregarded. Heijungs et al. (1992) therefore suggest integrating the emissions from the point where the waste is landfilled to infinity. Other researchers have used a time period of 100 years (Bez et al., 1998; Nielsen and Hauschild, 1998). Finnveden and Nielsen (1999) warn against neglecting the time period after 100 years and refute the assumption that an almost inert residue remains after 100 years, as suggested by Bez et al. (1998). Studies have shown that the emission of landfilled metals may be underestimated by a factor of more than a thousand if only emissions during the first century are considered (Finnveden and Nielsen, 1999).

A further complication is that the emission rate from landfills will vary as a function of time, but can not be assumed to decline exponentially with time (Finnveden et al., 1995). Finnveden (1995) suggests looking at two time perspectives: a short-term or surveyable time period, and a long-term or hypothetical, infinite time period. He suggests that the short-term period be defined as the time it takes to reach a pseudo-steady state. In this period, the chemistry is expected to be changing only slowly in response to external changes and the concentrations may be controlled by kinetic as well as equilibrium reactions. The surveyable time period is expected to correspond to roughly one century. The hypothetical, infinite time period is defined by the time it takes for complete degradation and spreading of all landfilled material (Finnveden et al., 1995).

5.2.2. LCA Studies Addressing Solid Waste Disposal

Although solid waste disposal has typically received little attention in LCA studies, it has been shown that the potential emissions from landfilling are, for some products, of importance to the final results, and it can lead to misleading results if they are neglected or underestimated (Finnveden et al., 1995). This has been duly noted, and more recently a number of models to predict emissions from landfill systems have emerged. These have focused on emissions from municipal solid waste landfills (Bez et al., 1998; Nielsen et al., 1998; Nielsen and Hauschild, 1998; Weitz et al., 1999). Consequently these studies tend to focus on the complicated task of allocating emissions from the landfill to the components within the landfill giving rise to them. This is not necessary for landfills or waste deposits associated with primary industries. These are usually fairly homogenous in composition, or if co-disposal takes place, all wastes usually arise from the product or process system under consideration.

There has tended to be an emphasis in LCA on post-consumer waste, with solid waste generated in other stages of the life cycle (e.g. mining and energy conversion) disregarded. A Swedish study shows the importance of solid waste generated at other stages in the life cycle, by calculating that approximately only 5% of the total waste produced annually in Sweden is of domestic origin, the balance being industrial, ash from energy conversion and sludge from waste water treatment (Finnveden, 1992). Although models of municipal waste landfills dominate the literature, Sundqvist et al. (1994) include a consideration of industrial wastes in their work. They considered coal ash landfill, hydroxide sludges and mine-tailings landfills. The very complex reactions taking place and the general lack of data meant that a very simplified modelling approach considering only three metals could be achieved. The emissions were integrated over a period of 100 years, as the developed model was too simple to be used to calculate a true surveyable time period. However the study was able to highlight that emissions from such landfills can be significant with respect to the overall life cycle, although predicted with a high degree of uncertainty (Sundqvist et al., 1994).

5.3. CONSIDERATION OF REGIONAL AND SITE-SPECIFIC IMPACTS IN LCA

Guidelines for the LCIA process have been laid out in the recent ISO standard (ISO, 2000a), although certain areas of the method are still undergoing vigorous research. Increasing calls have been made to develop the simple relative indicators used in LCIA to better reflect the actual impacts of processes. However, to a certain degree this is at odds with the fundamental methodology of LCA, and thus a consideration of local impacts has raised significant problems and debate. This debate is especially relevant to an assessment of resource-based processes, as the effects of these processes are often very localised. This is a particularly active area of LCA research, so the brief overview here does not attempt to be comprehensive, but rather aims to capture the main points of the arguments. This debate is primarily of relevance to the assessment of toxicity in LCA, specifically the inclusion of fate and transport into the assessment. An overview of the inclusion of fate and transport modelling into LCIA models is therefore given in appendix C.4.

Early LCA studies were predominantly product-related, taking place fairly far down the supply chain (e.g. paper vs. plastic packaging). A particular feature of LCA is that it aggregates effects all along this supply chain, so as to provide a comparative assessment of the products on an equal basis. The global nature of LCA is thus an important strength of the tool. However, it also presents some significant limitations to the inclusion of spatial and temporal effects in LCA, a consideration of which are needed if actual effects are to be predicted. As the use and development of LCA has progressed, there has been an increasing awareness of need to include impacts felt on a local and regional scale, and that not all systems can be adequately compared on global issues and resource consumption alone. A large number of impact assessment methods have been developed, each covering a range of areas of environmental concern, with varying degrees of comprehensiveness and levels of sophistication. However, in their review of impact assessment methods, Huppes et al. (1999) conclude that all the methods surveyed have severe limitations. Extensive research is thus still required to improve the accuracy and operability of these methods.

The SETAC-Europe second working group on LCIA (WIA-2) are of the opinion that this can best be done by increasing the environmental relevance of the indicators, and that the long term aim of LCIA should be to define all indicators at the level of endpoints (i.e. at the damage level) (Udo de Haes et al., 1999a). However, they recognise that, whilst making the indicator more environmentally relevant,

defining the indicator at the level of endpoint also renders it less certain in its relationship to the environmental interventions. Modelling at the endpoint level requires considerable additional data, much of it of a site-specific nature. This raises the possibility of generating spatially differentiated equivalency factors. The WIA-2 recommend that non-differentiated global factors for every impact category first be developed, and that spatially differentiated factors subsequently be developed for those factors where large variations of fate and exposure or of effect variables are observed (Udo de Haes et al., 1999a). They warn that the global factors must remain available, because attaining inventory data at an equivalent level of spatial differentiation may not always be possible. Such spatially differentiated equivalency factors have already been demonstrated for acidification by Potting et al. (1998).

However, there are also those that warn against the representation of local effects on a generic level. Barnthouse et al. (1998) warn that extrapolating to endpoints can give the false impression that the indicators represent absolute environmental processes (i.e. environmental concentrations), and stress that indicators can only ever be intended to represent relative overall emission loadings. In their opinion, only global, long-lived environmental processes are candidates for a reliable representation within LCIA, and that as the processes become more local and/or more transient, or are known to have thresholds or non-linear dose response curves, the reliability and representativeness of the indicators are increasingly lost. Thus, whilst Barnthouse et al. (1998) agree with Udo de Haes et al. (1999a) that some improvements in accuracy and relevance are possible by including spatial differentiation in the indicators, they view the threshold and dose-response assumptions as more intractable, i.e. these are not simple gaps in understanding that further research can address, but inherent constraints that force a consideration of unavoidable uncertainty in the use of generic indicators for the assessment of local effects (Owens, 1997a). A framework for assessing the accuracy of the impact categories typically considered in LCIA from a consideration of their inherent characteristics can be found in Owens (1996), and is summarised in Table 5-1.

Table 5-1 Potential accuracy of impact categories based on their underlying characteristics (Owens, 1996).

Impact categories	Spatial	Temporal	Dose Response	Threshold	Accuracy relevance
Global warming	Global	Centuries / decades	Linear (?)	No (?)	Good
Stratospheric ozone	Global	Centuries / decades	Linear	No	Good
Acidification	Continental / regional	Years	Non-linear	Yes	Fair
Eutrophication	Regional / local	Years	~ Linear	Yes	Fair-poor
Photochemical smog	Regional / local	Hours / day	v non-linear	No (?)	Poor
Ecotoxicity	Local	Hours-Years	Non-linear	Yes	Poor
Habitat loss	Regional (?) / local	Decades / years	Non-linear	Yes	Very poor
Biodiversity	Regional (?) / local	Years	Non-linear	Yes	Very poor

The concept of thresholds in LCIA methods has led to the “less is better” and “only above threshold” debate (Potting et al., 1999), i.e. whether the fact that many releases do not lead to any adverse environmental impact when present below a certain concentration should be taken into account. This debate is particularly relevant for the toxicity categories, where threshold play a critical role in the risk assessment methods. The underlying assumptions of the “less is better” approach is that emissions have intrinsically harmful properties, and that all sites have a similar sensitivity to the emitted substance (Potting et al., 1999), i.e. that all emissions are relevant and cause effects, whether or not this is true (Owens, 1997a).

The non-consideration of spatial and temporal differentiation in LCA (i.e. LCA’s practice of summing the emissions of a given substance throughout the life cycle) leads to an inability to deal with exposure and ambient concentrations, which, in turn, leads to an inability to deal with thresholds, critical loads, or dose-response characteristics (Owens, 1997a). Thus, those characterisation methods that do not consider spatial and temporal differentiation are inherently constrained to the “less is better” approach, in that they can only predict concentration increases and not the actual concentrations (Potting et al., 1999). The approach allows for a comparative analysis of emissions between alternative systems, but not of actual effects (Potting et al., 1999). However, it also allows the aggregation of seemingly negligible quantities of substances into totals that are implied to be significant (Owens, 1997a).

It is the opinion of certain researchers that LCIA methods are inherently limited to relative indicators that can never predict absolute concentrations, and are thus limited to the “less is better” approach (Barnthouse et al., 1998; Mackay and Seth, 1999; Owens, 1997a). From their perspective site-specific toxicity assessments are beyond the capabilities of LCA and are better handled by techniques oriented to an absolute approach, e.g. risk assessment (RA). This interface with RA is taken up by the “only

above threshold" methods, in which additional site-specific information is collected for those processes identified as having the largest share of emissions in the life cycle. Such methods as those used by RA are then used to predict whether threshold values will be exceeded, and only those emissions exceeding thresholds taken into account in the impact assessment (Potting et al., 1999). This increased accuracy in impact prediction has to be traded off against the extensive data requirements of a location-specific assessment. Also, the requirement that the actual location of emission be known means that this approach is not applicable to all decision contexts (see section 2.3). A further criticism of the "only above threshold" methods is the recognition that exposure to a combination of different substances with similar effect mechanisms can be additive, and thus lead to a toxic effect, even if the individual emissions are below the no-effect level (Potting et al., 1999).

Such an "only above threshold" approach has been demonstrated by Thiel et al. (1999), i.e. the use of RA at those processes identified by the LCA as having the largest share of emissions in the life cycle. Their study clearly shows the dramatically increased data requirements of such an integrated assessment, and the paper only presents an assessment of the impacts of NO_x emissions from a single site in the life cycle (that contributing the greatest to NO_x emissions). In addition, to calculate the actual environmental effects from this site, the environmental fate models had to be based on full production at the plant of interest, although the product of interest actually accounts for less than 10% of this production. This highlights the fundamental incompatibility between LCA and RA, i.e. the inability of LCA to provide an estimate of the absolute mass of releases. This results from two central methodological features of LCA, the functional unit calculation, and the allocation of resources and emissions in multi-product systems (Owens, 1997b).

The functional unit provides emissions normalised to a measure of performance of the system, which is crucial to the original purpose of LCA, i.e. a comparative analysis of products. However, there is no estimate of the absolute mass, which is required in a spatial context to address specific environmental impacts. Considering a single product in a multi-product system requires the materials and emissions to be allocated between the products, thus the LCA deals with only a limited portion of the overall releases in an operation. This again prevents a consideration of the absolute quantity of emissions (Owens, 1997a). The type of data included in an LCI is another feature of LCA recognised by Owens (1997a) as preventing the inter-linking of LCA and RA. LCI data is not necessarily compatible with that required by RA, such as LCI's use of aggregated data categories (e.g. total hydrocarbons, TDS

etc.), whilst data quality and availability, and the lack of data quality analysis, are also barriers to extension of LCA to RA.

Potting et al. (1999) assert that “less is better” and “only above threshold” do not have to be incompatible approaches, and conclude that the only difference between the approaches is the way which threshold values are dealt with. It is thus possible to combine the methods that predict above and below threshold exposures with the fate and transport models developed for LCA (e.g. by defining sub-categories for exposure situations above and below threshold, and weighting these sub-categories to indicate their relative importance). To unify these approaches additional factors are required, including an identification of source type (e.g. high/low point sources, indoor/outdoor sources), and for each source type, an estimation of the area of increased concentration, the exposure situation (e.g. above and below threshold on the basis of typical background concentrations), and the number of exposed people (Potting et al., 1999). However, this does not overcome the inability of LCA to predict absolute emission quantities. Owens (1997a) therefore disagrees that the “less is better” and “only above threshold” approaches can be compatible, and concludes that the inability to predict absolute emissions as well as the practice of system-wide data aggregation (that Potting et al. (1999) do address), largely removes spatial and temporal considerations. Thus thresholds for toxic effects do not exist under all circumstances, and all dose-response functions relate linearly to mass loadings. Thus, the inventory accounting procedures of LCA constrain LCIA to providing a relative indicator, able to provide a simplified directional perspective on environmental topics, and the capability of which varies considerably between impact categories (see Table 5-1) (Owens, 1997a).

An interesting study by Tolle et al. (2001) demonstrates conducting LCIA at increasing levels of detail and site-specificity. The study compares two equivalency factor approaches (“less is better” approaches) with simplified RA (“only above threshold” approach) for assessing the toxicity impact potential of emissions to air and water. The two equivalency factor methods, the PBT (Persistence, Bioaccumulation and Toxicity) method and the MFM (Multimedia Fate Model) method, incorporate toxicity benchmarks and data on potential exposure into a single equivalency factor. However, the MFM method incorporates the fate of the emission using Mackay Level III models (see appendix C for an overview of including fate and exposure models in LCIA). The simplified RA approach uses dispersion modelling with site-specific environmental data to determine the environmental concentrations of emissions and compare them with threshold values considered safe for humans and

wildlife. As in the study by Thiel et al. (1999), the extensive effort to conduct a site-specific RA limited the study to a single site in the life cycle.

Very different levels of effort were required for the three approaches, with the least effort required for the PBT method, intermediate effort for the MFM method, and considerably more effort required for the RA approach (approximately 24 times more effort than the PBT method, and 4 time more effort than the MFM method, even though fewer emissions were assessed with the RA approach due to lack of data, and because emissions were excluded if the release concentration was already below the threshold considered safe). Tolle et al. (2001) conclude that the selection of an appropriate method for LCIA needs to consider the number of sites and emissions involved in the life cycle, as well as the accuracy required from the assessment. The minor reduction in accuracy of the chemical ranking obtained with the PBT method may be justified by its considerable reduction in effort, and is the favoured method for an initial screening in a large LCIA where there is little or no previously compiled or site-specific information on environmental characteristics. As the number of sites and emissions decreases, the improved accuracy of the MFM method is likely to justify its additional effort. Where previously compiled data for regions and chemicals exist, the improved accuracy of the MFM method may be achieved with very little extra effort relative to the PBT method. The substantial additional effort required by the RA approach, and the lack of available information, suggests that this method is not justified unless there are only one or two sites included in the LCIA, which is unlikely given the comprehensive nature of most LCIA studies (Tolle et al. 2001).

From the above discussion one can conclude that LCA and RA are truly distinct tools, and that rather than trying to merge the two, LCA's ability to provide insights into hidden trade-offs and media shifts should be used in planning risk assessments or assisting risk managers to evaluate options (Owens, 1997b). It is clear that the specific data requirements of RA require that the system be well defined, and thus is not suitable for all decision contexts. It is best suited to those contexts where the actual location of the process is known (e.g. fairly far down the design "chain" for new processes, or small design changes within an existing process), and also where the product or process of interest is responsible for the full production at the site. From these considerations, it would appear that the relative indicator approach of LCIA is best suited to strategic/tactical (Type II) studies, and a combined LCA/RA approach to operational (Type III) decisions. It is also clear that a consideration of uncertainty is required when interpreting LCIA indicators, especially those using generic site-specific information to produce toxicity indicators.

5.4. POTENTIAL SOLUTIONS FOR ADDRESSING SOLID WASTE IMPACTS

It is clear from the above discussion that impact assessment is far from being a certain and well established component of LCA. The impacts from solid waste are particularly poorly assessed by current LCIA methods. This is to a large degree attributable to limitations in the inventory calculations, where the considerable complexities involved in predict the contaminants leaching from waste deposits, as well as methodological problems associated with the temporal nature of the emissions, means that emissions from waste dumps are typically not included in the inventory. If the substances emitted from the dump known, their fate and exposure can be estimated using the types of models discussed in appendix C.4, and an estimate made of their environmental impact via some sort of an effect model. However, the discussion in appendix C.4 shows that, in addition to a host of other problems, these models are ill-suited to a consideration of metal species, which are the most likely elements of concern in the context of resource-based industries (e.g. acid mine drainage, well recognised as a severe problem in the mining and processing of sulphidic ores, results in the emissions of heavy metals). These problems have led to this investigation into alternative solutions to obtain a relative indication of impacts from solid waste deposits.

5.4.1. Qualitative Indicators

In many studies, a fair amount of either qualitative or quantitative information may be known about the waste, although this is not in a format able to be included in an LCA study. For example, on-site monitoring programs may indicate emissions are occurring by measuring elevated concentrations in groundwater, although it is not possible to relate this to the mass emission required for LCA, or to separate out the emissions attributable to a particular site from background concentrations and other possible sources of emissions. Laboratory leach tests are able to give a relative indication of inherent toxicity or "leachability" of wastes, although the differences between laboratory and actual site conditions means that extrapolating this to actual emissions from a waste dump is not meaningful, or extremely complicated. However, a consideration of the potential impacts from waste dumps can be essential (see case study in chapter 7), and where the nature of the waste is affected by the decision, incorrect conclusions can result if the characteristics of the waste are ignored. In such cases, an inclusion of qualitative or semi-quantitative information (e.g. rating the waste on an ordinal scale), is better than assuming the waste is inert and only considering landfill area or volume.

5.4.2. Impacted Land Footprint

The fact that questionable toxicity data for only a few metal species exist (see appendix C.4), and the difficulty in quantifying the potential mobilisation of metal species from waste dumps, has led to the suggestion of a simplified proxy-indicator to characterise impacts from solid waste (Hansen et al., 2001b). This approach is quite distinct from the fate and exposure models developed for use in LCIA (see overview in appendix C.4). These methods use Mackay-type models to partition the contaminants into the various environmental media (soil, air, water etc.) using equilibrium partition coefficients, and then relate the calculated environmental concentrations to effects through toxicity benchmarks (e.g. LD₅₀, NOEC etc.). The approach developed by Hansen (2001) is different in that it looks at more sophisticated contaminant pathways into the environment (kinetics and equilibrium processes, non-ideal solution behaviour etc.), and considers the spatial distribution of contaminants. The result of leachate generation and subsequent plume dispersion is a potentially contaminated land area. If adequate toxicological models were available and the composition of the plume was known, the contaminants could be linked to actual environmental impacts. However, in the absence of reliable toxicity models and detailed knowledge of the contaminants in the leachate plume, the impacted land area serves as a useful midpoint indicator of environmental impact, where the land area is chosen as a quantity to which decision-makers are able to relate to, rather than, for example, a concentration limit.

The key requirements for the indicator are that it is able to be linked back to the nature of the waste, and that it is able to provide a meaningful relative indication of the impacts from solid waste deposits. The basic premise of the method is that solid wastes contain leachable constituents, such as heavy metals and salts, which may constitute an environmental risk if they remain mobile and bio-available. The indicator aims to capture purely this aspect of solid waste deposits (i.e. leachate emissions leading to groundwater contamination, and the potential of subsequent contamination of surface water and soils). Other aspects of solid waste disposal, such as land use, dust emissions and energy use, are dealt with separately, and their effects characterised by the relevant impact categories. In addition, this method has been developed for process wastes, particularly those from primary industries, which are characterised by large dedicated waste dumps. The impacted land footprint is also used to assess the impacts of stockpiles, as these generate similar impacts to waste deposits, and are also a potential source of leachate emissions.

To quantify the land area affected by the pollution plume, tools to predict leachate generation and pollution plume migration are required, as well as a methodology to define the boundary of the affected land footprint (Hansen et al., 2001b). ERA is used by the LCIA toxicity characterisation methods to compare predicted environmental concentrations with levels believed to cause effects, and thus enable an indication of risk (Mackay and Seth, 1999). It is therefore proposed that ERA can similarly be used to define the boundary between unacceptable and acceptable risk to the environment, i.e. to enable the boundary for the impacted land footprint to be drawn (Hansen et al., 2001b). The differentiation between acceptable and unacceptable risk is based on the environmental value to be protected, and can be translated into acceptable concentration limits of contaminants by a consideration of dose-response and other toxicological data. Alternatively such criteria as water quality guidelines can be used. Basing the indicator on the area of land affected rather than the actual toxicological effects is consistent with the practice of basing the assessment on easily measured midpoints, rather than the endpoint (the actual environmental damage). The use of water quality guidelines is consistent with basing the assessment on midpoints, as these have been worked back from the endpoints (i.e. to ensure protection of the environment) (Hansen et al., 2001b).

The extremely complex nature of the reactions taking place and the many species involved have limited the footprint approach to a consideration of salts. Salts are more readily modelled as the chemical reactions and mechanisms are better understood, whilst trace metal concentrations in leachates and soils are difficult to predict (their mobility is determined by aqueous and heterogeneous chemical reactions, particle coagulation and flocculation mechanisms, and are affected by pH, salinity, redox state and available ligands). Their inclusion would thus increase the modelling complexity significantly (Hansen et al., 2001b). Modelling the salts only is consistent with the level of inventory information available, so this is not seen as a particular limitation of the footprint approach. On the contrary, the ability of the method to characterise the solid waste on this reduced information set, is considered a significant strength of the method. The interrelation between salinity, metals mobility and bio-availability is complicated. With increased salinity, metals mobility increases due to inorganic salts competing with metals for adsorption sites and dissolved complexation with salt anions, whilst metals bio-availability decreases as there are presumably less metals in free ion form. The toxicity of trace metals can thus be said to be inversely related to salinity (Hansen, 2001).

Since salinity and metals mobility are related, it follows that trace metals will follow a salinity pollution plume. In the footprint model, salinity is measured as the ionic strength of the solution rather

than the more usual electrical conductivity, as this allows simpler modelling, and is also felt to be more meaningful in linking salinity measurements to toxic effects (Hansen et al., 2001b). The footprint “boundary” is calculated at two values. The first of these is a legislative limit for Total Dissolved Solids in drinking water, converted to ionic strength, and the second the “absolute extent” of the footprint, estimated as where the ionic strength becomes negligible. The latter is used in the case studies in chapters 7 and 8. Only a single footprint is presented, as the choice in basis did not change the relative rankings obtained.

The impacted land area is a function of the nature of the waste, the extent of the waste dump, the geohydrological and geochemical characteristics of the area, as well as the time over which the leachate has been generated (Hansen et al., 2001a). The degree to which site specific characteristics are taken into account is determined by the particular decision context. This is consistent with what has been proposed for other LCIA indicators, with generic indicators calculated first, and the possibility of later extending these to include spatial considerations, if applicable (Udo de Haes et al., 1999a). It has been argued above and in section 2.3 that spatially differentiated indicators are not always desirable (or feasible), and are limited to operational and perhaps tactical decision contexts. As with the other generic LCIA indicators, using generic information to model site specific processes introduces a certain amount of unavoidable uncertainty, and limits the indicator to a relative assessment (i.e. it has no absolute meaning) (see section 5.3). Significant and hidden value judgements are involved in the choice of generic conditions, and it should be checked that the choice of these do not change the outcome of the assessment (i.e. sensitivity analyses should be conducted). The temporal aspect of the impacted land footprint is included by taking “snapshots” of the footprint at various times (e.g. at closure, after 100 years etc.) (Hansen et al., 2001a). The footprint results used in the case studies in chapters 7 and 8 are given only for 100 years. This was because the relative rankings of the options were not affected by the duration after which the footprint was calculated. However, it is important that this be checked before a single time span is chosen.

A model for predicting leachate from a solid waste dump requires the following (Hansen et al., 2001a):

- A consideration of hydrodynamic behaviour, reflecting variations in degree of saturation corresponding to different disposal methods.
- Chemistry modelling, complex enough to consider dissolution and precipitation reactions, sorption and desorption, as well as kinetically controlled reactions, such as redox reactions and leaching.

- A consideration of solution thermodynamics, able to account for deviations from ideality, including high salinity.

Whilst the subsequent fate and transport of the contaminants in groundwater requires a consideration of advection (movement caused by the flow of groundwater), dispersion (movement caused by irregular mixing of waters during advection), and retardation (chemical or physical mechanisms which slow down the rate of contaminant migration, e.g. adsorption and precipitation). Details on the methods used can be found in Hansen (2001).

Uncertainty in model form is an inherent and unavoidable component of any model. Sensitivity analysis can be conducted to assess the effect of the choice of decision variables and model domain parameters (see Hansen 2001a), but the uncertainty in the choice of model form is less easily considered. One place where choice in model form has been considered is in the calculation of leachate composition. Each footprint is calculated at two extreme leachate compositions. The one assumes equilibrium concentrations in the pore water, and the other is a "straight through" value, assuming the concentration of the pore water does not change (e.g. due to channelling). It is likely that the actual pore water composition lies between these two extremes, so an average of the two footprints is used in the case studies in chapters 7 and 8. The variation in footprint due to these extreme leachate compositions was found to be a lot less than that due to an estimate of empirical variability (see section 8.2.5), so taking an average is not assumed to introduce appreciable additional uncertainty.

The impacted land footprint represents a significant step forward in operationalising the impacts from solid waste deposits, whilst working within the inventory and impact constraints imposed by these complex systems. An indicator based on the footprint approach is compatible with the other LCIA indicators and shares their general characteristics (i.e. whilst able to be site specific, are still meaningful in a generic relative assessment). There does not appear to be any barrier to summing the generic footprints across different sites and stages in the life cycle, so the approach is compatible with the global nature of LCA. The idea of a partially allocated impacted land footprint, whilst physically nonsensical, is no more so than the allocation of any of the other impacts routinely allocated in LCA. The use of the footprint in a global LCA context is therefore fine so long as the relative and generic nature of the indicator is borne in mind. The impacted land footprint only has an absolute sense when applied to a single site with site specific data, in which case, the absolute value is an estimate of the

area covered by the pollution plume. This feature is not unique to the footprint, and is a characteristic shared of all LCIA indicators (see section 5.3).

The technique is still relatively in its infancy, and adjustments are likely as the models are honed. At present the footprint is calculated in separate modelling platforms to the inventory assessment, and the analysis can only be conducted using single-point modelling. Potential improvements could be to allow better linking between the models to thus seamless data transfer, and ultimately to allow for inventory uncertainty to be carried through into the assessment of the footprint. Perhaps ultimately a range of generic indicator-multipliers based on pertinent waste characteristics could be developed, and thus enable the footprint indicator to be simply calculated in a manner consistent with the equivalency factor approach of the other LCIA methods.

CHAPTER 6

LCI OF SOUTH AFRICAN COAL BASED POWER GENERATION

South Africa's power industry is dominated by coal, a situation which will remain for the foreseeable future, given the cheap and plentiful supply available and the strong political drive to keep electricity prices low. Significant environmental, social and economic effects stem from this large-scale mining and combustion of coal. An aim of this study is to characterise the environmental effects of this collected activity from a life cycle perspective. Some background on South Africa's electricity industry is first presented, followed by a description of the model constructs used to generate an LCI for this essential industry. The chapter concludes with a presentation of the inventory.

6.1. THE POWER INDUSTRY IN SOUTH AFRICA

South Africa is in a period of critical development, accompanied by transformation of its industry base. Generation, transmission and distribution of electrical power have historically been in the hands of the state-controlled electrical utility, Eskom. Moves are currently afoot to deregulate this large enterprise to some degree, and split the line groups into separate entities (Energy, 1998; Enterprises, 2000). The SA Government's 1998 White Paper on Energy Policy supports taking gradual steps towards a competitive electricity market in the short term, while investigation into the desired form of competition is completed. It sets out the following policy objectives of the energy sector, of which the priorities are:

- Increasing access to affordable energy services
- Improving energy governance
- Stimulating economic development
- Managing energy-related environmental impacts
- Securing supply through diversity

Eskom is the fifth largest electrical utility in the world (rated according to both sales and capacity), with a nominal installed capacity of 40 GW (Eskom, 1999a). In 1998 it supplied 95% of South Africa's total available electricity, with the balance supplied by municipalities and industries that generate part of their electricity requirements. This translated to a total of 183 TWh (net) produced in Eskom stations, with a small percentage (1.4%) of the electricity sold imported from neighbouring countries (Eskom,

1999a). Eskom also dominates the Sub-Saharan electricity supply, and is responsible for 76% of installed capacity and 83% of production and trade of electricity in this region (Lennon, 1997). Eskom's capacity is heavily reliant on coal (see Table 6-1). 89% of the total nominal capacity is provided by coal-fired stations, which consumed 87.2 Mt of coal in 1998 (Eskom, 1999a). Eskom has notable excess capacity, with 11.5% of the total coal-fired capacity in reserve storage for possible future re-commissioning. The ability of the modern power stations to burn low-grade coals means that primary energy costs can be kept very low, making it extremely difficult for other fuels to penetrate the market. In addition, these coals are relatively low in sulphur compared to world averages, so the stations are run without flue gas desulphurisation (FGD) units.

The location of the power stations is highly regional, with all of the coal-fired stations to the North of the country, and only Koeberg nuclear station in the South (see Figure 6-1). The coal stations were all designed to be "mine-mouth", i.e. each supplied by an adjacent coal-mine. This severely restricts their location, but minimises coal costs. Six of the nine base-load stations, producing around 65% of the electricity, are clustered together in the Witbank region (top right of the inset box in Figure 6-1). This agriculturally productive region of the country is also the main coal-producing region, and is consequently heavily industrialised in parts. It is subject to significant environmental pressures. Water availability is critical, and air quality poor, particularly in winter, where local inversion layers trap aerial pollutants, resulting in very slow dispersion rates.

Table 6-1 Eskom generating mix (condensed from Tables 1, 2 and 3 in Eskom 1998 Annual Report)

Generating Plant	Total Nominal Capacity		Total Net Maximum Capacity ¹	Produced in 1998	
	MW	%		MW	GWh (net)
Coal-fired Stations: Total	35 539	89.1	33 605	165 473	90.4
In storage	4 150		3 871	0	
On order	2 130		2 007	0	
Gas Turbine Stations ²	342	0.9	342	3	-
Hydroelectric Stations ³	661	1.7	661	1 595	0.9
Pumped Storage Schemes ⁴	1 400	3.5	1 400	2 420	1.3
Nuclear Power Station	1 930	4.8	1 840	13 801	7.4
Total Eskom Stations	39 872	100	37 848	183 093	100

1. Difference between nominal and net maximum capacity reflects auxiliary power consumption and reduced capacity caused by age of plant and/or low coal quality.
2. Stations used for peaking or emergency supplies.
3. Use restricted to peaking, emergencies and availability of water.
4. Pumped storage schemes are net users of electricity. Water is pumped during off-peak periods to generate electricity during peak periods.

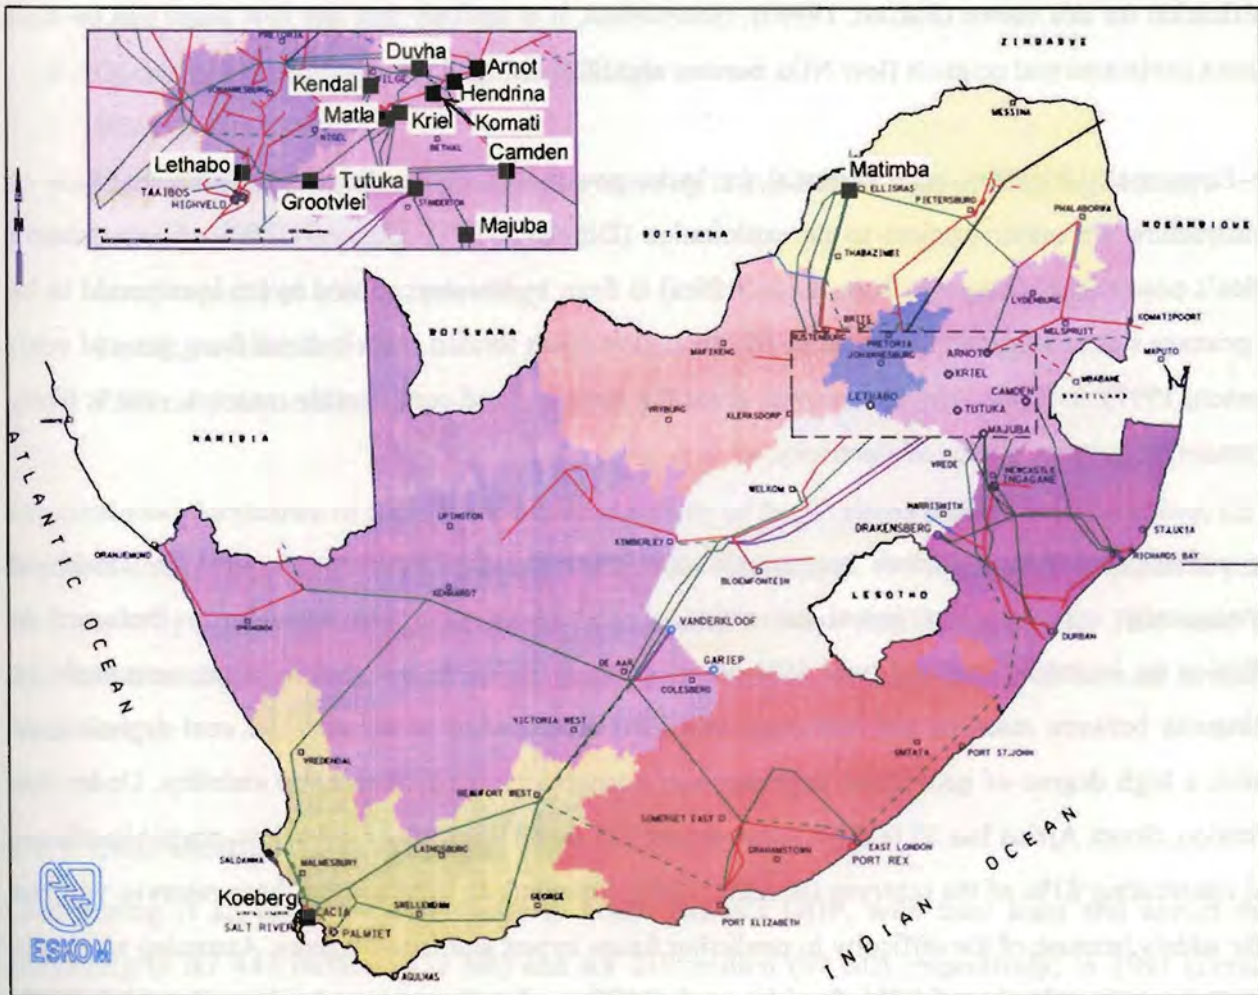


Figure 6-1 Electricity grid map of South Africa showing the highly regional placement of the power stations, with only Koeberg nuclear power station in the South of the country, and all but one of the coal-fired stations located within the box (source: Eskom).

This high regional density of stations has meant that water consumption is an important consideration in the design and maintenance of the power stations, and Eskom pioneered the first large-scale dry cooled station in response to these pressures. Air quality concerns forced CAPCO (chief air pollution control officer) to “cap” the air emissions in this region, forcing Eskom to explore coal-fields further afield, and to recognise that further exploitation of the reserves in this region will require a commitment to enhanced air pollution controls, notably around the fate of sulphur oxides and particulates. Eskom’s commitment to the government’s development programmes to reduce the real price of electricity and electrify homes precludes expensive capital expenditure to decrease emissions from existing stations (e.g. fitting flue gas scrubbers). Eskom cite the fact that electrification reduces air pollution in the region as a result of decreased biomass and coal burning for domestic heating and cooking as

justification for this stance (Eskom, 1999b). Nonetheless, it is unlikely that any new plant will be built without environmental controls (low NOx burners and FGD units).

Sub-Equatorial Africa has huge potential for hydro-power, although political instability and lack of infrastructure are major barriers to its exploitation (Dingley, 1995). Currently 79% of Sub-Saharan Africa's power (excluding that from South Africa) is from hydro-sources, and hydro is expected to be the primary source of growth in capacity for the region (with limited contributions from gas and coal) (Lennon, 1997). In South Africa, however, given the low cost, and considerable reserves, coal is likely to remain the primary source of electricity.

The estimates of South Africa's recoverable coal reserves are dependant on what is considered "economically minable". The initial estimation of 25 billion tons was dramatically increased to 58 billion by including coals of over 35% ash (De Jager, 1983). It is therefore more meaningful to distinguish between reserves and resources, where to be classified as a reserve, a coal deposit must exhibit a high degree of geological certainty, and economic and environmental viability. Under this definition, South Africa has 55 000 Mt reserves and 121 000 Mt resources, with low-grade bituminous coal constituting 81% of the reserves (Bredell, 1987). Opinions as to how long these reserves will last differ widely because of the difficulty in predicting future export and usage figures. Assuming an annual increase in coal production of 1.8%, Surridge et al. (1995) predict that coal production will peak in 2050, and then tail off over the next two centuries. The ability of the power stations to burn low-grade coals means that they will continue production long after coal exports have ceased.

Eskom's IEP (integrated electricity planning) process has predicted a 50% increase in energy demand between the years 2000 and 2015, assuming a long-term economic growth rate for South Africa of between 1.5% and 3.5% (Lennon, 1997). Under the current scenarios, utilisation of the excess capacity in the Southern African Power Pool (SAPP), and the implementation of demand side management measures, will result in no additional capacity likely to be required before 2010 (Lennon, 1997). Eskom are considering a number of established and new technologies to meet this demand, guided by the following factors (in approximate order of importance) (Lennon, 1997):

- Capital and operating cost
- Plant reliability and availability
- Access to indigenous, low cost fuel

- Lead times
- Operations flexibility (base load vs. peaking)
- Water availability
- Environmental considerations (likely to move up in importance with pending legislation)
- Security of fuel supply
- Local capacity to sustain technology
- Funding availability
- Political considerations

The continued dominance of coal is evident if the priority of these criteria remain as listed above, i.e. an emphasis on keeping costs low and using an indigenous fuel supply. A framework is required to judge the superiority of one technology over another according to these conflicting criteria. This study has contributed to such a process by providing baseline information around coal use to be used in such a framework.

6.1.1. Coal Mining and its Impacts

Coal mining is an important contributor to South Africa's GDP, with total local and export sales amounting to R7 442 million (157 Mt) and R8 210 million (57 Mt) respectively, in 1997 (Prevost, 1998). Electricity production dominates the local market, and accounted for 60% of the local saleable production in 1997 (Prevost, 1998). Over the years coal mining has been a cause of major environmental degradation in South Africa (Van Horen, 1996). Coal mining, particularly opencast mining, is associated with massive surface disruption. This results in changed land use and water catchment patterns, in addition to the noise, visual intrusion, dust and water contamination typically accompanying surface mining operations. Opencast mining accounts for 45% of South Africa's coal production (Prevost, 1998). Underground mining is generally less disruptive. Water contamination, surface subsidence and underground fires are some of the major impacts of underground mines (Wells et al., 1992). Only the environmental impacts of coal mining are considered here, although considerable social impacts also accompany coal mining, including the problems of migrant labour, and the high number of injuries and fatalities occurring in coal mines (Van Horen, 1996). A brief overview of the mining technologies employed by the South African coal mining industry, and their consequent impacts on the environment, can be found in appendix G.1.1.

Major environmental impacts stem from the high ash content of Southern Hemisphere coals, which require beneficiation to produce coals of acceptable quality for the world markets. Of the 279 Mt of run-of-mine (ROM) production in 1997, only 219 Mt was of saleable quality, resulting in 60 Mt of discard requiring disposal (Prevost, 1998). The resultant discard dumps are responsible for some of the most serious environmental effects of coal mining, including land sterilisation and groundwater contamination. On exposure to air and water the pyrites oxidise to form sulphuric acid, and iron oxides and hydroxides, which cause the pH to drop. The acid produced reacts with bases in the rock to form salts and consequently mobilises any heavy metals present. The resultant acid mine drainage (AMD) contains elevated levels of salts (mainly calcium and magnesium sulphates) and metals (predominantly iron, manganese and aluminium). The pyrite-rich discard is also susceptible to low temperature oxidation (so-called spontaneous combustion), and subsequent release of toxic air pollutants. AMD and spontaneous combustion can be minimised by preventing water and air getting to the pyrite and other sulphidic minerals. The power stations mostly burn ROM coal, so are not responsible for discard production, although some power stations are supplied by dual product mines. These mines maximise their coal production by producing a high-quality coal for export, and a medium quality power station coal, which contains a portion of the washing discard blended in with the ROM coal.

The location of South Africa's coalfields is significant. The Mpumalanga / Eastern Gauteng / Northern Free State region, where 65% of the reserves are to be found, has been extensively farmed, with little natural environment remaining. Coal mining therefore has little residual impact on natural ecosystems and land rehabilitation is usually able to restore the land to an acceptable state (Wells et al., 1992). However, from a water quality perspective, the coalfields occur in the worst possible location, since most mines are situated in the vulnerable upper reaches of Southern Africa's major river systems (Wells et al., 1992).

In response to the environmentally damaging nature of mining operations, the Minerals Act of 1991 was implemented, which requires any company wishing to mine, or those already doing so, to submit Environmental Management Programmes Reports (EMPRs). These form a legally binding contract between the mine and the Department of Mineral and Energy Affairs (DMEA) to show how they will protect the environment during the working life of the operation, and how they plan to rehabilitate the working areas after closure. An important feature of the legislation is the control over water resources. Mines require authority to abstract water and to discharge wastewater to designated watercourses, and

are responsible for ensuring that the discharged water complies with quality requirements. Contaminated water has to be adequately contained, and rainwater kept free from pollutants as far as possible (Baxter, 1993).

6.1.2. South Africa's Power Generation Technologies and Plant Mix

This study is limited to a consideration of the environmental effects of the coal-fired plants only, with the relatively small contributions from nuclear, gas and hydro not addressed. It is also limited to Eskom power stations, and does not include the small contribution (less than 5%) from old municipal power plants. In addition, the study only addresses generation, i.e. excludes transmission and distribution. An overview of the technologies employed in Eskom coal-fired power plants is given in appendix G.1.2, and covers the three broad aspects of the plants:

- *Boiler Plant* dealing with the "solids section" (milling and combustion of the coal).
- *Water Plant* dealing with the "water circuits" (production of potable, ultra-pure and cooling water).
- *Waste Management* where the "solids section" and "water circuit" overlap (co-disposal of combustion residues with liquid effluent).

South Africa generates electricity on an impressive scale. The most-recently constructed stations were built in a "six-pack" format (6 units of approximately 600 MW each). The generating capacity of the units reach as high as 686 MW, with one station (Kendal), the largest bituminous coal-fired power station in the world. Table 6-2 summarises some main features of Eskom's coal-fired stations.

All the stations employ pulverised fuel firing. The composition of the flue gas depends on the coal quality and conditions in the boiler, with carbon dioxide, water vapour, nitrogen oxides and sulphur dioxide the principal products. Thermodynamic equilibrium may not be reached and other compounds, such as methane, ammonia and carbon monoxide may be formed in small quantities. South African coals are generally low in chlorine and fluorine, although some HCl and HF is emitted, together with other volatile trace elements present in the coal, e.g. Hg and Se. Some of the less volatile elements condense onto the fly ash particles, where they also can be emitted e.g. As, Pb and Cd (Swaine and Goodarzi, 1995).

Table 6-2 Features of Eskom's coal-fired power stations. Condensed from Table 2 in 1998 Annual Report (Eskom, 1999a) and Eskom Statistical Yearbooks (Eskom, 1990; Eskom, 1997).

Eskom Power Stations in Commission at 31 December 1998							
Station Name	Number and capacity of generator sets	Number sets in storage	Total nominal capacity	Total net maximum capacity	Overall thermal efficiency	Specific water consumption	Generation load factor ² (1996)
			MW	MW	%	l/kWh	%
Amot	6 x 350	1	2 100	1 980	33.3	1.73	58.6
Camden ⁴	8 x 200	8	1 600	1 520	28.3	2.62	
Duvha	6 x 600		3 600	3 450	34.5	1.83	76.5
Grootvlei ⁴	6 x 200	6	1 200	1 120	28.7	3.23	
Hendrina	10 x 200		2 000	1 900	32.3	2.11	78.3
Kendal ³	6 x 686		4 116	3 840	34.3	0.11	65.4
Komati ⁴	5 x 100, 4 x 125	9	1 000	891	26.7	3.51	
Kriel	6 x 500		3 000	2 850	35.0	1.93	60.2
Lethabo	6 x 618		3 708	3 558	34.9	1.80	65.9
Majuba ¹	3 x 657		1 971	1 836			47.4
Matimba ³	6 x 665		3 990	3 690	33.5	0.14	76.3
Matla	6 x 600		3 600	3 450	35.5	2.10	77.0
Tutuka	6 x 609		3 654	3 510	35.3	1.92	47.0
Total / Average			35 539	33 605	34.4	1.38	67.2

1. Under construction. A further 3 units still to be commissioned.
2. kWh produced x 100 / (average net maximum capacity x hours in year)
3. Dry cooled stations.
4. Stations in storage.

The high ash content of the coals (ranging from 22 to 40%) results in extremely high particulate loading in the flue gas. Particulates are removed at high efficiencies (in excess of 98%, and as high as 99.9%) by electrostatic precipitators (ESPs) or fabric filters (FFs). At present, legislation only requires particulate monitoring in the stacks, by way of opacity meters. Each station is set a compliance limit by CAPCO. The limits are well within the stations' reach, with no stations violating their limit in 1998 (Eskom, 1999b).

A consequence of burning poor quality coal is also the huge volume of ash requiring disposal. Of the approximately 25 Mt of ash produced in 1998, just less than 5% was sold for use in cement products (Eskom, 1999b). Ash is either slurried with effluent and pumped out to the ash dam, or moistened with effluent and transported on conveyors to the ash dump or mined-out opencast workings. Considerable surface disruption therefore occurs, although this is mitigated to some degree by rehabilitating the dumps. Approximately 82% of disturbed land had been re-vegetated in 1998 (Eskom, 1999b). Ash

disposal is also associated with considerable potential for groundwater contamination (seepage from the ash dam, or leachate from the ash dump), and dust emissions, particularly from the dry ash dumps.

Eskom power stations are amongst the most water efficient in the world (Dutkiewicz and Gore, 1998). The high pressure on water resources in the locality of the stations has led to the development of advanced water saving technologies, which include dry cooling, dry ash disposal and sophisticated cascading effluent systems. Dry cooling can reduce a station's specific water consumption by more than an order of magnitude (see water consumption figures for Kendal and Matimba in Table 6-2, compared with the other stations). All of Eskom's newer stations strive to achieve a zero liquid effluent discharge (ZLED) policy. This means that all effluent generated, including storm water run-off, must be contained and evaporated on site. The water systems are thus designed to maximise water reuse and recycling. Cascading effluent systems ensure that the highest conductivity water is disposed of first, whilst lower conductivity effluents are placed back in the system. Bottom ash quenching and ash slurring in stations employing wet ash disposal, or ash conditioning and dust suppression in stations employing dry ash disposal, are the major effluent sinks available. This sink is smaller in stations with dry ash disposal, so the wet-cooled stations employing dry ash disposal require additional water treatment technology (e.g. a membrane plant) to be able to meet the ZLED requirement. Violations of the ZLED policy are required to be reported, with only four ash water spills reported in 1998 (Eskom, 1999b).

6.2. THE COAL-ELECTRICITY LIFE CYCLE

Electrical energy is an input into practically all processes, and the burdens associated with its provision have been found to significantly contribute to the overall burdens in many product- and process-life cycle studies. Energy data are thus indispensable in an LCA study, and consequently many life cycle inventories have been developed for energy-generating systems (see summary in Spath et al., 1999). LCA databases typically contain electricity data for many countries, built up by combining life cycle inventories of generating systems in the correct proportions to represent the fuel mix of that country, e.g. the IDEA and ESU databases (Frischknecht et al., 1994; Lubkert et al., 1991). Most LCA studies have tended to focus on the aerial emissions and energy requirements of electricity generation, with a lesser emphasis on solid waste and water emissions. Leachate from ash dumps and coal stockpiles is typically neglected, and water use often not included in the assessment (Spath et al., 1999). This current research study has a particular emphasis on water use and water-related impacts, because of the importance of these in the South African context.

In addition, the deficiencies of LCA with respect to solid waste have been identified (see section 5.2). There is thus a renewed interest in developing methodologies able to assess the implications of waste management in the LCA framework. This study contributes to this process, and suggests some methodological adaptations to ensure the localised impacts arising from coal mining and ash waste disposal are not overlooked (see section 6.2.6).

6.2.1. A Life Cycle Inventory Model of South African Coal-fired Power Generation

The life cycle presented here is that of coal; its extraction, processing, combustion and disposal of its residues. If viewed from the perspective of electricity, the life cycle is one of "cradle-to-gate" and not "cradle-to-grave", as the assessment ends at the generation of electricity. Transmission, distribution and use of the power are not covered. It thus provides an inventory of *undelivered* electricity. In addition, only process-related emissions are assessed. Burdens associated with offices, workshops etc. are not incorporated in the assessment. The primary aim of the study is to provide an inventory of the pollutant and resource flows associated with power generation in South Africa. This information is expected to be of value to other South African LCA studies, where relevant South African data can be used, instead of the European or American data commonly found in LCA databases. According to the application

types of LCA studies identified in section 2.3, this study is thus best defined by a Type I model, i.e. product declarations and studies providing information to be used in other studies. The considerations for flowsheet development, data requirements etc., as outlined in section 2.3.1, are thus applicable to this study.

A life cycle has both a spatial and a temporal dimension. This inventory is for an operating power station and mine. Building and commissioning the plants are not included, neither are the burdens associated with the materials of construction. Maintenance is also not included, and only consumable materials are included in the inventory. Emissions from waste management are mostly time-dependent, as the size of the waste dump changes over the life of the process. The temporal aspect of the life cycle can therefore not be avoided. Time dependent emissions are reported as averages for the time period corresponding to the functional unit, where the averages are calculated from the total potential emissions over the lifetime of the process (see appendix E for details).

The flowsheet for the coal-electricity life cycle is given in Figure 6-2. This shows the flow of coal through its life cycle, and the associated processes. The groupings of colours indicates the top level of breakdown in the flowsheet; mining and coal preparation, coal stockpile, coal combustion and particulate removal, water treatment, and ash and effluent disposal. These groupings or modules coincide with the major technology combinations available on the mines and power stations. This allows a single flowsheet to encompass all possible technology combinations by swapping in the relevant process module. These combinations are:

- Underground or opencast mining
- Stockpile at the mine or at the power station
- Boiler and particulate removal equipment type
- Wet or dry cooling
- Wet or dry ash disposal

These processes comprise the foreground system, and are broken down to the level of major unit operations. Details of the foreground / background split are discussed in the next section, followed by a discussion on the level of flowsheet breakdown and data sources used in this study. Necessary simplifications made and assumptions taken are also explained (details can be found in appendix E).

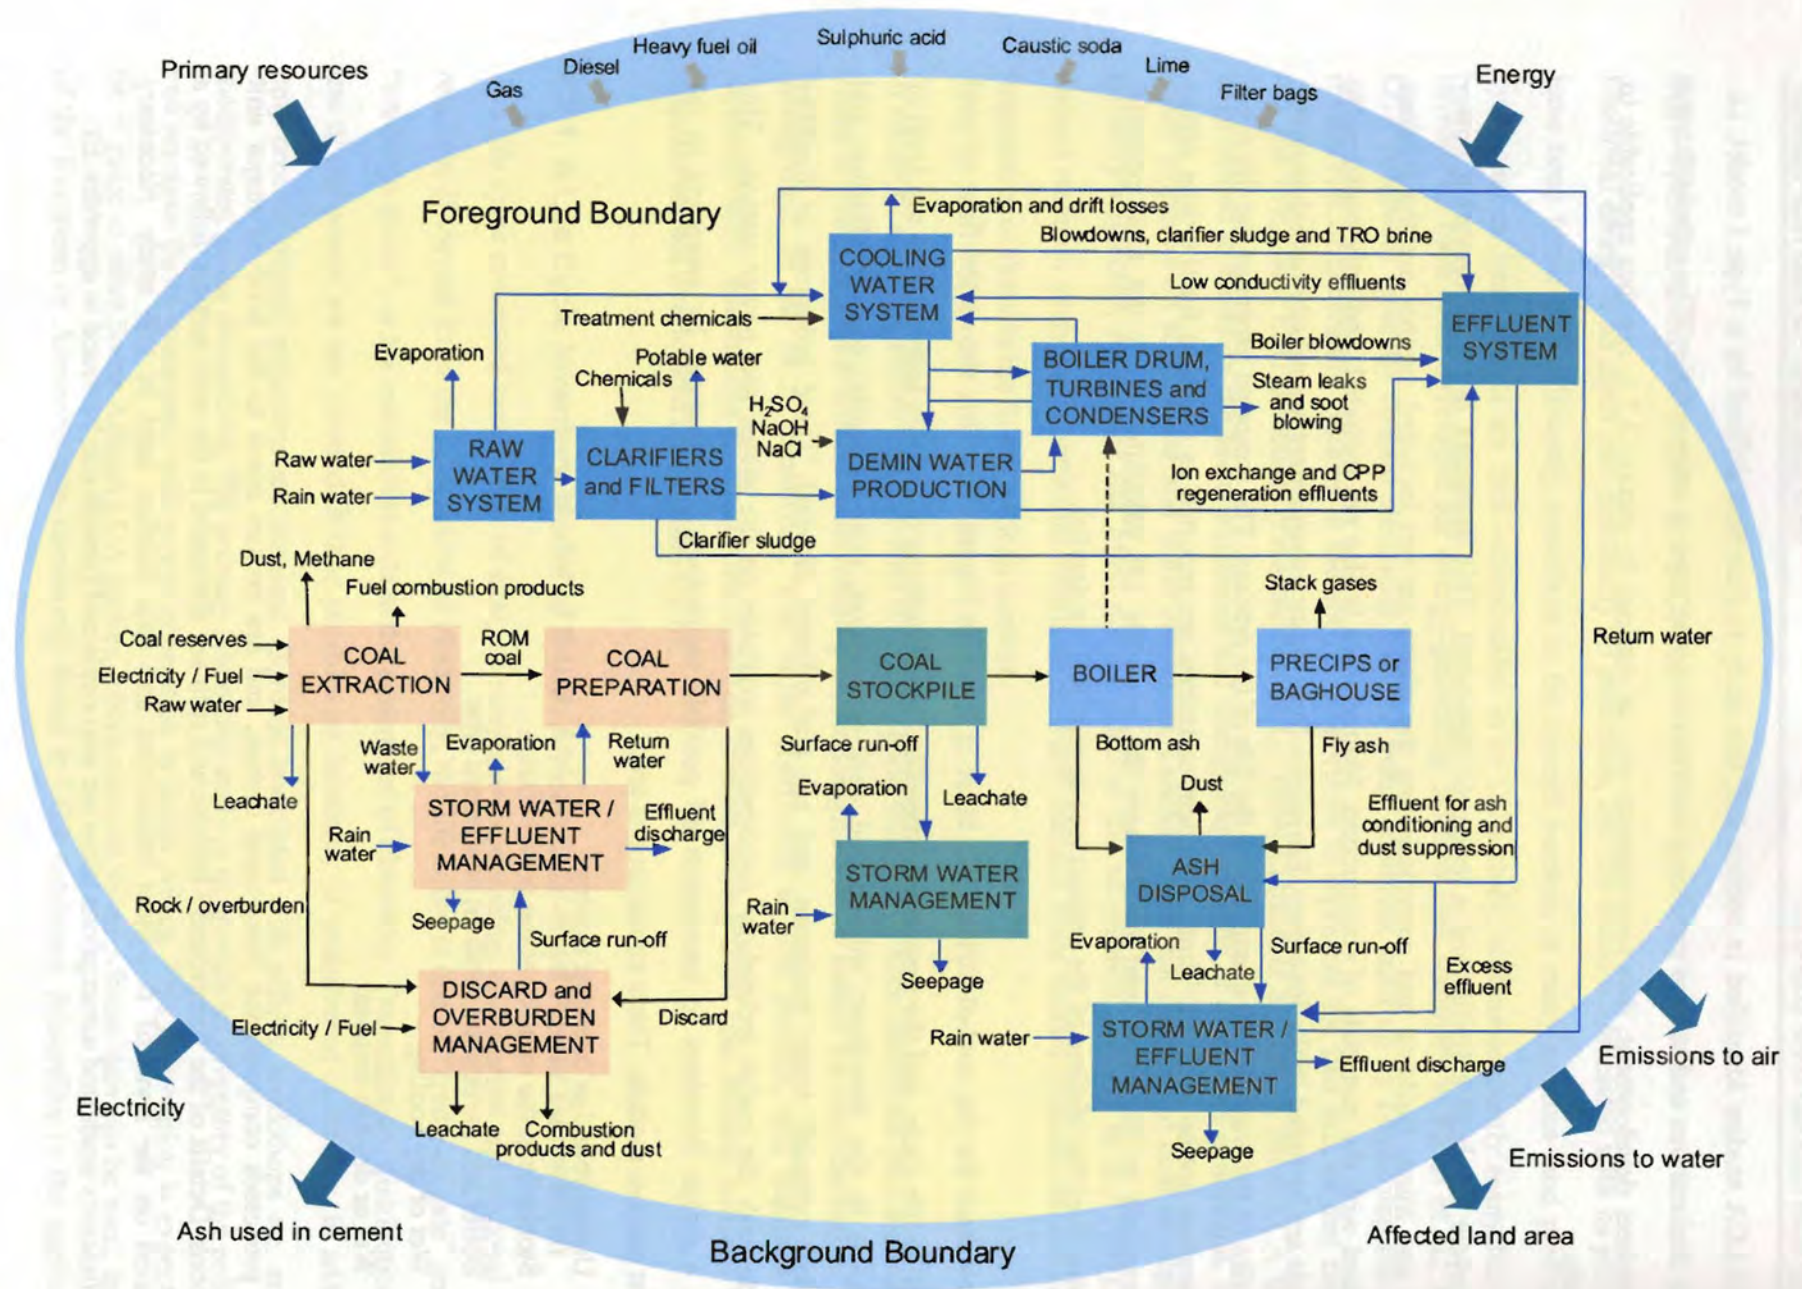


Figure 6-2 Schematic representation of the coal-electricity life-cycle.

6.2.2. Foreground and Background Systems

As discussed in section 2.1.3, it is useful to break the system into foreground and background sub-systems. The background system is defined as that over which you have no control other than the quantity of material (or magnitude of the function) input into the foreground system. An important consequence of this is that all associated burdens can be assumed to vary linearly with the quantity, thereby allowing the use of aggregated LCI data for that material. A notable distinction between the foreground and background systems is thus that inventory data for the background system can be taken from commercially available LCA databases (aggregated LCI data), whilst the foreground system requires data to be collected for the specific process at hand. LCI data for the foreground system is usually calculated using material balance models characterised by relevant process data.

The process blocks shown in Figure 6-2 comprise the foreground system. Mining could perhaps be seen to form part of the background system, i.e. the power station may well have no influence on the mine other than the quantity of coal purchased. However, mining is responsible for a significant proportion of the life cycle burdens. The marginal technology approach, and streamlining approaches, both warn against the use of average generic data for influential process blocks (Hunt et al., 1999; Weidema et al., 1999). Thus high quality inventory data is required around mining to accurately account for the burdens of the total system. In addition, the generic coal mining LCIs available in the LCA literature are not relevant to South African coals and conditions. Recourse is thus only made to generic and/or non-South African data where the equivalent colliery specific data is not available, and the mining LCI data generated at as disaggregated a level as possible. Furthermore, the nature of Eskom's contracts with some of the mines means that they do have some influence over the operation of the mine. Quality standards on the coal product limit how the mine can operate (e.g. coal preparation and blending restrictions), and material interactions between the mine and power station (other than the flow of coal) are evident in some cases, e.g. using ash to backfill opencast mining pits, using mine effluent in the power station cooling circuit (see case study in chapter 8).

Waste management processes are explicitly considered in the foreground. Whilst the mass of waste generated is quoted in the LCI, in order to be consistent with other LCI studies, this is not strictly appropriate as the waste dump is assumed to sit within the foreground system boundary (the mass flow does not cross the system boundary). Reporting the mass of waste generated is also problematic, as it is

not always straightforward whether the by-product material should be defined as a waste or as a potential future resource (e.g. coal washing discard). Regardless of its definition, it is the effects of containing the material that are of concern (i.e. the land occupied and the emissions from the dump), and not the mass of material. The dumps/stockpiles are treated as “pseudo” unit operations, in that they generate emissions (dust, leachate and stormwater run-off) which are either emitted or incorporated back into the process (in the case of surface run-off).

All ancillary materials comprise the background system. These include the production of fuels (those used “on site” and in the transport of materials), water treatment chemicals (lime, sulphuric acid etc.), and process ancillaries (filter bags, ion exchange resins etc.). “Cradle-to-gate” product inventories from LCI databases are used to supply the inventory data for these substances, as is consistent with the LCA methodology for processes consigned to the background system. These substances fit the definition of background processes, as only their input quantity is controlled by the foreground processes (which supports the use of aggregated data), and the quantities input into the foreground processes are small (which supports the use of generic data). The ancillary materials are consumed in the foreground system processes, so the emissions from their use are incorporated into the foreground system, although their production falls into the background system.

6.2.3. Flowsheet Construction

The aim of this case study is to provide inventory data on South African power generation (i.e. the Type I system, defined in section 2.3). As proposed in section 2.3.1, systems in this context require a flowsheet of sufficient resolution to allow the generation of high quality inventory data. Thus it should be detailed enough to allow for transparency (i.e. it should be clear what has been included), and for an assessment of uncertainty, either qualitative or quantitative. However, the primary driver of the system resolution is the level of aggregation in the available data. The highest quality data should be used, regardless of its resolution. For example, a metered overall electricity consumption figure may be known, but not the individual machinery requirements. The measured overall value should be used rather than the sum of estimated requirements of pumps, fans etc., as it is of higher quality. The fact that it gives no insight into which unit operations are consuming the most power is not relevant in the context of the study.

This is very much a feature of the decision context, and not of the inventory model. A greater level of flowsheet detail is required for different decision contexts (see considerations for Type II and III systems in sections 2.3.2 and 3). The level at which the data is specified is thus a function of the particular decision context, as well as the data availability. Case studies requiring different flowsheet considerations are presented in chapters 7 and 8. In these case studies, the aim of the inventory study is very different to that in this chapter (i.e. to provide insight and guide technology assessment), so a detailed disaggregated inventory of the processes under consideration is essential.

Presented in this chapter is the average grid inventory, calculated from a weighted combination of the inventories of each of the base-load stations. Detailed information was collected for three of the stations. These stations were chosen as together they represent Eskom's technology mix, i.e. wet cooling and wet ash disposal, wet cooling and dry ash disposal, and dry cooling and dry ash disposal. Inventory models were drawn up for each of these three stations, which were then used as generic power station models to model the other stations. The individual station models were adapted, where possible, with readily available station-specific information i.e. the majority of data inputs were kept the same across the same technology option, other than for certain key data, including the capacity and load of the station, the fuel and water quality, the thermal and electrical efficiency, the particulate removal efficiency, and the specific water consumption (see appendix F.1 for a full data listing).

The process blocks in Figure 6-2 are in most cases further broken down into major unit operations, as far as data availability permits. This level of flowsheet detail is necessary in order to perform the material balances to calculate the inventory data. Greater detail was generally available around power generation than mining, and the flowsheet is broken down to the level of major unit operations for the boiler plant and water plant. Not all information was available even at the broad level of the process blocks shown in Figure 6-2, e.g. only overall mine power use and fuel consumption figures were available. More detailed flowsheets, and detail on the modelling assumptions taken, are given in appendix E.

6.2.4. Some Methodological Considerations

The inventory results are presented per MWh power delivered to the grid. The inventories for the individual stations are calculated on the basis of an average year's production, the inventories summed, and the sum normalised to the total power sent out in the year. The functional unit is thus essentially the

net power generated (i.e. MWhSO) in a specified time period, in this case a single year. The period of operation is related to the power sent out by the load factor. This is the fraction of the total time available during which the power station is operating. A simplifying assumption made is that the power station is assumed either not to be operating or operating at full load.

The power sent out needs to be related to the time of operation because certain burdens cannot be directly related to the power sent out (or indirectly the flow of coal through the system), e.g. those emissions related to coal stockpiles, waste dumps and certain aspects of mining. The emissions generated from these processes act over the life of the mine and power station (and in some cases beyond) regardless of whether the station is operating for that specific time period or not. Specifying the functional unit as an operating period of interest, also enables accumulation or depletion of the coal stockpile to be taken into account.

The system boundary ends at generation, so the burdens associated with the transmission, distribution and use of the power are not included. The quantity of power used in an upstream process will not be the same as that sent out at the station, and should be inflated to account for power losses during transmission and distribution, so that the inventory reflect the power sent out by the station, rather than the power consumed by the system of interest. The bulk of the power losses occur during transmission, and are usually estimated per km of power line. The amount by which to inflate the electricity use can therefore be relatively easily estimated, provided the approximate distance between the generating source and the system of use is known. However, identifying the generating source in a power grid system is often not possible, although the fact that the power stations are all located fairly close together in South Africa makes a rough estimate possible.

A small portion of the ash produced ($\pm 5\%$) is used in cement production. Ash is not considered a by-product, but a waste, so a portion of the system burdens are not allocated to its production, nor is the system credited with the burdens avoided by replacing a portion of virgin material with ash in cement manufacture. The system does reflect the benefit of landfilling less ash, although the effect is negligible, since the ash sold makes up such a small percentage of the total ash produced. Eskom's stations are predominantly supplied by dedicated collieries, although three of the stations are supplied by dual product mines. These collieries produce a high quality coal product in addition to the power station coal, and by so doing, also produce a discard coal. A simple mass-based allocation rule has been

applied to these stations to apportion the mining burdens between the two products, i.e. the percentage of burdens allocated to the power station coal is equivalent to the percentage of total coal produced sold to the power station.

An important methodological decision to be made is whether discard coal is classified as a waste or a stockpiled resource. As a waste, it is not allocated any burdens of production, which it would be if considered a by-product (a future energy source). In the dual product mines, the power station coal is often blended with a portion of the washing discard. If the mine was only producing high quality coal, this discard would have been stored in a discard dump, with the associated environmental impacts, and waste of energy resources. In this study, discard is assumed a waste (i.e. assigned no mining burdens), in which case blending discard coal in with the power station coal not only avoids the burdens of stockpiling the discard, but also reflects the use of a "burden-free" energy source. The overall emissions of a dual product coal mine are therefore expected to be lower than from two separate mines, although the burdens of a dual product mine are greater than those from a mine only producing ROM coal. The simple allocation method used is not able to pick up the interplay between the two products, and thus is likely to overestimate the burdens attributed to the power station coal. However, the simple mass-based coal allocation is judged acceptable, since only a third of the power stations are supplied by dual product mines, making the effects slight. These considerations are investigated in more detail in the case study in chapter 7.

6.2.5. Data Sources

Very different data sources are used to characterise the foreground and background processes. To provide information for drawing up the foreground process models and to provide the data to characterise these models, a number of different sources were consulted. Background data were obtained from commercially available LCA databases. The sources of data used to calculate the inventory are discussed here, and listed in appendix D.

6.2.5.a Foreground System

Inventory data for the foreground processes are generated by simple "factor-based" models. The models are based on mass balance principles, and as such, rest on the assumption that all inputs must report to one or more of the output streams. Factors derived predominantly from process data essentially act as splitter functions that steer an input to its respective output streams. An important feature of the process

models is therefore that they are not based on any fundamental chemical or physical modelling, but on characterisation factors derived from process data. The models are thus very data intensive, and the accuracy of the calculated inventory data is very much dependent on the applicability of the data used. Where process data were not available, recourse has been made to data from process simulation models, literature and experimental data. A fair degree of uncertainty is introduced by using various sources of data (i.e. the accuracy of the source, its applicability to the case in hand etc.), in addition to the uncertainty caused by variability in the data. Uncertainty ranges are estimated for selected environmental interventions in section 6.3.2.

Process data were used preferentially, and an effort was made to tie into the reporting structures of the power station. Process data were collected for three representative Eskom stations and their associated mines on field visits and through email correspondence with station personnel. Data available were month end reports, as well as the raw monitoring data used to generate these reports. The most comprehensive information set was available for the water plant, where monthly water balance and chemical services (chemicals consumed) reports are generated. Monthly air quality reports to the regulatory body (CAPCO) were also a good source of data. Coal proximate analyses and water analyses were available for a number of years, although the varying frequencies of analyses, and the inconsistency in the number of components analysed, complicated the use of this data. Numerous interviews with power station personnel were undertaken, which were valuable in obtaining an understanding of the process, as well as providing process facts and figures. Interviewees generally provided process design information (i.e. unit efficiencies and plant specifications), with little indication as to what extent the operation of the plant deviated from this ideal. The mines run as separate entities, and were generally less forthcoming with information than the power stations. Copies of the mines' EMPRs and supporting documents were acquired, from which a good understanding of the operation could be obtained. Monthly water monitoring data were also acquired.

Information was also sourced at Eskom's research organisation and head office. Numerous interviews were held with Eskom research personnel and consultants, and past and present research projects investigated. Information gathering was hampered by the disparate nature of the personnel and research, and the lack of a centralised resource base. As with any large organisation, difficulties were encountered in connecting with the relevant personnel. General information was obtained from Eskom publications and Eskom-related papers in the open literature. Eskom publish an annual report, which

contains predominantly financial information, but also key performance data, and an annual environmental report, containing mostly qualitative assessments of Eskom's impact on the environment. The Statistical Yearbook, published every couple of years, is a good source of general operating figures for the various power stations. Technical Information Brochures have been published for the more recently built stations, which contain a description of the operation, as well as pertinent facts and figures. Power plant, mining and chemical engineering handbooks were used where the equivalent plant-specific information could not be found. A list of the documents and databases sourced, as well as a list of the personnel interviewed, can be found in appendix D.

6.2.5.b Background System

Generic LCI data were judged sufficient for background processes (see section 6.2.2). A fair number of LCA databases have been developed over the last ten or so years, and are probably most conveniently accessed through the many commercially available LCA softwares. The databases contain "cradle-to-gate" information for common materials and processes, the accuracy and comprehensiveness of which varies considerably between databases. Broadly speaking, the data is of European origin (or to a lesser extent North American), and is poorly documented regarding the technology on which the inventory is based, the age of the data, and comprehensiveness of the inventory model. Selecting the most applicable inventory and judging its uncertainty is therefore difficult.

The generic LCA data used in this study was obtained through the PEMS software (PIRA, 1996). Specific databases used were the ESU database (Frischknecht et al., 1994), the APME inventories (Association of Plastic Manufacturers in Europe, various reports from 1993-1998), the TENSIDE database and IDEA (Lubkert et al., 1991). Inventories were only available for the primary materials. The inventories for some ancillary inputs are therefore not complete, and only contain information on the materials and not the processing to form the final product. Falling into this category are filter bags, ion exchange resins and magnetite for dense medium separation of the coal. For these inputs, it seems likely that provision of the materials would form the bulk of the burdens in a complete product LCI. Complete inventories were available for the more common material inputs, e.g. the water treatment chemicals and fuels. LCI data for some materials were not available, notably synthetic fuels produced from coal.

6.2.6. Limitations of the Assessment of Waste Dumps and Coal Stockpiles

As discussed in section 5.2, waste dumps and coal stockpiles do not fit easily into an LCI assessment. Their non-steady state and imprecise nature ensure this. The size of the dump changes over time, so their potential emissions are a function of time, and the factors affecting these emissions are variable and mostly not able to be precisely calculated (e.g. rain ingress). Significant approximations are therefore inevitable if these processes are to be incorporated into an LCI assessment.

6.2.6.a Leachate Emissions

Ash and coal systems display complex interactions between their many components, so calculating leachate emissions is beyond the scope of most LCI studies. These difficulties have meant that often only the mass or volume of solid waste is reported, although this is inadequate, and often causes significant impacts of the system to be missed completely. A discussion on the limitations of LCA with regard to the prediction of impacts from solid waste systems is given in section 5.2. In the inventory presented here, a simplified “stop-gap” approach to calculating leachate emissions from dumps and stockpiles is used. The system boundary is taken to extend adjacently to the dump, so any leachate leaving the dump is assumed to enter the environment and so be included in the inventory. Attenuation in the soil, and the actual distance that the contaminants travel are features not accounted for. The fact that contaminants usually only “breakthrough” (i.e. are seen to be emitted), after the natural attenuative capacity of the underlying soil and rock has been depleted, is therefore not factored in, leading to an overestimate of the emission.

Average leachate rates and qualities are used to approximate the mass of waterborne salts emitted from the dumps, stockpiles and containment dams. The average leachate rates and qualities are extrapolated from available borehole monitoring data, or the results from experimental column leach tests are used in the absence of field data. The seepage from containment dams and pans is assumed to have the same composition as the water in the dam. This method is approximate, as it averages highly variable systems and severely simplifies the mechanisms involved. It also gives an overestimate, as in using average borehole analyses it is not possible to determine which salts are present as a result of coming into contact with the waste material, or which are “naturally” leached from the rocks and soils. To correct for this to some degree, a “background” salt load is subtracted from that estimated for the system. The background salt load is an estimate of what would have been the naturally occurring

leachate to groundwater before the area was disturbed, and is calculated from the former drainage patterns of the area and the typical undisturbed aquifer quality.

A considerable shortcoming of the simple method described above is that it allows no link to be established back to the sources of effluent and solid waste within the process, i.e. it is based only on monitored average leachate qualities, and is not able to respond to a change in waste composition or effluent quality. It is therefore inappropriate for technology assessment, where such a link is essential. However, in the context of this study, where the LCI is being used to characterise the overall environment effects of an operation, this rough assessment was deemed better than missing a portion of the impacts altogether. A method able to reflect changes in the nature of the waste is proposed in section 5.4.2, and demonstrated in the case studies in chapters 7 and 8.

6.2.6.b Surface run-off

Stormwater run-off emissions are similarly calculated from rough estimations of run-off volumes and average monitored compositions. The approximate volume of stormwater run-off is calculated by specifying the percentage of rainfall for the region resulting in run-off (i.e. that not evaporating or seeping into the ground) according to the particular type of land area on which the rain falls (rehabilitated or non-rehabilitated ash, discard dump, station terrace etc.). Run-off that is not collected in storm water drains or containment dams is assumed to enter a natural water course. The quality of the run-off is taken from the average water quality in the storm water drains, where monitoring data is available, otherwise the quality in the containment dam is used. This is more inaccurate, as the salt concentration is likely to have been diluted by rainfall or concentrated by evaporation. As with the estimation of leachate emissions, the monitored run-off composition is not able to differentiate between those salts “picked-up” from the waste material and those from the soils. Thus, surface run-off emissions are also taken as the difference of what would have been present in the run-off of the undisturbed area and what is currently predicted to occur in the transformed area.

6.2.6.c Water Consumption

Waste dumps and stockpiles have a significant effect on water catchment management, as do opencast mining operations. In most well-managed operations, surface run-off from “dirty areas” is collected and retained in evaporation dams, as it is mostly of too poor a quality to be discharged into the natural streams. The mine and power station typically use this water in their processes, wherever the use of

poor quality water is permitted (e.g. dust suppression, cooling water make-up etc.). In addition, highly compacted discard and coal stockpiles, and ash dumps that have undergone pozzolanic reactions, have very low permeabilities, which heavily reduce rain water ingress and subsequent recharge of the underlying aquifers. The impact of waste management on water catchment has been taken into account by adding the difference between what would have been the catchment yield of the area, and the subsequent yield of the transformed area, into the “water use” category of the inventory. Although not all this water has technically been used by the process, its use by any other process has been prevented. It can therefore be interpreted as having been consumed by the system.

6.2.6.d Land use

The land occupied by the containment of solid waste is a more relevant metric of the impact of solid waste than the mass or volume of the waste (although this is only one component of the assessment, as emissions from the waste also require consideration). In order to be consistent with the land use categories in the background LCIs, the land disruptions occurring during solid waste disposal have to be assigned to one of the following four land use categories used by the ESU database:

- Land transformation II-III (semi-natural to cultivated)
- Land transformation II-IV (semi-natural to built)
- Land transformation III-IV (cultivated to built)
- Land use (IV-IV) (built to built)

The mines and power stations all fall into roughly the same region of South Africa (see Figure 6-1). This area has been extensively farmed, and few natural areas remain (Wells et al., 1992), so all land transformations are assumed to be from a “semi-natural” state. Two distinct stages in the dump life can be defined, i.e. the unrehabilitated stage, while it is still being added to, and the rehabilitated stage, when the dump is re-vegetated. In most cases, the rehabilitation process is continuous, i.e. the dump is rehabilitated from the one side, whilst the other side is added to. The area of these land types therefore change continuously over the life of the operation, with a constant rehabilitated land area on closure. Unrehabilitated ash and coal waste dumps are defined as equivalent to “built” areas (i.e. assigned to the II-IV land category), because of the severely degraded nature of the land (i.e. sterile, and practically impermeable surface areas). However, once rehabilitated, the rehabilitated dump area is defined as “cultivated” (i.e. assigned to the II-III land category). For the ash dumps, the land transformations are only considered for the life time of the power station, i.e. it is assumed that once the operation is

complete, all dumps will be sufficiently rehabilitated so that alternative uses can be found for the land (e.g. grazing for cattle). The discard dumps are assumed to be fully rehabilitated at the end of the active mine life, but a further period of “occupation” is considered until the dump has been “re-mined”, i.e. an estimate is made of the number of years until the discard will be utilised.

6.3. INVENTORY OF SOUTH AFRICAN ELECTRICITY

A listing of the environmental interventions calculated for the South African coal-fired power generation LCI is given in Table 6-3. The inventory is presented for a typical generating mix of all base-load stations. The load factors used to calculate the generating mix and some operating details of the stations are given in Table 6-2, and a full listing of all data inputs (as well as a more detailed presentation of the inventory) can be found in appendix F.1. A more detailed analysis of the inventories than that following below is given in appendix G.2, where the relative contributions from the foreground and background systems, and from the life cycle stages, are presented. A historic inventory is of limited use, and to provide relevant LCI data for the design of new products and processes, the inventory needs to reflect the future emission profile (Frischknecht, 1997; Weidema et al., 1999). An inventory was therefore developed to approximate the generating mix in the near-term (± 10 years), and is presented in Table 6-5.

The development of a future inventory is inherently uncertain, as not only is historic emission data used to project emissions in the future, but many assumptions about future patterns have to be taken into account. The inventory presented in Table 6-5 represents some of this uncertainty by presenting a range of values rather than a single mid-point value. The range incorporates the uncertainty in future electricity demand, and in the composition of the future generating mix (see section 6.3.2).

6.3.1. South African Generating Mix

The inventories of three stations representing the spread of technology options in the base load stations are also presented in Table 6-3. The technology options these stations represent are given in Table 6-4.

Table 6-3 A listing of the key environmental interventions of South African coal-fired power generation. The inventory is presented for a typical combination of the base load stations' annual output. The inventories of three representative plants, from which the average overall inventory is derived, are also presented (see Table 6-4 for some details of the three plants).

		1996	Wet / Wet	Dry / Dry	Wet / Dry
per MWhSO		Generating	Station	Station	Station
		Mix			
process energy ^a (hard coal)	MJ	400	300	530	330
process energy (natural gas)	MJ	2.8	3.1	1.4	5.1
process energy (nuclear)	MJ	0.46	0.48	0.24	0.74
process energy (oil)	MJ	83	93	41	150
land transformation II-III ^b	m ² .a	26	38	1.7	37
land transformation II-IV ^b	m ² .a	18	21	7.0	25
gas reserves	kg	0.053	0.061	0.026	0.094
hard coal reserves	kg	390	380	380	500
oil reserves	kg	1.2	1.4	0.62	2.2
iron reserves	kg	0.30	0.39	0.005	0.81
limestone	kg	0.083	0.043	0.010	0.30
sulphur reserves	kg	0.039	0.041	0.012	0.027
water use	kg	1500	2200	180	1900
CO	kg	0.15	0.27	0.047	0.27
CO ₂	kg	1000	980	990	1000
HCl	kg	0.002	0.002	0.002	0.002
HF	kg	0.001	0.001	0.001	0.001
N ₂ O	kg	0.0004	0.0005	0.0002	0.0008
NO _x	kg	2.7	2.7	2.5	3.3
SO ₂	kg	8.0	9.1	11	7.2
methane	kg	0.056	0.050	0.12	0.038
non methane VOC	kg	0.016	0.018	0.007	0.030
As	kg	0.0005	0.0006	0.0005	0.0005
Cr	kg	0.008	0.009	0.008	0.008
Cu	kg	0.002	0.002	0.002	0.002
Mn	kg	0.009	0.008	0.009	0.009
Pb	kg	0.004	0.004	0.004	0.004
U	kg	0.0006	0.0007	0.0005	0.0005
Total suspended particulates	kg	0.94	0.93	0.85	0.83
Steam/water vapour	kg	1520	2100	190	2200
Waste water	kg	330	470	300	130
Al (waterborne)	kg	0.002	0.004	7E-05	4E-04
Fe (waterborne)	kg	0.11	0.28	0.014	0.0002
Mn (waterborne)	kg	0.017	0.043	0.002	0.0004
Na (waterborne)	kg	0.045	0.046	0.012	0.091
Chlorides (waterborne)	kg	0.049	0.071	0.018	0.069
Sulphates (waterborne)	kg	1.2	2.6	0.23	0.47
Sulphite (waterborne)	kg	0.0005	0.0006	0.0003	0.0010
Oils & greases (waterborne)	kg	0.001	0.001	0.001	0.002
TDS	kg	1.7	3.6	0.300	0.75
Solid waste	kg	180	250	130	200
Hazardous waste	kg	0.001	0.005	2E-09	9E-09

- a. "Process energy" refers to energy consumed by the process itself i.e. the energy to produce the power sent out, but not the energy transformed to the product. The source of the energy is given in brackets. All foreground electricity use is assumed to be derived from coal (i.e. nuclear contribution from background processes only).
- b. Semi-natural to cultivated" (II-III) includes grassed over areas, rehabilitated dumps and surface area of dams and pans. "Semi-natural built" (II-IV) includes the station terrace and coal preparation plant, coal stockpile, unrehabilitated mine spoils and ash dump.

Table 6-4 Key differences between the stations for which the inventories are presented in Table 6-3.

	Wet / Wet Station				Dry / Dry Station				Wet / Dry Station			
Ash disposal	Ash dam (wet)				Ash dump (dry)				Ash dump (dry)			
Cooling water circuit	Wet cooling				Natural draught dry cooling				Wet cooling			
Flue gas cleaning	ESP and FF				ESP with SO ₂ injection				ESP			
Coal quality	C.V.	(MJ/kg)	23	C.V.	(MJ/kg)	21	C.V.	(MJ/kg)	17			
	Ash	(%)	28	Ash	(%)	30	Ash	(%)	38			
	C	(%)	58	C	(%)	52	C	(%)	44			
	H	(%)	3.5	H	(%)	2.9	H	(%)	2.6			
	N	(%)	1.3	N	(%)	1.4	N	(%)	1.1			
	O	(%)	6.8	O	(%)	9.3	O	(%)	8.8			
	S	(%)	1.0	S	(%)	1.1	S	(%)	0.6			
	H ₂ O	(%)	1.3	H ₂ O	(%)	3.3	H ₂ O	(%)	4.8			
Mine type	Dual product, opencast mine.				Dedicated, underground mine, with small opencast section.				Dedicated, opencast mine.			
Coal preparation	Crushing and washing				Crushing only				Crushing and washing			
Date commissioned	1980-1984				1988-1993				1985-1990			

The Wet/Wet station represents the older plant, employing wet cooling and wet ash disposal. It consumes much more water than either of the other two plants, and fairs worst in most of the water emissions, as the wet ash disposal method is associated with high salinity effluent seeping from the ash dam. In addition, this station is supplied by an opencast mine, which further exacerbates its poor scores in water emissions (see comparative mining inventories in Table 6-6).

The Dry/Dry station represents the newest plant to be built. It employs both dry cooling and dry ash disposal, which dramatically decreases its water requirements and water-related impacts. It is also supplied by a predominantly underground mine, so it fairs best in practically all emissions to water.

The Wet/Dry station, supplied predominantly by an opencast mine, falls somewhere between these two extremes. Although it consumes less water than the wet ashing station, it fairs only marginally better in most, but not all, water emissions. This is because the dry ashing method provides less of an effluent sink than the wet ashing method, and spills result if excess effluent occurs.

Table 6-5 Inventory of possible near-term future generating mix. The median assumes a 25% increase in coal-fired electricity demand, met by the current generating mix (assuming 1996 values) and re-commissioning in-storage capacity. The minimum/maximum range represents the 90% confidence interval over the scenarios explored to investigate the uncertainty in the future coal-fired energy demand and the generating mix.

	per MWhSO	Future Generating Mix		
		Median	minimum	maximum
Process energy (hard coal)	MJ	400	380	- 430
Process energy (natural gas)	MJ	2.9	2.6	- 3.2
Process energy (nuclear)	MJ	0.47	0.43	- 0.50
Process energy (oil)	MJ	87	79	- 95
Land transformation (II-III)	m ² .a	27	21	- 30
Land transformation (II-IV)	m ² .a	18	16	- 20
Hard coal reserves	kg	400	390	- 410
Oil reserves	kg	1.3	1.2	- 1.4
Iron reserves	kg	0.31	0.26	- 0.37
Limestone	kg	0.089	0.076	- 0.095
Water	kg	1500	1200	- 1600
CO	kg	0.15	0.13	- 0.20
CO ₂	kg	1000	980	- 1000
NO _x	kg	2.7	2.6	- 2.8
SO ₂	kg	7.9	7.6	- 9.2
Methane	kg	0.055	0.050	- 0.069
Non methane VOC	kg	0.016	0.015	- 0.018
As	kg	0.0005	0.0005	- 0.0005
Cr	kg	0.008	0.008	- 0.008
Mn	kg	0.009	0.008	- 0.009
Pb	kg	0.004	0.004	- 0.004
U	kg	0.0006	0.0006	- 0.0006
Total suspended particulates	kg	0.99	0.88	- 1.1
Steam/water vapour	kg	1500	1200	- 1600
Waste water	kg	340	300	- 360
Fe (waterborne)	kg	0.12	0.093	- 0.16
Na (waterborne)	kg	0.046	0.040	- 0.048
Chlorides (waterborne)	kg	0.052	0.045	- 0.058
Sulphates (waterborne)	kg	1.3	1.0	- 1.6
TDS	kg	1.8	1.4	- 2.2
Solid waste	kg	180	170	- 200

Since all three stations are mine-mouth stations, the transport emissions are predominantly from transporting non-coal fuels and materials to the mine and power station. Underground mining is assumed to use predominantly electrical energy, so has far lower diesel requirements than opencast mining. The Dry/Dry station thus has the lowest contribution from background processes, and consistently scores lowest in the interventions where the background processes dominate (see Table G-1 in appendix G.2.1). The Wet/Dry station has the highest ancillary material requirements, both in mining and power generation, and thus has the highest fuel requirements and background manufacturing burdens.

Table 6-6 A comparison of selected inventory items for underground and opencast coal mines.

per ton coal product		Underground Mine	Opencast Mine
Process energy (hard coal)	MJ	590	300
Process energy (natural gas)	MJ	0.38	6.3
Process energy (nuclear)	MJ	0.055	0.79
Process energy (oil)	MJ	12	210
Land transformation II-III	m ² .a	2.5	53
Land transformation II-IV	m ² .a	0.36	37
Gas reserves	kg	0.007	0.12
Hard coal reserves	kg	710	810
oil reserves	kg	0.18	3.1
Iron reserves	kg	0.001	1.3
Limestone	kg	0.010	0.054
Water use	kg	280	200
CO	kg	0.016	0.23
CO ₂	kg	57	73
NO _x	kg	0.16	0.6
SO ₂	kg	0.45	0.32
Methane	kg	0.26	0.059
Total suspended particulates	kg	0.045	0.56
Waste water	kg	580	200
Mn (waterborne)	kg	0.0006	0.0006
Na (waterborne)	kg	0.015	0.14
Chlorides (waterborne)	kg	0.005	0.095
Sulphates (waterborne)	kg	0.013	0.75
TDS	kg	0.002	1.2
Solid waste	kg	4.1	6.5

Although the Dry/Dry and Wet/Dry stations have higher auxiliary energy requirements (particularly the Dry/Dry station), their large unit sizes enable thermal efficiencies comparable with that of the Wet/Wet station to be achieved. The Wet/Wet station still produces less CO₂ per MWhSO as a result of the better quality coal burnt at this station. This is a consequence of the more recently commissioned stations being designed to burn poor quality coal, thereby allowing the exploitation of low quality coal reserves, and freeing up the high quality reserves for export. The older stations therefore perform better with respect to CO₂ and particulate emissions, and generate lower solid waste volumes, as a consequence of their burning lower ash and higher C.V. coals. Although this is balanced by their lower thermal efficiencies, and the fact particulate emissions depend not only on the ash volume, but also on the efficiency of the particulate collection device.

The older stations generally have less efficient ESPs, and although the worst of these have been replaced by FFs with very high collection efficiencies, they still tend to have higher particulate emissions than the newer stations. The Dry/Dry station is fitted with SO₃ injection to improve the

collection efficiency of the ESPs, so achieves low particulate emissions even with a high ash coal. Mining particulate emissions affect the overall particulate loading to a lesser degree, with open-pit mining generating significantly more dust than underground mining. The ash disposal method also affects particulate emissions, albeit to a much smaller degree, with dry ash disposal creating more of a dust problem than wet ash disposal.

The volume of solid waste generated also does not follow the ash content of the coal exactly, but is influenced by the thermal efficiency, the particulate removal efficiency and any ash sales occurring. The Dry/Dry station, although burning a slightly higher ash coal, produces slightly less ash requiring disposal than the Wet/Wet station, as both it and the Wet/Dry station sell a small portion of their ash. The high mass of solid waste for the Wet/Wet station reflected in Table 6-3 stems from high discard coal volumes, and its ash volume is in fact between that of the Wet/Dry and Dry/Dry station. The Wet/Dry station is burning an extremely high ash coal, and thus has the highest ash volume requiring disposal.

The power station coals are typically supplied without coal preparation (other than crushing), although the mine supplying the Wet/Dry station is mining such poor quality coal that a simple de-stoning wash is required. This has the side-effect of producing a low sulphur coal, which accounts for the lower SO₂ emissions of this station. The discard/shale produced from this wash is not included in the waste volume as it is used to backfill the opencast mine. The mine supplying the Wet/Wet station is a discard product mine and is producing both a discard that is used as backfill (not reflected in the inventory) and one which is dumped, causing the high waste volumes seen in Table 6-3. The allocation of this discard to the power station coal is questionable, and results from the simple mass-based allocation method used to allocate mining burdens between the two coal products (see section 6.2.4).

6.3.2. A Consideration of Uncertainty

The aim of this study is to provide an inventory of South African electricity that can be used in the life cycle design and analysis of South African products and processes. The study uses a combination of recent process-specific data, literature data and generic LCI data to produce life cycle inventories for three representative South African power plants. These inventories are then extrapolated to cover the potential future South African generating mix. This extrapolation is expected to introduce considerable uncertainty, in addition to the uncertainty in the inventory itself, i.e. that due to uncertainty in

variability in the data used to generate the three representative inventories. Additionally, uncertainty is associated with predicting the future generating mix. The latter is investigated in detail here. The extent of this uncertainty is then compared with that arising from uncertainty in the empirical parameters used to characterise the inventory models, where this is inferred from the uncertainty analysis presented in chapter 8.

6.3.2.a Uncertainty in Modelling the Future Generation Mix

The further into the future one tries to predict, the more uncertain the prediction generally becomes. This study is thus limited to a relatively short- to mid-term prediction of ten years. In this time-frame, coal is expected to still dominate the generating mix, the existing plants still be in operation and the coals used of similar quality to what they are now. The inventories generated from current data will therefore still be applicable.

A key determinant of modelling the future generating mix is the estimate of additional capacity that will be required. Coal-fired electricity demand is estimated to increase approximately 25% over the next ten to fifteen years, resulting in a most likely estimate of an additional 20 000 GWh required by the year 2010. The assumptions made to arrive at this estimate, and the estimated uncertainty in this value, are given in appendix G.3.1. The next step is to estimate how this additional power demand is to be met. Eskom has considerable excess capacity in storage, which is most probably sufficient to meet this demand, e.g. for the range in power demand estimated in appendix G.3.1, there is a 75% probability that no new capacity will be required if 80% of the capacity in storage is re-commissioned (see Figure G-3). Whether to meet the demand with a new plant, or to refurbish an old one, will depend on the environmental, economic and social pressures at the time, and is thus difficult to predict, and requires investigation in a scenario analysis.

The inventory presented in Table 6-5 is built up by combining three inventories representing:

- currently installed capacity;
- capacity brought out of storage; and
- new coal-fired capacity.

These are combined in varying proportions to approximate the future generating mix. Figure 6-3 shows selected results of the combined inventory calculated at three different percentages of the available in-

storage capacity brought back on line. The inventories used to represent the three capacity "types" are kept the same (the generating mix of 1996 is used to represent currently installed capacity, a slightly amended inventory of the Wet/Wet station to represent the capacity brought out of storage, and the inventory of the Dry/Dry station to represent new coal fired capacity). At a low percentage of the in-storage capacity re-commissioned, new capacity will most likely be required (e.g. if only 20% of the capacity in storage is re-commissioned, 60% of the predicted capacity shortfall will have to be supplied by a new plant). At higher percentages of the in-storage plants re-commissioned, it is likely that no new capacity will be required (e.g. re-commissioning 60% of the plants in storage is sufficient to supply the predicted capacity shortfall (most likely estimate), but taking the uncertainty of this prediction into account, there is a 25% probability that 35% of the capacity will have to be supplied by a new plant. If all the in-storage capacity is re-commissioned, there is only a very low probability (less than 10%) that new capacity will be required at all for the time frame investigated.

The inventories used to represent the three capacity "types" are investigated in Figure 6-4, which varies the inventories chosen to represent the installed, re-commissioned and new plants independently of each other in a sensitivity analysis. The main features of these figures are highlighted below, and are explained in greater detail in appendix G.3.2. Box and Whisker plots are used in Figures 6-3 and 4 and these are able to display the range in the data caused by the uncertainty in the prediction of the capacity required in the mid-term, as well as the most likely value (the median). The "boxes" give the interquartile range (the range spanned by the mid 50% of the data), whilst the "whiskers" give the total range of possible values.

The majority of inventory items is surprisingly insensitive to the percentage of plants in storage re-commissioned, with only those environmental interventions that differ very markedly between the old and new plants significantly increasing the uncertainty range over that due to the future energy demand (as shown by the range spanned by the "whiskers" in Figures 6-3).

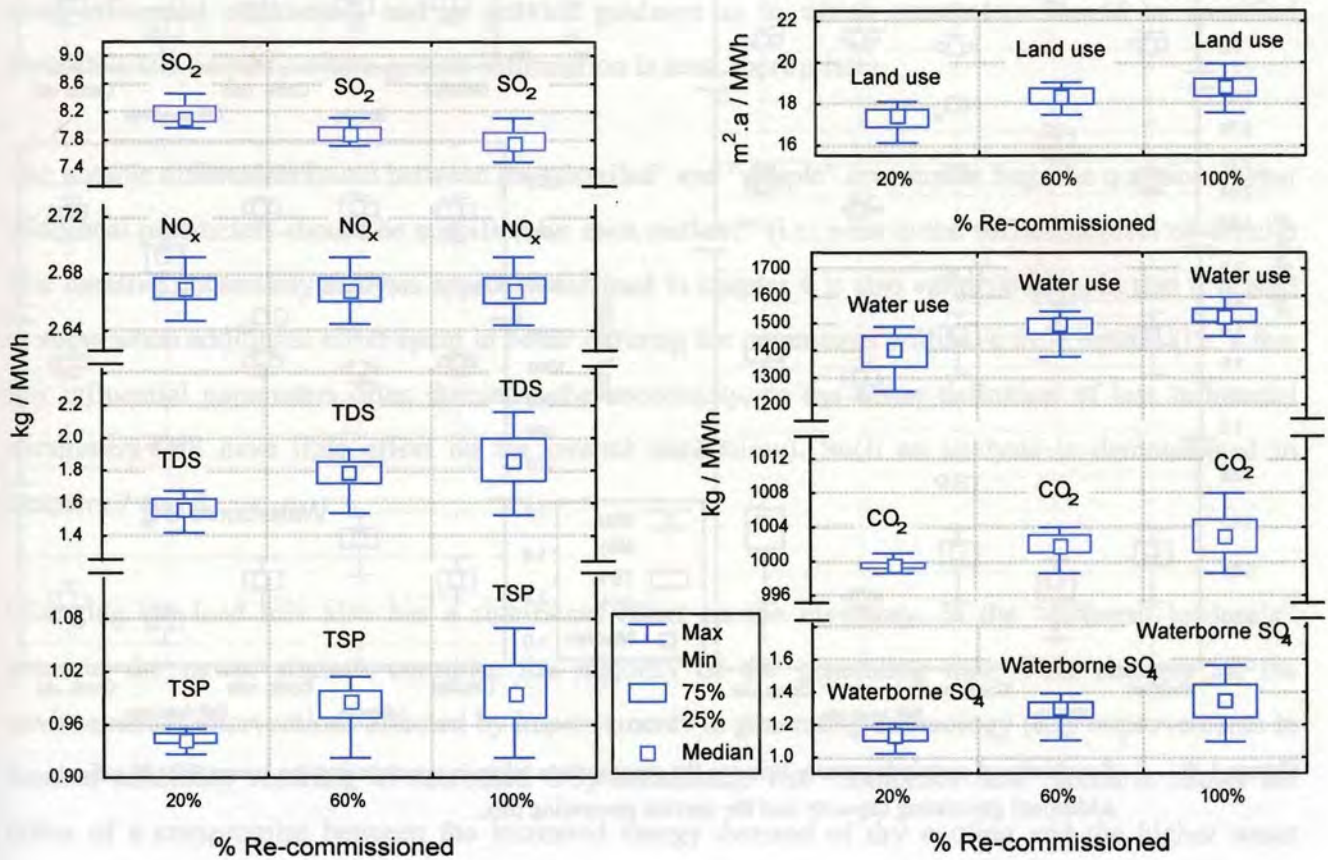


Figure 6-3 Sensitivity of selected environmental interventions to the % of plants in storage re-commissioned.

The sensitivity to the choice of inventory representing the installed capacity is shown by the difference between the “detailed”, “simple” and “different load mix” scenarios. “Detailed” represents the most likely scenario, with the grid inventory constructed from as detailed inventories as possible of each currently operating station. In the “simple” scenario, the grid is approximated by a combination of three generic plant-types, with no plant-specific details of the stations. The “different load mix” scenario takes the same individual station inventories as in the “detailed” scenario, but combines these according to a hypothetical future load mix, placing more emphasis on the newer plants (i.e. less contribution from the older plants, relative to the 1996 load mix).

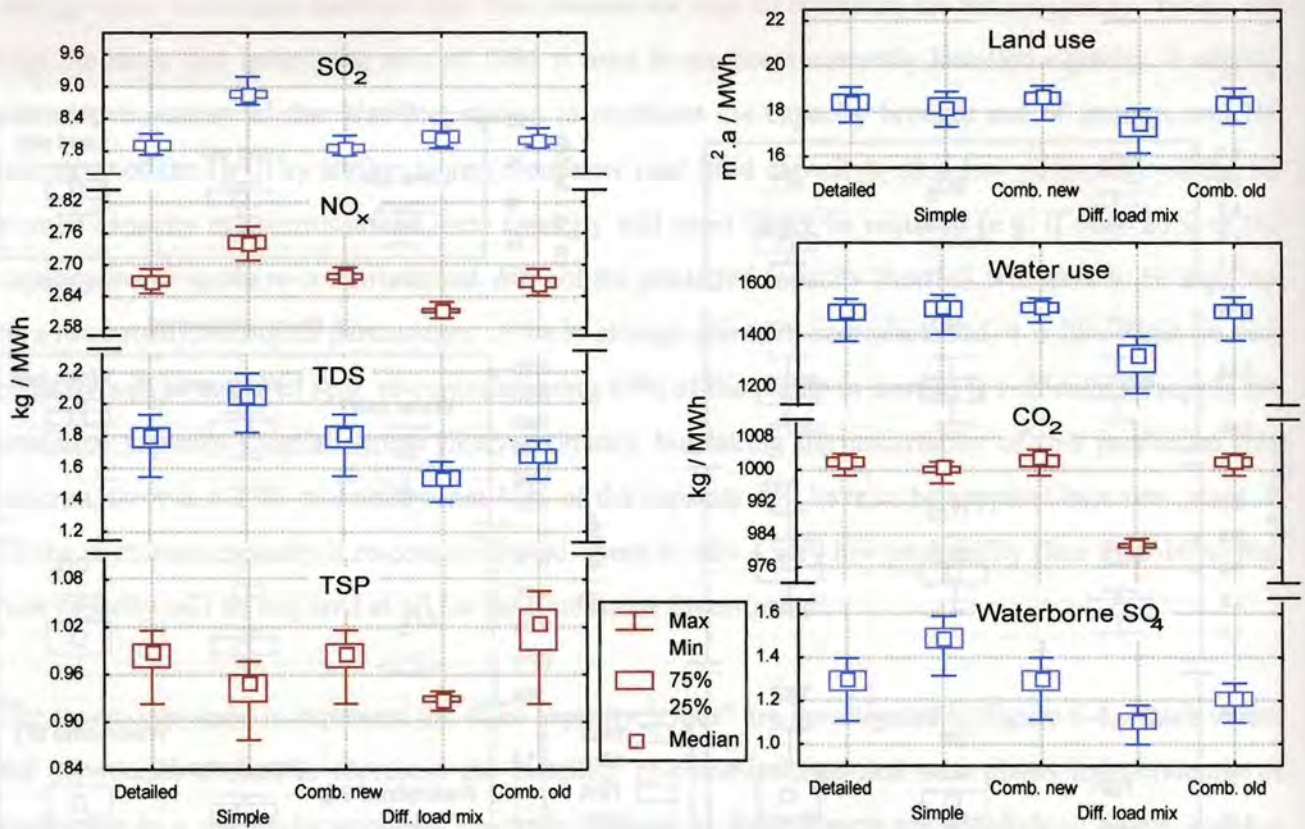


Figure 6-4 Sensitivity of certain inventory items to the assumptions taken in constructing the inventories to reflect additional generating capacity and the current generating mix.

The sensitivity to the choice of inventory used to represent new capacity is shown by the “combined new” scenario, which assumes a combination of dry/dry and wet/dry plants are built to fulfil the capacity shortfall, instead of the single Dry/Dry inventory, as in the other scenarios. The sensitivity to the choice of inventory representing the re-commissioned capacity is shown by the “combined old” scenario, which uses a combination of the inventories of the older wet/wet plants to get a better mix of coal qualities and mining methods, rather than the single Wet/Wet inventory used in the other scenarios.

The inventory of the future grid mix is most sensitive to the inventory representing the installed capacity, since this comprises the largest share of the total capacity. The differences between the “detailed” and “simple” inventory mixes show the significant differences obtained through more detailed modelling, and suggests that, at least, some plant specific data should be considered. However, this sort of “backwards” sensitivity analysis is of limited benefit, because other than giving insight into the uncertainty of the system, the effort in conducting a more detailed inventory has already been

expended. Of greater value is an uncertainty analysis, such as that described in chapter 4 (i.e. a stochastic analysis coupled with an uncertainty importance analysis). This can be used to identify the most influential parameters, and so provide guidance as to which parameters should be specified accurately (i.e. identify where generic information is least appropriate).

The notable differences found between the “detailed” and “simple” inventories begs the question, “what additional parameters should be specified for each station?” (i.e. what is the sufficient level of detail?) The iterative uncertainty analysis approach outlined in chapter 4 is also valuable here, in that it is able to show when additional effort spent in better defining the parameters will have little benefit (i.e. a few key influential parameters often dominate the uncertainty, so the better definition of less influential parameters will have little effect on the overall uncertainty). Such an analysis is demonstrated in chapters 7 and 8.

Changing the load mix also has a significant effect on the inventory. In the “different load mix” scenario the newer stations comprise the majority of the generating mix. This changes all the environmental interventions affected by improvements in generating technology (e.g. improvements in thermal efficiency resulting in decreased CO₂ emissions). The “combined new” scenario shows the effect of a compromise between the increased energy demand of dry cooling and the higher water requirements of wet cooling. This has a negligible effect for the particular scenario explored, and will have a more significant effect when a greater portion of the capacity required is assumed to come from new plants (e.g. a prediction further into the future, when all in-storage capacity has been brought back on line). The choice of inventory used to represent the re-commissioned plants is also less important. Although, generally more significant than the choice of inventory representing new capacity, it has less effect in most of the intervention categories than the choice of installed capacity inventory.

The choices governing the compilation of the future generating mix inventory are shown to introduce more uncertainty than that due to the uncertainty in energy demand, i.e. the range spanned by the scenarios in Figure 6-4 is greater than the range within the individual scenarios (as given by the “box” and “whiskers”). The choice of installed capacity load mix is predominantly responsible for this dominance, which is to be expected in a prediction of a relatively near-term generating mix.

6.3.2.b Data Uncertainty

The South African coal-fired electricity LCI was calculated for mid-points only, but an estimate of data uncertainty is achieved by extrapolating the uncertainty range found for a single power station. A full description of the empirical data uncertainty analysis conducted on the single power station can be found in chapter 8. The single station inventory is calculated using the same type and mix of data sources as the individual station inventories combined to form the South African (SA) grid LCI. Thus it can be assumed that they will have similar data uncertainty.

The total uncertainty range estimated for the future SA grid LCI is given in Table 6-7. The relatively high empirical parameter uncertainty of the individual station inventory results in data uncertainty dominating that arising from the choice of generating mix and the estimation of future energy demand. This can be seen by comparing Tables 6-5 and 6-7, where the range of values in the 90% confidence range is found to exceed that calculated previously for all environmental interventions, other than water use and sulphur dioxide emissions. This increase in range is especially significant for those interventions that are calculated from very uncertain data (few values, or non station-specific data) (e.g. trace metals), or from generic data (e.g. non methane VOCs).

The empirical parameter uncertainty of the SA grid LCI is approximated by fitting Gamma distributions to each inventory element, specified by the CV found for that environmental intervention in the single station empirical parameter uncertainty analysis, and the median value in Table 6-5. The Gamma distribution is chosen as it gives the closest fit to the uncertainty distribution of the simulated output, although for very uncertain interventions (i.e. those with high CVs), it tends to overestimate the uncertainty. This slight exaggeration of the data uncertainty is deemed preferable to an underestimate, as it accounts to some degree for the higher data uncertainty in those power station inventories which contain a greater proportion of generic data. The Gamma distributions are fitted to the combined grid inventory, and not to the individual stations, so the averaging effects of combining the individual inventories are not incorporated, also leading to a likely overestimate of the actual data uncertainty. The range given in Table 6-7 can thus be seen as a conservative estimate of the total uncertainty.

Table 6-7 Range of values calculated for inventory of possible near-term future generating mix, expanded to include data uncertainty. The confidence interval reflects the empirical data uncertainty, the uncertainty in future coal-fired energy demand and the future coal-fired generating mix.

	per MWhSO	Future Generating Mix	
		minimum	maximum
Land transformation (II-III)	m ² .a	18	- 33
Land transformation (II-IV)	m ² .a	12	- 25
Hard coal reserves	kg	340	- 440
Oil reserves	kg	0.31	- 3.2
Iron reserves	kg	0.086	- 0.72
Limestone	kg	0.025	- 0.20
Water	kg	1200	- 1700
CO	kg	0.046	- 0.33
CO ₂	kg	850	- 1100
NO _x	kg	1.3	- 4.5
SO ₂	kg	6.6	- 9.2
Methane	Kg	0.023	- 0.11
Non methane VOC	kg	0.003	- 0.049
As	kg	0.0002	- 0.001
Cr	kg	0.001	- 0.027
Mn	kg	0.002	- 0.024
Pb	kg	0.001	- 0.010
Total suspended particulates	kg	0.45	- 1.6
Sulphates (waterborne)	kg	0.23	- 3.4

6.4. CASE STUDIES

To conclude this chapter two applications of the South African inventory are briefly presented.

6.4.1. Comparison with European LCA data

The graphs in Figure 6-5 show a comparison between the inventories calculated by this study and four European electricity inventories, taken from Ökoinventare für Energiesysteme (Frischknecht et al., 1994). The average South African generating mix is compared to a European average mix and a UK average mix. To obtain insight into the coal-fired components of these averages, the inventories of two stations representing the “best” and “worst” of South Africa’s plants are compared to two European coal plants, a modern station in France, and an older “poor” station in ex-Yugoslavia. For ease of presentation, impact assessment results are presented here for only two well accepted impact categories, acidification and global warming. The inventory for South Africa includes only the coal component of the generating mix (i.e. the burdens from the relatively small nuclear and hydro contributions are not reflected).

The dominance of coal in South Africa’s generating mix is clearly visible in the comparison, with the South African inventory showing the highest utilisation of coal reserves, with negligible gas and oil consumption. The reliance of South Africa on coal is also shown by the relatively high contribution to acidification and global warming. The UK’s generating mix contains a higher fossil-fuel component than the average European mix, and thus a higher potential contribution to acidification and global warming. Even inflated by incorporating catchment losses into the inventory, South Africa’s commitment to water efficient power generation is clearly visible in the comparison.




Figure 6-5 Comparison of South Africa’s electricity grid (SA) with that of Europe and the UK, and a comparison of two South African power plants (an older station with wet cooling and wet ash disposal, and a modern station with dry cooled and dry ash disposal) with two European stations (a modern French plant, and an older Eastern European plant). Data for the comparisons is from the ESU database (Frischknecht et al., 1994), as documented in the PEMS LCA software (PIRA, 1996).

Surprisingly, given the high ash content of the Southern African coals, the analysis shows marginally lower total TSP emissions for the South African grid than the UK average inventory. This can probably be attributed the use of limestone sorbent for FGD in the UK power stations, as limestone mining is associated with high particulate emissions. Unfortunately no TSP emissions are reported for the European stations, so the individual coal-fired stations can not be compared across this category. South Africa's high contribution to land transformation is attributed to the high volumes of ash produced, and the consequently large areas required for ash dumps. The increasing number of land-intensive opencast coal mines also contribute to this high land use. The high land use of opencast mines is shown by the marked difference between the two South African stations.

On an individual coal-fired power station level, the potential contribution of South Africa's stations to global warming and acidification is very competitive to that of Europe. The very large discrepancy between the modern and old European station is probably due to the modern station's incorporating FGD. The South African stations do not have FGD units, but the relatively low sulphur Southern African coals results in their releasing only marginally more sulphur dioxide per MWhSO than a plant with FGD. The South African stations show a marginally better performance in global warming potential than the modern European plant. This is attributed to the low methane emissions in South African coal mines, as the South African inventory actually shows marginally higher CO₂ emissions than the modern European plant. Also, the South African system has very low transport emissions, as all of the stations are mine-mouth, further reducing their hydrocarbon emissions. The relatively high contribution to global warming from the "poor" European station is most probably due to the older station's low thermal efficiency.

The South African stations are shown to consume less coal reserves than their European counterparts. This is not because they are more thermally efficient, but is a consequence of their burning near-discard quality coal. In the South African inventory, the mass of coal reserves utilised is calculated by dividing the energy extracted by a reference calorific value of coal reserves (taken as 29 MJ/kg). The comparison is somewhat questionable, as it is not known how the reserves in the ESU database are calculated. The actual mass of material burnt is not a meaningful indicator of reserve depletion, as the South African stations would then show far greater values than the European stations, and not the fact that by burning such poor quality coal they are freeing up higher quality coal for other uses. The higher oil reserve consumption observed for the European stations is attributed to higher coal transport distances, since

the South African power stations are all mine-mouth. Both of the South African power stations are more water-efficient than the two European stations, with the considerable water savings of dry cooling very evident.

The South African inventory shows disproportionately high TDS emissions relative to Europe and the UK. This is most probably due to underestimation of TDS emissions in their inventories, as they show far higher emissions than the South African inventory in other water emission categories, particularly chlorides, sodium and nitrates. Their sulphate emissions are also anomalously low, which may also be due to an error in the database, as the individual coal-fired stations show fairly high sulphate emissions. No TDS values were reported for the individual European stations, although they show higher emissions than the South African stations in most water categories. This can be attributed to the ZLED policy of the South African power stations and many of their tied collieries, which requires that all effluent and storm water run-off be collected and retained on site, where it is either re-used or evaporated in containment dams. Thus very low surface emissions to water occur, with most of the emissions reflected in the inventory due to leachate from the dumps, and seepage from evaporation dams.

6.4.2. Electricity Input into a Zinc Refinery

The high electricity consumption of minerals refining offers some interesting opportunities to explore the application of the SA electricity LCI in Type I decision situations, i.e. as background information into a separate decision system. In particular there is the opportunity to explore the possibility of whether the electricity LCI used could change the outcome of the study. In addition to their high electricity requirements, minerals processing technologies are interesting in that they encounter some classic “problem shifting” situations. Newer technologies often employ electricity rather than burning coal on site, which shift burdens to the site of electricity production, whilst innovative technologies, such as bio-processing, change the nature of the wastes, and can potentially shift burdens from one release medium to another. This section examines a small aspect of a study presented by Stewart (1999). The purpose of the case study was to select the optimal technology for a proposed zinc refinery in the Eastern Cape in South Africa.

The current excess generating capacity in South Africa has led to Eskom providing incentives to encourage local electricity-intensive projects. An example of such a project is the Hillside Aluminium

Smelter in Richards Bay, in which the cost of electricity is linked to the selling price of the metal, thus essentially fixing the profits of the process. The electrowinning of base metals is a similarly electricity intensive process, although the lower throughput of material means that their total electricity requirements are not nearly as high as aluminium purification. Nonetheless, a zinc refinery was proposed, and a detailed and extensive EIA process ensued. Although the project did not go ahead, the technology assessment study showed the kind of considerations that needed to be taken on board in selecting the optimal technology combination for the process.

The initial phase of the technology assessment identified the possible flowsheet options shown in Figure 6-6. Both a single technology route, and a combination of technologies were considered. Goal programming, considering a number of economic, social and environmental criteria, was used to determine the optimum flowsheet. Details of the technologies considered and the criteria used in their selection can be found in Stewart (1999). The analysis found that a combination of technologies was always preferred, with the best option a 44:56% split between pressure leaching and roasting, with all of the roast product treated in the Imperial Smelting Process (ISP). A considerable drop in performance resulted when the single technology route was chosen. When considered singly, the pressure leach/electrowin process was favoured, with the roast/leach/electrowin and roast/ISP processes scoring very similarly for the attributes considered (roast/leach/electrowin was marginally favoured).

The contribution of electricity provision to the total burdens for four key environmental interventions is shown in Table 6-8. The model in Stewart (1999) only considers electricity use in the electrowinning process, so no electricity burdens are reflected for the ISP route. In addition, only the major flows through the zinc refining process are considered (feed extraction and processing, ancillary materials, and transport are not included), so the roaster is the only process, other than electricity production, for which emissions to air are reflected. The contribution to the total emissions from electricity production are surprisingly small. Only water consumption shows a notable contribution, since the zinc refining technologies are not terribly water intensive. The very high emissions of CO₂ and SO₂ from the roasting process completely dominate those from electricity production. It is thus only in the pressure leach process that CO₂ and SO₂ emissions from electricity provision are significant (all are from power generation), although the actual quantities are very small compared to those from the roasting process. The high mass of solid waste produced by zinc refining means that electricity provision is insignificant with respect to the total mass of solid waste generated.

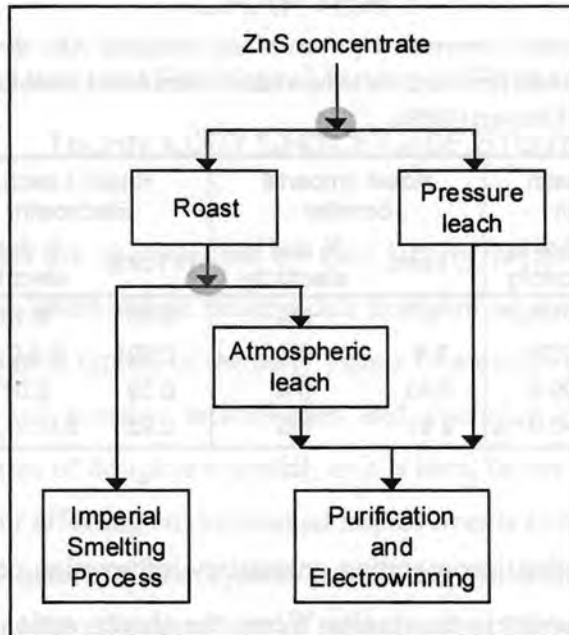


Figure 6-6 Flowsheet options for zinc refining. The shaded circles show where possible stream splits can occur (Stewart, 1999).

The contribution of electricity burdens to each refining process was calculated for the range of values given in Table 6-7, reflecting a range of possible SA electricity LCI combinations, and the uncertainty underlying them. Surprisingly, given that the range in Table 6-7 is fairly extensive, this was found to have no effect on the total values found for the environmental interventions considered. Even for these interventions where the contribution from electricity production is high (water consumption, and CO₂ and SO₂ for the pressure leach process), the quantities are sufficiently low that at the level of reporting (two significant figures), the effect of the range is insignificant. The uncertainty in electricity production is thus not sufficient to change the outcome of the technology selection process. However, the uncertainty in the modelling of the zinc refining process is not investigated. A single value of electricity consumption by the electrowinning process is considered (3.6 kWh/ton zinc produced). If this value were to be higher, uncertainty in electricity production may well be shown to be significant.

Table 6-8 Selected environmental interventions per ton of zinc produced. Also shown is the percentage of the total arising from electricity provision, for the particular intervention considered. Calculated using the models and assumptions of Stewart (1999).

	Pressure Leach, Electrowin		Roast, Imperial Smelter		Roast, Leach, Electrowin		Optimum Technology Combination	
	t / t zinc	% due to electricity	t / t zinc	% due to electricity	t / t zinc	% due to electricity	t / t zinc	% due to electricity
Water consumed	0.05	10-12%	0	0%	0.05	9-12%	0.03	10-13%
Carbon dioxide	0.004	100%	3.4	0%	0.90	0.3-0.4%	1.5	0.1-0.2%
Sulphur dioxide	2.8e-5	100%	0.40	0%	0.39	0.01%	0.18	0.01%
Solid waste	0.98	0.06-0.07%	0.91	0%	0.93	0.07-0.08%	0.95	0.04%

The above brief study shows that incorporating uncertainty information on the background LCI used allows for a greater sense of surety in the results. Where the chosen option can be shown to be robust across the known range of background information, the choice of LCI to use becomes less significant (if a mid-point analysis is to be pursued). However, not only should the sensitivity to the background LCI be checked, but variations within the foreground should be considered as well (i.e. the operating range and the empirical variability within the process(es) under consideration), since these affect the quantity of background product/service input into the foreground. The following two chapters explore case studies where both foreground and background uncertainty are taken into account.

CHAPTER 7

LCA MODELLING FOR DECISION MAKING IN PRIMARY INDUSTRIES: TECHNOLOGY CHOICE CASE STUDY

Strategic and tactical studies are characterised by their lack of precise information, as both require modelling of future systems, where at best, process data from similar systems can be extrapolated to the future system. This constraint is typical of the early phases of a design study, where the first step is to choose among a number of possible technologies and operating conditions. An environmental component in the early phases of design is essential, as it is here, before the design becomes fixed, that the greatest opportunities for affecting environmental improvements exist. However, in the early stages of design the unspecified nature of the system means that full-scale LCA is severely restricted (Azapagic, 1999; Keoleian, 1993). The emphasis of the LCA therefore needs to be on scenario analysis and uncertainty assessment.

This chapter looks at the use of LCA to inform strategic decision making in primary industries, and presents a case study on selecting the optimal technology and operating conditions for re-powering an in-storage power station. The case study is broken into two parts, and represents the progression from a strategic to a tactical decision level (and ultimately down to an operational level) in a design study. The first phase of the study is typical of a strategic technology scanning study, where all relevant technologies are investigated. This is consistent with a screening or first-pass LCA, and requires the evaluation of a large number of options, about which little specific information is known. The inventory is therefore calculated at a relatively low level of resolution, incorporating predominantly generic LCI data. The next phase of the study is typical of the early phases of a design study, where the choice has been narrowed, and the chosen options explored in some detail. Here the system is much more precisely defined, which allows for more detailed data sets to be collected, and a more detailed inventory to be constructed, i.e. it focuses on a particular technology, with the options narrowed to the possible operation of the chosen technology, instead of spanning the “typical” operation of a number of technologies. This phase is typical of that in the early stages of design, where the broad outline has been drawn, although the actual configurations of the system are yet to be decided.

7.1. DECISION CONTEXT

The features of tactical or strategic decisions are discussed in section 2.3.2. These features are consistent with those demonstrated by this case study, as the study is conducted without a specific implementation date or location in mind, and involves elements wider than the primary sphere of influence of the decision maker, i.e. it involves aspects not directly under the control of the decision maker (in this case the power supplier) (e.g. the supply, availability and quality of future discard coal sources). The case study system therefore shares the properties of tactical/strategic decision systems described in section 2.3.2.

Studies to support tactical and strategic decisions are distinct from operational and “historical-type” studies primarily with respect to the quantity and quality of data available to them. Tactical/strategic studies require data to characterise a system different to that from which the data is collected, and one to be implemented some time in the future. The other two study types primarily use data from the actual system of interest to characterise a current system, or one where changes occur in the short term. Tactical and strategic studies thus require data to be predicted for the relevant mid or long term marginal technology, and are associated with a significant amount of inherent uncertainty, especially as strategic decisions covering long time horizons often require that developing technologies be incorporated into the assessment. Coupled with this high empirical uncertainty, is the high uncertainty in model parameters (decision variables) typically encountered in tactical and strategic studies, as a result of the, as yet, loose definition of the system. A feature of tactical studies is thus their emphasis on scenario analysis, and the exploration of possible operating states.

Although the case study is conducted without a specific power station in mind, some site specificity is introduced, as the locations of the power stations currently in storage are known. The older stations that are no longer operating, but have the potential for being reconditioned and brought back into service, all fall into a fairly localised region of the country, which mirrors the localised coal-producing region (see Figure 6-1). Site specific considerations therefore play a more significant role than usually found in tactical studies in the selection of the relevant impacts to be considered, as well as in the importance attached to these impacts.

7.2. PRELIMINARY TECHNOLOGY SCANNING STUDY

The aim of this study is to evaluate technologies able to reduce SO₂ emissions from coal-fired power generation. A preliminary technology screening takes place very early on in a design study, and is used to narrow down the choices to two or three options, which can then be designed in detail, and on which full-scale LCA studies can be conducted. The study is thus required to look broadly at all possible manners of sulphur-removal, relevant to all possible plants and locations.

Technologies developed to reduce SO₂ emissions demonstrate well the dilemma of choosing between disparate technologies offering relative degrees of improvement in environmental performance, and which result in a trade-off between environmental impacts. They range from 'end-of-pipe' control strategies, such as FGD systems, to modifying the combustion conditions, or burning a "cleaner" coal. LCA provides a systematic methodology to organise information about the technologies on a consistent basis, enabling their incorporation into an optimisation strategy, where together with economic and social information, the optimal technology for a particular situation can be identified. The study is thus intended to provide an environmental argument to the usual economic assessment in the preliminary technology-screening phase of a design study.

A full account of the study can be found in appendix H-1, with only a brief overview presented here.

7.2.1. System Definition

The objective of the study is to identify feasible technologies to reduce SO₂ emission from a coal-fired electricity plant, and to evaluate the relative performance of these technologies. A life cycle basis is required to evaluate the technologies because all shift the sulphur burden within the life cycle, i.e. all the technologies result in lower emissions of SO₂ from the power station stack, but at the expense of higher solid and/or liquid waste volumes.

As comprehensive a range of technologies as possible was considered. Three possible intervention "sites" to effect sulphur-removal in the coal-electricity conversion chain were identified, and two representative technologies chosen for each of the three sites, selected according to their applicability to the South African coals and conditions, as well as their level of demonstration. The technologies were required to be equally applicable as a retro-fit on an existing plant, or built with a new plant, i.e.

they had to be compatible with the existing technology employed on the South African power stations.

The following technologies were investigated:

- Before combustion: Conventional dense medium coal washing (at medium and high density), and microbial desulphurisation of the coal.
- During combustion: Combustion in a fluidised bed boiler with in-situ desulphurisation (atmospheric bubbling bed, and circulating fluidised bed).
- After combustion: Flue gas desulphurisation (wet limestone process, and a lime spray dryer process).

The inventories developed are for typical operating conditions, and are intended to give a measure of average operating performance. They are therefore calculated using mid-point values, representing the “most likely” performance of the technologies. The major assumptions and data sources are given in appendix H.1.2.

A significant problem with modelling data deficient systems is the danger that emissions that cannot be quantified are simply ignored. This is a concern when dealing with systems for which the data is not of the same level of completeness, as this leads to the under-specified system appearing better than it is, and the more comprehensive system being prejudiced by its more complete inventory. It is therefore essential that the inventories be specified to the same degree of completeness, i.e. if an environmental intervention is specified for one system, it must be specified for all the systems being compared. A quantitative analysis is therefore inappropriate for interventions which are known for only some of the options, and the comparison should rather be limited to a qualitative assessment of that intervention.

An alternative approach, is to estimate the intervention for the systems where it is not known, and place high uncertainty on this estimate. However, a full quantitative data uncertainty assessment is not feasible at an early screening phase of a design study, so uncertainty management of an estimated data input has to be via a qualitative uncertainty assessment. The simplest of these is placing a “flag” on the uncertain data element (the so-called “post-it-note” approach). However, this approach was not pursued because of problems with aggregating data “flags” through the assessment, and the danger that the estimated data give a false sense of completeness. Environmental interventions that are not known for all options are therefore excluded from the quantitative analysis, and only considered qualitatively. A qualitative data uncertainty matrix, using a pictorial representation developed by Graedel and Allenby (1995), is used to provide an indication of the quality of the data used in the quantitative component of

the study. An example of the matrix is shown in Table H-6 in appendix H.1. The overall “blackness” of the matrix gives an indication of the overall data quality according to two DQIs, completeness and confidence level. The “completeness” indicator highlights data gaps, whilst the “confidence” indicator encompasses the applicability of the data to the study in hand and the reliability of the data source.

The significant restrictions on data availability at a technology-screening stage of design mean that quantitative LCA may not be sufficient to fully characterise the impacts of the system, and a qualitative assessment is required to fill in the gaps. For example, the quantity and quality of leachate generated from waste landfills is highly site-specific and depends on complex characteristics of the waste and the landfill site. A quantitative evaluation of landfill emissions is therefore not feasible at this level of analysis, and consequently only the volume of waste produced is quantified. The nature of the wastes is especially relevant in this study, as the sulphur burden is transferred from an air-borne emission to a solid waste in all the technologies considered. The quantitative comparison alone is thus not sufficient for these systems, and a qualitative assessment is required to incorporate the impacts from solid waste into the assessment.

A qualitative assessment is therefore used in conjunction with the quantitative LCIA. Although more sophisticated semi-quantitative LCA approaches have been developed (Graedel, 1998), a qualitative approach developed by Graedel et al. (1995), using the same pictorial representation as the qualitative data uncertainty matrix, is used. This method is chosen because it keeps the quantitative and qualitative components of the study distinct from each other (i.e. no confusion can arise between a semi-quantitative matrix score and the inventory results). The qualitative impact assessment matrix encompasses a more comprehensive range of impacts than the quantitative LCIA, and estimates the certainty of an impact occurring in a number of broad environmental impact categories, together with the level of concern attached to the impact. The “certainty” indicator gives an estimate of the uncertainty of the prediction, whilst the “concern” indicator is informed by both the predicted extent of the impact (e.g. volume of emissions), and the degree of significance attached to the impact. An example of the qualitative impact matrix is shown in Table 7-1 for two of the systems considered. Where possible, the qualitative indicator is informed by the quantitative inventory and LCIA results.

7.2.2. Selected Results and Conclusions

Comparisons between the technologies at each intervention site were first conducted, from which the best technology from each intervention site was chosen (see appendix H.1.3). Figure 7-1 shows the chosen sub-set of technologies: coal treated in a high-density washing process, PF boiler replaced with a circulating fluidised bed boiler (CFBC), and flue gases treated in a lime spray dryer process (SDA). A combined system is also presented, in which the discard produced by the medium density wash system is burnt in the CFBC system, thereby avoiding the waste of useful energy and the impacts from stockpiling the discard. The systems are compared on a basis of the annual power output from a typical South African power plant, and are all defined as if retro-fitted to a refurbished older station.

From this first-pass analysis the FBC system appears to be the technology of choice. It shows a marked decrease in acidification and eutrophication potential, whilst showing minor improvements over the other systems in all other categories, except dust emissions and landfill volume. These minor improvements are principally due to its better efficiency relative to the older PF boiler, and a different relative profile is obtained if compared to a modern or extensively refurbished PF system (see Figure H-6 in appendix H). The CFBC system is also favoured in a qualitative assessment of the impacts from the different wastes produced by the systems. A detailed interpretation of the results can be found in appendix H.1.3.d.

A particular strength of the CFBC technology is its ability to burn low quality fuels, whilst the most significant disadvantage of the washing technology is its production of discard. This suggests the use of the combined system shown in Figure 7-1, which results in a more efficient system utilising all the coal mined, and replacing the high risk discard waste with a lower risk CFBC waste. The performance of the combined system falls between that of the medium density wash system and the CFBC system for most of the impact categories. It falls closest to the medium density wash system, since this system supplies the bulk of the power in the combined system (the volume of discard produced is much smaller than the volume of coal produced, so its combustion adds only a small portion of the power produced by the combined system). Even burning a very poor quality discard coal, the “improving” effect of the CFBC system can be seen by the lower contribution to acidification and eutrophication of the combined system relative to the medium density wash system alone.




Figure 7-1 Comparative performance of the most promising technologies to reduce sulphur emissions, relative to an older PF station with no additional environmental controls, and burning ROM coal.

The most significant feature of this combined system is the removal of the discard and the “avoided” burdens which result. This highlights the considerable limitations of reporting only the quantity of waste, and not the potential risks associated with its disposal, as these “avoided” burdens are not reflected at all in the quantitative analysis. Table 7-1 presents a qualitative comparison of the medium density washing system and the combined system. The qualitative assessment is essential, as it reflects the positive effects of converting the polluting discard dumps to the more benign FBC waste, and not merely the benefits of utilising an otherwise wasted energy source in a cleaner combustion technology.

The above analysis shows that quantitative LCA alone can not reliably inform decision making in information deficient systems. At low levels of detail it can potentially give misleading results, especially if an influential aspect of the systems is omitted, or if one system is more complete than the others. In addition, the inability of LCIA to assess the risk of site-specific impacts results in misleading comparisons of technologies in which waste management impacts are important. A combination of a quantitative and qualitative approach is therefore recommended for technology scanning studies of primary industries, where a qualitative assessment is backed up by whatever quantitative information is available.

Table 7-1 Qualitative assessment of the effects of the conventional coal washing system (coal washed at medium density, the coal burnt in an older PF boiler, and the discard stockpiled) and the combined system (coal washing system as before, but the discard burnt in a CFBC system with limestone desulphurisation).

	Medium Density Coal Wash				Combined Washing and FBC System			
	Resource extraction	Feedstock preparation	Power generation	Waste management	Resource extraction	Feedstock preparation	Power generation	Waste management
Resource consumption								
Local air impacts								
Atmospheric impacts								
Water contamination								
Soil contamination								
Land degradation								
Human toxicity								

Explanation of symbols used in impact tables

Concern / Certainty	Negligible	Minor	Moderate	Significant
25%				
50%				
75%				
100%				

The qualitative assessment matrix highlights the problem areas of the system, and can be used to identify sites in the life cycle system which could possibly “swing” the results of the study. This identifies the sites where qualitative reasoning must be used to determine the relative impacts of the systems, or if the scope of the study allows, where quantitative information needs to be generated. In this study, the need for a quantification of the impacts associated with disposal of the various sulphur-containing residues has been identified.

Although inventories generated for strategic level decision support are highly uncertain, quantitative uncertainty analyses are infeasible because of the large number of options requiring consideration, and the low level of data availability. The qualitative uncertainty matrices used here are therefore more practicable, with the qualitative impact matrix able to provide an indication of the uncertainty of the

impact occurring, and the data uncertainty matrix aiding interpretation of the quantitative results. Whilst it may be possible to obtain an estimate of the empirical data uncertainty of the inventory, especially should it become more commonplace for generic LCI databases to be published with uncertainty estimates, this is likely to still be inadequate, since the uncertainty due to the choice of decision variables is likely to dominate in a strategic-level study. An assessment of model parameter uncertainty is not feasible in a study with many loosely defined options, as the number of scenarios to consider becomes unmanageably large. The general lack of information at this level of study also makes it difficult to define feasible operating ranges over which the model parameters should be specified in a sensitivity or parametric analysis. The determination of all options at their “typical” operation states, and a qualitative estimate of the uncertainty therefore appears to be all that is possible at this level of analysis.

A systematic uncertainty analysis is therefore postponed until the next phase of the design, where the smaller number of options and the increased level of information make this possible. The rest of this chapter considers such a subsequent step in the design process, and presents the challenges to including a full quantitative uncertainty analysis into technology assessment in primary industries. It also aims to address the most significant gaps in the inventory, notably around solid waste disposal.

7.3. AN INVESTIGATION INTO THE COMBUSTION OF DISCARD COAL IN FLUIDISED BED BOILERS AS AN OPTION FOR RECONDITIONING OLD POWER STATIONS

The potential of FBC as a technology for reducing SO₂ emissions is identified in the preliminary technology scanning study. Of particular importance is its potential for burning colliery discards, thereby avoiding the significant environmental impacts of stockpiling the discard. However, burning discard in a fluidised bed results in a trade-off between impacts, primarily between the water and air pollution from the stockpiled discard, and the emissions to air and ash waste produced when the discard is burned. This case study investigates the combustion of discard coal in a reconditioned power station in considerably more detail than the preliminary investigation. The study again takes place at an early stage in the design process, and the operating conditions of the system are not yet fixed, although the environmental profile of the system is expected to vary considerably with differing operating conditions. The study has therefore been reformulated to some degree, and aims to determine the conditions under which the system need operate to achieve a particular level of certainty in net environmental benefit, rather than characterising the environmental performance of a “typical” operation. The key question to be answered has therefore been formulated as:

“Under what operating conditions is it environmentally beneficial to re-power an old PF station with an FBC boiler burning discard coal?”

This particular formulation of the question focuses the study on a parameter analysis, rather than trying to pre-define a number of tightly defined scenarios. In this way, significant decision variables (e.g. the quality of the discard sourced, the type of sorbent used etc.) are investigated as part of the model parameter uncertainty analysis, rather than specified in discrete scenarios “up front”.

The primary system to be modelled is an old station with reconditioned boiler units burning discard coal. Also requiring definition are the scenarios for comparison. To determine whether the FBC system is “environmentally beneficial” requires comparison against the power generating options it is displacing. Different scenarios are possible, as the comparative basis changes depending on the driver for the project. Most likely is that the station is being re-commissioned because the additional capacity is required for the grid, i.e. if the station is not re-powered using FBC technology, the capacity will be supplied from some other source. Possible other sources include re-commissioning the station as a

conventional PF station, building a new PF station, or operating the existing stations at higher loads. The first scenario is complicated by the fact that the extent to which the station is refurbished will significantly influence its environmental performance, and the comparison will have to take this into account. Alternatively, the driver behind the project may not be the requirement of new capacity, e.g. the primary aim of the project could be to remove the discard dumps, or political drive to demonstrate a “clean coal” technology. In this case, re-powering the station would displace existing power off the grid, and a relevant comparison would be between the re-powered FBC station and the average grid mix.

7.4. METHODOLOGICAL FRAMEWORK

This case study follows the methodological framework proposed for the systematic investigation of uncertainties in LCA models, summarised in section 4.4. The overall approach is schematically represented in Figure 4-1, which illustrates the layered approach, in which empirical uncertainties, model uncertainties and uncertainties in model form are investigated in a looped fashion. An outline of the key steps in the uncertainty analysis is given in Figure 4-3.

A necessary first step is the definition of the scenarios to be investigated. A suitable model must then be constructed to determine the environmental profile of the scenarios. This encompasses the usual decisions taken during the goal definition and scoping phase of LCA. Since the study involves the comparison of alternative systems, it is also necessary to define criteria against which they can be rated (e.g. impact categories or selected environmental interventions). This section presents the major considerations in the definition of these structures.

7.4.1. Scenario Definition

The scenarios to be investigated are derived from the problem statement. Investigated here is the first scenario outlined in the problem statement above (i.e. the extra grid capacity is required), and the decision is whether to refurbish the plant as a PF or an FBC system. The systems to be modelled are thus an old plant re-commissioned with FBC boiler units, and the plant re-commissioned with the original PF boiler units.

The case study only considers discard sourced directly from the coal washing plant, and not reclaimed from a dump. Both discard sources should be considered for sensitivity, although the differences are

expected to be slight, as both essentially reflect removing a portion of the discard dump (either in practice or “virtually” by way of avoided burdens). The reclaimed discard should also reflect the burdens associated with the energy used to reclaim the dump. The study is limited to a consideration of discard produced within a fairly close radius of the power station (approximately 12km). This constraint is set so as to determine the number of units that could be brought back on line. For the in-storage stations, there is, on average, sufficient discard in this radius to produce a maximum of 400 MW_e. Other likely limiting factors (e.g. water availability) are not considered. Power output is therefore not included in the model parameter analysis, although considering both washing plant discard and reclaimed discard, or increasing the transport radius, would allow more units to be brought back on line.

The same number of units are assumed to be re-commissioned for the PF refurbishment. This assumes that if not provided by the FBC system, the same grid capacity will be provided by the PF system. This may not strictly be true, since this will depend on coal availability in the area. However, the functional unit chosen (see below) means that the choice of PF capacity only has an effect on those burdens which are not directly related to the power production (see appendix E.1). The effect on these time-dependent burdens is slight, especially since the power output of the PF and FBC systems are not expected to differ widely. The assumption of equivalent power output is thus judged to be acceptable. The coal supply is assumed to be from dual product mines, since the study considers a re-conditioned station (i.e. its dedicated colliery is assumed to have stopped operating and it is supplied by an existing nearby mine adjusting its coal product to also produce a power station feed, or the dedicated mine subsequently started producing a high-quality coal product to sustain itself when the power station ceased operating). This introduces the problem of allocating the mining burdens between the power station and high-quality coal product (see section 7.4.2.c).

7.4.2. System Definition

To allow meaningful comparisons across the scenarios, consistent definitions of system boundaries, functional unit, data structures etc. are required. The same inventory models as presented in appendix E are used in this case study. The discussion in section 6.2 covers the reasons for the selection of the system boundaries (foreground and background), functional unit, level of flowsheet breakdown and data sources, and some methodological adaptations. As the same reasoning is applied here, the discussion is not repeated, and this section only considers elements particular to this case study.

7.4.2.a Functional Unit

In addition to the removal of the discard dump, a key aspect of the FBC system is the energy savings resulting from producing useful energy from an otherwise wasted energy source. The functional unit needs to reflect this, as well as compare the systems on an equivalent basis. As before, a dual functional basis is required to relate the power sent out to the duration of operation, so that those processes not able to be directly related to the production of power are incorporated (see appendix E.1). A dual time and product basis is therefore taken, in which the burdens calculated for an average year's operation are normalised to the total power produced in the year. The normalisation is necessary, because without it a more efficient system merely reflects increased burdens and not an increased energy product.

7.4.2.b Flowsheet Construction

The key requirements for inventory models to be able to incorporate uncertainty assessments is discussed in section 4.2.1, primarily that the LCI model have explicit variable inputs at as disaggregated a level as possible. The inventory model used here is specifically built up on these considerations, with the emphasis on the exploration of model parameters, whilst allowing for empirical parameter uncertainty to be incorporated. Unlike the case study presented in chapter 6, where the aim was to provide quality LCI information, this study is primarily a comparative and explorative analysis. The flowsheet therefore needs to be constructed in sufficient detail to allow the differences between the systems to be explored, and a detailed flowsheet and disaggregated data set is thus required where the systems differ. The systems are therefore modelled at the greatest level of resolution the data allows, which is at the major unit operation level for the foreground processes.

7.4.2.c Allocation Issues

The use of discard presents some difficult methodological considerations. If the discard is defined as a waste it is allocated no burdens, and the system burning it is essentially "credited" with a free energy source. The system therefore reflects the benefits of not having to mine virgin coal. However, the system does not reflect the potential burdens avoided by removing the discard dump, unless the comparative system boundary is expanded to include the discard dump. Equivalently, the FBC system can be credited with the burdens avoided by no longer having to stockpile or dispose of the discard. It can also be argued that discard is a future energy supply and not purely a waste. This is especially valid where discard is taken directly from the coal preparation plant, where the line between a low quality coal product and a discard becomes very blurred. The allocation procedure developed by Weidema

(2001) is useful here, in that the distinction between wastes and co-products is not important for the selection of the allocation procedure. The output, whether defined as a waste or co-product, can be regarded as a dependent co-product, and the applicable allocation procedure chosen accordingly. Where the dependent co-product is not utilised fully (as is the case with the colliery discard), Weidema's procedure recommends all burdens be allocated to the determining co-product (in this case the coal product), other than the burdens arising from the intermediate treatment of the dependent co-product (e.g. the burdens associated with reclaiming the discard dumps), which are allocated to the dependent co-product (i.e. the discard). In addition, the product system in which the dependent co-product is used, is credited for the avoided waste treatment of the dependent co-product (Weidema, 2001) (i.e. the FBC system is credited with the avoided disposal/stockpiling of the discard).

Another consideration is that the mine could permit a lower washing efficiency, and thus the production of more discard at a higher quality, when it is known that a market is available for the discard. In this case, the use of the discard does have an effect on the operation of the mine, and it can no longer be regarded as a dependent co-product. An allocation on the basis of the value of the product is not meaningful, as the discard is essentially worthless if defined as a waste, or extremely difficult to value as a "future" resource. A simple mass-based allocation is therefore reasoned to be more valid than a value-based allocation. The fact that the discard production can be increased whilst the coal production stays constant (although the mass of coal seam mined increases and coal and discard qualities change), means that the allocation is most meaningfully based on a marginal analysis, i.e. by modelling the consequences of increasing the volume of discard produced (Weidema, 2001). A marginal analysis is therefore used to determine the portion of mining burdens to allocate to the discard, other than those burdens arising from discard disposal, which are again "credited" to the discard product to reflect the "positive" burdens achieved by avoiding its stockpiling.

The allocation of burdens to discard reclaimed from discard dumps is simpler, as it can be assumed that the dumps were constructed sufficiently long ago that their reclamation in no way influences the current operation of the mine. They are thus not allocated any burdens, other than the positive effects of decreasing the size of the dump (i.e. the "avoided" burdens), and the burdens associated with their reclamation. The "avoided" burdens of the reclaimed discard are essentially identical to those of the washing plant discard, except that rather than removing discard from the dump, the use of washing plant discard avoids its placement there in the first place.

In addition to discard allocation for the FBC system, the PF system requires a suitable allocation of the mining burdens between the power station coal and the mine's existing coal product. As with the discard allocation, a marginal or mass-based approach is appropriate as the two products are better regarded as combined production than joint production (i.e. the product volumes are independent of each other, and depend on the relative demand for each product). An alternative approach is to avoid the allocation altogether by extending the system to a consideration of both energy products (i.e. coal and electricity). With this approach, a mining only system is also required, so that the removal or avoidance of the discard dumps can be reflected (i.e. "before" and "after" scenarios are required to reflect the removal/prevention of the discard dump). Whilst such a basis was initially pursued for this case study, problems with the unequal nature of the energy products, and the inclusion of an additional scenario not actually under consideration for the decision, clouded the analysis to such an extent that the problems outweighed the benefits of the simpler allocation (i.e. the normalisation/functional unit uncertainties and valuation uncertainties introduced were greater than the allocation uncertainty removed). The uncertainties arise in the dual product approach because the dual functional unit, based on the total "effective" energy of the combined product, is not able to capture anything of the relative benefits of coal and electrical energy. The system was therefore judged to be better investigated on the basis of a single function (i.e. the electricity produced), and the effect of the coal and discard allocation investigated by way of sensitivity analyses.

7.4.2.d Uncertainty Propagation

The propagation of empirical parameter uncertainty through the inventory model is calculated in a simulation using Latin Hypercube sampling. Latin Hypercube sampling has been shown to be always as good as, and never worse, than the other commonly implemented sampling routines (see section A.1.1.b). A sampling size of 1000 was calculated to be sufficient for a minimum of 95% confidence that the median lies within 3 estimated percentiles. The actual precision achieved may be greater than this estimate, as the equations used to estimate the required sampling size for the specified accuracy only apply for Monte Carlo sampling, and Latin Hypercube sampling is likely to achieve the same or a better precision (see section A.1.1.a). This precision was chosen primarily because of the need to keep the time taken to complete each simulation to a minimum. The precision achieved with 1000 samples is judged sufficient, given the high empirical uncertainty in the system, and the rough distributions used to define the input parameters.

single index is therefore not appropriate. The difficulty of simultaneously interpreting probabilistic output across multiple criteria, whilst retaining the insight not possible with a single weighted index, is overcome by plotting the options on a principal component plane (see section 7.5).

It is not in the scope of this thesis to review the many impact assessment methods in use. The EI 99 method was chosen to illustrate the transformation from inventory data to impact indicators, primarily because data on the uncertainty of the equivalency factors are available for most of the impact categories (see appendix C.2 for details on the EI 99 method). In tactical/strategic type studies there are likely to be a wide range of stakeholders, so a comprehensive set of impact categories is required to be sure all stakeholder concerns are addressed. This study therefore starts by a consideration of the full EI 99 set of impact categories. For the reasons stated above, the EI 99 method is not followed all the way through to a single index, i.e. the system is assessed at the damage category level.

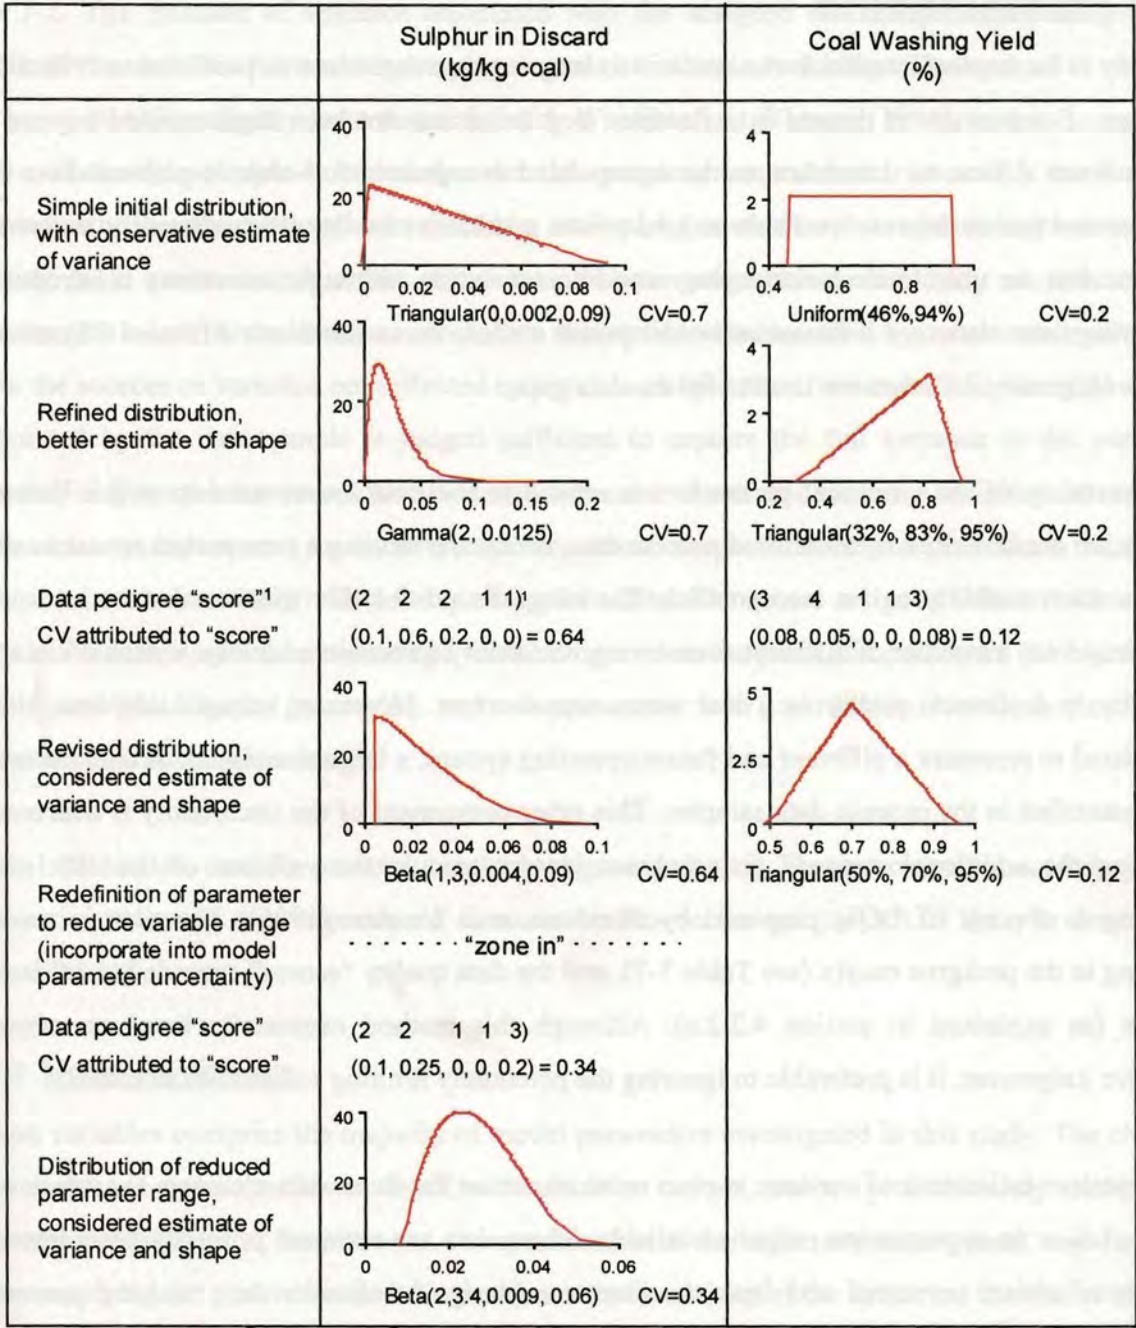
Two additional criteria reflecting burdens considered important in the context of the study, but not captured by the EI 99 damage categories, are included. Water use is included as a criterion because of its known importance in the regional context of the study. This includes raw water purchased by the mine and power station, as well as the water “consumed” by the process by virtue of its impact on the catchment area (see section 6.2.6.c). Also of importance in the regional context of the study is the effect of the systems on regional water quality. A key aspect of the study is the trade off between the emissions from the discard dump, and those from the ash/gypsum dump. The level of detail at which the mine and power station inventories are constructed is not sufficient to capture the water quality impacts as reflected in the EI 99 toxicity categories, so the use of waterborne sulphate emissions as a “proxy” indicator of water-related impacts is suggested as a stop-gap measure (see section 5.1.3). A better solution is provided by the impacted land footprint (see section 5.4.2 for details on the calculation of the impacted land footprint), which gives an indication of the potential of solid waste deposits to contaminate water bodies through leachate generation. It incorporates the land occupied by the dump and its leachate potential into a single indicator.

7.4.4. Definition of Parameter Uncertainty

The parameters input into the model are defined either as empirical or model parameters, which determines how their uncertainty is addressed (see section 3.1.1.a for reasons and definitions of the parameter types). The guidelines given in section 4.2.2 are followed to characterise the uncertainty of the empirical and model parameters (Figure 4-3 summarises the steps).

The full set of data inputs specifying the scenarios is given in appendix F.2. The model parameters are defined by a single value, which differs according to the model parameter scenario state, whilst the empirical parameters are assigned probability distributions. Uniform and triangular distributions are used for the initial runs of the analysis, and the distributions of those parameters returned as important from the rank-order correlation analysis are subsequently updated (see section 4.2.3 for details on prioritising parameter uncertainty importance using rank-order correlation coefficients). Gamma, beta, normal or lognormal distributions are used for the updated distributions, depending on which is the most applicable according to the data element on hand (see section A.2 for guidelines on the choice of probability distribution). Discrete conditional probability distributions are used to characterise the compositions of composite quantities (i.e. coal, sorbent, ash and water qualities). For these quantities it does not make sense to vary the individual components independently of each other, as they are highly correlated (e.g. the carbon content of the coal can not be varied independently of the C.V., or the nonsensical situation of a low carbon coal having a high C.V. will result).

An example of the process of defining successive distributions of increasing relevance (discussed in section 4.2.2) is shown schematically for two example parameters in Figure 7-2. The estimate of “Washing Yield” is only supported by point estimates from various sources, so the shape cannot be better characterised than by a triangular distribution, whilst “Sulphur in Discard” is supported by a large data set (discard quality across many mines, and over a number of years), from which a distribution shape can be estimated. The full variable range of sulphur in discard is extensive, and the variance of the parameter is unable to be reduced further without redefinition. A subset of various discard blends is therefore defined, each with substantially lower variability, and each blend incorporated into the model parameter analysis.



1. DQIs in pedigree "score": certainty, completeness, temporal correlation, geographical correlation and further technological correlation (see Table 3-7).

Figure 7-2 Schematic of process of successive definition of empirical parameter uncertainty for those parameters returned with high uncertainty importance by the correlation analysis (see section 4.2.2.a for discussion on assigning CV estimates to data quality scores, and the choice of distribution shape).

7.4.4.a Empirical Parameters

The study is for implementation in the medium to long term, so data must be predicted to reflect future conditions. Combustion of discard in a fluidised bed boiler has not been demonstrated beyond pilot scale in South Africa, so this data must be extrapolated to represent full-scale implementation. Fairly extensive and recent data are available around power generation for the plants currently in operation, and these data are used to characterise the reconditioned station, although uncertainty is introduced in transferring these data to a different and older power station. As in the South African LCI case study (chapter 6), generic LCI data are used to fill the data gaps.

The uncertainty of the empirical parameters is related to the data source used to define them. The majority are characterised by monitored process data, taken over as long a time period as was available, to allow for variability to be incorporated. The range found for this monitored data incorporates uncertainty from a number of different sources, e.g. variability (process variability, seasonal variability, variability in feedstock quality etc.) and measurement errors. However, because this data must be extrapolated to represent a different and future operating system, a large component of data uncertainty is not quantified in the process data samples. This other component of the uncertainty is addressed by estimating the additional variance (i.e. that not incorporated in the variance of the data sample) according to the set of DQIs proposed by Weidema and Wesnæs (1996). The data is evaluated according to the pedigree matrix (see Table 3-7), and the data quality “scores” extended to estimates of variance (as explained in section 4.2.2.a). Although this method necessarily involves substantial subjective judgement, it is preferable to ignoring the potentially limiting sources of uncertainty.

The subjective estimation of variance is even more important for those data elements for which only a single value or an approximate range is available. These data are obtained primarily from interviews with power station personnel and from the literature (design handbooks etc.), and are assumed to represent “design” or “most likely” values. The matrix scores assigned to these data elements are poorer than those assigned to the monitored data, as typically little is known about how representative the estimate is. Conservative CV estimates are applied to these data types, as they have no quantitative variability information to guide the conversion from qualitative matrix “score” to quantitative variance estimate. Where possible, the range of values available in the literature, or the range for a related quantity, is used to guide this estimate.

An example of the matrix scoring and equivalent CV estimate is given for the two parameters in Figure 7-2. The estimate of variance associated with the assigned score depends on the particular variable, as the variables have widely varying degrees of intrinsic variability. The manner in which the DQIs are used differs slightly from that defined by Weidema and Wesnæs (1996), in that the measured variance of the data sample is included under the completeness indicator, when it is not known to which indicator this observed variance should be attributed, i.e. whether the variance in the data sample is arising from measurement errors, from temporal or geographical variability et cetera. If the parameter gets a low “completeness score”, then this variance is increased by an amount estimated to include the sources of variance not reflected by the data sample, but if it scores well, then the variance incorporated by the data sample is judged sufficient to capture the full variance in the parameter. Appendix F.2 gives the scores and the associated estimates of variance for those variables returned with high uncertainty importance, i.e. those parameters whose conservative initial definitions required re-definition to more accurately reflect the quality of the data. The scores assigned reflect the source and quality of the data, whilst the estimates of variance depend on the particular parameter (i.e. the inherent variability exhibited by the parameter), and two parameters with a particular score do not necessarily have the same estimate of variance.

Generic LCI data are used for processes in the background system. The approximate uncertainty in this data is estimated using the ranges found in recent studies on the quality of LCI data, as described in section 4.2.2.a.

7.4.4.b Model Parameters

Decision variables comprise the majority of model parameters investigated in this study. The choice of these variables is under the direct control of the decision makers, and are predominantly related to the degree of refurbishment of the plant (i.e. the specification of the boiler and water plant), and the choice of mine supplying the coal or discard. To determine the sensitivity of the system to the choice of model parameter, the feasible limits of operation and the most likely value have to be defined. Not all decision variables can be defined as feasible ranges, and instead require a discrete choice that involves the choice of a number of related parameters, e.g. wet or dry ash disposal. In these cases, the choice of one parameter determines the choice of others, e.g. the choice of dry ash disposal requires a certain level of refurbishment on the water plant to limit the effluent volume to that able to be handled by the ashing system. Also, it does not always make sense to vary the parameters independently. In these instances,

run-off collected and the containment dam capacities are defined together as “good”, “average” or “poor” storm-water management. This streamlines the assessment by simultaneously varying groups of parameters that are likely to change together.

The minimum, maximum and most likely values are then used in a sensitivity analysis to screen the parameters for their uncertainty importance (i.e. the parameter value varied to its extremes while all other model parameters are held at their most likely values). A parameter is labelled as “significant” if the median of the output sample calculated with the parameter at either extreme, lies outside the 50% confidence interval of the sample calculated with the most likely value. The parameter is labelled as “moderately” or “less significant ” if the median value of the extreme output sample changes that of the most likely sample by greater than 10% or 5% respectively, but there is a large overlap between the samples (the median of the extreme sample lies within the 50% confidence interval of the most likely sample). Parameters changing the median less than 5% at their extreme values are not considered in the parametric analysis, and are kept at their most likely value in all scenarios.

The sensitivity analyses are required because the large number of decision variables at a tactical decision level make a full parametric analysis impractical. However, sensitivity analyses do not pick up possible extreme values resulting from combinations of parameters, so there is the danger that a potentially important parameter (or combination of parameters) be excluded from the parametric analysis. This introduces an unavoidable amount of uncertainty as a result of using a less rigorous method, but to minimise this, a knowledge of the system is used to do a qualitative check on the parameter before it is excluded from further analysis, i.e. the effect of each parameter on the parameter of concern is evaluated, and if the possibility of an extreme value is identified, the output is calculated for both variables set to their extremes. For example, the choice of ESP or FF is shown not to be significant where an average discard blend is used, but may be so for the choice of a poor discard blend, as at high ash contents even a small change in particulate collection efficiency may be important.

Another significant simplification of the model parameter analysis is that the important parameters are grouped into “best”, “worst” and “most likely” scenarios, rather than a full combinatorial analysis conducted. The results from the sensitivity analyses are used to assign the minimum and maximum values of each significant parameter into either the “best” and “worst” model parameter scenario states. However, a choice of parameter does not always cause the system to be “worse” or “better” in all the impacts considered, e.g. an increase in the percentage of SO₂ removal in the boiler, causes a decrease in

acidification potential, but an increase in particulate emissions. As discussed in section 4.2.2.b, a possible solution is to establish a priority order of the criteria. The “best” and “worst” scenarios are thus established according to the highest priority criterion. The choice of parameters to define the three limit scenarios are given in Table 7-2. For the parameter groupings considered here, water-related impacts (water use, followed by waterborne sulphate emissions), are considered first (as these are a particular driver for the project), after which local air quality impacts are considered (SO₂ emissions, followed by particulates). Thus, if the parameter at its extreme value causes a greater than 5% increase in the median value of water use, this value is used in the “worst” scenario, and its opposite extreme value used in the “best” scenario, even if it causes a decrease in another impact.

Table 7-2 Choice of parameter across the three extreme scenario states for those decision variables found to be significant in a sensitivity analysis.

	Worst (Little refurbishment)	"Most likely"	Best (Significant investment)
Maximum unit capacity (MW)	125	200	200
Number of generator sets	4	2	2
Load factor	64%	80%	80%
Sorbent type	dolomite	limestone	limestone
% SO ₂ removed	30%	60%	90%
Particulate control	ESP	ESP	FF
Mine type	opencast	underground	underground
Mine power source	adjacent station	grid	grid
% coal bypassing washing plant	15%	20%	25%
Station, stockpile and mine life (yrs)	35	20	15
Stormwater and effluent management	poor	average	good
Distance coal transported (km)	25	10	3
Method of coal transport	rail	conveyor	conveyor
Transport distances (km)	850	500	150
Transport mode	road	road	rail
Ashing method	wet	wet	dry
Water plant configuration	un-optimised	optimised	optimised, and necessary adjustments for dry ashing
Stockpile size (reserve time)	3yrs	3mths	3wks

7.5. RESULTS AND INTERPRETATIONS

The uncertainty analysis process is an iterative one, where successive revisions are made until an acceptable resolution is obtained in the results. Figure 7-3 shows the final output from this process. The results presented here are primarily concerned with the outcome of this process, although reference is made to the key factors driving this progression in the sections that follow.

Figure 7-3 clearly shows the benefits of burning discard in a fluidised bed boiler. The figure shows a principal component representation of the output data, where the output samples have been projected onto a best-fit plane, so as to obtain a two-dimensional (planar) representation of the data, maximising the information able to be displayed in two-dimensions. A plot of the PC loadings (the coefficients of the eigenvectors corresponding to the PCs) is overlain on the plot of the transformed output samples. These are plotted as arrows (or “stressor vectors”) which represent the distance from the origin to the best-fit plane, and give information on the strength and independence of the criteria (see section A.3.2 for information on interpreting PC plots).

The strongest differences between the PF and FBC systems are in their fossil fuel use and their impacted land footprint (as seen by the magnitude of their stressor vectors). The length of the footprint vector reflects the considerable benefits of removing the discard dumps, i.e. the ash/gypsum dump has a consistently smaller footprint than the discard dump. The considerable savings in fossil fuel resources reflect the use of a “waste” energy source to generate power. The discard is defined as a waste from the mining system, and as such, is not allocated any mining burdens other than the “avoided” burdens resulting from the removal of the dumps. It therefore does not reflect any fossil fuel resource consumption, as all fossil fuels consumed and extracted during mining are allocated to the coal product. Also caused by the avoidance of mining burdens are the lower contributions to climate change, water use and summer smog of the FBC system relative to the PF system. This is less marked for climate change, because if compared on the basis of the power station alone (i.e. without the effects of mining), the FBC system has a slightly higher climate change burden than the PF system (caused predominantly by the use and transport of limestone in the FBC system).




Figure 7-3 Representation of FBC and PF operating “spaces”, given by ratios between the output samples and the PF “most likely” scenario for the chosen impact categories, transformed on the 1st and 2nd PC plane (see Table H-12 for corresponding PC loadings).

The lower contribution to the combined effect of acidification and eutrophication from the FBC system stems from the far lower NO_x emissions from the FBC boiler. The effect is much more marked for a consideration of NO_x alone, as SO_2 emissions are not always lower in the FBC system (they depend on the degree of desulphurisation). This is shown in Figure 7-4, which presents the same results as Figure 7-3, but on the level of environmental interventions rather than impacts. The high sulphur content of the discard burnt in the FBC system relative to the coal burnt in the PF system, means that only at high levels of desulphurisation does the FBC boiler have lower SO_2 emissions than PF boiler. The SO_2 emissions of the “worst” FBC scenario (with only 30% SO_2 removal) are higher than those of the PF scenarios, whilst those of “best” FBC system are significantly lower. This causes the strongly negative trend in SO_2 emissions in Figure 7-4, where the maximum variance in SO_2 emissions is between the FBC scenarios and not between the PF and FBC systems. This is a consequence of the PC plots’ ability to only show the variance incorporated in the first two PCs. PCA maximises the variance able to be viewed in two dimensions, and thus picks up the greater source of variance in SO_2 emissions

(i.e. between the FBC scenarios), whilst the variance between the two systems is picked up by a lower order component. From Figure 7-4 it appears that FBC systems are always worse than PF systems with respect to SO_2 emissions, although Figure H-10, which considers only the “best” and “most likely” scenarios, shows this is not the case.

Winter smog is shown to be ambivalent between the PF and FBC systems (i.e. neither system plots strongly with or against the winter smog vector). The winter smog vector falls between those of the SO_2 emissions and TSP emissions, with the direction it takes dependent on the relative magnitude of these two emissions (see Figures H-9 and 10 or Figures 7-3 and 4). Although the FBC system always causes an increase in particulate emissions, the volume of SO_2 emitted can “swing” winter smog to being significantly worse or slightly better than the PF system. The toxicity categories also do not show a strong tendency to be better or worse in either of the systems. The high variability in the output of these categories leads to the PCA being unable to distinguish any clear trends between the systems.

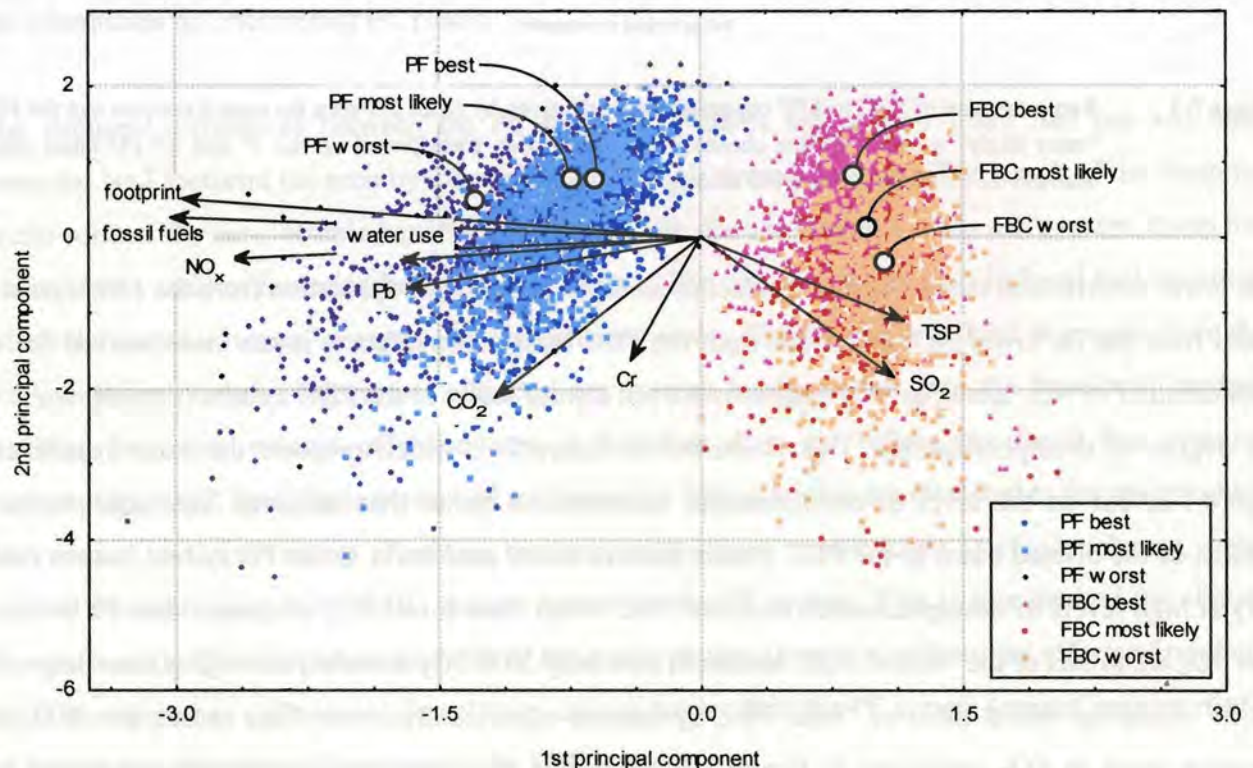


Figure 7-4 Representation of FBC and PF operating “spaces”, derived from the results for selected environmental interventions and the impacted land footprint (in original units), transformed on the 1st and 2nd PC plane (see Table H-13 for corresponding PC loadings).

7.5.1. System Definition

The above discussion shows that a number of the positive effects of the FBC system are a result of the difference in mining burdens between the PF and FBC systems, i.e. the higher PF values for many environmental interventions arise from the mining burdens associated with the coal fuel, since the FBC system is assigned no mining burdens other than being “credited” with avoiding the disposal or stockpiling of the discard. This is the most methodologically defensible allocation approach where the discard is not fully utilised and can be regarded as a dependent co-product of producing high-quality coal (Weidema, 2001) (see discussion in section 7.4.2.c). This is likely to be the case for most power station/discard coal interactions, and is the scenario for this case study, where the scale of discard-based power generation proposed is such that not all the discard produced will be utilised. However, should discard-based power generation be proposed on a much larger scale, i.e. where all the discard produced is utilised, an alternative allocation procedure based on the processes displaced or avoided because of using discard should be investigated. An alternative discard allocation is investigated in section 7.5.5, for a different scenario where the production of a relatively high-quality discard is assumed to have some affect on the mining operation (i.e. where it is not a truly dependent co-product).

The significance of the mining burdens with regard to the differences seen between the PF and FBC systems means that the allocation of mining burdens to the PF system is also important. The PF system is assumed to be supplied by a dual-product mine (see section 7.4.1), so mining burdens need to be allocated between the power station coal and the high-quality coal product. The two products can be regarded as combined production rather than joint-production (i.e. they can be varied independently of each other), which means that the allocation is most meaningfully based on the consequences of varying the quantity of power station coal produced (Weidema, 2001). The relevant portion of mining burdens to allocate to the power station coal is therefore calculated by varying the percentage of coal bypassing the washing plant. The calculation of the marginal allocation burdens is rendered less meaningful by the fact that the simple mining model used relates power use on the mine to the flow of coal produced, and so the change in burdens predominantly reflects the change in power consumption, other than for land occupation, dust emissions, and waterborne sulphates, which are not related to the flow of coal.

A mass-based allocation method is also applied for sensitivity, where the mining burdens allocated to the power station coal are based simply on the percentage of the total coal produced sold to the power station. This apportions fewer burdens to the power station coal than the marginal allocation, and thus improves the PF system's performance on all impacts. This has the effect of moving the PF system closer in performance to the FBC system, and so decreases the confidence with which the differences between the systems can be discerned. This can be seen in Figure 7-5, where the sample "cloud" of the PF "most likely" scenario with mass allocation plots closer to the FBC scenarios than the other PF scenarios. This causes greater overlap between the PF and FBC operating "spaces", i.e. there is less separation between them in Figure 7-5 than 7-3, and thus less certainty with which differences between them can be discerned. However, the overlap is not sufficient to change their relative performance. This can also be seen in Figure 7-5, where the orientation of the PF mass allocation scenario relative to the FBC scenarios has not changed with respect to the criteria stressor vectors, although the degree of separation has changed. For example, the PF mass allocation scenario still plots in the direction of the criteria vectors for all but the toxicity criteria, representing a decrease in performance relative to the FBC scenarios, although it plots a lot less strongly in the direction of the arrows.

A benefit of removing discard dumps is eliminating the risk of "spontaneous" combustion. This high consequence, but low probability, event is excluded from the analysis, because it introduces so much uncertainty into the output, that no conclusions are able to be drawn from the results. "Spontaneous" combustion is therefore better incorporated into the analysis by way of a scenario analysis. Figure 7-5 shows the results from such an analysis, which, in addition to the previous scenarios, includes a scenario in which 20% of the discard dump burns over its lifetime. The effect of the discard dump burning is to enhance the performance of the FBC system relative to the PF system with respect to its contributions to climate change, winter smog and acidification and eutrophication. This is because the "avoided" burdens approach "credits" the FBC system for removing the disposal/stockpiling of discard, and thus credits it with avoiding the emissions associated with spontaneous combustion. This effect is particularly marked for winter smog, since the prevention of significant emissions of CO, SO₂ and TSP impacts strongly on this category. Including the effects of "spontaneous" combustion into the analysis thus only improves the performance of the FBC relative to the PF system. Not including them therefore underestimates the relative benefits of the FBC system over the PF system, with the significance of this oversight dependent on the level of risk afforded to "spontaneous" combustion occurring.




Figure 7-5 Extension of the FBC and PF operating “spaces” to include case when discard is burning, and a different coal allocation basis. Derived from ratios between the output samples and the PF “most likely” scenario transformed on the 1st and 2nd principal component plane (see Table H-14 for corresponding PC loadings).

7.5.2. System Evaluation

This section examines the relative importance and independence of the criteria chosen to summarise the impacts of the systems, with particular reference to the use of PCA to facilitate this analysis. Discussion on the uncertainty of the criteria, and thus the relative value of keeping them on an inventory level rather than an impact level, is postponed to section 7.5.3. Figure 7-3 shows the final set of criteria considered. The length and orientation of the vectors give an indication of the strength and independence of the criteria (see appendix A.3.2 for information on interpreting PC plots). The impacted land footprint and the extraction of fossil fuels clearly show the strongest differences between the PF and FBC systems. The combined effect of acidification and eutrophication also shows a strong degree of difference between the systems, with the remaining categories more or less equally influential, other than the toxicity categories (carcinogenic effects on humans and ecotoxicity), which show a lesser degree of influence.

The different units and large differences in magnitude of the various criteria mean that the PCA has to be based on a correlation matrix and not the actual data set (see appendix A.3 for a discussion on the basic principles of PCA). This standardisation of the data essentially places each criteria on an equal footing in the analysis. For example, the standardisation allows land use in $\text{m}^2 \times \text{yr}$, mass flows in kgs (ranging from very high flows of major components to extremely small flows of trace emissions), and energy flows in MJs to be compared as to the relative strength of the differences between the systems. However a problem with this is that the relative significance of the impacts is no longer reflected. The length of the arrows on the PC plot thus do not indicate the significance of the criteria with respect to each other, but merely the magnitude of the differences between the systems. For example, if two systems shows a 50% increase in both winter smog potential and ecotoxic emissions, resulting from a difference of 0.05 DALYs, and 500 PDF.m².yr for winter smog and ecotoxicity respectively, the impacts are regarded as equally important, as they each increase the contribution to the impact category by 50%. However, with the normalisation factors given in the EI 99 method, a change of 0.05 units of winter smog is roughly 14 times more significant than a change of 500 units of ecotoxic emissions. Thus when interpreting the length of the arrows on the PC plots, it must be remembered that they do not reflect the significance attached to the impacts in the context of the study, nor the absolute magnitude of the impacts (e.g. the impact may show a very significant increase over the base case, but the base case impact may be so low that the actual effect is negligible).

A similar “levelling” effect on the criteria results from the need to base the analysis on a ratio or relative difference to remove the correlation of uncertainties between the scenarios. This is particularly important for the analysis based on impact categories, as the uncertainty associated with the equivalency factors considerably increases the overall uncertainty of the scenarios, but is common to all the scenarios. Thus for case studies where a considerable number of the uncertain inputs are common to both systems (as in the following case study in chapter 8), and/or where the results are based on impact categories, the analysis is most meaningfully based on ratios between the output samples and a suitable base case scenario (since this cancels out the uncertainty common to both systems, provided the output samples of both systems are calculated from identical random samples). This effective “standardisation” of the data causes the same problems as are discussed in the previous paragraph for basing the PCA on a correlation matrix (as any absolute value given to the potential effects by a normalisation step is cancelled out by taking a ratio of the options) .

Basing the analysis on ratios of the output samples is generally preferred to basing it on the percentage difference between each scenario and a base case scenario, since this does not increase the number of outliers (extreme combinations) as calculating the percentage difference does. The choice has very little effect on the PCA since it is based on the correlations between the data points (other than to marginally increase the scatter of the uncertainty “clouds”). However, cumulative probability plots based on the percentage difference between the systems enable these plots to provide a clear representation of the degree of confidence in the results (see section 7.5.3.c). Thus the analysis is based on the percentage difference between the systems where cumulative probability plots are also required as part of the analysis.

The full set of EI 99 categories, other than “extraction of mineral resources” and “land occupation” are retained, as each category is shown to reflect unique information about the systems (i.e. the impact vectors do not plot on top of each other in Figure 7-3). “Extraction of mineral resources” is excluded because its effect is very small (see Table H-11 and its accompanying discussion in appendix H.2.1), whilst “land occupation” is excluded because it is incorporated into the “impacted land footprint” (see appendix H.2.1).

The PCs from which Figure 7-3 is derived are given in Table H-12. The relatively low contribution of the first two PCs to the total variance (49%) reflects the complexity of the systems underpinning the results, and the many sources of variance being picked up by the PCA. For the two-dimensional PC representation to be useful, the first two PCs need to account for a high percentage of the total variance. The first two PCs reflect the differences between the systems (i.e. between the FBC system and PF system), as these show the greatest contribution to the variance. The differences between the scenarios within each system (e.g. the “best”, “most likely” and “worst” scenarios) are less well represented in the first two PCs, and it may be necessary to view the lower order components to get a better understanding of the interplay between the different scenarios (see section 7.5.5).

The toxicity categories are associated with considerable empirical parameter uncertainty, as is discussed in the following section. The benefit of including such highly uncertain quantities as criteria is questionable, since at such high uncertainty it is never possible to obtain a high degree of confidence in the perceived differences between the systems (as is discussed in section 7.5.3.c and shown in the analysis in appendix H.4). The PCA therefore does not find any significant differences between the systems for ecotoxicity and carcinogenic effects, and they are relatively weakly reflected in the first

two components. However, they are responsible for a high proportion of the variance in the 3rd and 4th PC (see Table H-12), where the major cause of the variance within these categories is attributed to a different source (in this case, their high empirical variability). If the variance attributed to the toxicity categories is discounted (given that at such high uncertainty they are not adding anything to the analysis), the effective variance accounted for by the first two PCs increases to around 70%, which is much closer to the 75% guideline given by Murtagh and Heck (1987) for the two-dimensional PC plot to provide a meaningful representation of the data. Important insights can therefore be derived from the PC tables, and the systems should not be interpreted on the PC plots alone, especially if the first two components account for a relatively low proportion of the overall variance.

The systems are also compared on a selection of inventory data. The environmental interventions displayed in Figure 7-4 are chosen according to their uncertainty importance, i.e. the interventions that contribute most significantly to each impact category are identified according to the magnitude of their rank-order correlation coefficient (see tables in appendix H.3). This selection procedure assumes the analysis will be extended to an impact level. If the analysis stops at the intervention level, some other selection criteria would be required (e.g. the overall uncertainty importance, or the perceived significance of the intervention). The plots give insight into the impact categories for which the contributing environmental interventions show opposing trends, e.g. winter smog (TSP and SO₂) and the combined effect of acidification and eutrophication (SO₂ and NO_x). Although the move from inventory to impact data is associated with a significant increase in empirical parameter uncertainty (see Table 7-3), this is not reflected in the PC plots. As discussed above, the equivalency factor uncertainty common to both systems cancels as the results are based on ratios calculated from the same random sample. Section 7.5.3.b explores inventory versus impact uncertainty in greater detail.

7.5.3. Empirical Parameter Uncertainty

The empirical parameter analysis follows an iterative procedure, as explained in section 7.4.4 and shown schematically in Figure 7-2. Even after several rounds of iterations, during which the rough distributions of those variables identified as important were refined, empirical uncertainty remains high. This is seen by the extensive scatter of sample points in Figures 7-3 and 7-4, and in Table 7-3. The variance calculated for each criterion is fairly similar between the options, with the “worst” option showing a greater degree of scatter because of the greater uncertainty in predicting the effects of opencast mining.

for the “waterborne sulphate” category, which shows the typical features of the analysis. Initially a high reduction in variance is obtained, and a more careful definition of the top two or three key parameters yields high returns. This quickly tapers off, and additional effort in redefining variables yields little benefit, especially if the uncertainty in the top variables is not able to be significantly decreased through better parameter characterisation (i.e. they are high contributors to the uncertainty within the current scope of the study, even after their definitions are revisited, and require significantly increased data collection or modelling effort for their uncertainty to be substantially reduced).

Rank-order correlation coefficients are a simple and effective means to order each parameter’s contribution to the overall empirical parameter uncertainty (i.e. to determine its uncertainty importance) (see section 4.2.3). However, the rank-order correlation analysis was found to underestimate the uncertainty importance of those parameters defined by conditional probability distributions, where these are used to specify the uncertainty of those parameters known to be highly correlated to another parameter (see section 7.4.4). The problem arises because of the complexity in defining continuous conditional probability distributions, and their being approximated by discrete probability distributions. These are limited to only a few sample points, so can not be effectively analysed by a correlation analysis. The uncertainty importance of the relatively few parameters defined by discrete conditional probability distributions therefore has to be investigated separately. This was done by computing the output with each of these parameters set at their mid-point value (i.e. the analysis repeated for each of the parameters with their variance removed), and the change in variance of each of the criteria evaluated (a large decrease indicates the parameter as significantly contributing to the overall variance of that criterion).

Table 7-4 Parameter intervention strategy to reduce variance in waterborne sulphates.

Parameter returned with high uncertainty importance	Intervention	Measure implemented	CV ¹
			1.68 ²
Rainfall	change from monthly to annual variation	gamma distribution (CV=0.98) to normal distribution (CV=0.27)	1.05
Mine surface run-off area	split into underground and opencast variability	from combined CV of 0.41 to individual CVs of 0.14 (uniform distributions)	0.70
Dirty dam seepage volume	more relevant distribution and estimate of CV	uniform distribution (CV=0.57) to gamma distribution (CV=0.35)	0.69
Dirty dam seepage sulphates	more relevant distribution and estimate of CV	uniform distribution (CV=0.57) to triangular distribution (CV=0.38)	0.66
Power use per kg coal (underground mine)	more relevant distribution	uniform distribution (CV=0.57) to lognormal distribution (CV=0.62)	0.64

1. overall CV in waterborne sulphates

2. initial CV estimate, prior to any revisions.

Very high uncertainty yields meaningless results, since no differences between the systems can be discerned with clarity. An investigation of the results of the uncertainty importance analysis yields where interventions can best be made to improve the resolution. Parameters exhibiting high variability are found to offer the best opportunities for variance reduction in strategic studies, as these are often due to the unspecified nature of the system, and can therefore be reduced by better scenario definition. Incorporating the full variability of the parameters across all possible operating states yields such high uncertainty as to produce un-interpretable results. Some “zoning in” therefore needs to occur, in which the full range of variability is broken down into possible operating states. For example, considering all possible sources of discard coal yields too great a variability in discard quality. Particular blended discard products with much smaller quality ranges are therefore assumed, and a range of blended discard products incorporated into the model parameter analysis. Similarly, considering all sorbent sources yields very high variability in their reactivity, composition and transport distance. The sources are thus divided into limestone and dolomite, each with greatly decreased variability, and the choice of sorbent type incorporated into the model parameter analysis.

Incorporating the full range of variability in parameters specifying the mining system similarly results in very high variance in the output. Mining was therefore split into underground and opencast operations, each defined as separate, less variable states (e.g. instead of including the full range of methane emissions encountered in South African mines, this can be divided into two far more contained underground and opencast ranges). It was also found not to be meaningful to incorporate highly variable “risk” events. The extensive variance these introduce obscures any trends that would otherwise have been able to be discerned. Such low probability, high impact events are thus more meaningfully incorporated as “what if” scenarios, i.e. the event occurring is compared as a discrete scenario against the standard case of it not occurring (e.g. the scenario including spontaneous combustion in the discard dump in Figure 7-5).

7.5.3.b Inventory versus Impact Uncertainty

Table 7-3 shows the considerable increase in uncertainty accompanying the move from inventory to impact data. The averaging effect of aggregating the interventions into potential impacts is in most cases overshadowed by high estimates of uncertainty for the equivalency factors, representing the high uncertainties in impact assessment models. This presents an interesting dilemma, in that as the relevance of the information increases (impact rather than inventory), so the ability to distinguish between the systems decreases.

The increase in uncertainty from inventory to impact data is shown in the figures in appendix H.3, which plot the cumulative probability of the impact category output samples and the environmental interventions contributing most to that category. The extent of variance in the output is shown by the shape of the curves. The typical “S” shaped curve is observed for those categories with low to medium uncertainty, which changes to the steep curves of highly skewed distributions for the more uncertain categories (e.g. compare the curve for SO₂ emissions with that of TSP emissions in appendix H.3.3). The figures in appendix H.3 show that for the highly uncertain impact categories, the cumulative probability plots are useless for discerning any information about the systems, whilst the degree of separation and order of the systems can be discerned from the curves based on inventory data.

This suggests that for the impact categories with very high model and empirical uncertainty, it may be more meaningful to base the analysis on inventory rather than impact results. This observation is rendered less significant when the analysis is based on a ratio or difference between the systems being compared. In this case, the uncertainty common to both systems (i.e. that due to the equivalency factors) is cancelled out, provided the same random uncertainty sample is used in the calculation of both systems. However, it is not always desirable or possible to base the analysis on ratios, or for the results to be generated from identical uncertainty samples. Taking a ratio removes information on the relative magnitude of the effects the changes represent (see section 7.5.2). In addition, taking a ratio implies a pair-wise comparison, although the choice of a “base case” may not always be straightforward in a study comparing more than two systems. Finally, the requirement that the systems are calculated from identical random samples implies that the systems have been generated from consistent models, although this is unlikely to be achieved for comparisons with literature systems or those generated by a different person or group.

Thus there are likely to be many cases where the additional impact uncertainty will not be able to be simply removed as common to both systems. In these cases, cumulative probability plots (such as those in appendix H.3), can be used to assess whether inventory-level information may be preferable to impact information. The following section looks at this in greater detail.

7.5.3.c An Uncertainty Audit Trail

The most important aspect of an uncertainty analysis is not merely in the quantification of uncertainty, but how it directs efforts back into the analysis. The overall aim of the analysis is to support decision-making, and thus the ability to direct the decision maker to the point of intervention most able to improve the clarity of the analysis is therefore valuable. The uncertainty importance analysis is able to identify those parameters contributing the greatest uncertainty to the results, but the extent to which their variance needs to be reduced is also required. On the PC plots, this can be loosely judged by eye, by the degree of separation between the system “clouds”. The degree of overlap tolerated represents the degree of risk the decision makers are willing to take, whilst non-overlapping “clouds” represent significantly different systems. The iterations (identifying and refining the definitions of key parameters) can be continued until the desired degree of separation is achieved.

A more precise quantification of the significance of the differences between the systems may be achieved by plotting the cumulative probability of the ratio or difference of the result samples. The relative difference, defined by Coulon et al. (1997), is used in this study as it is thought to allow the clearest interpretation of the cumulative probability plots. The results are calculated as a percentage difference between each scenario (*B*) and the most likely scenario (*A*):

$$\left(\frac{A - B}{A} \right)$$

For system *B* to be significantly better than system *A* the ratio must be substantially positive. The y-intercept represents the degree of confidence that this is so, i.e. the probability that *A* is larger than *B*. A drawback of these plots is that they require a pair-wise comparison of the options. Each criteria is presented on a separate plot, so they quickly become tedious for a large number of options and criteria. Nevertheless, these plots are useful if a numerical target for the reduction in variance is desired, and are used to investigate the reduction in variance required to see significant differences between the systems. This process is shown schematically in Figure 7-6, using a comparison between the most likely PF and FBC systems’ contribution to the combined effect of acidification and eutrophication as an example. The process is explained in greater detail in appendix H.4, where a more in depth example is presented.

Extremely high variances are present for certain interventions and impact categories (see Table 7-3). At high variances (CVs in excess of 1), even considerable differences in the mean values, as well as

stochastic dominance, is not sufficient to discern differences between the systems with a high degree of confidence (i.e. greater than 60%) (see example in appendix H.4). This is important for the choice of selection criteria, as it is not useful to base the analysis on criteria for which it is not possible to obtain a high degree of confidence in the perceived differences between the systems. This is especially important in deciding whether to base the analysis on inventory or impact level criteria, as transforming already highly uncertain inventory data to estimates of potential impacts, may introduce so much uncertainty that significant differences between the systems can no longer be discerned. The analysis of variance outlined in Figure 7-6 provides guidance as to when inventory data rather than the more uncertain impact data should be used, by identifying when high confidence levels will never be obtained for a particular equivalency factor variance and degree of separation between the systems. This is only relevant for a comparison between systems calculated from non-identical random uncertainty samples, in which case the common impact assessment uncertainty cannot merely be cancelled out (see previous section).

The decision of whether inventory or impact level information is most meaningful for decision-making will depend on the extent of the uncertainty, and the degree to which a few key environmental interventions contribute to the impact category. In the example in appendix H.4, it is found that no useful information can be discerned from a consideration of carcinogenic effects, but that the performance of the scenarios relative to each other is discernible for Cr emissions. The fact that Cr emissions dominate the category to such a high degree (see Table H-20) implies that better decision support can be obtained using Cr emissions as a proxy for carcinogenic effects. However, the decision is less straightforward for impacts with a large number of environmental interventions of equal significance. Logistically there is a limit to the number of interventions that can be considered (i.e. that can be simultaneously interpreted in an analysis), and a trade-off between an ability to better distinguish between the systems, and missing the effect of some interventions will have to be made.

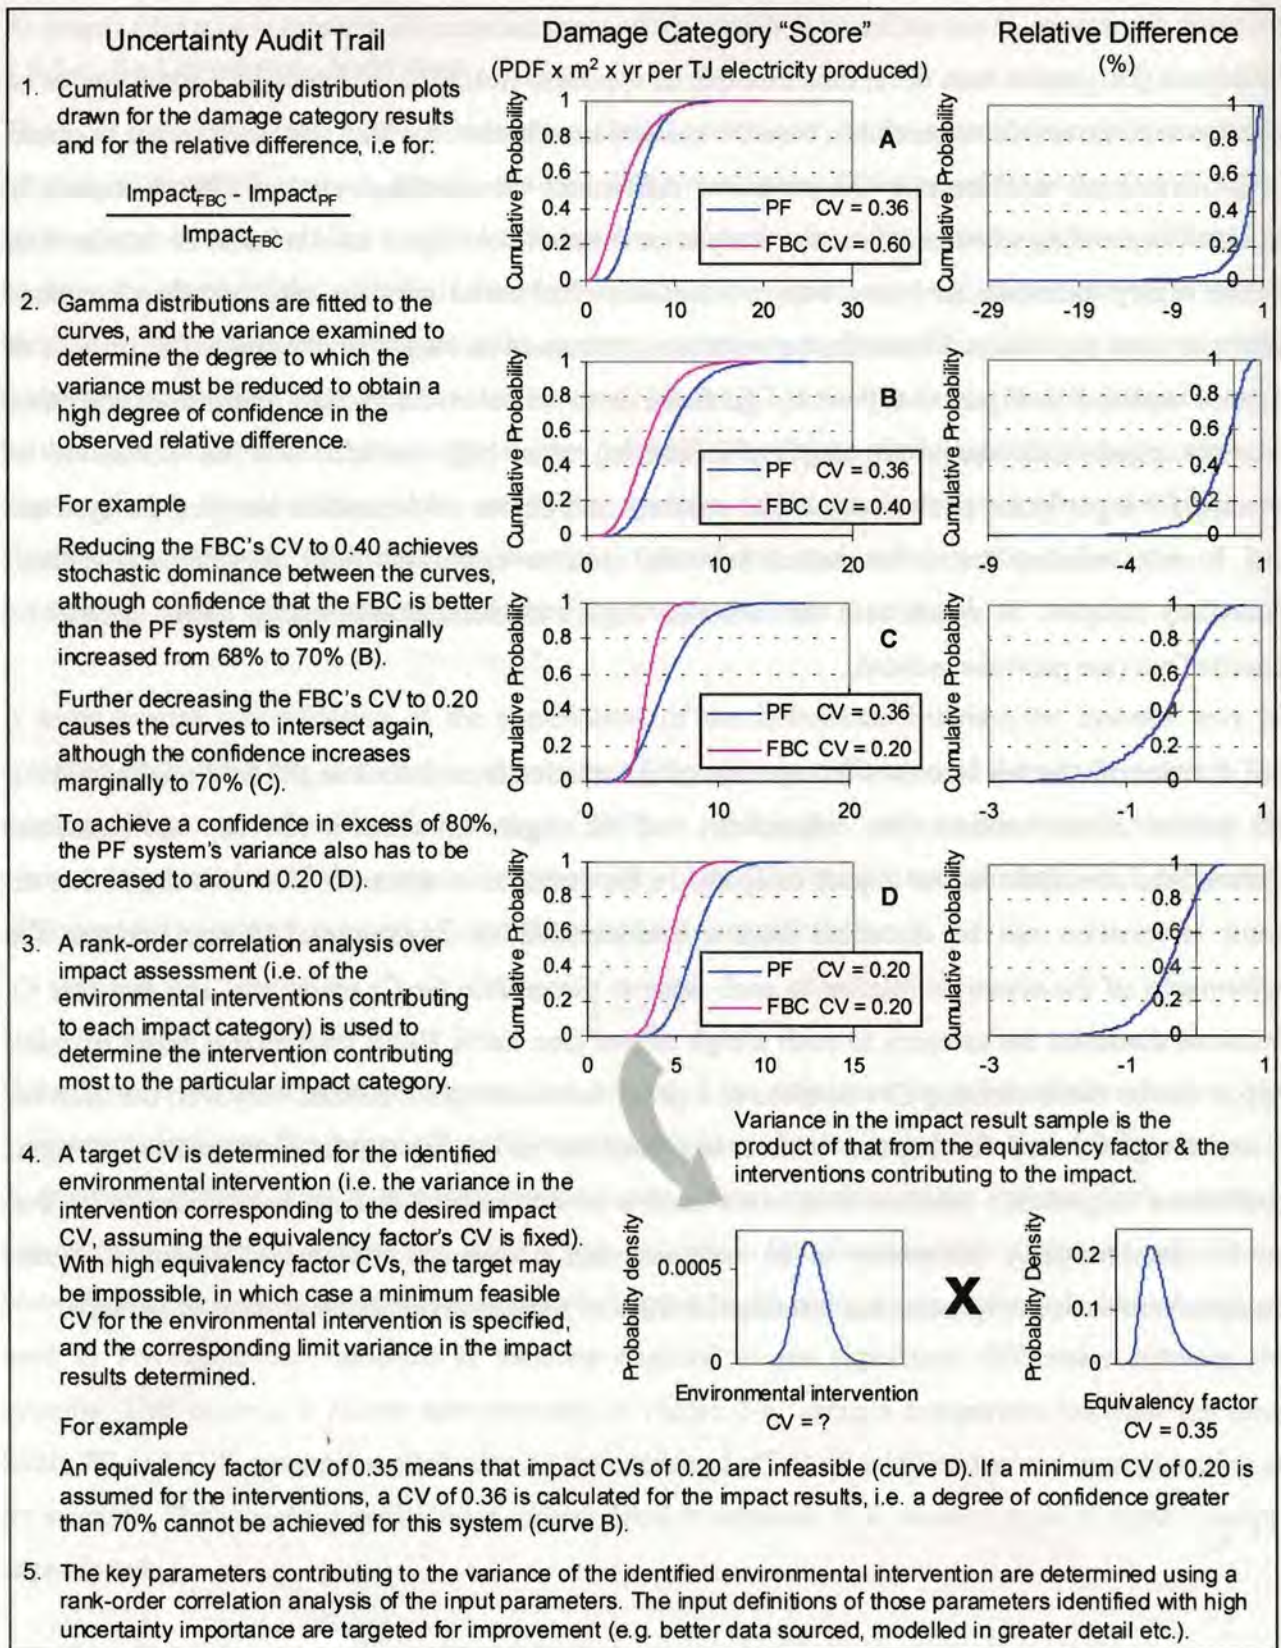


Figure 7-6 Plan and schematic of "top down" approach to direct effort in reducing uncertainty.

A significant finding of this study is that at the high levels of uncertainty encountered in these types of systems, it may never be possible to achieve high levels of confidence unless the absolute difference between the systems is increased, i.e. for certain criteria the systems are too uncertain and too similar to say with a high degree of confidence that one system is always better than another. An analysis such as that outlined above can therefore prevent valuable time being wasted in reducing data uncertainty with little return. A high degree of stochastic dominance is a necessary criteria for achieving 90% or greater confidence levels, but is not sufficient. The mean values of the uncertainty samples also have to be sufficiently distinct from each other, i.e. there must be a fair degree of separation between the systems. A generalisation of such analyses as that presented in Figure 7-6 and appendix H.4 yields the following “rules of thumb” for systems with CVs of around 0.2:

- A difference of at least 30% in the mean values is required to achieve 90% confidence that the one system is preferred over the other.
- A maximum confidence of 80% is achievable for systems displaying less than 20% difference in their mean values (and 65% for less than 10%).

It is the finding of this case study that for primary industries at this level of analysis (strategic/tactical) CVs less than 0.2 are unlikely to be encountered. In fact CVs in excess of this are far more likely. The preceding “rules of thumb” therefore give useful minimum guidelines for setting the “acceptable” degree of confidence controlling the number of iterations required in the empirical parameter analysis (see Figure 4-1).

7.5.4. Model Parameter Uncertainty

A specific requirement of the study is to investigate the operating “space” of the FBC system, having defined it by previous sampling so that it can be seen under what conditions, and to what degree, it performs better than a conventional PF system. This is done through an explicit consideration of model parameters. In a system with a large number of model parameters requiring consideration, the range over which the system performs is first identified, which allows the model parameter analysis to be streamlined into a consideration of the extreme options (see Table 7-2 and accompanying discussion). A more detailed parametric assessment is possible at the later stages of the decision process, when the number of parameters requiring consideration has been reduced (see following section). Decision variables comprise the majority of the model parameters incorporated into this assessment. Model choice parameters are less well addressed, but their effect is expected to be less significant in a

comparative assessment where the same models have been used for both systems (i.e. their choice is expected to introduce a common bias, rather than affect the comparative analysis).

The effect of the choice of model parameters can be seen by the separation between the three extreme options in Figure 7-3. Although the high empirical parameter uncertainty makes it difficult to discern the range covered by each option (the sample “clouds” largely plot on top of each other), a definite “shift” in the colours can be seen. The range is seen to be more extensive for the PF system, with the wide separation between the “best” and “worst” PF options attributable to the large differences in the environmental impact of opencast and underground mines. The smaller difference between the “best” and “most likely” options primarily reflects the differences caused by the degree of refurbishment of the power station, with the separation between the “most likely” and “best” PF scenarios reflecting the change in water management at the power station (i.e. the move from wet to dry ash disposal). The range covered by the FBC systems is smaller, and driven primarily by their different SO₂ removal efficiencies (the sorbent type and the degree of desulphurisation). These affect not only the SO₂ emissions, but also the volume of sorbent used, and all the burdens associated with its production and transport.

The three extreme options span the overall operating “space” of each system, which is represented by the total area covered by the FBC and PF “clouds” in Figure 7-3. This is shown to be extensive, and the relative performance of the FBC and PF systems changes depending on the points in the “cloud” at which the comparison is taken. The choice of decision variables is therefore extremely important, as it determines in what region of the operating space the system falls, and thus the relative performance between the systems. For example, in Figure 7-3, a comparison between the lower end of the FBC region (around the “worst” mid-point) and the PF “most likely” point yields that the FBC system is marginally worse than the PF system on climate change, combined acidification and eutrophication, and toxicity effects, and strongly worse on winter smog, whilst a comparison with the top end of the FBC “cloud” (around the FBC “best” mid-point) shows the FBC system to be better than the PF system on all these impacts.

7.5.5. Further Analysis: Choice of Desired FBC Operating Region

From the preceding results it can be determined that FBC burning discard is preferable to a conventional PF system for re-powering old power stations, but that the operating conditions chosen will determine to what extent, and with what confidence, this can be predicted. The next phase of the assessment is thus to select a number of discrete operating states that can be taken further into a subsequent analysis, where the options are compared against environmental, economic and social criteria, using such methods as multi-attribute utility theory/multi-attribute value theory (Basson and Petrie, 2001). The ability of the PC plots to allow the interpretation of results in a solution space is particularly useful here, as it is possible to interpret the effects of many variables simultaneously. They are thus ideal for use as a screening tool for scenario selection. Whilst cumulative probability plots of pair-wise comparisons yield more precise comparative information on the two systems, these have to be evaluated sequentially. The overview able to be achieved with principal component plots is indispensable in enabling the selection of the relevant options for a subsequent pair-wise comparison.

This section briefly presents the first step of such a subsequent analysis, which represents the next step down the decision “continuum” from tactical to operational decision contexts, i.e. a few more design choices have been taken, and the system is better defined than before. The specific station to be refurbished has been chosen, so the comparison is against a single PF option and specific mine. The model parameter uncertainty is thus considerably reduced. The empirical parameter is also reduced due to lower system variability, i.e. variability across a single plant, rather than across a number of possible plants and technologies.

Figure 7-7 presents the investigation into some key parameters. For clarity, these are split between two plots (i.e. half the options are hidden on each plot). An additional option not considered in the preceding analysis is that of dry cooling, which was initially thought to be out of the scope of the project because of the extensive changes that would be required to the plant. However, it was subsequently brought back into the analysis because of the serious competition for water in the area chosen for the refurbishment.

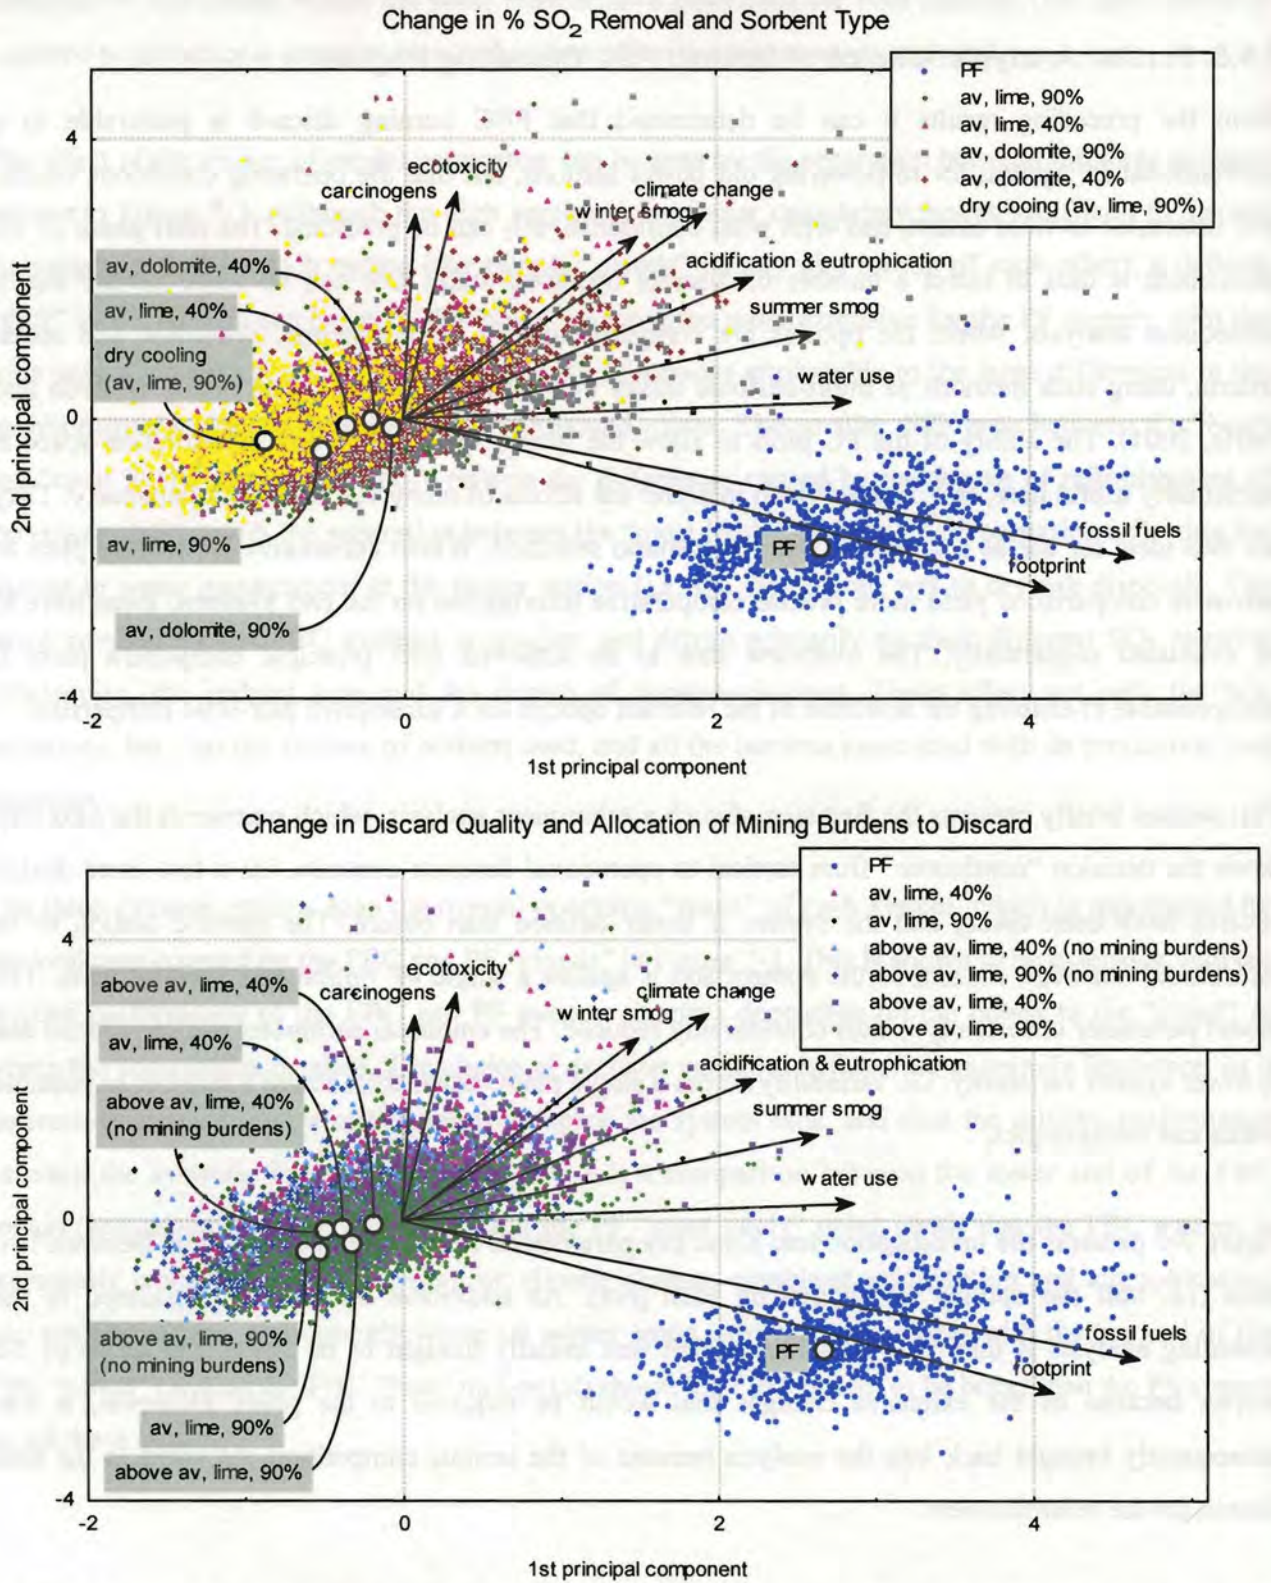


Figure 7-7 PC plots investigating possible FBC operating conditions, including sorbent and discard source, degree of desulphurisation and cooling technology. Derived from the % difference between the scenarios and the PF scenario (see Table H-15 for corresponding PC loadings).

The FBC options plot fairly close together in Figure 7-7, with the large empirical uncertainty “clouds” as before. The variance between the PF and FBC systems is much greater than that between the FBC options, and therefore falls under the first PC, with that accounting for the differences between the FBC options incorporated into the lower order PCs. To get a clearer picture of the interplay between the options it is thus necessary to also view the lower order components. Figures H-11 and 12 give the plots for the 3rd and 4th PCs, respectively. The plots each emphasise a different set of criteria, as shown by the length of the arrows in the plots, or by the PC loadings in Table H-15. The choice of which lower order PCs to represent thus effectively involves choosing the visualisation that emphasises the criteria of interest, although its contribution to the overall variance may be over-emphasised. For example, the PC table shows a high contribution to the 4th PC from the combined effect of acidification and eutrophication. Figure H-12 therefore emphasises the differences between the FBC options, which have significantly different contributions to acidification and eutrophication due to their different degrees of desulphurisation. However, this representation under-emphasises the differences between the PF and FBC systems (their “clouds” plot closer together in Figure H-12).

A high degree of sulphur removal, using lime as a sorbent, gives the best overall performance. This configuration with dry cooling gives an even better performance. A switch from lime to dolomite significantly increases both winter and summer smog, because of the higher sorbent volume requiring production and transport (the mass of sorbent required to achieve a high SO₂ removal with lime is approximately equivalent to that to achieve a low SO₂ removal with dolomite). Whether or not a higher quality discard blend gives a better performance is dependent on whether mining burdens are allocated to it. With no mining burdens allocated, it shows marginally better performance than the equivalent option burning an average quality discard. However, a fairly marked increase in a number of impacts results if it is reasoned that the discard should reflect a portion of the mining burdens (see section 7.4.2.c). The system is relatively insensitive to this choice in terms of the differences between the PF and FBC systems, but it is important in determining the ranking of the particular FBC options. Simple marginal allocation factors were calculated for the discard by varying the volume of discard produced. As with the allocation of coal, the factors are inaccurate because they are developed by simply varying the efficiency of the coal washing plant. In particular, they overestimate the portion of burdens ascribed to the discard, because the inventory model relates mining power use to the flow of coal produced.

Whilst the PC plots clearly represent the trade-offs between the options, if a particular confidence limit is to be set, an analysis based on the cumulative probability of the options is required. Cumulative

probability plots of the relative difference between the PF option and the various FBC options are given in appendix H.5. The order and distance between the curves reinforce the trends able to be discerned on the PC plots, whilst the y-axis intercept for each option represents the degree of confidence that the PF option performs better than the FBC option on that criterion. The FBC options all show a very high probability that they will significantly reduce the impacted land footprint, waterborne sulphates, land occupation and fossil fuel extraction. The very high uncertainty in the toxicity categories means that no trends can be discerned with confidence for human carcinogenic effects or ecotoxicity. Whether or not the FBC option is better or worse than the PF system on its contribution to winter and summer smog, and the combined effect of acidification and eutrophication is strongly dependent on which sorbent is used, and the degree of desulphurisation in the boiler, with water use and climate change less dependent. The only FBC option that is never worse than the PF system on these impacts (i.e. its y-intercept never above 50%) is that using lime with 90% SO₂ removal. This is also the only option that shows greater than 90% confidence of decreasing the combined effect of acidification and eutrophication.

7.6. CONCLUSIONS

Combustion of discard coal in a fluidised bed boiler is shown to be an environmentally preferable option for re-powering an in-storage power plant, compared with a conventional PF combustion system. However, for certain impact categories (most notably winter smog), this is only shown to be the case for certain plant configurations. Whilst the FBC system's better performance in fossil fuel resource consumption and impacted land footprint (incorporating both land transformation and leachate from the waste) is predicted with high probability, its degree of improvement in the other categories is less certain, and largely dependent on the manner in which mining burdens are allocated to the discard. This is because in these categories the benefits of the FBC system relative to the PF system are predominantly a result of the discard being defined as a waste (i.e. the fact that the discard does not bring any mining burdens into the FBC system, as the coal does into the PF system). The allocation basis thus introduces a fair degree of model uncertainty into the analysis. The effect of this can be assessed via a sensitivity analysis, with the degree of difference between the systems obviously decreasing as the proportion of mining burdens allocated to the discard increases. For a fairly small percentage of the mining burdens allocated to the discard (as determined from a marginal analysis), the choice does not significantly affect the relative performance between the PF and FBC systems, but is important to the choice of FBC operating parameters.

High uncertainty in both empirical and model parameters is encountered in strategic/tactical decision systems. To a large degree, this uncertainty is inherent in the nature of these studies, and thus, the emphasis is rather on the management of uncertainty, than on its reduction. The management of empirical variability and model parameter uncertainty cannot be taken in isolation, as a large component of the empirical parameter variability is also due to the unspecified nature of the system. Only by defining discrete operating “states”, corresponding to particular groupings of model parameters and the associated empirical parameters, could sufficient resolution be obtained in the results to yield conclusions as to the relative performance of the systems.

The efficacy of the iterative approach to uncertainty assessment proposed in chapter 4, is demonstrated in this case study. The high empirical uncertainty was found to be largely irreducible given the scope and nature of the study, and thus could not be reduced further following 3-5 iterations of parameter refinements and uncertainty importance analyses (depending on the particular impact category). This forces a different slant on the analysis, and suggests a bottom-up analysis of uncertainty, in which the variance of the system is analysed to see whether or not a desired level of confidence can ever be achieved for a particular criterion. The high level of uncertainty in this study meant that for many criteria, a high degree of confidence can only be achieved by increasing the degree of separation between the systems, e.g. for it to be stated with greater than 90% confidence that the FBC system performs better in terms of its combined effect of acidification and eutrophication, at least 90% SO₂ removal is required in the fluidised bed boiler.

The high degree of uncertainty introduced by the choice of decision variables can be seen by the fact that the difference between the model parameter extreme operating states, i.e. between the “best” and “worst” options, lies outside the spread of empirical uncertainty. The importance of the model parameter uncertainty analysis is demonstrated by the fact that, for many criteria, the relative performance of the systems is dependent on where in the operating space the comparison is taken. An indispensable part of a strategic/tactical study is therefore to investigate the effect of model parameters and the range of empirical variability encountered. PC plots are shown to be particularly useful for this, as they allow the results to be interpreted in a solution space, thereby making it possible to interpret the effects of many variables simultaneously. This is especially valuable for uncertainty analysis results, where the large volume of results would otherwise have to be evaluated sequentially. Whilst cumulative probability plots of pair-wise comparisons of scenarios yield more precise comparative information on the two systems, the overview able to be achieved with PC plots is indispensable in enabling the

selection of the relevant scenarios for pair-wise comparison, i.e. it is particularly useful for strategic/tactical decisions as it can be used as a screening tool for scenario selection. The plots are also able to display where overlaps are occurring between solution spaces, and are thus useful in an iterative analysis, to show when a distinct solution has been obtained. However, the use of the PC plane to display multi-dimensional results in two dimensions is only valid if the first two PCs are responsible for a high proportion of the variance in the results. This limitation of PC plots needs to be kept in mind when interpreting them, and a consideration of the corresponding PC tables is required to prevent possible misinterpretation of the relative effects of the systems.

The strength with which the relative performance of technologies can be discerned is determined by the amount of information available to qualify the systems. This relates to the degree to which the systems have been specified, which in turn relates to where in the decision process the study falls (i.e. where it falls on the strategic / tactical / operational decision context continuum). This study demonstrates the far greater strength of conclusions able to be drawn as a study progresses, notably the ability to include elements into the study that had not been feasible at an earlier level of analysis (e.g. the impacts associated with solid waste dumps), and the ability to include a quantitative uncertainty assessment. At an early stage of analysis, qualitative assessments may be all that are possible for certain information deficient aspects of the study. Whilst this is certainly preferable to not considering them at all, this study show the relative power of quantitative results.

CHAPTER 8

LCA MODELLING TO INFORM OPERATIONAL DECISIONS: CASE STUDY

The third decision context identified in section 2.3 is that of operational decisions, in which the decision to be taken is entirely within the decision maker's sphere of influence. This chapter presents a case study set in such an operational context.

8.1. DECISION CONTEXT

Operational decisions are those in which the decision is internal to the production system. These decisions are typified by short time scales and well-specified systems, and are typically concerned with the optimisation and running of existing processes (see section 2.3.3). The case study presented here is typical of such a decision type, and investigates incorporating a poor quality water stream into an existing power station water plant. The water is to be incorporated within the confines of the existing plant, and only operational changes can be made (e.g. stream splits, treatment chemical dosages etc.). Refurbishment of certain existing equipment may be required, but no new equipment or changes to the layout of the plant (i.e. the ordering of the unit operations) is to be considered. As with moving from a strategic to a tactical decision level, there is a continuum in moving from a tactical to an operational decision level. The decision type encountered here is purely operational, as it applies wholly within an existing system. However, should the need for extensive refurbishment or ultimately new or additional water treatment processes be identified, the decision moves towards that of a tactical level, i.e. the time frame of the decision increases and the number of external factors requiring consideration increases.

Site specific considerations play an important role in operational decisions (since they apply to existing systems, the site of implementation is obviously known). Operational decisions are therefore more often addressed by site-specific environmental tools, such as ecological risk assessment (ERA), than LCA. However, this case study shows that LCA does allow some useful insights into these types of decisions, as well as exploring the opportunity to interface an LCA approach with a risk assessment approach. Splitting the LCA system boundary into foreground and background sub-systems allows for such an interface, where a disaggregated foreground inventory can provide information to risk assessment models. This is demonstrated here with the use of the "footprint" analysis, which although used here as a generic indicator, so as to be compatible with the other generic LCIA indicators, could as

easily have been used to predict the actual risk of leachate generation by characterising the models with site-specific information. For operational decisions, especially those applied to primary industries, an assessment based purely on global effects would be inappropriate, as the effects of such decisions are often felt predominantly on a local level.

8.1.1. Problem Statement

A power station's tied colliery is struggling to contain and dispose of its excess effluent. Two distinct effluents can be identified. The first is from flooded old workings which need to be drained before opencast mining can commence, and is not of too poor a quality (around 400 mg/l sulphate). At present this water is being discharged into a nearby river. There is a finite quantity of this water available, as the source will be depleted once the workings have been completely drained. The second effluent is of very poor quality (around 2000 mg/l sulphate) and is arising because the open-cast pit has created an artificial water gradient in the area. Large volumes of water are therefore draining into the pit. The mine uses as much of this water as possible (e.g. for dust suppression), contains what it can on site in evaporation pans, and discharges the balance into a nearby river in a seasonal release scheme. This water will be available beyond the duration of the mine life, until the pit has been filled and the water gradient for the area re-established.

The mine and power station both fall within a generally water-scarce region, so water consumption and the increased salinity of natural water bodies are issues of critical concern. The adjacent power station consumes huge volumes of water in its cooling towers (± 140 Ml/day), so there is the potential for the power station to decrease its raw water consumption by using a portion of the mine effluent, whilst also reducing the salinity effluent load for the combined mine/power station system. However, incorporating the poor quality water into the power station's water circuits increases the need for chemical and energy intensive water treatment processes, and also changes the concentration and volume of the effluents produced. These effluents are used to condition the ash, and for dust suppression on the ash dump, and thus have the potential to change the nature of the waste. The question to be addressed is thus:

“What volume of water should the power station be taking, and where in the water circuit should this water be placed, in order for there still to be a net environmental benefit to the mine/power station system?”

The power station is not considering modifying its water plant, so only operational changes can be made to accommodate the poorer quality water feed, i.e. only the operation of the various units adjusted, with no changes to the configuration of the plant. The extent of the changes required will be dictated by existing water quality constraints within the water plant, and the quality of the mine water.

The power station consists of a number of inter-connected water circuits. An overview of the water and ashing plants given is appendix G.1.2. Mine water would most likely be used in the cooling water circuit, but could also be be used to produce ultra-pure water for the boiler circuit. A schematic of these processes is given in Figure 8-1. In both of the options trade-offs will have to be made between reduced raw water consumption and mine water emissions on the one hand, and increased chemical and energy use in water treatment, as well as increased water plant effluent and potential for leachate formation from the ash dump, on the other.

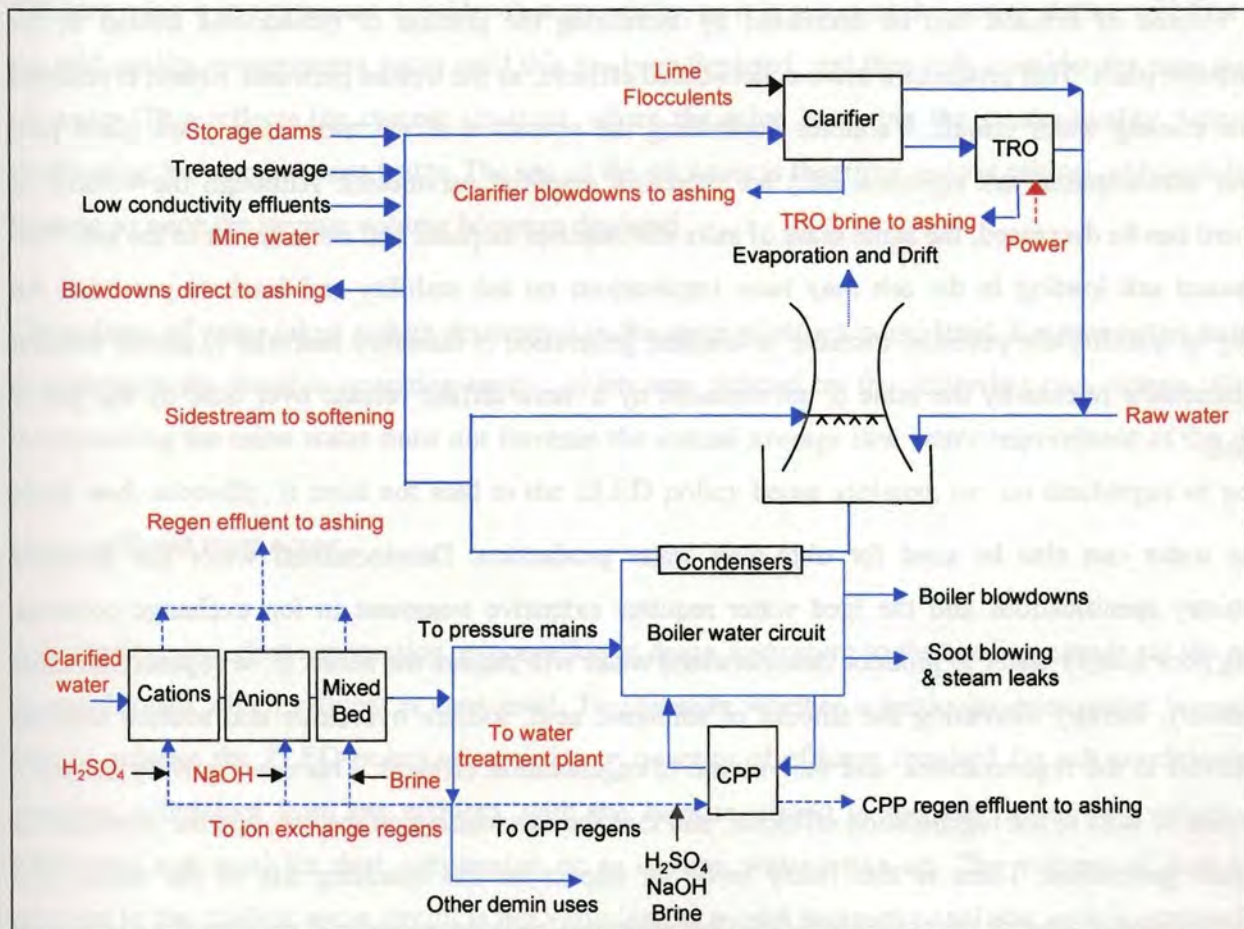


Figure 8-1 Schematic of cooling water circuit and ultra-pure water production, with flows affected by incorporating mine water shown in red.

The ZLED policy of the power station places a major constraint on the volume of effluent than can be generated from the water plant, as this may not exceed the effluent “sink” capacity. The primary processes consuming effluent are ash conditioning and dust suppression (effluent is used to moisten the ash before it is conveyed to the ash dump, and is sprayed on the unrehabilitated areas of the dump for dust control). Stormwater run-off and excess effluent are collected in dams. The water from these dams is used for dust suppression, returned to the cooling water circuit, or lost to evaporation. The moisture retentive capacity and volume of the ash produced, as well as the capacity of the storage dams, are therefore significant parameters, as they dictate the volume of effluent that can be disposed of without violating the ZLED policy.

If mine water is used as cooling water make-up, more frequent blowdowns are required from the cooling circuit, i.e. fewer cycles of concentration can be achieved. This consequently increases the volume of effluent requiring disposal, as well as the amount of salts requiring entrainment in the ash. The volume of effluent can be decreased by increasing the portion of blowdowns treated in the membrane plant. This produces a more concentrated effluent, as the treated permeate stream is returned to the cooling water circuit. Variables controlling the operation of the membrane plant (flow rate, power consumption, salt rejection etc.) are therefore essential parameters. Although the volume of effluent can be decreased, the same mass of salts still requires disposal and entrainment in the ash. This increased salt loading in the ash may have implications on ash stability and leachate potential. An ability to quantify the potential increase in leachate generation is therefore essential to assess whether an immediate release by the mine is not replaced by a more diffuse release over time by the power station.

Mine water can also be used for ultra-pure water production. Demineralised water has stringent chemistry specifications and the feed water requires extensive treatment in ion exchange columns. Using poor quality water to produce demineralised water will require the resins to be regenerated more frequently, thereby increasing the amount of sulphuric acid, sodium hydroxide and sodium chloride consumed in the regenerations, and the volume of regeneration effluent. This consequently increases the mass of salts in the regeneration effluents, and thus the salt loading on the ash, and the potential for leachate generation. There is also likely to be an impact on the operating life of the resins. The frequency and chemical consumption of the ion exchange resin regenerations are therefore important parameters.

8.2. METHODOLOGICAL FRAMEWORK

Whilst the same general approach is taken in this case study as in the previous case study, the emphasis of the assessment is very different. Operational problems are typically more constrained, and apply to well defined systems, so this case study has considerably less emphasis on scenario analysis. Other differences in system boundary definition and parameter characterisation are also evident.

8.2.1. Scenario Definition

The case study investigates the deviation from standard operating performance when a poor quality water stream is incorporated into the water plant. The basis for comparison is thus the typical operation of the plant (i.e. current operation without mine water). The incorporation of the two different qualities of mine water is investigated relative to the annual average operation of the water plant. The two different mine water types are considered sequentially, as it is assumed the power station will first take the mid-quality compartment water until this has been depleted, and then only consider the poor quality pit water. This reflects the current situation, where the mine is storing the poorer quality water and discharging the compartment water. The use of the pit water is therefore not yet critical, although it will become so once the storage volume becomes depleted.

The volume of water taken and its destination in the water plant are considered in a parametric analysis to determine the feasible operating range, which was defined by the following two criteria. Firstly, incorporating the mine water must not increase the annual average raw water requirement of the water plant, and, secondly, it must not lead to the ZLED policy being violated, i.e. no discharges of power station effluent must occur.

Ash conditioning, dust suppression, evaporation in dams, and return to the cooling circuit are the major means through which effluent is consumed. To calculate whether a particular mine water scenario is able to achieve the ZLED policy, the minimum quantity of effluent required for ash conditioning is assumed withdrawn from the effluent, with the balance placed in the ash dams, from which it is withdrawn and used for dust suppression or as cooling water make-up. The volume of dam water returned to the cooling water circuit is not varied in the model parameter analysis, as it is assumed that the return system will continue being operated at full capacity (i.e. the volume of water returned stays constant, and is the maximum possible). Where the volume of effluent generated exceeds the dust

suppression requirements and the storage/evaporative capacity of the dams, the volume of effluent used for ash conditioning is increased until the excess is absorbed. The percentage moisture of the conditioned ash is thus not varied independently as a parameter, but is dependent on the volume of effluent. The requirement to maintain the ZLED policy is thus translated into a maximum ash moisture content. This is set by when the ash becomes too wet to be able to be handled by the existing conveyor/stacker system. However, this limit is not attained, as the other constraint is found to be limiting first, i.e. before the maximum moisture content of the ash is reached, raw water use no longer shows a decrease relative to when no mine water is used.

8.2.2. System Definition

A feature of operational-type problems is that they require a detailed foreground model, as they typically involve small changes within the system, which would not be picked up by a model at a high level of aggregation. The relative placement of the foreground/background systems is therefore important, as the same level of detail is not necessary throughout the system. The exact choice of functional unit, system boundary, allocation method etc. will depend on the study in hand, although the choice is generally simpler for the more contained operational-type studies than for multi-faceted tactical or strategic studies.

8.2.2.a Functional Unit

This is taken as an average year's operation of the water plant, i.e. the cooling function and demin water required by the average annual operation of the power station. The increased power requirements of the treatment plant are assumed not to affect the operation of the power station or colliery, i.e. the increased auxiliary power requirements do not affect the annual average load of the power station. The increased power consumption is calculated to be at most 0.5% of the total power output, which supports this assumption. The changes in the system burdens arising from incorporating mine water are very small in comparison to the total burdens of the mine/power station system. The output is therefore presented as a percentage change relative to the system not using mine water.

An average year's operation of the water plant is chosen as this is most compatible with the data available for the study, although a daily or monthly operating window is equally valid. The time period cannot be less than the monitoring time interval of the data. For this study, the data is predominantly available in monthly reports, so a time period of a year chosen in order to be sure that the variance of

the data samples provides a sufficient estimate of annual variability (i.e. the monthly variation will provide an overestimate of annual variability, since the latter is typically lower due to averaging effects). The study is therefore not able to address possible “spikes” in the operation of the plant (e.g. a sudden increase in volume or decrease in quality of the mine water). These effects are better assessed in an operability study, where more sophisticated process models can assess the effect of such occurrences on the unit operations (e.g. damage to the membrane plant). The aim of this study is to capture the possible range in environmental effects, and make explicit the trade-offs that will have to be taken, given that the actual volume of water taken will be driven by operability and cost studies.

8.2.2.b System Boundary

The same system boundary is relevant here as in the previous case studies (see section 6.2). The boundary incorporates the mine, power station, associated waste and effluent management processes, and the provision of major ancillary materials. However, the relative placement of the foreground and background processes differs in this case study (see Figure 8-2, compared with Figure 6-2). Only a relatively small part of the system is directly affected by the decision in hand. These affected processes comprise the foreground system, and include the water plant (excluding potable water production), the ash and effluent disposal systems, and effluent management at the colliery. The remaining mine and power station processes fall into the background system, where they contribute to the background burdens through the provision of power to the water plant. Other background burdens considered are the provision of water treatment chemicals and materials, and the production of fuels for the transport of these materials to the power station.

8.2.2.c Flowsheet Construction

As in the previous case study, a key requirement of the system model is explicit parameter inputs and a sufficient level of flowsheet detail to explore the effects of the proposed changes. In this study it is only the water plant and effluent management processes that are directly affected by the proposed changes, and thus require detailed system models. The water plant is modelled at a unit operation level, although some units are omitted or combined (e.g. pumps are considered by their combined energy consumption). The same inventory models as described in appendix E.2 are used to provide inventory data for the balance of the power station and mine. Although not required to investigate changes to the system, the detailed models are required to provide relevant background data to the foreground system. The remainder of the background processes (provision of chemicals, fuels and auxiliary materials) are incorporated at a high degree of aggregation from available LCI databases.

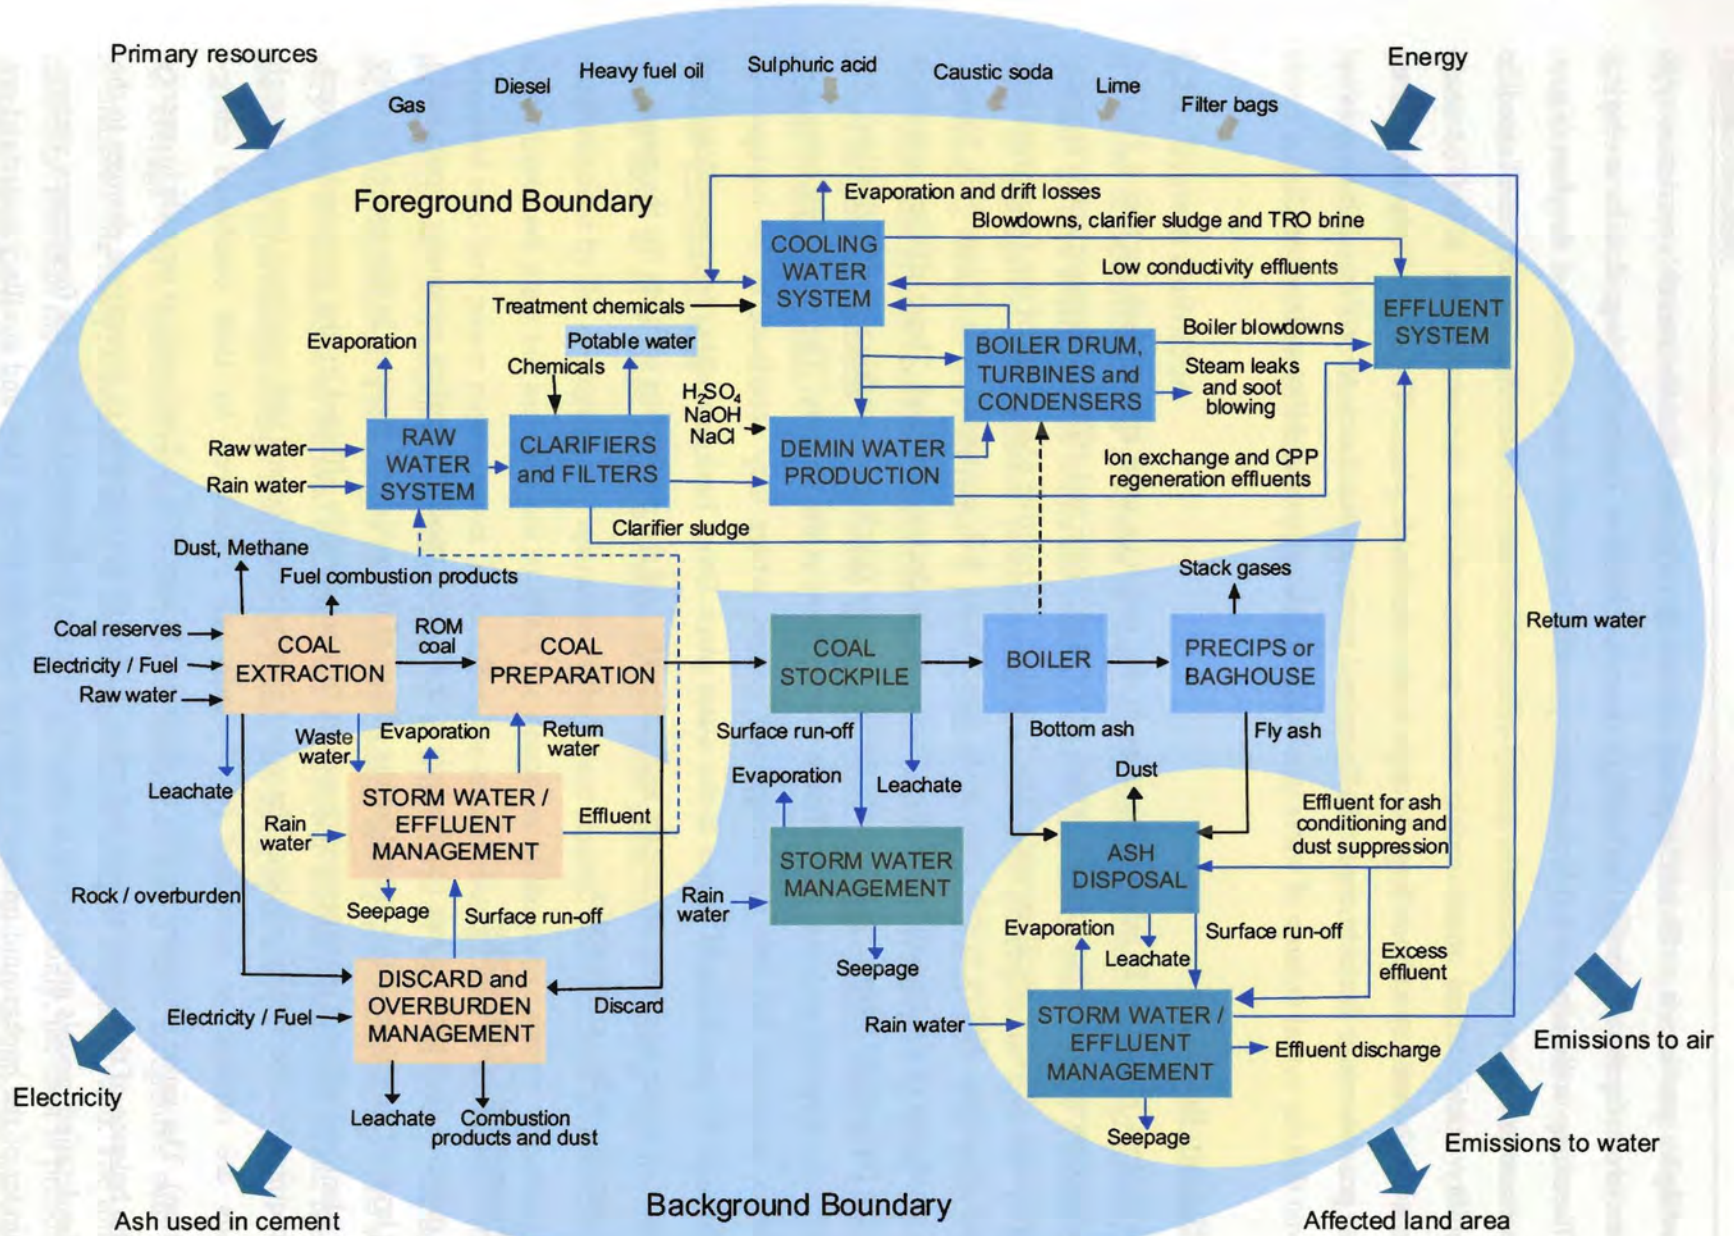


Figure 8-2 Schematic of the coal-electricity life-cycle, showing foreground boundary for mine water case study.

8.2.2.d Allocation Issues

Allocation problems do not arise in this case study. The system boundary incorporates both the mine and power station, which enables the positive effect of reducing the effluent discharged from the mine to be reflected. The colliery supplying the power station is a dedicated supplier, so the allocation of mining burdens is not an issue.

8.2.3. System Evaluation

A relevant set of indicators needs to be defined against which the scenarios can be compared. Operational type decisions typically involve only a small group of stakeholders, so a relevant subset of indicators particular to the study in hand is more likely to be agreed upon, than in tactical studies, where a wide range of impacts incorporating a wide variety of stakeholder concerns require consideration.

From the perspective of the decision makers in this study (i.e. the power station management), a key issue is the impact of the decision on the surrounding water bodies, both in terms of the volume of water used, and the increased salt loading of surface and underground water bodies as a consequence of mine and power station discharges. There is a potential that immediate mine water releases to the river may merely be shifted to a slow release of salts over time from the ash dump. The major impacts of the system will therefore only be felt on a local level, as evidenced by increased salinity in the natural water bodies surrounding the power station and colliery. High salinity is associated with a number of negative environmental impacts (Carlson and Adriano, 1993), although these are missed by the global-type indicators typically used in LCIA. The use of an “impacted land footprint” is suggested to address this gap (Hansen, 2001b), which provides a semi-site specific indicator for comparing the leachate potential of solid waste deposits (see section 5.4.2). The footprint is assumed to encapsulate emissions to groundwater and those to surface water bodies via leachate migration. Waterborne sulphate emissions are used as an indicator of direct salinity emissions (spills and discharges to surface water bodies) in the absence of an impact-level indicator (see section 5.1.3).

At the same time, there is the concern of increased background emissions associated with the increased consumption of auxiliary power and chemicals. These require minimisation from both a cost and an environmental perspective. The analysis presented here does not attempt to minimise the background

emissions, but merely to give an indication of their sensitivity to the choices made. Impacts arising from the increased production and transport of treatment chemicals, and power use in the water plant, are able to be assessed either with a “typical” set of LCIA impact categories (e.g. the EI 99 set) or with relevant environmental interventions (e.g. CO₂ emissions). The minimum criteria found suitable to capture all aspects of this system are chosen with the aid of PCA, as explained in section 8.3.3.

8.2.4. Definition of Parameter Uncertainty

As in the previous case study, a quantitative analysis of uncertainty is undertaken. Section 4.2.2 explains the methodology to characterise the uncertainty of empirical and model parameters (Figure 4-3 summarises the steps), so the discussion here is limited to the particular features of the case study. The parameters and the distributions used to characterise them are given in appendix F.3.

8.2.4.a Empirical Parameters

Empirical parameter uncertainty tends to be simpler to characterise in operational case studies than tactical/strategic studies, since the decision is applied over a short time scale to an already operational system. It is therefore possible to utilise existing process data, without the difficulties of subjectively estimating the uncertainties involved in extrapolating data from related processes or to future operating conditions. Empirical parameter uncertainty is therefore captured by the first two study-independent indicators identified by Weidema and Wesnæs (1996) (reliability of source and completeness), with little or no uncertainty being added by the three study-dependent indicators (the temporal, geographical and further technological correlation indicators). Obviously this is only relevant to parameters for which process data are available from the specific process under consideration. Where literature data, or data from similar processes, are used, the full suite of indicators is applicable.

Data from the actual system of interest monitored over at least a year is available to characterise the majority of foreground parameters in this case study. The empirical uncertainty assigned to these parameters is based simply on the variance exhibited by the data sample. This is assumed to be sufficient to capture the uncertainty arising from the reliability and completeness of the data, i.e. it incorporates fluctuations arising from measurement errors (other than systematic errors), and from inherent variability in the system. The variance of the data sample at the level at which it is monitored (daily or monthly) is used initially. This represents an overestimate, since the system is calculated for an annual output, and the annual variability would be lower due to averaging effects. However, the

higher variability is used to define the initial estimate, so as to incorporate any missed effects (e.g. deviations caused by the operational change, or biases in the data). If the parameter is found to have high uncertainty importance (see section 4.2.3 for the method by which this is determined), its definition is revisited and its variance adjusted to better reflect annual variability (so long as other possibly effects contributing to the uncertainty can be discounted).

The temporal indicator is discounted because there is less than 3 years difference between the year for which the process data are available and the year of the study (which is within the time period specified by Weidema and Wesnæs (1996) for the lowest temporal indicator “score”). The geographical and technological indicators are always redundant for data monitored from the actual process of interest, except where the operational change involves a change in technology or different operation of the technology. In this case the “further technological correlation” indicator is important, and the additional variance introduced in extrapolating data from the existing process to model the change will have to be estimated subjectively. The amount of variance added will depend on the extent of the difference between the proposed and existing technology. For this case study, the variance of the relevant parameter samples was increased 5-15% to reflect the additional uncertainty in applying data monitored from the water plant without mine water to the case where mine water is used.

Simple triangular distributions are used to characterise the majority of parameters, as these show the central tendency of the data sample, without creating a false impression of a precise knowledge of the distribution shape. Gamma distributions are used for those parameters that show a central tendency but have a low probability of extreme values, whilst uniform distributions are used for those parameters for which little is known other than their range, or have been shown by the rank-order correlation analysis to have low uncertainty importance (i.e. their first-pass rough distribution can be retained). The variability in water and coal quality is given by discrete conditional probability distributions. Discrete analyses representative of the range spanned by the data are selected from the water and coal quality monitoring data, and each analysis assigned a probability of occurring from a frequency analysis of the data (see values assigned for coal and water qualities in appendix F.3).

8.2.4.b Model Parameters / Decision Variables

The case study is applied to a particular and well defined system (i.e. to a specific power station and mine). Model parameter uncertainty is therefore low and a consideration of model parameters can be limited to those specifically associated with the operational change in hand, i.e. a large number of

scenarios do not have to be generated to cover a fuzzy definition of a future system, as in tactical and strategic studies. Only those parameters specifically controlling the proposed change require investigation as decision variables, with the bulk of the parameters able to be specified as empirical parameters, whose definitions reflect the variability in the operation of the system. Also, in contrast to the previous case study, where the large number of model parameters required that they be grouped into parameter or operational “states”, the fewer model parameters under consideration here mean that the model parameters can be treated as continuous variables, which together define the operating space of the system. This space is investigated in a systematic parametric analysis.

The model parameter analysis can be limited to four main parameters:

- those controlling the mine water taken, namely its volume and source (quality);
- the destination of the mine water in the water plant (the split between cooling make-up and ultra-pure water production); and
- the operation of the cooling water circuit (the volume of water treated in the membrane plant).

All other operational parameters are assumed to be dependent on these parameters (e.g. the treatment chemical dosage rates depend on the volume and quality of the water treated, and the operation of the effluent management system depends on the volume of effluent generated by the water plant).

8.2.5. Model Uncertainties

Simple mass balance models are used to evaluate the options (flowsheets and explanations of the models used to approximate the processes are given in appendix E.2). The models require factors derived from process data to specify them, and there are thus large model uncertainties associated with simplifying the complex processes into systems able to be specified by mass balance factors. These model uncertainties are unable to be quantified, as an alternative model is not available for comparison (e.g. one modelling the complex water chemistry involved). However, the generation of such a model is not only beyond the scope of the study, but also of limited relevance given the goals of the study. However it is important that the limitations of the model be borne in mind when considering its output. Specifically the fact that the output from the model is only ever as good as the data put into it. Possible “knock on” effects on the operation of the system are in danger of being missed (such as a decrease in ion exchange resin or membrane lifetimes). The uncertainty analysis can highlight where simplifications may potentially lead to inaccurate results by placing high uncertainty on the relevant empirical

parameter inputs. It can thus identify whether the simplification is significant and so direct where more detailed modelling or better data collection is required. However, this is obviously only relevant where a problem is suspected, and unforeseen effects are still in danger of being missed by the analysis.

Model uncertainties due to data gaps and omissions are only able to be assessed qualitatively. Data gaps identified in this study include certain water plant materials, notably the osmosis plant membranes and the sand filter beds. Production and transport of these materials are not included, as no data are available on their usage rate, or no suitable LCI data on their manufacture could be found. The effect of increased membrane plant usage to treat the additional blowdown from the cooling water circuit is reflected only in the additional power required to run the membrane plant. No effects on the membrane plant itself are included, although a significant decrease in membrane service lifetime is likely. Maintenance procedures are also not included in the model, so the possible increase in maintenance caused by the use of mine water is not reflected. The concentration limits set for the cooling water circuit, which determine the volume of cooling water blowdowns and the volume withdrawn for lime softening in the clarifier, are assumed to be set sufficiently high so as not to increase corrosion or scaling problems.

The most significant model uncertainties are therefore around applying the model based on the existing process, to the process with the proposed change. Treatment chemical dose rates, ion exchange resin regeneration rates etc. are derived from the existing process, and assumed to still apply in the changed system, i.e. the dose rates are assumed to vary linearly with the volume of water treated or, in the case of ion exchange resin regenerations, with the mass of salts requiring removal. The proposed changes are reasoned to be small enough to make such assumptions valid. However, the effect of large deviations from typical operating conditions can be investigated by a sensitivity analysis (as is demonstrated by an investigation of salt rejection in the membrane plant in appendix I.4.4). The effect of the change can be incorporated to a degree by translating the model uncertainty into an empirical uncertainty, i.e. by assigning higher uncertainty to the input parameter when it is applied to the changed system (increasing the variance attributed to the technological correlation DQI). In this way possibly limiting simplifications can be highlighted by the uncertainty assessment, thus directing where additional modelling/data collection is required to achieve a particular level of certainty in the output.

In addition, simplifications have been made in modelling the relation of the foreground system to the background system. The foreground system is the focus of the operational decision, and incorporates

the water plant and the effluent management systems of the mine and colliery (see Figure 8-2). The remainder of the power station and colliery processes fall into the background system, where their input into the decision is purely through the provision of power to the foreground system. The proposed changes to the water plant cause an increase in its power consumption, through increased use of the membrane plant and additional pumping requirements. However, these additional auxiliary power requirements are small in comparison to the total power output of the power station (less than 0.5%), so the changes they cause in the background system are assumed to be negligible. The average load factor and mass of coal consumed in the power station are therefore assumed to remain the same, i.e. no additional mining impacts are considered, and the ash volume and dump size do not reflect the slight increase in power consumption.

Extending the inventory data to a measure of the impacted land footprint is far more complex than the equivalency factor approach used to calculate the other impact indicators. The impacted land footprint is calculated in a separate deterministic model, and is therefore not reported as an uncertainty sample, but as a single “most likely” value. Its uncertainty therefore has to be estimated separately, and it is the only criterion for which the uncertainty estimate is not based on a propagation of empirical parameter uncertainties in a simulation analysis. Instead, the range in values is estimated by calculating the footprint from two extremes sets of data inputs, representing the 90% confidence interval of the empirical parameter inputs. A distribution was then fitted to the calculated mid-point value to reflect this range, with the distribution shape matched to the dominant source of empirical uncertainty (the variation in rainfall). The footprint is calculated for two extreme reaction states, modelled as equilibrium and “straight through” (see section 5.4.2). This was found to cause only a 6-10% variability across the size of the footprint, much less than that predicted to arise due to empirical variability ($\pm 50\%$). The consideration of reaction state therefore did not change the uncertainty estimate.

8.3. RESULTS AND DISCUSSION

The following section presents the results of the analysis into the inclusion of mine water into the water circuit of a large, modern, wet cooled power station. The investigation starts broadly by identifying the feasible operating range, determined according to the criteria given in section 8.2.1. This range requires identification so that the limits of the model parameters can be set, and the effect of the uncertainty they introduce can be considered in a model parameter analysis. As in the previous case study, PCA is used to provide insight into the structure of the complex data sets, and to provide guidance as to the selection of the criteria against which the options can be compared. The PC plots provide a visual overview of the empirical uncertainty, which is subsequently presented more precisely using cumulative probability, and box and whisker plots.

8.3.1. Identification of Operating Range / Selection of Options

The volume of water taken and its destination in the water plant are considered in a parametric uncertainty analysis over the feasible operating range, in which the parameter's operating ranges are discretised and "stepped through" to approximate the continuous operating space. A fairly coarse step interval is chosen, as the greater accuracy of a smaller interval was considered unnecessary given the high empirical parameter uncertainty (see section 8.3.3). The volume of water taken is considered at 1 Ml intervals, up to the maximum available for the compartment water case, and up to the maximum volume able to be incorporated before violating the constraints set on the system for the pit water case.

This model parameter uncertainty analysis should not be confused with an optimisation analysis. The former aims to investigate the range in potential effects caused by using mine water, whilst the latter would be required to identify a particular operating point to best meet a particular set of objectives. The "extreme" options selected below and taken forward in the analysis are chosen because they represent the edges of performance (i.e. assuming a continuous operating space, the performance spanned by the options is from the base case to the extremes identified). To determine the optimum operating point would require an objective function to be defined, and a suitable optimisation routine to be chosen (e.g. non-linear programming). A multi-objective optimisation process would be required, as the problem involves the contrasting objectives of minimising costs and maximising environmental performance. Whilst a cost objective can be easily defined, a suitable environmental objective is more complex to define. This could either be to simply minimise water use, or the salinity load to natural water bodies,

or it could involve a number of environmental criteria, such as those in Figure 8-9. The current analysis is assumed to be prior to such an optimisation analysis, where the tradeoffs between the criteria and options are explored so that an informed choice of criteria can be made in an optimisation process.

Figure 8-3 shows the operating range identified for the medium quality mine water (compartment water) and Figure 8-6 shows that for the poor quality mine water (pit water). As a three-dimensional system is being represented on two-dimensional plots, one of the parameters have to be held constant in each plot, whilst the other two are “stepped through” their range at discrete intervals. The value of the parameter held constant is chosen so as to maximise the operating ranges of the other two parameters. The points on the graphs represent the feasible operating states at the considered resolution of the model parameters. The points are joined as the parameters are in fact continuous variables, and any point along a given line can be considered a feasible operating state. Similarly, the volume of water taken is a continuous variable. The lengths of the lines represent how wide an operating range is spanned for that particular combination of parameters.

In the simple inventory model used, the changes arising from the use of mine water are assumed to be proportional to either the volume of the water taken, or to the mass of salts in the mine water (see section 8.2.5). This results in the linear relations seen in the figures dominating the analysis. For example, the number of resin regenerations, and so the additional demineralised water required to effect the regenerations, is assumed proportional to the mass of salts requiring removal, and consequently to the volume of mine water used to supplement raw water make-up to the demin plant. The volume of mine water to the demin plant is thus linearly related to the raw water requirements of demineralised water production. Similarly, the membrane plant is modelled with a constant percentage salt rejection and permeate recovery. The decrease in cooling water make-up as the throughput of the membrane plant increases, is thus linearly related to the volume of water treated in the membrane plant.

The membrane plant is assumed to be able to operate at an average annual maximum of 80% of its total capacity, to allow for maintenance. Where this limit is not reached, the capacity of the membrane plant is determined by the volume of the cooling circuit blowdowns (the volume treated cannot exceed the volume of the blowdowns), and the effluent requirements of the system (sufficient effluent must be produced for ash conditioning, dust suppression and return to the cooling system), e.g. in Figure 8-3, for 2 Ml water incorporated into the cooling circuit, the membrane plant can not be operated at more than 40% capacity or insufficient effluent to meet the requirements of ash disposal results.

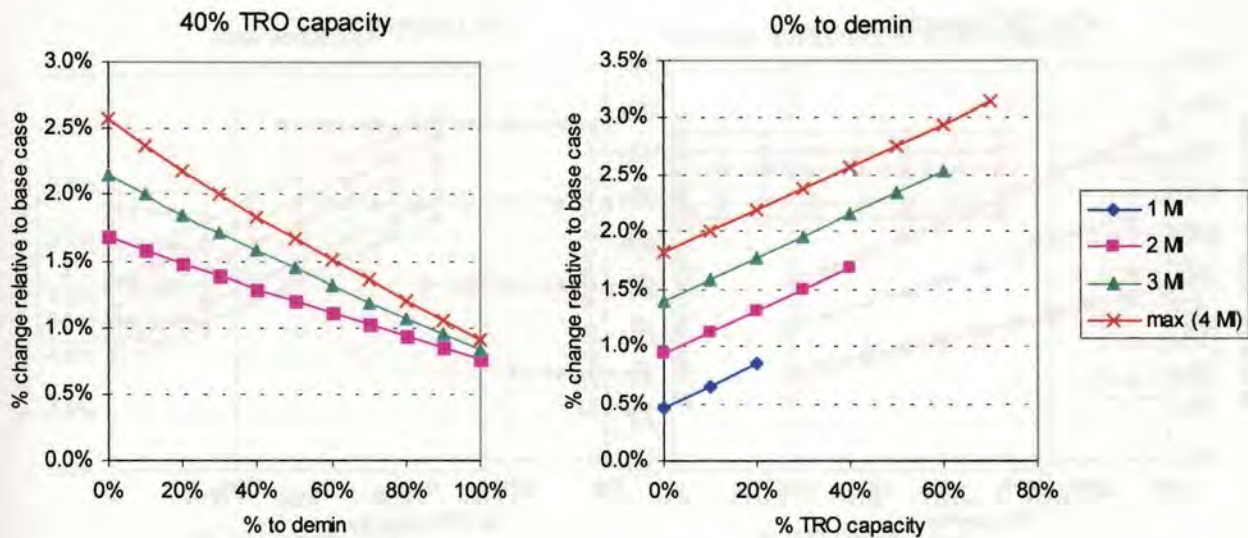


Figure 8-3 Raw water use for various volumes of compartment water taken and different plant configurations.

Compartment water is able to be incorporated across practically the full operating ranges of the decision variables (curtailed only by the maximum membrane plant capacity), whilst the pit water case is seen to have relatively small operating ranges, particularly as the volume of water taken increases, i.e. both minimum and maximum capacity limitations for the membrane plant, and maximum demin plant/cooling water splits (e.g. in Figure 8-6, for 1 Ml mine water make-up in the cooling circuit, the membrane plant can only be operated between 20 and 50% capacity).

There is a limited volume of compartment water available from the mine, and Figure 8-3 shows that this maximum volume can be incorporated into the water plant whilst easily meeting the criteria (i.e. the maximum compartment water can be incorporated without violating the ZLED policy, and results in a net water saving to the plant). At low volumes of compartment water, the membrane plant needs only to be run at low capacity. Even at the maximum volume of compartment water the membrane plant does not have to be run at its maximum capacity, although the minimum raw water usage occurs at the maximum membrane plant capacity (see Figures 8-3 and 5). Using poor quality water to produce demineralised water increases the salt loading on the ion exchange resins, requiring them to be regenerated more frequently, and thereby consuming additional demineralised water and chemicals. The greatest water savings therefore always occur when all the mine water is placed in the cooling water circuit (i.e. 0% to demin). Reductions in sulphate emissions also show this trend, with the maximum reductions possible when all water is placed in the cooling circuit (see Figure 8-4).

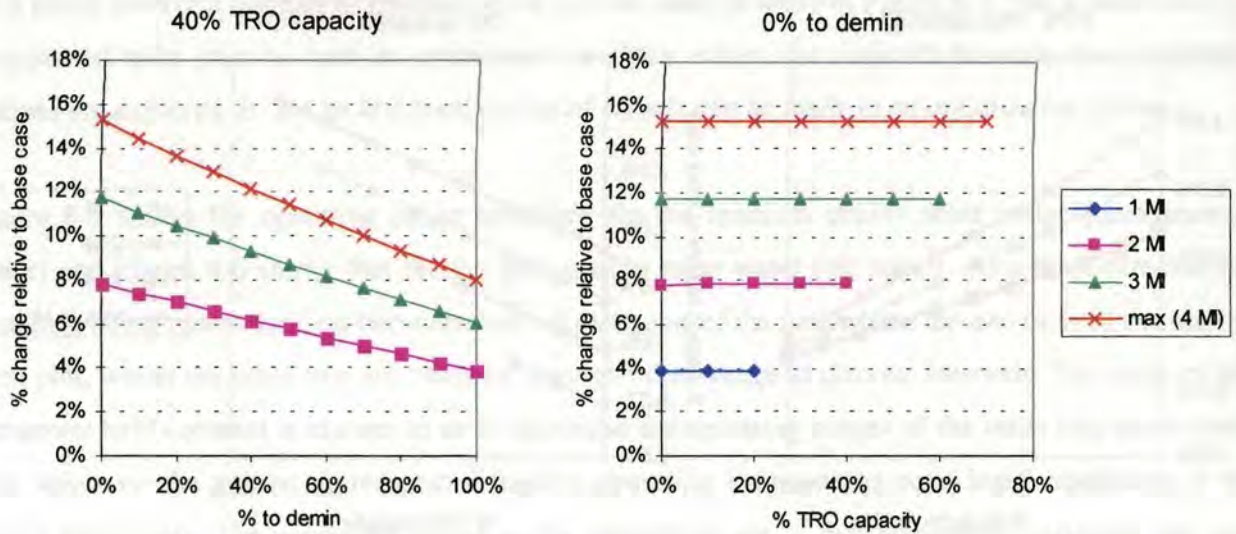


Figure 8-4 Sulphate emissions to surface water for various volumes of compartment water taken and different plant configurations.

The operating range reflected in Figure 8-4 is for those conditions that meet the operating criteria (i.e. no spills from the power station effluent management system). The mining effluent is therefore the only notable source of waterborne sulphates reflected in Figure 8-4, and the mid-points of the sulphate emissions essentially remain constant over the given operating range. However, the operation of the membrane plant affects the volume of effluent produced, and therefore the risk of spills from the effluent system. This feature is captured by the empirical uncertainty analysis (see section 8.3.3).

Figures 8-3 to 5 show that for maximum water savings and reduction in waterborne sulphate emissions, the maximum compartment water should be taken, with all of it used as cooling water make-up, and the membrane plant operated at its maximum possible capacity. This extreme case is investigated in greater detail in the following sections.

As with the compartment water case, the maximum water savings and reductions in waterborne sulphate emissions are achieved for all pit water incorporated into the cooling water circuit, and the membrane plant operated at its maximum capacity (see Figures 8-6 to 8). The trend of decreasing water savings for increasing volumes of mine water used for demineralised water production is even more marked for use of pit water, because the very poor quality water causes considerably more frequent resin regenerations.

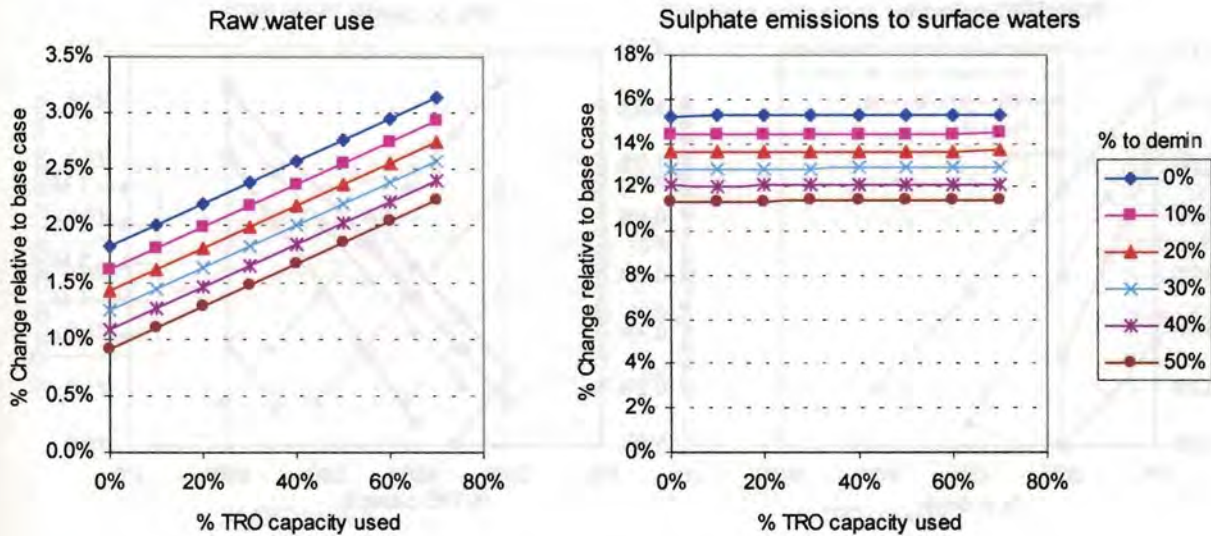


Figure 8-5 Raw water use and sulphate emissions for various plant configurations and maximum volume of compartment water taken.

The opposite trend to the compartment water case is seen for increasing volumes of pit water incorporated into the cooling circuit. For the use of compartment water, the maximum water intake causes the greatest raw water savings, whereas the raw water savings decrease as the volume of pit water taken is increased. This is a result of the very poor quality water considerably decreasing the number of cycles of concentration that can be achieved in the cooling water circuit, thereby greatly increasing the volume of cooling water blowdowns.

For the same reason, utilising compartment water enables greater raw water savings than pit water (annual average of 2.5% compared with 1.3%), although the use of pit water shows far greater reductions in sulphate emissions to surface water bodies (60% compared with 15%). There is also a difference between the two water types in the certainty that this reduction in sulphate emissions will occur. For high volumes of pit water this reduction may not always be achieved, because the high effluent volumes resulting from the use of pit water carry the risk of spills from the effluent system. In addition, they increase the risk of leachate from the ash. These issues are explored in detail in section 8.3.3.

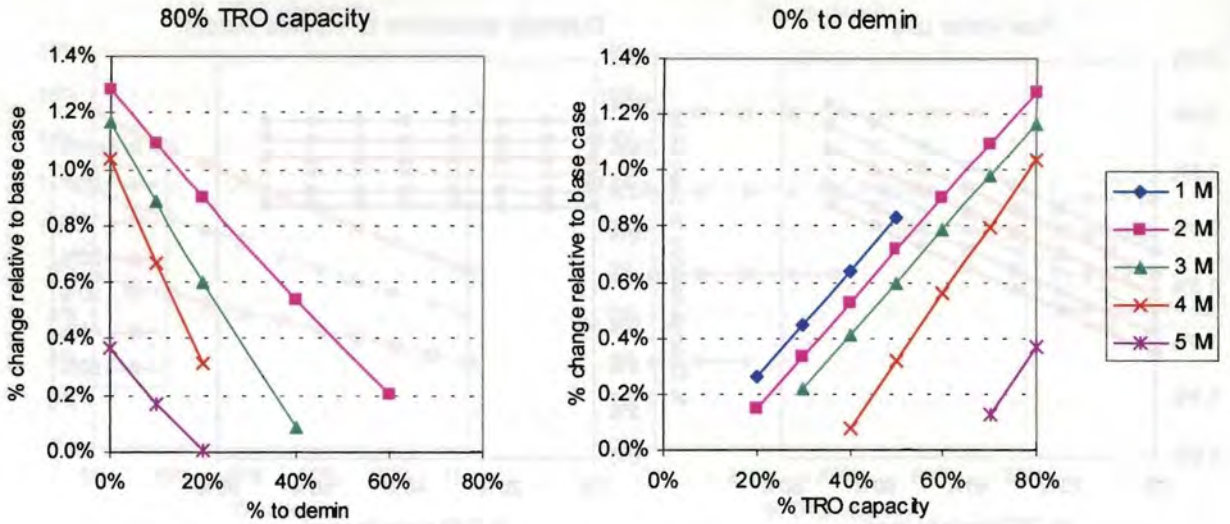


Figure 8-6 Raw water use for various volumes of pit water taken and different plant configurations.

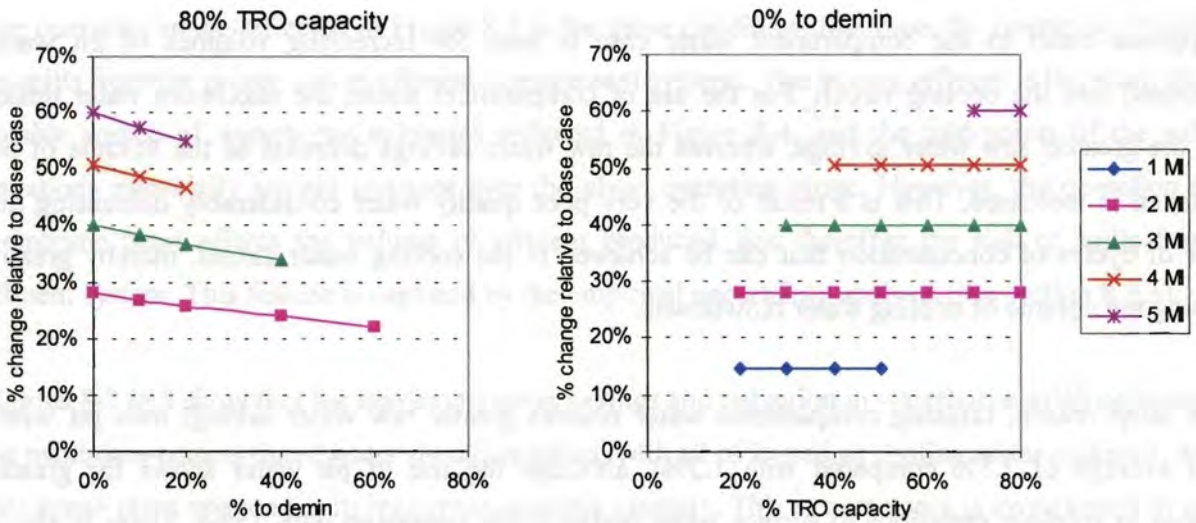


Figure 8-7 Sulphate emissions to surface waters for various volumes of pit water and different plant configurations.

Two extreme operating states are identified from the operating ranges shown in Figures 8-6 to 8. The first is if the maximum water savings are to be obtained. Figure 8-6 shows this to be for 2 Ml/day pit water utilised with maximum membrane plant operation and all water used in the cooling circuit. However, if the aim is rather to achieve maximum reduction in sulphate emissions to surface waters, Figure 8-7 shows that the same plant operation but using 5 Ml/day pit water should be used. These two extreme cases are investigated in greater detail in the following sections.




Figure 8-8 Raw water and sulphate emissions for various plant configurations and 2 Ml/day pit water taken.

8.3.2. Selection of Criteria for Comparison

A significant problem with studies considering uncertainty is the large data sets that result from the analysis. PC plots are suggested as a convenient representation of uncertain data where simultaneous interpretation across a number of dimensions (environmental interventions or impact categories) is required. Two such plots are shown in Figures 8-9 and 10, for the analysis conducted for a set of impact categories, and the key environmental interventions contributing to these categories, respectively (see section A.3.2 for an explanation on interpreting principal component plots). Raw water use, sulphate emissions to surface waters and the impacted land footprint are also considered in both analyses (see section 8.2.3).

For the two sets of criteria, the first two PCs account for 78% and 73% of the total variance, respectively. This high proportion of total variance supports the use of the PC plots to interpret the systems, as little information is being lost in the two-dimensional representation. The first PC reflects the changes between the systems resulting from the additional auxiliary power consumption of the water plant, and the second the mine water related effects (i.e. the decrease in waterborne sulphate emissions and water use, and the increase in chemical manufacture and transport). The principal component tables and a brief discussion on their interpretation are given in appendix I-1.

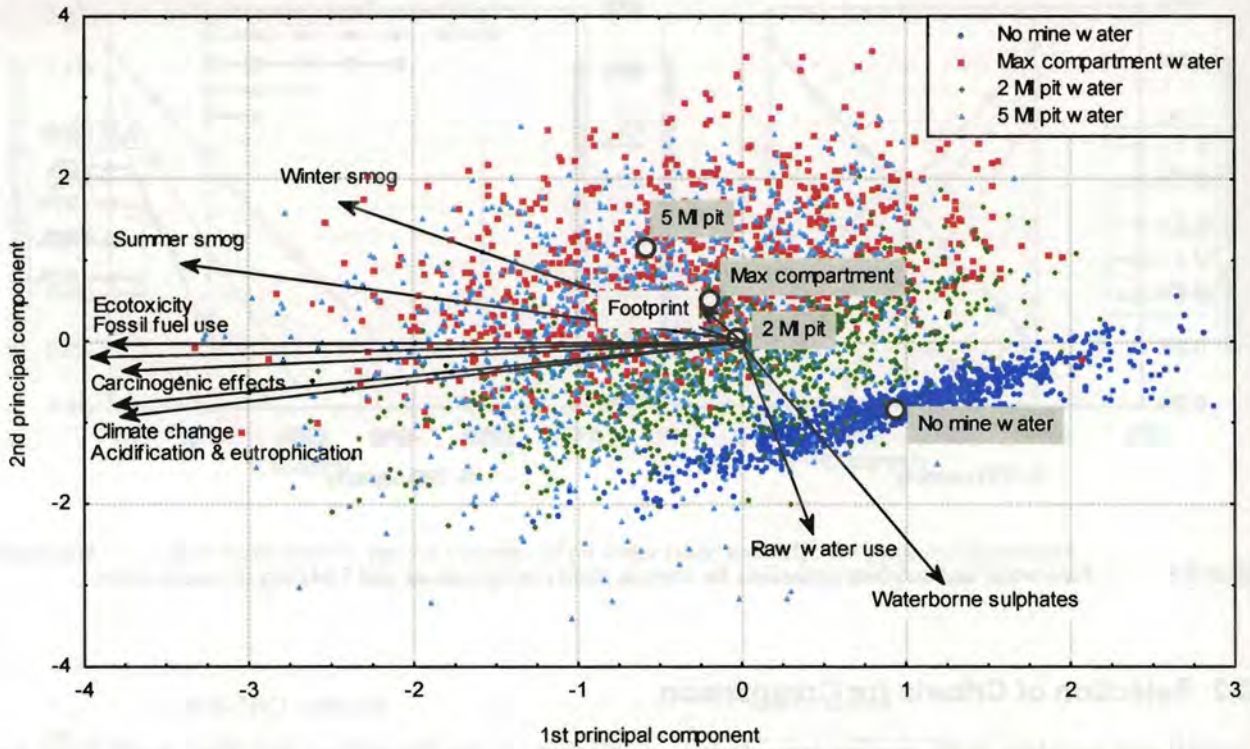


Figure 8-9 Output samples of the three mine water scenarios selected above (representing the edges of performance) transformed on the 1st and 2nd PC plane. The analysis is based on the % difference between the options and the no mine water case for the EI 99 set of impact categories, raw water use, sulphate emissions to surface water and the impacted land footprint (see Table I-1 for corresponding PC loadings).

The “strength” of the criteria (the extent and certainty of the differences between the systems for each impact or intervention) is represented by the magnitude of the coefficients of the eigenvectors corresponding to the high-order PCs (i.e. the PC loadings in Tables I-1 and I-2), or the length of their vectors in Figures 8-9 and 10 (see section A.3 for theory on calculating and interpreting PCs). The criteria are shown by the analysis to be more or less equally good at discriminating between the options, other than those for water use and the impacted land footprint. The criteria related to the increase in auxiliary power consumption (most strongly indicated by the CO₂, SO₂, and NO_x emissions, and their related impacts) are responsible for the strongest differences between the options, with waterborne sulphates and winter smog/TSP emissions showing the next greatest contribution.




Figure 8-10 Output samples of the three mine water scenarios selected above (representing the edges of performance) transformed on the 1st and 2nd PC plane. The analysis is based on the % difference between the options and the no mine water case for key environmental interventions and the impacted land footprint (see Table I-2 for corresponding PC loadings).

As discussed for the previous case study (see section 7.5.2), the fact that impact categories and environmental interventions are in various units means that the PCA has to be based on a correlation analysis. This standardisation of the data removes the relative significance of the criteria, as the analysis is based on the correlations between the options (i.e. on the magnitude of the differences between the options), and does not take the relative magnitude of the criteria into account. In addition, the fact that the options have many processes in common (only the water plant and ash disposal are changing, whilst the mine, other power station processes and background processes are common to all options), as well as the fact that impact assessment uncertainty is common to both processes, requires the analysis to be based on a ratio or percentage difference between the options. This is required to remove the correlations of uncertainties between the options (see section 7.5.2), and has a similar “levelling” effect on the criteria as the PCA does (i.e. removes any consideration of the relative magnitude of the criteria).

The “strength” of the criteria determined by the PCA thus only reflects the relative magnitude of the observed changes in the system, and not the actual amount emitted, nor their significance relative to the total amount emitted by the wider system or region. Table 8-3 reflects this, and shows the change in burdens both as percentage change for the system considered (the water plant and effluent disposal system), and for the total mine/power station system. This shows that the change in the auxiliary power requirements of the water plant is notable across the options for the system considered, although this is relatively insignificant in terms of the total power output of the power station, i.e. the change in the energy related emissions arising from power use in the water plant is very small in comparison to the emissions from the whole mine and power station system.

Table 8-3 shows that those environmental interventions with a high percentage contribution from the system of interest to the total mine/power station system are under-emphasised when the background levels are not taken into account, i.e. the significance of the impacted land footprint (which falls fully into the system of interest, since only the ash dump footprint is considered in the analysis), raw water use, and waterborne sulphates are under-emphasised by the analysis (as seen by comparing their values in the 1st column with those in the 3rd column of Table 8-3). This accounts for raw water use and the impacted land footprint being shown as only weak criteria by the PCA (see Figure 8-10 and Table I-2), since the correlation matrix on which the analysis is based does not take into account the relative significance of the criteria (i.e. as if based on the 1st and not the 3rd column of Table 8-3).

Although the definition of a reduced foreground system boundary is useful from a modelling perspective, as it allows only those processes directly affected by the decision choices to be modelled in detail, Table 8-3 shows that a functional unit based on that sub-system may distort the significance of the burdens from the sub-system. The consideration of a smaller sub-system is also of value in that it allows changes to be discerned which would otherwise be lost in the background “noise” of the system, which may be essential for an operational decision that only affects a small part of a larger operation. However, when interpreting the significance of the changes, the magnitude of the emissions from the total system need to be borne in mind.

Table 8-3 Significance of the changes observed for the maximum compartment water case relative to the no mine water case when the total mine/power station system is considered rather than the water plant alone.

Burdens / Impacts	% Change relative to no mine water case (sub-system) ¹	System burdens of base case as a % of the total burdens ²	% Change relative to no mine water case (total burdens) ³
Carcinogenic effects on humans	-10%	0.6%	-0.1%
Summer smog	-11%	2.3%	-0.3%
Winter smog	-10%	2.6%	-0.3%
Climate change	-9%	0.6%	-0.1%
Ecotoxic emissions	-10%	1.1%	-0.1%
Combined effect of acidification and eutrophication	-9%	0.6%	-0.1%
Extraction of fossil fuels	-11%	1.0%	-0.1%
Impacted Land Footprint	-2%	100%	-1.6%
Water use	3%	79%	2.4%
CO ₂	-9%	0.6%	-0.1%
NO _x	-10%	0.6%	-0.1%
SO ₂	-9%	0.6%	-0.1%
Methane	-11%	1.9%	-0.2%
Non methane VOC	-11%	2.3%	-0.3%
Total suspended particulates	-10%	10%	-1.2%
Waterborne Sulphates (surface)	15%	69%	9.1%

1. % Change between systems for functional basis considered, i.e. water plant and ash and effluent management (see Figure 8-2).

2. Burdens arising from sub-system considered as a % of the burdens arising from the total mine/power station system.

3. % Change between systems considering the total burdens of the mine/power station system.

The lengths and clustering of the vectors in the PC plots provide a visual representation of the information contained in the PC tables, and their interpretation thus follows that given for the tables in appendix I.1. The plots clearly show the separation out of energy-related effects, chemical/material related effects and water-related effects. The subsequent analysis can thus be significantly streamlined by basing the analysis on a representative environmental intervention for each of the distinct areas, thereby decreasing the number of selection criteria, but not losing any information on the relative effects of the options. Basing the analysis on a few chosen interventions is possible in this case study because of the clear correlation between the interventions, and the fact that the analysis is based on the relative difference between the systems, i.e. the change in energy consumption can be as well represented by CO₂ emissions as by climate change. However this may be less applicable where the ability to interpret the relevance of the criterion in a subsequent valuation step is important.

CO₂ emissions were chosen to represent the increased energy consumption of the water plant (cluster of CO, NO_x, SO₂ and CO₂ burdens in Figure 8-10, and toxicity, fuel use, climate change and combined effect of acidification and eutrophication in Figure 8-9). The relatively low emissions of methane and NMVOCs from the boiler plant, relative to the other emissions (i.e. NO_x, SO₂ and CO₂), means that

these are the only transport emissions not dominated by boiler emissions. Thus the trends caused by increased transportation and chemical consumption are able to be distinguished by the PCA for methane (and NMVOCs), i.e. their vectors plot in an orientation distinct to those of the auxiliary power related-burdens (methane and NMVOCs in Figure 8-10, and summer smog in Figure 8-9). Methane emissions are therefore chosen to represent increased transportation and chemical consumption. TSP emissions (winter smog in Figure 8-9) are retained as they plot distinct from the other criteria. These primarily represent the increase in chemical consumption, as they are dominated by the emissions from the production of lime for water treatment (the background LCI used for lime production has very high particulate emissions).

8.3.3. Model and Empirical Parameter Uncertainty

Three different representations of uncertain output are included here, as each provides different insights into the system. The cumulative probability plots are able to quantify the degree of overlap of the options, whilst box and whisker plots display the relative magnitude of model and empirical parameter uncertainty. The PC plots clearly display any tradeoffs that have to be made, as they are able to consider all indicators simultaneously.

The results presented in Figures 8-9 and 10, and in the cumulative probability and box and whisker plots, are for the three extreme scenarios identified in section 8.3.1. These scenarios are chosen as they represent the extremes of the operating ranges, thereby allowing the plots to be interpreted as representing a continuous operating space, where the full range of performance spanned by introducing mine water is from the base case to the extreme scenarios identified. An additional scenario, in which the membrane plant is refurbished and the plant is operated more tightly within its design specifications is also considered in the cumulative probability and box and whisker plots. The refurbished membrane plant considers a significantly different way of running the water plant, and demonstrates some different trends, so is included for comparison with the selected extreme options (see Figure I-8 and accompanying discussion in appendix I.4.4). The refurbishment allows higher salt rejections to be obtained, whilst the more controlled operation of the plant means that these occur over a smaller band of variability. To ensure that the inlet concentration to the membrane plant does not exceed the maximum design specification for guaranteed salt rejections, the maximum allowable concentration limits for the circulating cooling water have to be decreased. These lower limits decrease the number of cycles of concentration that can be achieved and thus increase the volume of cooling water blowdowns.

8.3.3.a Principal Component Plots

In addition to the insights obtained into the interplay between the criteria, the PC plots provide a visual representation of the uncertainty spread in the output samples. Instead of a single mid-point value, each scenario is associated with a “cloud” of points, the spread of which represents the empirical uncertainty present (see Figures 8-9 and 10).

Although the output samples (the uncertainty “clouds” on the PC plots) are for the three extreme scenarios (the mid-points of these scenarios are shown on Figures 8-9 and 10), the solution space to the problem is in fact continuous. The plots are therefore better interpreted as showing the range of both empirical and model parameter uncertainty encountered in the study. The empirical parameter uncertainty is reflected in the scatter of points on the plots, and the model parameter uncertainty in the spread of the options. The model parameter uncertainty can be interpreted as forming a continuum away from the base case (blue dense scatter on the PC plots), in the direction of the colour bands, i.e. in the direction of the “waterborne sulphate” and “water use” vectors, but bent to the left (in the direction of the negative impact arrows, e.g. winter and summer smog, climate change etc.) as the volume of mine water taken increases. Thus depending on the volume and quality of the water incorporated in the cooling water circuit, an option could fall anywhere between the base case and the scatter of the high volume options.

Figures 8-9 and 10 show that as the volume of mine water taken increases, so the reductions in sulphate emissions and raw water use increase (the options plot further against the direction of these vectors). However, as the volume of mine water taken increases, so the chemical and energy consumption also increases, as reflected by the “swing” of the options in the direction of vectors representing these effects (e.g. climate change, winter and summer smog etc.). The opposing directions of these vectors show that the raw water savings and decrease in sulphate emissions always come at the expense of increasing the other potential impacts and emissions represented on the PC plots.

The increasing uncertainty in the systems as the volume of mine water taken increases is shown by the increasing scatter of the output samples. The system incorporating no mine water shows fairly little uncertainty spread, as this is a well characterised system based on good data. The subsequent incorporation of mine water increases the sample spread, as it introduces an element of uncertainty into the modelling of the water plant, e.g. uncertainty in chemical usage, material service lifetimes etc. As

the operation becomes more extreme (greater volumes of water incorporated), so the uncertainty increases (departs further from the normal operation conditions for which the data samples apply). The large uncertainty around the high-volume mine water options is primarily due to predicting waterborne sulphate emissions from spills from the effluent management systems. Their uncertainty cannot be reduced, as these emissions are only roughly quantified in the inventory model using approximate seepage and run-off rates. They are also to a large degree inherently uncertain, as the largest contributors to their uncertainty are rainfall and evaporation, which are intrinsically variable quantities.

8.3.3.b Probability Plots

A more precise investigation of empirical parameter uncertainty can be achieved using cumulative probability plots (see Figure 8-11 for that of raw water use, whilst plots for the other selected criteria can be found in appendix I-2).

In the cumulative probability plots the options are compared as ranges of probable output, which give an indication not only of whether the option is better (or worse) than the base case, but of how probable its better (or worse) performance is. The results are presented as a percentage change relative to the base case (no mine water taken), i.e. by:

$$\frac{Burden_{\text{base case}} - Burden_{\text{mine water option}}}{Burden_{\text{base case}}}$$

with the “burden” for each option calculated for an average year’s operation of the water plant and effluent management system, except for the footprint, which is calculated as the extent of the land affected after 100 years. The ratio is calculated this particular way round so that a beneficial change is reflected by the degree to which the sample falls in the positive range on a plot of its cumulative probability. For the criteria selected, less is always better, so a sample falling in the positive range indicates that that option decreases the emissions relative to the base case. The degree of confidence in the relative performance of the options is given by the y-intercept, which gives the probability that incorporating no mine water is better than the option, i.e. options plotting in the positive range alone show 100% confidence that they are better than the no mine water system.




Figure 8-11 Cumulative probability of % decrease in raw water consumption relative to the base case, in which no mine water is used.

For example, in Figure 8-11 a y-axis intercept of 12% shows there is a 12% probability that using 2 Ml/day of pit water will not decrease the raw water use (i.e. an 88% probability that the option is better than the base case). To be fully confident that using 2 Ml/day pit water will always decrease the raw water use requires refurbishment of the membrane plant, i.e. for this option, the y-axis intercept decreases to 2%, indicating a 98% probability that incorporating 2 Ml/day pit water in the refurbished plant decreases the raw water use relative to the base case. Incorporating a high volume of pit water (5 Ml/day) is associated with significantly more uncertainty. Whilst the median lies above zero (a decrease of 0.5% is evident at 50% probability), a large proportion of the sample spread lies below zero (i.e. only 60% confidence of decreasing the raw water requirements). Figure 8-11 also shows that incorporating the maximum compartment water introduces no risk that the raw water benefits would not be achieved (i.e. its curve lies fully above the zero intercept on the y-axis).

The degree of separation between the probability curves for each option gives an indication of the model parameter uncertainty, i.e. the extent of the differences between the curves represents how significant an effect the operating conditions have on the performance of the options. For example, in appendix I-2, the plot for CO₂ emissions shows the curves to all fall more or less on top of each other. This indicates that model parameter uncertainty has little effect on CO₂ emissions, whilst in Figure 8-11, distinct curves can be seen for each option, which suggest that the operating conditions are significant in terms of the raw water use. Model parameter uncertainty is covered in section 8.3.3.d.

The probability plots for the other criteria are presented and discussed in appendix I-2. As expected, there is a high probability that methane, CO₂ and TSP emissions will be increased relative to the base case, with the degree related to the volume of mine water taken. An important subsequent analysis would be to determine the significance of these emissions, given the extremely high volumes emitted from the power station as a whole, i.e. how much weight should be given to a relatively small increase attributed to the water plant. As discussed above, methane emissions are being used to indicate the increase in the manufacture and transport of ancillary materials. The significance of the impacts arising from these processes should be investigated, to enable the significance of the indicator to be judged (i.e. the weight that should be placed on methane emissions relative to the other criteria).

The compartment water case is the only option which is certain always to reduce sulphate emissions to surface water bodies. Whilst incorporating high volumes of pit water achieves the greatest reductions in sulphate emissions, it also has a small probability of causing the greatest increase in sulphate emissions. It can not be stated with greater than 70-82% confidence that incorporating pit water will decrease the waterborne sulphate emissions. This is because of the high volumes of effluent resulting from the use of pit water, and thus the pressure these place on the power station's effluent management system (i.e. they create a small probability that spills may occur).

The reductions in waterborne sulphate emissions are compared against their potential to increase the impacted land footprint. This is critical, as there is little value in merely shifting the point and duration of the emission. The impacted land footprint is calculated in a separate deterministic model, and thus its uncertainty is estimated differently to the other criteria (see section 8.2.5). The high uncertainty estimated for the impacted land footprints, and the fact that differences between the footprint sizes of the options are slight, result in no significant differences between the systems able to be distinguished. Nonetheless the analysis shows that using mine water always increases the potential for leachate formation. The increase is most pronounced for high volumes of pit water taken, although with the high uncertainty estimates, this is only shown with a 65% probability of occurring. Whilst the degree of certainty is relatively low, there is a small probability that this increase could be extremely large.

8.3.3.c Box and Whisker Plots

An alternate representation of the uncertainty samples, that of a box and whisker plot, is shown for raw water usage in Figure 8-12. Plots for the other criteria are given in appendix I-3. These plots enable the empirical uncertainty within each option to be visualised (the extent of the box and whiskers for a

chosen confidence level), as well as displaying the degree of overlap between the options. Box and whisker plots are a sort of “uncertainty enhanced” bar chart, and allow a comparison of several options and criteria at once, although many charts result for a large number of criteria and options.

Figure 8-12 and the graphs in appendix I-3 are drawn on the same basis as the probability plots, i.e. the percentage change in the water plant and effluent system relative to the case where no mine water is taken. As with the probability plots, the degree to which the sample falls in the positive range suggests how confident the decision maker can be for that selection criterion. On the box and whisker plot this can be judged by eye, with it easily able to be seen whether the full range spanned by either the box, or the whiskers as well, falls above the zero line. The confidence ranges displayed as the extremes of the box and whiskers relate to the degree of confidence that the decision makers are comfortable with. This need not be the same for all criteria, and will depend to some degree on the perceived significance of that criterion. The 50% and 90% confidence intervals are displayed by the box and whiskers, respectively, in Figure 8-12 and the figures in appendix I-3.




Figure 8-12 Box and whisker plots of raw water use for the various mine water extreme cases considered. The boxes span the 50% confidence range, and the whiskers the 90% confidence range.

The same trends as those discussed above and in appendix I-2 for the probability plots are shown by the box and whisker plots. Although they give little information on the shape of the probability sample, in some ways the box and whisker representation allows a clearer interpretation of the probabilistic output sample. This is particularly true of criteria with highly skewed distributions, or those with a large number of outliers. By plotting only the confidence range of interest, the box and whisker plots focus the decision only on the relevant portion of the results, resulting in a clearer picture of the trends. An advantage of the box and whisker representation over the cumulative probability plots is that they clearly display the degree of overlap between the options, and are thus a better representation of model parameter uncertainty, i.e. the degree to which the range spanned by the options expands the empirical range for each option can be easily assessed. Model parameter uncertainty is discussed in the following section.

8.3.3.d Model Parameter Analysis

A feature of operational decisions is that they require the consideration of fewer model parameters (the choice of model parameters is limited to those directly affected by the decision in hand). This allows a systematic parametric analysis to be conducted, without creating an unreasonably large set of parameter combinations. The fact that this study could be limited to a consideration of only three model parameters, allowed such an analysis to be conducted. The trends and ranges caused by the choice of model parameters for the selected criteria are presented and discussed in appendix I-4 (Figures 8-3 to 8 show the graphs for raw water use and waterborne sulphates).

Whilst the trends are of importance in that they direct the choice of model parameters to achieve the desired outcome, of equal importance to the trend is the degree to which a change in the model parameter affects the performance of the system, i.e. the choice of a parameter which causes only a small change in system performance, even when varied over its entire feasible operating range, does not introduce significant model uncertainty. Whether or not a change in system performance can be judged to be significant or not needs to be interpreted in relation to the extent of the emission or impact. For example, over the possible range of membrane plant capacities, an increase in CO₂ emissions from a 4.5% to a 10% change relative to the base case scenario appears fairly significant (see Figure I-7). However, the very high background CO₂ emissions means that this 5.5% range in performance from the water plant is only a 0.03% range for the mine/power station system as a whole (considering the information in Table 8-3). The choice of membrane plant capacity therefore does not introduce significant model parameter uncertainty into the prediction of CO₂ emissions.

The importance attached of the range in performance caused by model parameter uncertainty also needs to be assessed relative to the empirical uncertainty. If the model parameter varied to its extremes does not notably increase the range in performance determined from a consideration of empirical uncertainty, the choice of parameter can be viewed as not significant for that particular criterion. This can be seen in Figures I-6 and 7, where the very high range in performance arising from uncertainty in the empirical parameters, as given by the error bars in the figures, dominates that arising from the model parameters (as given by the space between the curves, and the range covered by each curve). This can perhaps be more clearly seen in the box and whisker plots in Figure 8-12 and appendix I-3. Here the empirical and model parameter uncertainty are represented by the spread of the box and whiskers (i.e. the range spanned within each option), and the range spanned by the options, respectively. A high degree of overlap of the options indicates a dominance of empirical parameter uncertainty (e.g. as for CO₂ emissions in Figure I-1). The overall uncertainty range (due to both model and empirical parameters) can be read as the overall top and bottom bounds of the whiskers (e.g. in Figure 8-12, this can be read as ranging from -3% to 5% for raw water use).

Both the box and whisker and the cumulative probability plots show that raw water use has the most significant model parameter uncertainty between the options, whilst for CO₂ emissions, this is practically non-existent. Methane and TSP emissions demonstrate fair model parameter uncertainty. The emissions resulting from the background system (i.e. the provision of energy and chemicals) exhibit high empirical uncertainty. This is attributed to the fairly high uncertainty specified for the operation of the water plant when high volumes of mine water are incorporated (e.g. increased pumping requirements and chemical dosage rates), as well as the high uncertainty placed in the aggregated non-South African background LCI data. For these emissions, the high empirical uncertainty means that a high degree of meaning can not be attributed to the relatively slight differences between the options, especially in light of the fact that they contribute so little to the total emissions from the wider mine/power station system. Similarly, the high empirical uncertainty assumed for the footprint overshadows the effect of the model parameters, although a general downward shift is evident as the volume of mine water incorporated is increased.

Model parameter uncertainty in sulphate emissions is significant, not so much because of the additional range it introduces (which is actually quite small because of the high empirical parameter uncertainty), but because of the shift it produces from the fully confident range (compartment water) to a low probability of actually increasing sulphate emissions (see Figure I-2). For all criteria, other than water

use, the refurbished and non-refurbished 2 Ml/day options plot extremely close, demonstrating that refurbishing the membrane plant introduces surprisingly little model parameter uncertainty. However, the fact that it shifts the raw water use into the 100% confidence range is significant, as it suggests that for full confidence, some degree of refurbishment would be required.

From an analysis of the trends in the model parameter analysis, it can be concluded that the box and whisker plots represent the full uncertainty range (both empirical and model) for the inclusion of mine water into the cooling water circuit (with the exception of TSP emissions, which show a very slight increase over the range predicted for the extreme scenarios). The uncertainty ranges predicted are only for the case where all the mine water is used as cooling water make-up. Utilising a portion of the mine water for demin production causes more extreme effects in all criteria considered, and thus extended ranges would be seen in the box and whisker plots. The significant increases in the negative impacts of the system (power and chemical consumption, footprint), with little or no benefits in the positive (water use and waterborne sulphate reduction) mean that utilising mine water for demineralised water production is unlikely to be considered (see discussion in appendix I.4.2).

8.4. CONCLUSIONS

From the analysis above and those in the appendices, it can be concluded that incorporating mine water does not always have the positive benefits one would expect. This is particularly true for the very poor quality pit water, which if the full range of uncertainty is included, has a low probability of both increasing the raw water consumption and salinity emissions to surface and underground water bodies (the exact opposite of the desired effect). However, the maximum compartment water can be incorporated with confidence, as this was shown to always decrease raw water consumption and salinity emissions to surface waters, whilst only marginally increasing the leachate potential from the ash dump. However, the use of compartment water does increase the energy and chemical requirements of the water plant, which translates to an increase in potential global and regional impacts (see Figure 8-9). A trade-off between these increases and the water and sulphate benefits thus has to be found, which will depend on the significance placed on each of the indicators. The model parameter analysis shows that these benefits will still be achieved for the full range of membrane plant capacities and volumes of water.

For certainty to be achieved that the maximum water-related benefit can be attained for the 2 MI/day pit water case, some degree of refurbishment of the membrane plant is required. As with the compartment water case this system shows a relatively small increase in leachate potential from the ash, and increases in energy and material related emissions. To obtain certainty that sulphate emissions to surface waters may not actually be increased due to spills from the ashing system, some amendments to the ashing system may be required (e.g. additional spare dam capacity). This could also be achieved by mixing a higher percentage of effluent with the ash, although this then trades the potential for surface emissions with that of leachate emissions. The effect of increased moisture content in the ash has been shown to be potentially extremely significant for the size of the impacted land footprint, so this would not be recommended.

The incorporation of higher volumes of very poor quality mine water would appear not be worth the risks. This carries a substantially higher potential of leachate from the ash dump, as well as increased risks of spills from the power station effluent management system. Thus, although associated with high median sulphate reductions, the high variability in the system could result in these systems actually increasing sulphate emissions through spills and leachate. These systems also show the highest increase in energy and material consumption, which would appear not to be worth the comparatively slim chance of raw water savings or reductions in sulphate emissions.

The results of the study have been significantly enhanced by investigating the options as probabilistic ranges instead of mid-points. A particular benefit of the uncertainty analysis is its ability to recognise the risks of incorporating the high volumes of pit water. These would be missed if only a mid point analysis were conducted, which would only reflect the benefits of high waterborne sulphate removal, and not the possibility of spills from the ashing system. Also key to warning against the use of high volumes of pit water is the inclusion of the impacted land footprint indicator. This picks up the potential of significant increases in leachate from the ash dump, which would otherwise be missed from the assessment.

The enhanced interpretation of the system that can be achieved through PCA is again demonstrated by this case study. In particular, the ability of the PC plot to portray a continuous decision space, incorporating both empirical and model parameter uncertainty, against a number of criteria is notable. However, some limitations of the assessment were also encountered; that of their inability to represent the significance of the changes in the criteria with respect to each other. The three representations of

uncertainty used in the study complement each other, and are each able to enhance different aspects of the uncertainty of the system, i.e. the PC plots clearly present the trade-offs involved and the “spread” of the operating space, whilst probability plots clearly display the level of confidence able to be held in each criteria for each option. Box and whisker plots are a good representation of model parameter uncertainty, able to show the degree of shifting between the options, and the full range over which the systems act.

This case study demonstrates the markedly different approach and methods suitable to the analysis of continuous operational type decisions, as opposed to tactical type decisions. These differences are particularly notable for the type and quantity of the data available to characterise the system models, and the levels of foreground/background detail required. In addition, in operational type decisions, the few model parameters allow a definite and contained operating space to be defined, whereas tactical type decisions require the investigation of a number of possible scenarios spanning a large and undefined operating space. Model parameter uncertainty thus plays a dominant role in tactical studies, whilst for operational studies, although required to define the operating space, model parameters tend to contribute less to the overall uncertainty in the system, i.e. they are generally dominated by empirical parameter uncertainty, even though this empirical uncertainty is typically much lower than that present in tactical decision systems. This is clearly shown in this case study, where the choice of model parameter generally caused a small shift in the operating space of the option relative to the range spanned by its empirical uncertainty.

CHAPTER 9

CONCLUSIONS

This thesis has investigated the tools and methods to enhance environmental decision making in resource-based industries. This development has been informed by the definition of different decision contexts, recognising that the type and quantity of information governing the decision-system has a number of significant consequences for the decision-making process. The nature of primary industries raises two central themes of this thesis, i.e. the need for suitable indicators of their environmental impacts, and an ability to quantify the high uncertainty and variability in their performance.

9.1. SIGNIFICANT FINDINGS

This section draws together insights on the inter-linking themes mentioned above, acquired through an investigation of case studies in coal-fired power generation.

9.1.1. Modelling the Environmental Performance of Resource-Based Industries

Resource-based processes typically involve large-scale disruption of the environment, with imprecise processes and diffuse emissions. The modelling of their environmental performance therefore raises significant challenges. For prospective assessments (i.e. to predict the effect of process changes on future performance), monitoring data alone are typically not sufficient to assess the environmental performance, and models are required to provide the link between monitored data and process performance, and to predict the products and emissions from the process inputs. In the construction of such models, many disparate sources of data, available at different levels of aggregation, and over various time intervals, have to be brought together into a coherent assessment. A single point model, reconciling the different types of monitored data and reconciling monitored data with approximations filling in data gaps, as well as reconciling data from the different sub-processes, is infeasible at the level of data availability, and would create a false sense of accuracy. An “uncertain” definition of the system is therefore much more meaningful, in which variables are defined over ranges of values to cover inconsistencies and imbalances in the system.

In addition, mining processes have inherently high variability, which further supports resource-based industries being modelled as ranges of potential performance, rather than “typical” operations. The high natural variability between mining processes means that in most cases these processes can not be modelled as averages, as the range in performance between different operations is so high as to obscure any meaningful information in the system. The relevant process of interest therefore needs to be identified (e.g. particular mining method, ore grade, water management features etc.), and the variability within the particular process incorporated in the assessment.

9.1.2. Use of LCA for Environmental Decision Support

LCA is the structured approach to environmental decision making chosen to organise the large data sets into useful information able to support process-related decisions in resource-based industries. As explained in section 1.1, LCA is uniquely qualified as an environmental decision support tool where the overall goal of the decision-making process is sustainability. Another significant benefit of the LCA methodology is that it provides demarcated steps of problem structuring, enabling the compilation of comparative information sets. It is thus able to structure complex systems into a contained number of environmental indicators, against which the systems can be compared on a consistent basis.

However, whilst its problem structuring ability is viewed as a significant strength of LCA, it is also one of its most significant weaknesses. This is because it is not able to incorporate stakeholder considerations into the problem structuring, and forces a particular frame of reference on the decision. It is therefore clear that if LCA is to be able to be tailored to specific decision contexts, it needs to move away from its original conception as a prescriptive analytical tool to assess products, to application as a decision-making process to promote learning and support environmental decision making. This allows the interface between LCA and other environmental decision-support tools to be explored, and a more flexible and adaptable LCA tool to be developed.

9.1.2.a Inventory Models for Resource Based Processes

The need to model the variable and imprecise resource-based processes as ranges of probable performance rather than single mid-points is identified above. This requires the inventory model be developed at a disaggregated level, so that independently varying the inputs in a stochastic model does not result in meaningless output (see section 4.2.1). Mass-balance models using “splitter” functions to apportion material inputs into their respective output streams were found to be the most appropriate for

the level of assessment required and the type of data available. Emission factors were used for those emissions that could not be related directly to the input flow responsible for them. The requirement of LCI to include all process within the material life cycle results in the need to quantify the inputs and outputs from a wide array of processes (e.g. from reactors for which complex process models are available, to waste dumps which are approximated by simple run-off and ingress factors). The factor-based inventory models are built up at a consistent level of detail for all processes, although the data sources and models upon which the factors are based differ considerably between the processes in the life cycle. The source and accuracy of the factors are able to be reflected in the inventory models through the uncertainty assigned to the input parameters.

The large surface disruption that occurs with resource-based processes (e.g. from mining operations and solid waste disposal), results in these processes often having very significant impacts on the water catchment of the area. Coupled with their significant potential for diffuse water pollution, this means that these processes are capable of water-related impacts of far greater significance than merely the process water consumption typically reflected in an LCI. The impact of the system on the water catchment of the area is therefore included in the assessment by reflecting the change in the quantity and quality of surface run off and seepage to groundwater in the inventory. However, these are inherently approximate and variable quantities to determine, which again points to the need for an uncertain representation of the system.

9.1.2.b Impact Assessment and Resource-Based Industries

The LCIA impact categories in commonly used assessment methodologies were found to be deficient for a full assessment of primary industries, partly because of the level of inventory detail available, and partly because of LCIA's inherent limitation with respect to predicting site-specific impacts (see section 5.3). This is of particular consequence to resource-based industries, since a significant portion of their impacts are felt on a local scale. The LCIA impact categories are sufficient to assess the potential impact of the emissions to air, although those impacts acting on a local scale, and to a lesser extent a regional scale, are associated with high uncertainty. Impacts related to emissions to water are less well assessed. This results from a disparity between the level at which the emissions are reported in the inventory, and that for which the water-related equivalency factors are developed. The former is constrained by the monitoring detail available, which for the case study systems investigated in this thesis, was found to be at a high level of aggregation due to a reliance on water quality "indicators",

such as conductivity or TDS, and major salt concentrations (e.g. sulphates and chlorides). Whilst equivalency factors for emissions to water are typically only available for toxic metal species (as in the EI 99 method), and no account is taken of the link between a high dissolved salt content and the potential for elevated concentrations of metal species, or indeed that high salinity effluents themselves can constitute an environmental problem.

The extremely large dedicated waste deposits of resource-based processes are assessed particularly poorly by LCA. Often only the mass or volume of the waste is identified, and possibly a qualitative rating assigned (i.e. whether the waste is toxic or not). Mostly the waste is assumed inert (or the dump linings assumed infallible), and only the land area required for containment of the waste considered an issue. These simplifications are taken because of the difficulties of predicting time-dependent emissions from complex waste systems, even though waste dumps, including those of ostensibly non-toxic wastes, have the potential of significant low-grade pollution (i.e. salinity rather than toxicity problems). The above problem of the disparity between the level at which emissions to water are predicted, and that for which toxicological impact factors are available is relevant here as well, in that the complex nature of predicting leachate emissions from waste materials means that predicting individual metal species is likely to be beyond the scope of most studies. A method based on an easier to calculate “indicator” component is thus more likely to be workable. This stance is strengthened by the fact that the toxicological factors developed for metal species are of questionable accuracy. The “impacted land footprint” is demonstrated here as essential to evaluate the relative impact of high volume, low-grade waste deposits. Material stockpiles are similarly overlooked in their potential for groundwater contamination, and thus the “impacted land footprint” is also relevant for their assessment.

9.1.3. Decision Contexts of Resource-Based Industries

Recognition that different levels of information are available for different decision contexts, allows the definition of three different decision models, based on the flow of information to characterise the decision system (see Figure 2-2):

- Type I: Decision context in parallel to production system
- Type II: Production system within decision context
- Type III: Decision context within production system

The production system is the particular sphere of influence of the person generating the inventory information. In the Type I model, the decision-maker merely imports information from another system

(e.g. background LCI data), so has no influence over its construction. In the Type II model, the decision-system is wider than the production system, and contains elements external to the system generating the information, whilst in the Type III model, the decision-system falls wholly within the system generating the information. The Type II and III models are perhaps better interpreted as representing the extremes of a single decision type, i.e. decisions of this nature fall somewhere along the progression from strategic, to tactical, to operational decisions, with the Type II model representing those of a strategic or tactical nature, and the Type III model representing decisions of an operational nature.

Since information flow is the defining feature of the decision model types, the choice of decision model has significant implications on the quantity and quality of data required of the study, and consequently on the uncertainty present in the system. The decision context has particular implications with regards to how well defined the problem to be addressed is, and thus the relative importance of the various sources of uncertainty. At the very lowest level of information availability (e.g. strategic studies requiring the assessment of a large number of options for implementation a considerable distance into the future), only a qualitative assessment may be possible. This is especially likely to be the case if different levels of detail are available for the systems under comparison, where an inconsistent quantitative analysis is likely to give misleading results. Even if a quantitative inventory can be attempted, a full quantitative assessment of uncertainty is unlikely to be feasible at this level of assessment. A qualitative uncertainty assessment is probably more meaningful, and although simultaneously interpreting quantitative and qualitative data is much more difficult, the inclusion of qualitative representation of uncertainty or impacts, is certainly preferable to missing certain impacts, or over-estimating the significance of differences between highly uncertain systems (see section 7.2).

In addition to suggesting the magnitude and relative importance of model and empirical parameter uncertainty, the decision type has implications for the selection of the criteria used to evaluate the systems. As the decision moves more towards a strategic nature, the less likely it is that the location of the processes within a particular product or process system will be known. Thus undue emphasis on local considerations in the evaluation of the system is inappropriate. Local impacts can only be predicted on a generic level, and will consequently be associated with high uncertainty. Operational decisions, however, apply to existing processes or product systems, whilst tactical decisions apply to those systems further along the decision-making process, by which time the location of the processes

within the system are more likely to have been chosen (e.g. specific suppliers will have been chosen). Regional and local impacts therefore do play a role in the evaluation of these decisions. Operational decisions have the greatest potential for the LCA decision process to interface with site specific environmental analysis tools (such as ERA), since these processes have the greatest potential to furnish the considerable site-specific information required for such assessments.

9.1.4. Uncertainty Analysis

LCA is an inherently uncertain procedure in that it combines data sources of varying accuracy and representativeness, and employs subjective judgement in applying this data to future operating systems. Subjective judgements are also present in the definition of the systems, and in the modelling choices determining the accuracy and complexity of the inventory and impact models used. Nonetheless, LCA results are most often presented as single values, which in a comparative analysis, gives the often incorrect impression that one system is always better or worse than another system. The procedure developed in chapter 4 therefore offers a marked improvement in LCA's decision support capabilities, where stochastic models enable the degree of confidence in differences between comparative systems to be taken into account, and parametric analyses and/or sensitivity analyses allow an explicit investigation of the operating space.

Different types of uncertainty cannot all be meaningfully assessed by the same type of analysis. A comprehensive structure, able to incorporate an assessment of the various sources of uncertainty, is therefore required. The framework in Figure 4-1 is proposed to supply this structure, where empirical parameter uncertainty, model parameter uncertainty, and uncertainty in model form are investigated in a looped fashion. The innermost loop assesses empirical uncertainty in an iterative probabilistic analysis. Model parameter uncertainty is assessed next, by a parametric analysis, or by a combination of sensitivity analyses and a parametric analysis, if a large number of model parameters require consideration. The top-most layer is an assessment of model form, in which alternative model forms are investigated in a sensitivity analysis. A simulation approach using Latin Hypercube sampling of the uncertain input distributions is chosen as the most appropriate method for a probabilistic assessment of empirical parameter uncertainty, and a rank-order correlation analysis the chosen method to determine the relative uncertainty importance of the parameters input into the model. These are shown to fulfil their functions well, and are capable of delivering accurate results, within reasonable constraints of computing power and time (see sections 7.5 and 8.3).

An uncertainty analysis is a valuable tool to direct effort back into the inventory analysis to reach the desired level of confidence in the results, or, where this is not possible, to highlight where the desired level of confidence can not be achieved within the current framework of the study. The empirical parameter uncertainty analysis takes an iterative approach, in which the input distributions of parameters shown to have high uncertainty importance are refined until an acceptable variance in the output sample of each criterion is achieved. This allows a “quick and dirty” first round definition of input distributions, with the probability distributions of only those parameters subsequently shown to be important requiring a rigorous definition of uncertainty. Attaining a desired level of variance in the results will most probably require that data collection be revisited, i.e. more accurate or applicable data sourced for the study. If this is not possible given the particular context of the study, the limit variance is determined for the particular data available, i.e. after a few iterations, in which the distributions of the top contributing parameters to a particular criterion are defined as accurately as the data allows, the limit variance is reached, and subsequent iterations and refinement of parameters of lower uncertainty importance yield little, or no, reduction in the output variance.

This thesis aims to develop LCA as a decision-making process. The placement of the uncertainty analysis within this decision-making process is shown in Figure 9-1, where a condensed version of Figure 4-1 is overlain on the key steps of the decision-making process. The uncertainty analysis provides input into criteria setting by highlighting where a chosen selection criterion is inappropriate. This is an iterative process, where the initially comprehensive range of impacts considered, is subsequently refined by removing those criteria which are not helpful to the analysis, i.e. those criteria which have such high empirical uncertainty within the data quality constraints of the study, that no differences between them can be discerned with any degree of confidence, or where the differences between the systems are too slight to see significant differences even at low variances in the output.

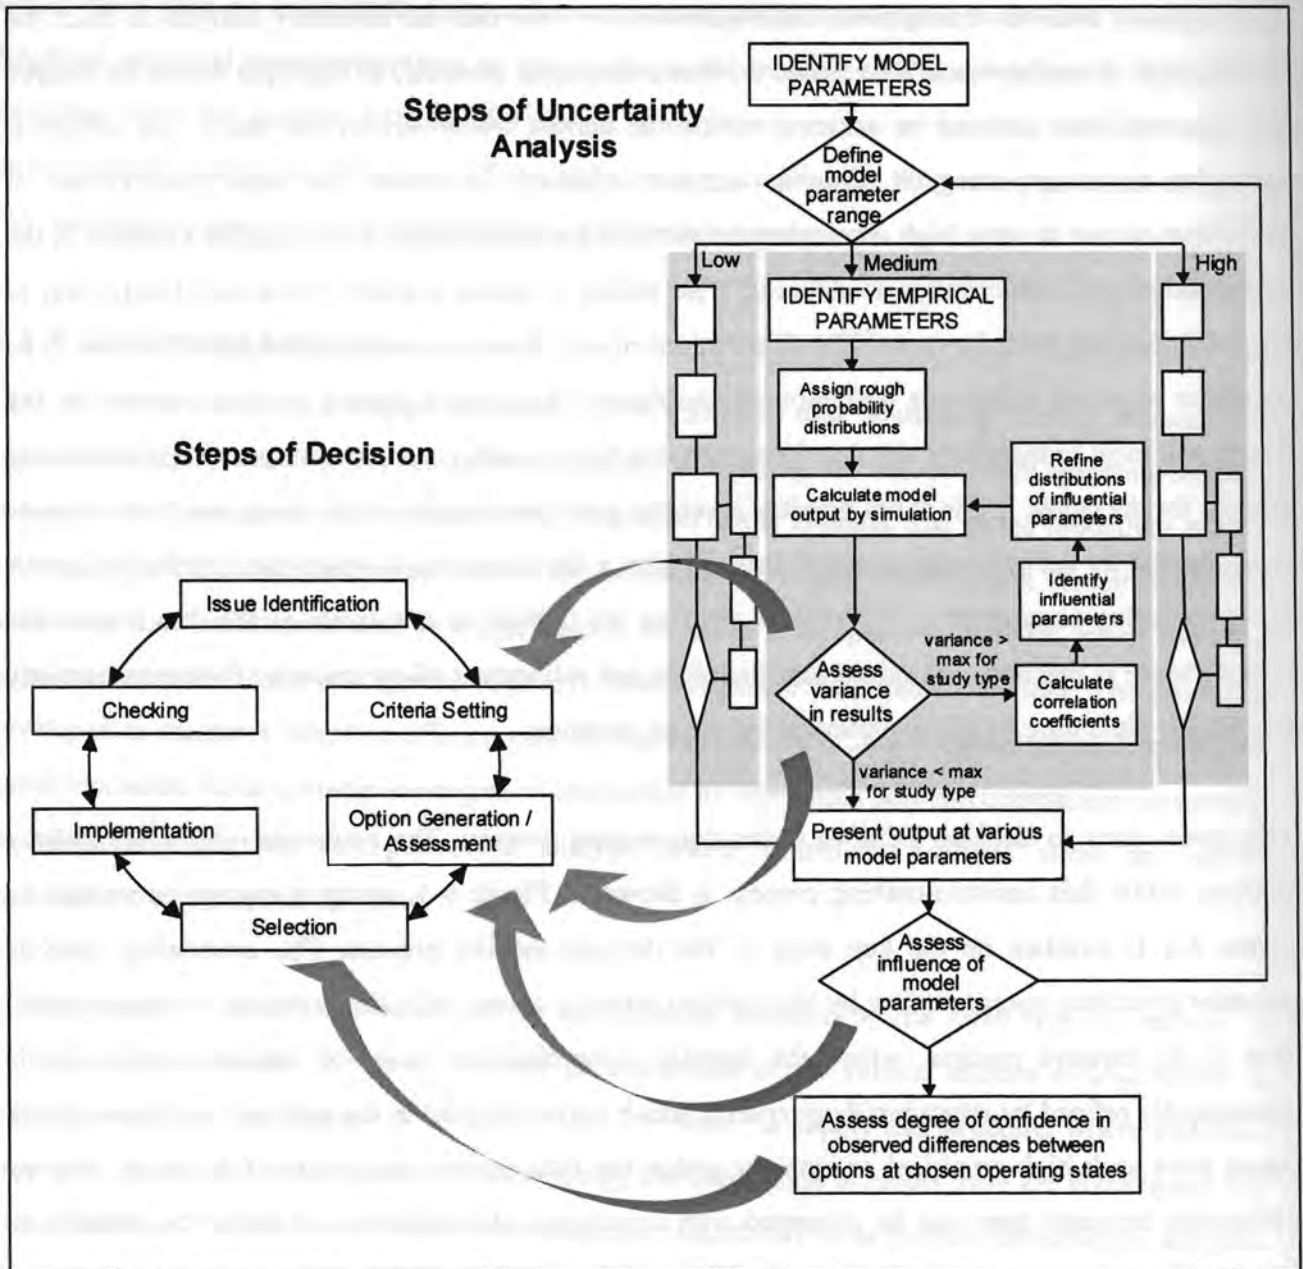


Figure 9-1 Placement of the uncertainty analysis in the context of the overall decision-making process.

The uncertainty analysis is also extremely useful at the stage of generating options to be considered in the decision. The systematic model parameter analysis forces a consideration of all decision variables, thereby allowing an exploration of the full solution space of the system. This allows an informed selection of the best operating states (appropriate combination of operating parameters) to be taken further in the decision. The iterative empirical parameter assessment guides the definition of the systems by highlighting where a tighter definition of the system is required, i.e. by highlighting those

parameters contributing significantly to the variance in the output with high geographical, temporal or technological variability (where variability is defined distinct from inherent variability or randomness, as in Figure 3-1).

The uncertainty analysis considerably enhances the selection of the preferred option, as it enables an estimation of the probability that a perceived difference between the options will always occur (i.e. an estimate of the degree of confidence that can be held that the one option performs worse or better than the other for that particular criterion). There is considerable feedback between the different steps of the decision-making process, and thus between the various aspects of the uncertainty analysis. For example, the options under consideration are refined until significant differences between them can be seen (i.e. their output variance reduced and/or their operating parameters adjusted until the required degree of confidence is obtained), or if this is not possible within the constraints of the study, the criteria are refined until a meaningful set is obtained for the options under consideration.

Less emphasis is placed on an assessment of uncertainty in model form in this thesis, as to a large degree this is constrained by the methodology of LCA, i.e. the standardised methodology of LCA forces particular model constructs. Working within the current structure of LCA it is possible to evaluate particular methodological choices by determining the output for different system choices (e.g. choice of allocation method). However, far greater inherent model uncertainties remain (e.g. uncertainties arising from aggregating over process life-times and across locations). This again returns to the dilemma of defining LCA as an analytical tool or decision-making process, where it needs to be recognised that taking advantage of LCA's structured approach forces a particular framing of the decision, and thus a certain degree of inherent model uncertainty.

9.1.4.a The Influence of the Decision Context on the Uncertainty Analysis

The implication of the choice of decision model type on the management of uncertainty is particularly marked for the consideration of model parameter uncertainty, which relates to the definition of the system. The more "strategic" the decision, the greater the number of elements outside the decision-maker's sphere of influence, and consequently the greater the number of parameters that can only be addressed through scenario analysis. The model parameter analysis therefore plays a dominant role in these studies, and the emphasis of the uncertainty analysis is on exploring the possible operating states of the system. Operational decisions, on the other hand, are defined as those in which the decision

context falls wholly within the system of interest. Thus external factors out of the control of the decision makers do not feature in a truly operational problem, and a consideration of model parameter uncertainty can be limited to those parameters directly affected by the decision in hand.

For the Type I model it is necessary to distinguish between the perspective of the analyst generating the information, and the analyst using the information, since the defining feature of the Type I model is that these are not connected. For the former (i.e. from the perspective of the person generating the information), there is essentially no model parameter uncertainty since the system merely describes an existing operation (i.e. the so-called "historical" inventory). However, from the perspective of the user of the information, there is typically extremely high model parameter uncertainty, since generic "background" LCIs, as available in LCA databases, are typically poorly described, and therefore likely to be inappropriately applied. Choosing to model a process with a generic LCI implicitly suggests a compromise between ease of modelling and coping with high uncertainty, since high uncertainty is inherent to using average data rather than the applicable marginal data. An assessment of model parameter uncertainty is not possible for "background" LCIs as currently available in LCA databases, as these are presented as aggregated systems. In fact, choosing to model with average, aggregated data, can be seen as essentially removing model parameter uncertainty, replacing it with the high empirical uncertainty of applying average and generic data.

The decision type also has significant implications for empirical uncertainty assessment, particularly for determining the number of internal iterations in the empirical analysis, i.e. on the level of variance in the output judged to be "acceptable". As the decision moves from an operational to a strategic level, so the time horizon of the decision increases. Thus, to support decisions on a strategic level, it is necessary to model the performance of future operating systems. A certain amount of inherent uncertainty is consequently involved in determining the values of the empirical parameters characterising these systems. The "acceptable" variance in the output is therefore typically high, and most likely limited by data uncertainty and availability, rather than by the scope of the study (set by the decision-makers). Operational decisions are associated with lower empirical parameter uncertainty, as parameters can predominantly be taken from data for the actual process of interest, and they do not require extrapolation to represent future conditions since operational decisions cover relatively short time-frames. The degree of confidence in these studies is therefore more likely to be set by the decision-

makers (by matching the degree of certainty required to the level of risk able to be tolerated, and to the resources able to be spent on data collection and evaluation), than by the inherent nature of the data.

For the Type I models it is again necessary to distinguish between the perspective of the analyst generating the information, and the analyst applying the inventory. From the perspective of the analyst generating the information, these systems potentially have low empirical parameter uncertainty, since they describe existing systems. Equally though, they may have high empirical uncertainty if they represent highly averaged processes (e.g. average production over a number of technologies will have high variability) or if many simplifications are taken (e.g. non-specific data is used, even though applicable process data is available). From the perspective of the user of the Type I models, these typically have high uncertainty, especially since they are most often available as poorly documented single-point inventories. This uncertainty will only be able to be reduced if “background” LCIs, as available in LCA databases, are better documented and reported with uncertainty information (most usefully as ranges of probable output, rather than single mid-points).

9.1.4.b Presenting and Analysing Uncertain Results

Principal Component Analysis (PCA) was found to provide useful information on the underlying structure of the result sample, particularly with respect to the strength and independence of the criteria chosen to evaluate the systems (see sections 7.5.2 and 8.3.2). In addition, the PC plots enable a powerful “graphical summary” of the results, in which the full output sample can be plotted on a single graph. However, to fully interpret the output samples, two commonly used representations of probabilistic samples were also used, namely “box and whisker” plots and cumulative probability plots. The three representations of uncertainty complement each other, as each are able to enhance different aspects of the results (see section 8.3.3).

The PC plots present clearly any trade-offs that have to be made between the criteria, and the “spread” of the operating space (due to both empirical and model parameter uncertainty), whilst cumulative probability plots display the level of confidence able to be held in each criterion for each option. Box and whisker plots are good at representing the relative importance of model and empirical parameter uncertainty, and show the degree of shifting between the options, and the full range over which the systems act. Whilst the probability and box and whisker plots are useful representations where a quantitative estimate of the relative uncertainty between systems is required (i.e. the degree of overlap

or confidence), they become extremely tedious (many pair-wise combinations) and difficult to interpret when a large number of options are involved. In this case, the PC plot is invaluable, as it enables the full data set to be displayed on a single plot, provided a sufficiently high percentage of the overall variance is displayed by the first two components. This all-encompassing view makes them particularly useful for a preliminary screening of options, enabling the selection of a smaller number of options for further analysis.

9.1.4.c Characterising Input Uncertainty

The most challenging step of the uncertainty analysis process is the definition of the probability distributions for the input parameters. Although the uncertainty matrix developed by Weidema and Wesnæs (1996) provides a structure for the estimate of variance for each parameter, the value assigned is unavoidably subjective. Subjectivity is involved in the choice of DQI score, the extrapolation of this score to an estimate of the variance it introduces, and the choice of a relevant distribution shape. The greatest subjectivity is introduced in the extrapolation of score to variance, as the guidelines given by Weidema and Wesnæs (1996) for scoring the data quality minimise the former, whilst the fact that different distribution shape defaults can be applied to different classes of parameters according to their general characteristics, minimises the latter.

The definition of default estimates of variance corresponding to a particular DQI score does not appear feasible because of the wide range of potential variability within different types of parameters. The choice thus has to be based on a consideration of the particular parameter in hand, informed where possible from the range demonstrated for similar parameters, or for the same parameter in a related system. It is therefore possible to draw on the experiences of other studies, and as uncertainty analyses hopefully become more commonplace in LCAs, it may be possible to develop default scores and corresponding CVs for particular parameter types and application areas. However, the very large number of parameters in an LCA study, and the very wide variety of application areas of LCA, mean that this is unlikely to be possible for all parameters. Thus there will inevitably always be some degree of subjectivity in the method. Nonetheless, there is still considerable value in the ability of the method to communicate the subjective judgement of data quality in a comprehensive way.

The degree of subjectivity depends to a large degree on the type of decision being supported, with strategic decisions requiring the greatest degree of subjective judgements. The empirical parameter

uncertainty of operational-type decision systems is simpler to characterise than strategic/tactical-type systems, because the decision is applied over a short time scale to an already operational system. The study is thus able to utilise existing process data to characterise most of the parameters, and so avoids the difficulties of subjectively estimating the uncertainties involved in extrapolating data from related processes, or to future operating conditions. Where the process data is available over a suitably long time interval, the subjectivity is minimised further, as the shape of the distribution can be obtained from the data sample, and the variance of the sample can be assumed to incorporate all potential sources of variability.

9.1.4.d Inventory versus Impact Uncertainty

The LCIA models are associated with high uncertainty, particularly those predicting local toxicological effects. In these categories, the additional uncertainty from the equivalency factors overshadows the reducing effect of aggregating the inventory data, and a significant increase in the variance of the impact score can be seen over that of the inventory data giving rise to it. This implies that for systems with already high inventory uncertainty, extending the inventory data to a consideration of impacts yields such high variance in the results, that no significant differences between the systems will be discernable. Thus for studies supporting decisions of a more strategic nature, an assessment based on inventory-level criteria is more appropriate for those emissions potentially giving rise to impacts acting on a local scale. This further supports an earlier conclusion that a consideration of local impacts is inappropriate for strategic decision systems, as their geographical location is not yet known, and their prediction is associated with irreducible uncertainty. However, the use of inventory-level criteria will result in a trade-off between lower uncertainty in the selection criterion, but higher valuation uncertainties.

In a comparative LCA, the impact assessment uncertainty is common to all options. The considerable increase from inventory to impact uncertainty is thus less noteworthy where the results are based on a ratio or percentage difference between the options, since a significant portion of the impact assessment uncertainty cancels out between the options. Where the options share a number of identical sub-processes, basing the analysis on a ratio is similarly required to remove correlations in the output. Thus where options have a considerable number of common elements, it is more meaningful to base the analysis on a ratio or percentage difference of the output with the output of a suitably selected base case, since the misleading effects of correlation of uncertainty among the options can be removed. A

significant constraint is that this is only possible if the uncertainty samples of the options are generated from identical random samples in the simulation analysis, i.e. it is not possible in comparing options generated from different studies.

9.1.5. Key Features of the Case Studies

The comparison of the inventory developed for South African power with an average European and UK grid mix highlights the features of South Africa's electricity industry, notably its reliance on coal, and its commitment to water-efficient power generation (see section 6.4.1). The dominance of coal in South Africa's energy mix shows in its relatively high contribution to acidification and global warming, since the European and UK mixes against which it is compared have lower fossil fuel components in their generation mix. However, on a power station level, the South African power stations compare favourably with their European counterparts, with marginally lower contributions to global warming, and only marginally higher contributions to acidification than a European plant with FGD. The former is attributed to the lack of coal transport in the South African stations (they are all fed by "mine-mouth" collieries), whilst the latter occurs because of the relatively low sulphur content of South African coals. The emphasis on water management at the South African power stations is clearly visible, with markedly lower water consumption figures and lower emissions of waterborne salts than their European counterparts. The comparison is hampered by the lack of documentation on the LCIs in the reference source of the data, rendering it less valid for certain environmental interventions, as it is not certain how these have been defined, or whether the category is complete (e.g. land transformation, coal reserves, TDS etc.)

The preliminary screening study of technologies to reduce sulphate emissions from coal-fired power stations showed that quantitative LCA alone cannot reliably inform decision making in information deficient systems. At low levels of detail it can potentially give misleading results, especially if an influential aspect of the systems is omitted, or if one system is more complete than the others. In addition, the inability of LCIA to assess the risk of site-specific impacts resulted in misleading comparisons of technologies in which waste management impacts were important. A combination of a quantitative and qualitative approach is therefore recommended for technology scanning studies of resource-based industries, where a qualitative assessment is backed up by whatever quantitative information is available. Such a combined assessment pointed to the potential of fluidised bed

combustion with in-situ desulphurisation as a technology for reducing sulphate emissions (see section 7.2.2).

The performance of fluidised bed combustion relative to conventional PF combustion was chosen for rigorous assessment in a subsequent case study, where the combustion of discard coal in a fluidised bed boiler was shown to be the environmentally preferable option for re-powering an in-storage power plant. However, for certain impact categories (most notably winter smog), this was only shown to be the case for certain plant configurations, where the degree of desulphurisation in the FBC boiler was shown to be a particularly important parameter. The determination of the FBC system as the preferred technology was strongly dependent on having an adequate indicator of the potential impact of the solid waste produced by the two systems, since the FBC system essentially results in the replacement of one type of solid waste dump with another. The impacted land footprint, which provides a relative indicator of the potential for groundwater contamination, was thus shown to be essential for an assessment of these technologies.

The FBC system's superior performance in fossil fuel resource consumption and impacted land footprint was predicted with high probability. However, its degree of improvement in the other impact categories was less certain, and largely dependent on the manner in which mining burdens were allocated to the FBC system. This is because in these categories, the benefits of the FBC system relative to the PF system are predominantly a result of the discard being defined as a dependent co-product (i.e. the fact that the discard does not bring any mining burdens into the FBC system, as the coal does into the PF system). Also critical to the evaluation of these systems is reflecting the "avoided" burdens of removing or reducing the volume of discard coal requiring disposal or stockpiling (depending on whether discard is regarded as a waste or a future resource). The allocation of mining burdens was thus shown to be important when considering different fuel sources. This is especially significant for South African power stations, where there is an increasing trend to burning near discard-quality coal in the stations, thereby "freeing-up" the higher quality reserves for higher revenue markets. The poorer quality coals show higher generation impacts, thus the benefit of choosing these coals rests on their being associated with lower mining burdens. The accuracy of the mining data and the relevance of the allocation method are therefore crucial to determining the environmentally preferable coal source. The need for higher quality mining data, and more detailed mining models to accurately model the consequences of varying the product volumes of the dual-product mines, are therefore identified.

The operational case study shows that appropriate conclusions can not be drawn around decisions in coal-fired power generation without an assessment of the site-specific impacts of effluent and waste management. The case study also shows that it is essential the high uncertainty and variability in these processes be taken into account, as incorrect decisions can be made if the analysis is based on a mid-point analysis (in the case study, the option showing the greatest potential for environmental improvement from a consideration of the mean values, also shows the greatest potential for negative impacts, albeit at a lower probability).

9.2. THE WAY FORWARD

Limitations to the use of LCA in decision-making is its prescriptive approach, as well as the perceived arbitrariness of its results. The former is being addressed through a greater emphasis on LCA's application as a decision-making process rather than an analytical tool. However, greater demonstration is required of the use of LCA as a decision-support process, where the emphasis is on the exploration of the decision space, and the generation of meaningful criteria and options, rather than generating a single environmental profile of questionable relevance. Greater meaning can be attached to LCA results by including an explicit assessment of the effect of the underlying choices, and the quality of the data from which they are generated. Including a representation of empirical uncertainty through stochastic modelling is a well developed technique. However, characterising the uncertain input parameters is an inherently subjective process, which will only be able to achieve greater acceptance through a greater demonstration of the method. Through increased experience, and an opportunity to learn and "borrow" from data ranges in similar studies, the level of effort and subjectivity in assigning data quality "scores" and the corresponding estimates of variance, will be able to be reduced.

Including an analysis of uncertainty undeniably increases the time and effort required of an environmental assessment. However, a far more robust and defensible assessment results, and insights into the systems are obtained that would not be possible with a mid-point analysis. Ways to minimise the considerable extra effort required therefore need to be devised. It is suggested that an uncertainty analysis become an integral part of the assessment, as it will always be a tedious and time-consuming process if viewed as an "add-on" to a mid-point study. This removes the need to determine average or "typical" parameter values, and the focus of data collection from the start should be on incorporating the full range of data values encountered. These can be refined if the parameter is subsequently found to

be of high uncertainty importance. Including an uncertainty analysis as an integral part of the assessment also forces inventory models to be developed at as low a level of aggregation as possible. This is required for meaningful parameter definition, as the uncertainty range can be based on the actual measured data and not estimated for an aggregated quantity, and to avoid meaningless output from randomly varied inputs in a simulation analysis. This also has the potential for vastly increasing the transparency and usefulness of inventory models, where these are available at a disaggregated level, rather than a single column of numbers.

It is important to note that the choice of an LCA approach forces a particular framing or evaluative paradigm on the assessment. The implicit judgements that this implies introduce an inherent amount of uncertainty over and above that which can be calculated from stochastic or parametric analyses. This study (and others addressing uncertainty in LCA) has primarily focussed on assessing the effect of empirical parameter uncertainty and the choice of decision variables. More research is required to assess the effect of model domain parameters and the choice of model form or structure. The ability to explore these within the framework of LCA is limited to the familiar methodological choices of the LCA method (e.g. impact assessment method), and does not allow an assessment of the potentially far more significant “paradigmatic” uncertainty. This is significantly more challenging as it requires exploring completely different perspectives from which to view the problem, and is only relevant where LCA is defined as a process rather than a tool, since it will likely require significant adaptation of the LCA process.

In certain decision-making situations (e.g. those concerned with predominantly site-specific impacts), LCA is inherently associated with high uncertainty. To deliver results with lower uncertainty, LCA therefore needs to be used in a complementary manner to other assessment tools that may be more applicable to the particular decision context (e.g. site-specific assessment tools, such as EIA). Ways to facilitate and formalise the incorporation of such complementary assessments into LCA, and thus the formulation of LCA as a flexible decision-making process need to be investigated.

Limitations to the assessment of resource-based industries are their lack of data, and the fact that high variability between processes makes it difficult to infer data for one process from another. This is particularly true for mining processes. There is the need to move away from “reactive” environmental monitoring towards obtaining measures of environmental performance that can be linked to the

processes responsible for them, and thus to opportunities to reduce them. The fact that the emissions from mining processes typically vary considerably over the life of the operation further complicates their assessment, and requires that a time-dependent emission profile be developed. Further work is therefore required to better characterise mining processes. The case studies in coal-based power generation show the significance of the impacts arising from waste management in decision systems where a high importance is placed on water-related impacts. This is likely also to be the case for other mining and minerals processing operations, so similar models are required for other high volume mineral processing wastes.

The high uncertainty found in resource-based industries shows that, for many criteria, a mid-point analysis results in a misrepresentation of the system (i.e. implies a change between the systems, although the change is not statistically significant). This suggests that an uncertainty analysis is an essential component of an environmental assessment of resource-based industries, especially for those studies of a more tactical/strategic nature. Greater work therefore needs to go into characterising the uncertainty and variability of these processes.

This study has provided base-line information on coal-fired power generation in South Africa. Similar information needs to be determined for other potential generating options (e.g. gas and nuclear), so that environmental considerations can explicitly be taken into account in planning South Africa's future generating mix. However, it appears that coal-based power generation will continue to provide the bulk of South Africa's power for the foreseeable future, and the best that can be done is to minimise its impacts. The framework provided by this study can be used to evaluate future innovations in coal-based power generation, and to guide the design of these processes, as has been demonstrated for the FBC system. This study has looked at undelivered power. Work on including the impacts of transmission and distribution needs to be done, so that the relative siting of plants can be taken into account, and coal-fired electricity can meaningfully be compared with potential distributed generation options (e.g. solar panels).

The most significant contribution of this work is the development of an approach to support prospective environmental decision-making. A procedure to guide the assessment of loosely defined option sets is developed, along with an ability to incorporate the often high empirical data uncertainty of future systems. The outcome of the analysis is an explicit consideration of the full decision-space in which the

future systems could act. The powerful data analysis and presentation features of PCA provide an enhanced interpretation of the large information sets. In particular, they allow an exploration of the full operating space, clearly identifying the trade-offs between the various options. A stochastic modelling approach considerably increases the relevance of the LCA results, by allowing a degree of confidence to be attached to the perceived differences between the options. The iterative uncertainty analysis procedure developed facilitates the incorporation of empirical data uncertainty by using rank order correlation analyses to focus the analysis on those parameters with high uncertainty importance. The combined model parameter and empirical parameter uncertainty analysis therefore considerably enhances the decision-support capabilities of the LCA process.

This work also has significance for the environmental management of coal-based power generation. In particular, the need for a comprehensive assessment has been demonstrated, incorporating all aspects of power generation and fuel procurement. The case studies show that firm conclusions can not be drawn around technology choice either on the boiler plant or water plant without a consideration of the impacts of solid waste disposal. The quantitative evaluation of the leachate potential of the solid waste dumps, provided by the “impacted land footprint”, is therefore invaluable. Finally, the enhanced decision-support possible with the uncertainty analysis procedure outlined above is highly significant for environmental management decisions in coal-based power generation. The case studies show that incorrect conclusions could be drawn without a consideration of the high uncertainty and variability in these processes. The model parameter analysis allows the operating parameters to be determined for which the system has a high probability of meeting the desired objective. For example, the analysis allowed the particular combination of decision variables to be selected for which a definitive statement could be made that FBC burning discard can be considered the environmentally preferable option for re-powering an in-storage power station.

REFERENCES

- Albritton R. G., Derwent R. G., Isaksen I. S. A., Lal M., and Wuebbles D. J. (1995) Trace Gas Radiative Forcing Indices. In *Climate Change 1994*. Intergovernmental Panel on Climate Change, Cambridge University Press.
- Azapagic A. (1999) Review Article: Life cycle assessment and its application to process selection, design and optimisation. *Chemical Engineering Journal* **73**, 1-21.
- Azapagic A. and Clift R. (1994) Allocation of Environmental Burdens by Whole-System Modelling - the Use of Linear Programming. In *Allocation in LCA* (eds. G. Huppes and F Schneider), pp. 54-60. SETAC.
- Azapagic A. and Clift R. (1998) Linear Programming as a Tool in Life Cycle Assessment. *International Journal of Life Cycle Assessment* **3**(6), 305-316.
- Azapagic A. and Clift R. (1999a) Life Cycle Assessment as a tool for improving process performance: A case study on Boron Product. *International Journal of Life Cycle Assessment* **4**(3), 133-142.
- Azapagic A. and Clift R. (1999b) Allocation of Environmental Burdens in Co-Product Systems: Product Related Burdens (Part 1). *International Journal of Life Cycle Assessment* **4**(6), 357-369.
- Azapagic A. and Clift R. (2000) Allocation of Environmental Burdens in Co-Product Systems: Product Related Burdens (Part 2). *International Journal of Life Cycle Assessment* **5**(1), 31-36.
- Barnthouse L., Fava J., Humphreys K., Hunt R., Laibson L., Noesen S., Norris G., Owens J., Todd J., Vigon B., Weitz K., and Young J. (1998) Life-Cycle Impact Assessment: The State-of-the-Art. *Report of the SETAC Life-Cycle Assessment (LCA) Impact Assessment Workgroup, SETAC LCA Advisory Group*. SETAC.
- Basson L. and Petrie J. G. (2001) Multi Criteria Approaches to Valuation in Life Cycle Impact Assessment. *submitted to the International Journal of Life Cycle Assessment*.
- Baumann H. (1996) LCA Use in Swedish Industry. *International Journal of Life Cycle Assessment* **1**(3), 122-126.
- Baumann H. and Rydberg T. (1994) A Comparison of Three Methods for Impact Analysis and Evaluation. *Journal of Cleaner Production* **2**(1), 13-20.
- Baxter B. (1993) Cleaning up the act. *SA Mining, Coal, Gold and Base Minerals* (September 1993), 14-15.
- Bez J., Heyde M., and Goldhan G. (1998) Waste Treatment in Product Specific Life Cycle Inventories, An Approach of Material Related Modelling, Part II: Sanitary Landfill. *International Journal of Life Cycle Assessment* **3**(2), 100-105.
- Bredell J. H. (1987) South African Coal Resources Explained and Analysed. Geological Survey of South Africa.
- Carlson C. L. and Adriano D. C. (1993) Environmental Impact of Coal Combustion Residues. *Journal of Environmental Quality* **22**, 227-247.
- CHAINET (1999) Summary of Guidebook on Analytical Tools for Environmental Design and Management in a Systems Perspective (Draft). CML and IPTS.
- Chevalier J.-L. and Le Téo J.-F. (1996) Life Cycle Analysis with Ill-defined Data and its Application to Building Products. *International Journal of Life Cycle Assessment* **1**(2), 90-96.
- Clift R., Frischknecht R., Huppes G., and Weidema B. P. (1998) Towards a Coherent Approach to Life Cycle Inventory Analysis. SETAC-Europe.
- Consoli F., Allen D., Boustead I., Fava J., Franklin W., Jensen A. A., de Oude R., Parrish R., Perriman R., Postlethwaite D., Quay B., Seguin J., and Vigon B. (1993) *Guidelines for Life-Cycle Assessment: A "Code of Practice"*. Society of Environmental Toxicology and Chemistry (SETAC).

- '97 on Coal Science and Technology. The Fossil Fuel Foundation, Randburg, South Africa.
- Coulon R., Camobreco V., Teulon H., and Besnainou J. (1997) Data Quality and Uncertainty in LCI. *International Journal of Life Cycle Assessment* 2(3), 178-182.
- Cowell S. J. (2001) SETAC Working Group on LCA and Decision-Making. *Presentation at SETAC Europe Conference*. Madrid, May 2001.
- Cowell S. J., Hogan S., and Clift R. (1997) LCANET Theme Report: Positioning and Applications of LCA. European Network for Strategic Life-Cycle Assessment Research and Development.
- Crofts C. (1990) Emissions Reduction in the UK: Accommodating Waste Production From Sulphur Abatement Systems. *Energy and Environment* 1(1).
- CSIR (1991) The Situation of Waste Management and Pollution Control in South Africa (Report to the Department of Environmental Affairs). CSIR Programme for the Environment.
- De Jager F. S. J. (1983) Coal Resources of the Republic of South Africa - an Evaluation at the End of 1982. Geological Survey of South Africa.
- De Smet B. and Stalmans M. (1996) LCI Data and Data Quality: Thoughts and Considerations. *International Journal of Life Cycle Assessment* 1(2), 96-104.
- Department of Minerals and Energy (1998) White Paper on Energy Policy of the Republic of South Africa. Government Printer.
- Detz C. M. and Barvinchak (1979) Microbial Desulphurisation of Coal. *Mining Congress Journal* (July 1979), 75-84.
- DMEA (1985) South African Discard and Duff Coal - National Inventory 1985. Department of Mineral and Energy Affairs.
- Dingley C. (1995) An electricity grid for sub-equatorial Africa? *Africa* 3(2).
- Diwekar U. M. (1999) Process Design for Environmental Problems. In *Tools and Methods for Pollution Prevention* (eds. S. K. Sikdar and U. Diwekar), pp. 277-297. Kluwer Academic.
- Dubreuil A. and McGeer J. C. (1999) Do We Have Adequate Toxicity and Fate Models for Metals Emissions? *3rd APEC/GEMEED Environmental Co-operative Workshop on Sustainable Development of Mining Activities*. Cairns, Australia.
- Dutkiewicz R. K. and Gore B. (1998) Emission Comparisons for Eskom and Some International Utilities. *Journal of Energy in South Africa* 9(3), 79-83.
- Ekvall T. (1994) Principles for Allocation at Multi-Output Processes and Cascade Recycling. In *Allocation in LCA* (eds. G. Huppes and F Schneider), pp. 91-100. SETAC.
- EPRI (1983) Economic Evaluation of FGD Systems. Volume 1: Throwaway FGD Processes, High- and Low-Sulphur Coal.
- EPRI (1989) The Current State of Atmospheric Fluidized-Bed Combustion Technology. World Bank.
- Eskom (1999a) Annual Report 1998. ESKOM Corporate Communication Department.
- Eskom (1999b) Environmental Report 1998. ESKOM Corporate Communication Department.
- Fava J., Consoli F., Denison R., Dickson K., Mohin T., and Vigon B. (1993) A Conceptual Framework for Life-Cycle Impact Assessment. SETAC.
- Fava J., Denison R., Jones B., Curran M. A., Vigon B., Selke S., and Barnum J. (1991) *A Technical Framework for Life-Cycle Assessment*. SETAC foundation.
- Fava J., Jensen A. A., Lindfors L., Pomper S., De Smet B., Warren J., and Vigon B. (1994) Life-Cycle Assessment Data Quality: A Conceptual Framework. *Proceedings of a workshop in Wintergreen, VA, October 1992*. SETAC and the SETAC Foundation for Environmental Education.

- Finnveden G. (1992) Landfilling - a Forgotten part of Life Cycle Assessments. In *Product Life Cycle Assessment - Principles and Methodology*, Vol. Nord 1992:9, pp. 263-280. Nordic Council of Ministers.
- Finnveden G. (1994) Some Comments on the Allocation Problem and System Boundaries in LCA. In *Allocation in LCA*, (eds. G. Huppes and F Schneider), pp. 65-71. SETAC.
- Finnveden G. (1996) "Resources" and Related Impact Categories. In *Towards a Methodology for Life Cycle Impact Assessment* (ed. H. A. Udo de Haes), pp. 39-48. SETAC-Europe.
- Finnveden G. (1997) Valuation Methods Within LCA - Where are the Values? *International Journal of Life Cycle Assessment* 2(3), 163-169.
- Finnveden G., Albertsson A., Berendson J., Eriksson E., Hogland L., Karlsson S., and Sundqvist J. (1995) Solid waste treatment within the framework of life-cycle assessment. *Journal of Cleaner Production* 3(4), 189-199.
- Finnveden G. and Lindfors L. (1997) LCANET Theme Report: Life-Cycle Impact Assessment and Interpretation. European Network for Strategic Life-Cycle Assessment Research and Development.
- Finnveden G. and Lindfors L.-G. (1998) Data Quality of Life Cycle Inventory Data - Rules of Thumb. *International Journal of Life Cycle Assessment* 3(2), 65-66.
- Finnveden G. and Nielsen P. H. (1999) Long-Term Emissions from Landfills Should Not be Disregarded. *International Journal of Life Cycle Assessment* 4(3), 125-126.
- Frischknecht R. (1997) LCANET Theme Report: Goal and Scope Definition and Inventory Analysis. European Network for Strategic Life-Cycle Assessment Research and Development.
- Frischknecht R. (1998) Life Cycle Inventory Analysis for Decision-Making: Scope-Dependent Inventory System Models and Context-Specific Joint Product Allocation. Doctor of Technical Sciences, Swiss Federal Institute of Technology Zurich.
- Frischknecht R., Hofstetter P., Knoepfel I., Dones R., and Zollinger E. (1994) Ökoinventare für Energiesysteme. Schlußbericht, Laboratorium für Energiesysteme, ETH.
- Funtowicz S. O. and Ravetz J. R. (1990) *Uncertainty and Quality in Science for Policy*. Kluwer Academic Publishers.
- Goedkoop M. and Spriensma R. (1999) The Eco-Indicator 99: A Damage Oriented Method for Life Cycle Impact Assessment (Methodology Report). PRé Consultants B.V.
- Graedel T. E. (1998) *Streamlined Life-Cycle Assessment*. Bell Laboratories, Lucent Technologies, Prentice Hall.
- Graedel T. E. and Allenby B. R. (1995) *Industrial Ecology*. AT&T, Prentice Hall.
- Graedel T. E., Allenby B. R., and Comrie P. R. (1995) Matrix Approaches to Abridged Life Cycle Assessment. *Environmental Science and Technology* 29(3), 134A-139A.
- Grisel L., de Beaufort A., Wrisberg N., Coelho-Schwitz V., Glavind M., van Dam A., Carlson R., Schmidt A., Schmidt W.-P., van den Berg N., Meyer U., Bretz R., Pajula T., McKeown P., Hochfeld C., and Kreissig J. (1997) LCANET Theme Report: Databases and Softwares. European Network for Strategic Life-Cycle Assessment Research and Development.
- Grobbelaar C. J., Maake R., and Asamoah J. K. (1997) Coal in South Africa's energy economy and the Impacts of National and International Environmental Pressures. *Southern Africa Coal Conference '97*. Sandton, South Africa
- Guinée J. (1993) Data for the Normalisation Step within Life Cycle Assessment of Products. CML.
- Guinée J. (1994) Review of Classification and Characterisation Methodologies. In *Integrating Impact Assessment into LCA* (ed. H. A. Udo de Haes, A. A. Jensen, W. Klöpffer, and L. Lindfors), pp. 51-60. SETAC.

- Guinée J. and Heijungs R. (1995) A Proposal for the Definition of Resource Equivalency Factors for Use in Product Life-Cycle Assessment. *Environmental Toxicology and Chemistry* 14(5), 917-925.
- Guinée J., Heijungs R., Udo de Haes H. A., and Huppes G. (1993a) Quantitative life cycle assessment of products, 2: Classification, valuation and improvement analysis. *Journal of Cleaner Production*. 1(2), 81-91.
- Guinée J., Heijungs R., van Oers L., van de Meent D., Vermeire T., and Rikken M. (1996a) LCA Impact Assessment of Toxic Releases. Generic Modelling of Fate, Exposure and Effect for Ecosystems and Human Beings with Data for About 100 Chemicals. CML, RIVM.
- Guinée J., Heijungs R., van Oers L., Wegener Sleswijk A., van de Meent D., Vermeire T., and Rikken M. (1996b) Uniform System for the Evaluation of Substances (USES): Inclusion of Fate in LCA Characterisation of Toxic Releases. *International Journal of Life Cycle Assessment* 1(3), 133-138.
- Guinée J., Udo de Haes H. A., and Huppes G. (1993b) Quantitative life cycle assessment of products, 1: Goal definition and Inventory. *Journal of Cleaner Production*. 1(1), 3-13.
- Hansen Y. (2001) Assessing the Environmental Impact Associated with Solid Waste Management in the Coal Based Power Generation Industry. Ph.D., University of Sydney (in progress).
- Hansen Y., Notten P. J., and Petrie J. G. (2001a) The Environmental Impact of Ash Management in Coal-Based Power Generation. *submitted to Applied Geochemistry*.
- Hansen Y., Notten P. J., and Petrie J. G. (2001b) Incorporating the Environmental Impact of Ash Management into Life Cycle Assessment. *submitted to the Journal of the South African Institute of Mining and Metallurgy (SAIMM Journal)*.
- Hanssen O. J. and Asbjørnsen O. A. (1996) Statistical Properties of Emission Data in Life Cycle Assessments. *Journal of Cleaner Production* 4(3-4), 149-157.
- Heijungs R. (1994) A Generic Method for the Identification of Options for Cleaner Products. *Ecological Economics*(10), 69-81.
- Heijungs R. (1995) Harmonization of Methods for Impact Assessment. *Environ.Sci. & Pollut.Res.* 2(4), 217-224.
- Heijungs R. (1996) Identification of Key Issues for Further Investigation in Improving the Reliability of Life-cycle Assessments. *Journal of Cleaner Production* 4(3-4), 159-166.
- Heijungs R. and Guinée J. (1994) Using Environmental Multi-media Models in LCA: The Flux-Pulse Problem, Paper 18. CML.
- Heijungs R., Guinée J., and Huppes G. (1997) Impact Categories for Natural Resources and Land Use, Paper 138. CML.
- Heijungs R., Guinée J., Huppes G., Lankreijer R. M., Udo de Haes H. A., Wegener Sleswijk A., Ansems A. M. M., Eggels P. G., van Duin R., and de Goede H. P. (1992) *Environmental Life Cycle Assessment of Products - Guide and Background*. CML, TNO, B&G.
- Hertwich E. G. (1999) Toxic Equivalency: Accounting for Human Health in Life-Cycle Impact Assessment, Ph.D., University of California, Berkeley.
- Hertwich E. G. and Hammitt J. K. (2001) A Decision-Analytic Framework for Impact Assessment. Part 1: LCA and Decision Analysis. *International Journal of Life Cycle Assessment* 6(1), 5-12.
- Heule R. and van den Berg N. (1998) Spy Plots: A Method for Visualising the Structure of LCA Data Bases. *International Journal of Life Cycle Assessment* 3(2), 114-118.
- Hofstetter P. (1998) *Perspectives in Life Cycle Impact Assessment; A Structured Approach to Combine Models of the Technosphere, Ecosphere and Valuesphere*. Kluwer Academic Publishers.
- Holmburg J., Karl-Henrik R., and Eriksson K.-H. (1996) *Socio-ecological Principles for a Sustainable Society*. The Natural Step Foundation.

- Le Téno J.-F. (1999) Visual Data Analysis and Decision Support Methods for Non-Deterministic LCA. *International Journal of Life Cycle Assessment* 1(4), 41-47.
- Lennon S. J. (1997) Clean Coal Technology Choices Relating to the Future Supply and Demand of Electricity in Southern Africa. *Journal of Energy in South Africa* (May 1997), 45-51.
- Lindeijer E., Müller-Wenk R., Baitz M., Köllner T., Klöpffer W., Renner I., I Canals L. M., Ten Houten M., and Weidema B. P. (2000) SETAC Europe Taskforce on Resources and Land, Impact Assessment of Resources and Land Use, 4th draft, April 2000. *SETAC Globe* 1(5), 38-41.
- Lindfors L.-G., Christiansen K., Hoffmann L., Virtanen Y., Juntilla V., Hanssen O. J., Ronning A., Ekvall T., and Finnveden G. (1995) Nordic Guidelines on Life-Cycle Assessment. Nordic Council of Ministers.
- Lubkert B., Virtanen Y., Muhlberger M., Ingman J., Vallance B., and Alber S. (1991) An International Database for Ecoprofile Analysis (IDEA), WP-91-30. International Institute for Applied Systems Analysis (IIASA).
- Mackay D. and Seth R. (1999) The Role of Mass Balance Modelling in Impact Assessment and Pollution Prevention. In *Tools and Methods for Pollution Prevention* (eds. S. K. Sikdar and U. Diwekar), pp. 157-179. Kluwer Academic Publishers.
- Maurice B., Frischknecht R., Coelho-Schwartz V., and Hungerbühler K. (2000) Uncertainty Analysis in Life Cycle Inventory. Application to the Production of Electricity with French Coal Power Plants. *Journal of Cleaner Production* 8, 95-108.
- McKone T. E. (1993) CalTOX, A Multimedia Total Exposure Model for Hazardous-Waste Sites. Lawrence Livermore National Laboratory.
- Meier M. A. (1997) Eco-Efficiency Evaluation of Waste Gas Purification Systems in the Chemical Industry. In *LCA Documents*, Vol. 2 (ed. W. Klöpffer and O. Hutzinger), pp. 271. Eco-Infoma Press.
- Meier M. A. and Hungerbühler K. (2000) Uncertainty Analysis and Toxicity Classification in Life-cycle Assessment Using the Case-study of Gas Purification Systems. *International Journal of Life Cycle Assessment* 5(2), 75-77.
- Menke D., Davis G. A., and Vignon B. W. (1996) Evaluation of Life Cycle Assessment Tools. Environment Canada.
- Miettinen P. and Hamalainen R. P. (1997) How to Benefit from Decision Analysis in Environmental Life Cycle Assessment. *European Journal of Operational Research* 102, 279-294.
- Ministry of Public Enterprises. (2000) An Accelerated Agenda Towards the Restructuring of State Owned Enterprises: Policy Framework. Ministry of Public Enterprises Republic of South Africa.
- Morgan M. G. and Henrion M. (1990) *Uncertainty. A Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis*. Cambridge University Press.
- Müller-Wenk R. (1998) Depletion of Abiotic Resources Weighted on the Base of "Virtual" Impacts of Lower Grade Deposits in Future. Universität St. Gallen.
- Murtagh F. and Heck A. (1987) *Multivariate Data Analysis*. Reidel Publishing Company.
- Nash J. and Stoughton M. D. (1994) Learning to Live with Life Cycle Assessment. *Environmental Science and Technology*. 28(5), 236-237.
- Nielsen P. H., Exner S., Jorgensen A.-M., and Hauschild M. (1998) Product Specific Emissions from Municipal Solid Waste Landfills. Part II: Presentation and Verification of the Computer Tool LCA-LAND. *International Journal of Life Cycle Assessment* 4(3), 225-236.
- Nielsen P. H. and Hauschild M. (1998) Product Specific Emissions from Municipal Solid Waste Landfills. Part I: Landfill Model. *International Journal of Life Cycle Assessment* 3(3), 158-168.

- NRCan (1997) Sustainable Development of Minerals and Metals (prepared in connection with Canada's participation at the meeting of the United Nations Commission on Sustainable Development). Natural Resources Canada.
- Öko-Institut (1993-98) Environmental Manual for Power Development Version. 1.1. GTZ (Deutsche Gesellschaft für Technische Zusammenarbeit).
- Ott W. R. (1995) *Environmental Statistics and Data Analysis*. Lewis Publishers (CRC Press LLC).
- Owens J. W. (1996) LCA Impact Assessment Categories: Technical Feasibility and Accuracy. *International Journal of Life Cycle Assessment* 1(3), 151-158.
- Owens J. W. (1997a) Constraints on Moving from Inventory to Impact Assessment. *Journal of Industrial Ecology* 1(1), 37-49.
- Owens J. W. (1997b) Life-Cycle Assessment in Relation to Risk Assessment: An Evolving Perspective. *Risk Analysis* 17(3), 359-365.
- Pesonen H.-L., Ekvall T., Fleischer G., Huppel G., Jahn C., Klos Z. S., Rebitzer G., Sonnemann G. W., Tintinelli A., Weidema B. P., and Wenzel H. (2000) Framework for Scenario Development in LCA (SETAC-Europe LCA Working Group 'Scenario Development in LCA'). *International Journal of Life Cycle Assessment* 5(1), 21-30.
- Pesso C. (1993) Life Cycle Methods and Applications: Issues and Perspectives. *Journal of Cleaner Production*. 1(3-4), 139-142.
- Petersen J. (1998) Assessment and Modelling of Chromium Releases in Mineral Processing Waste Deposits. Ph.D., University of Cape Town.
- PIRA (1996) PEMS (PIRA Environmental Management Software) Version 4. PIRA International.
- Pistikopoulos E. N., Stefanis S. K., and Livingston A. G. (1994) A Methodology For Minimum Environmental Impact Analysis. *AIChE Symposium Series Pollution Prevention and Product Modifications* 90(43), 139-150.
- Potting J., Hauschild M., and Wenzel H. (1999) "Less is Better" and "Only Above Threshold": Two Incompatible Paradigms for Human Toxicity in Life Cycle Assessment. *International Journal of Life Cycle Assessment* 4(1), 16-24.
- Potting J., Schöpp W., Blok K., and Hauschild M. (1998) Site-Dependent Life-Cycle Impact Assessment of Acidification. *Journal of Industrial Ecology* 2(2), 63-87.
- Prevost X. M. (1998) Coal. In *South Africa's Mineral Industry 1997/1998*. Department of Minerals and Energy, Mineral Economics Directorate and Mineral Bureau.
- Raimondo J., Van Zyl H., and Bonner D. (1996) Environmental Cost Internalisation in the South African Coal Mining Industry. *A joint project of UNCTAD/UNEP - "Environmental Cost Internalisation in the Commodity Sector*. African Environmental Solutions.
- Rossiter A. P. (1994) Process Integration for Pollution Prevention. *AIChE Symposium Series* 90, 12-20.
- Russell C. (1991) Coal Mining: Custodian or Culprit. *Mining Survey* (1), 7-14.
- Sorin M. and Paris J. (1997) Combined Energy and Pinch Approach to Process Analysis. *Computers Chem Engng* 21(supplement), S23-S28.
- Shonnard D. R. and Hiew D. S. (1999) Environmental Assessment and Optimisation of Chemical Process Designs. *AIChE 1999 Annual Meeting Abstracts*, 207d.
- Spath P. L., Mann M. K., and Kerr D. R. (1999) Life Cycle Assessment of Coal-fired Power Production. National Renewable Energy Laboratory.

- Spengler T., Geldermann J., Hahre S., Sieverdingbeck A., and Rentz O. (1998) Development of a Multiple Criteria Based Decision Support System for Environmental Assessment of Recycling Measures in the Iron and Steel Making Industry. *Journal of Cleaner Production* 6, 37-52.
- StatSoft Inc. (1999) Electronic Statistics Textbook. StatSoft (<http://www.Statsoft.com/textbook/stathome.html>)
- Steen B. (1997) On Uncertainty and Sensitivity of LCA-based Priority Setting. *Journal of Cleaner Production* 5(4), 255-262.
- Stefanis S. K., Livingston A. G., and Pistikopoulos E. N. (1995) Minimising the Environmental Impact of Process Plants: A Process Systems Methodology. *Computers Chem Engng* 19(supplement), S39-S44.
- Stevens A. (1997) Moving Companies Towards Sustainability Through Eco-design: Conditions for Success. *Journal of Sustainable Product Design* (October 1997), 47-55.
- Stewart M. (1999) Environmental life cycle considerations for design-related decision making in minerals processing. Ph.D., University of Cape Town.
- Stewart M. and Petrie J. G. (1996) Life Cycle Assessment for Process Design - The Case of Minerals Processing. In *Clean Technology for the Mining Industry* (eds. M. A. Sanchez, F. Vergara, and S. H. Castro). University of Concepcion.
- Sundqvist J., Finnveden G., Albertsson A., Karlsson S., Berendson J., Eriksson E., and Hogland L. (1994) Life Cycle Assessment and Solid Waste, AFR-REPORT 29. Swedish Waste Research Council (AFR).
- Surrige A. D., Grobbelaar C. J., and Asamoah J. K. (1995) On South African Coal Reserves/Resources and Their Utilisation. *Colloquium: Coal Processing, Utilisation and Control of Emissions*. South African Institute of Mining and Metallurgy, Randburg, South Africa.
- Swaine D. J. and Goodarzi F. (1995) *Environmental Aspects of Trace Elements in Coal*. Kluwer Academic Publishers.
- Tavoulareas E. S. and Charpentier J.-P. (1995) Clean Coal Technologies for Developing Countries. *Energy Series*. World Bank.
- Thiel C., Seppelt R., Müller-Pietralla W., and Richter O. (1999) An Integrated Approach for Environmental Assessments - Linking and Integrating LCI, Environmental Fate Models and Ecological Impact Assessment Using Fuzzy Expert Systems. *International Journal of Life Cycle Assessment* 4(3), 151-160.
- Tolle D., Hesse D. J., Chadwell G. B., Cooper J. S., and Evers D. P. (2001) Comparison of Two Equivalency Factor Approaches with Simplified Risk Assessment for LCIA of Toxicity Impact Potential. *International Journal of Life Cycle Assessment* 6(2), 96-105.
- Trinius W. and Le Téo J.-F. (1999) System Boundaries According to Decision Scope: A Concept of Focal Zones. *International Journal of Life Cycle Assessment* 4(5), 253-256.
- Tukker A. (1998) Uncertainty in Life Cycle Impact Assessment of Toxic Releases. Practical Experiences - Arguments for a Reductionalistic Approach? *International Journal of Life Cycle Assessment* 3(5), 246-258.
- Udo de Haes H. A. (1993) Applications of Life Cycle Assessment: Expectations, Drawbacks & Perspectives. *Journal of Cleaner Production* 1(3-4), 131-137.
- Udo de Haes H. A., Jolliet O., Finnveden G., Hauschild M., Krewitt W., and Müller-Wenk, R. R. (editorial committee of SETAC-Europe: Second Working Group on LCIA) (1999a) Best Available Practice Regarding Impact Categories and Category Indicators in Life Cycle Assessment: Background Document for the Second Working Group on Life Cycle Impact Assessment of SETAC- Europe (WIA-2): Part 1. *International Journal of Life Cycle Assessment* 4(2), 67-74.
- Udo de Haes H. A., Jolliet O., Finnveden G., Hauschild M., Krewitt W., and Müller-Wenk, R. R. (editorial committee of SETAC-Europe: Second Working Group on LCIA) (1999b) Best Available Practice

- Regarding Impact Categories and Category Indicators in Life Cycle Assessment: Background Document for the Second Working Group on Life Cycle Impact Assessment of SETAC- Europe (WIA-2): Part 2. *International Journal of Life Cycle Assessment* 4(3), 167-174.
- UNEP-IE (1996) Life Cycle Assessment: What it is and how to do it. United Nations Environment Programme Industry and Environment.
- USEPA (1995a) Guidelines for Assessing the Quality of Life-Cycle Inventory Data. U.S. Environmental Protection Agency, Office of Air Quality Planning and Standards.
- USEPA (1995b) Life-Cycle Impact Assessment: A Conceptual Framework, Key Issues, and Summary of Existing Methods. U.S. Environmental Protection Agency, Office of Air Quality Planning and Standards.
- Van Horen C. (1996) *Counting the social costs. Electricity and externalities in South Africa*. UCT and Elan Press.
- van Leeuwen C. J. and Hermens J. L. M. (1993) *Performing Ecological Risk Assessments*. Lewis Publishers.
- WCED (1987) *Our Common Future*. World Commission on Environment and Development, Oxford University Press.
- Weidema B. P. (1993) Quality Criteria for Product Life Cycle Assessment. In *Environmental Assessments of Products: A Textbook on Life Cycle Assessment* (ed. B. P. Weidema). UETP-EEE.
- Weidema B. P. (1998a) Application Typologies for Life Cycle Assessment. *International Journal of Life Cycle Assessment* 3(4), 237-240.
- Weidema B. P. (1998b) Multi-User Test of the Data Quality Matrix for Product Life Cycle Inventory Data. *International Journal of Life Cycle Assessment* 3(5), 259-265.
- Weidema B. P. (2001) Avoiding Co-Product Allocation in Life-Cycle Assessment. *Journal of Cleaner Production* 4(3), 11-33.
- Weidema B. P., Frees N., and Nielsen A.-M. (1999) Marginal Production Technologies for Life Cycle Inventories. *International Journal of Life Cycle Assessment* 4(1), 48-56.
- Weidema B. P. and Wesnæs M. S. (1996) Data Quality Management for Life Cycle Inventories - An Example of Using Data Quality Indicators. *Journal of Cleaner Production* 4(3-4), 167-174.
- Weitz K., Barlaz M., Ranjithan R., Brill D., Thorneloe S., and Ham R. (1999) Life Cycle Management of Municipal Solid Waste. *International Journal of Life Cycle Assessment* 4(4), 195-201.
- Wells J. D., Van Meurs L. H., and Rabie M. A. (1992) Terrestrial Minerals. In *Environmental Management in South Africa* (eds. R. F. Fuggle and M. A. Rabie). Juta.
- Wenzel H. (1998) Application Dependency of LCA Methodology: Key Variables and their Mode of Influencing the Method. *International Journal of Life Cycle Assessment* 3(5), 281-288.
- Wenzel H. (1999) Life Cycle Assessment in Pollution Prevention: Trends in Method Development and Simplifications. In *Tools and Methods for Pollution Prevention* (ed. S. K. Sikdar and U. Diwekar), pp. 119-129. Kluwer Academic Publishers.
- Wenzel H., Hauschild M., and Alting L. (1997) *Environmental Assessment of Industrial Products*. Chapman & Hall.
- Wisberg N. and Gameson T. (1998) CHAINET Definition Document. CHAINET (European Network on Chain Analysis for Environmental Decision Support).

APPENDIX A

STATISTICAL METHODS

Summaries and definitions of the various statistical techniques used and referenced in this thesis.

A.1. PROPAGATION AND ANALYSIS OF UNCERTAINTY

An overview of statistical methods developed for the analysis of uncertainty are presented here. All are well established statistical methods and details can be found in the extensive statistical literature. The Monte Carlo sampling method is expanded upon in greater detail, as this is the method chosen for use in this study. This brief overview is based on that of Morgan and Henrion (1990).

The analysis of uncertainty involves measuring the degree to which each input contributes to the uncertainty in the output. Consider function f , with two uncertain inputs, x_1 and x_2 , and one output, y :

$$y = f(x_1, x_2)$$

the analysis of uncertainty is the degree to which x contributes to y . Methods to quantify this are termed *measures of uncertainty importance*, and are denoted by $U(x, y)$.

The simplest measure of uncertainty importance is *sensitivity*. This is the rate of change of the output y with respect to variation in an input x , i.e. the partial derivatives of output y with respect to each input. The derivatives are evaluated at the nominal scenario, where the nominal scenario is defined as the vector of initial “best guess” or “most likely” values for the inputs. Sensitivity is thus defined as:

$$U_s(x, y) = \left[\frac{\partial y}{\partial x} \right]_{\mathbf{X}^0}$$

$$\text{where } \mathbf{X}^0 = (x_1^0, x_2^0)$$

and x_1^0 and x_2^0 denote the nominal input values (usually the mean, median or mode of the input probability distribution). A problem with using sensitivity for comparing the uncertainty importance of different inputs is that it depends on the scale of x and y , i.e. on the units of measurement. To avoid this, the sensitivity is often normalised, and the changes in x and y defined in relative terms as a fraction of their nominal values, e.g. the percent change in y induced by a 1 percent change in x . This normalised sensitivity, defined as the ratio of the relative change in y induced by a unit relative change in x , is termed *elasticity*:

$$U_E(x, y) = \left[\frac{\partial y}{\partial x} \right]_{\mathbf{X}^0} \times \frac{x^0}{y^0}$$

Sensitivity and elasticity consider only the slopes of the response surface and do not consider the degree of uncertainty in each input. An input that has a small sensitivity but a large uncertainty may be just as important as an input with a larger sensitivity but smaller uncertainty. *Gaussian or first order approximation* is the simplest uncertainty analysis that considers both sensitivity and uncertainty. In this approach, a variable’s uncertainty importance is measured as the product of its sensitivity and uncertainty, i.e. the product of the partial derivative and standard deviation:

$$U_G(x, y) = \left[\frac{\partial y}{\partial x} \right]_{x^0} \sigma_x$$

Gaussian approximation can be used directly to measure the uncertainty propagation, i.e. to estimate the uncertainty of the output. The variance of the output, $\text{Var}[y] \equiv \sigma_y^2$, is estimated as the sum of the squares of the contributions from each input (where $\text{Var}[x] \equiv \sigma_x^2$):

$$\text{Var}[y] \approx \left(\left[\frac{\partial y}{\partial x_1} \right]_{x^0}^2 \text{Var}[x_1] \right) + \left(\left[\frac{\partial y}{\partial x_2} \right]_{x^0}^2 \text{Var}[x_2] \right)$$

Gaussian approximation is a local approach in that it considers the behaviour of the function only in the vicinity of the nominal scenario. It is thus fairly accurate for smooth functions and small uncertainties, but is likely to produce misleading results for complicated functions and large uncertainties. In such cases, a global approach is required, that evaluates the function for scenarios distant from the nominal scenario. The *nominal range sensitivity* method computes the effect on the output of varying each input from its low to high value, while keeping the other inputs at their nominal values:

$$U_R(x_1, y) = f(x_1^+, x_2^0) - f(x_1^-, x_2^0)$$

$$U_R(x_2, y) = f(x_1^0, x_2^+) - f(x_1^0, x_2^-)$$

where $[x_1^-, x_1^+]$ and $[x_2^-, x_2^+]$ denote the bounds of plausible variation for the inputs.

The nominal range sensitivity is more than a local approach, as it evaluates the model for extreme values of each input, but is not truly global, as it holds all the other inputs at their nominal values when looking at the effect of each input. The effect of one input may depend on the values of other inputs, so an approach is required that evaluates the effect of each parameter for several values of the other inputs. This can be obtained by a *parametric analysis* that evaluates y for a sequence of different values for each input, holding the others constant. Figure A-1 shows an example of such an analysis.




Figure A-1 Projection in two dimensions of the three dimensional response surface resulting from a parametric analysis of y with respect to x_1 .

A *scenario tree*, an example of which is shown in Figure A-2, is a useful way to represent possible combinations of inputs. Each node represents an uncertain quantity or event, and each branch from the node one of its possible outcomes. Each path through the tree represents a sequence of event outcomes determining a specific scenario. The number of scenarios increases exponentially with the number of uncertain inputs, so the computational effort to evaluate every scenario rapidly becomes infeasible with increasing numbers of inputs, as does the ability to display and analyse the results. For this reason, often only a few special interest scenarios are examined, e.g. “most likely”, “best case” and “worst case” scenarios, where the input parameters are set to their nominal, best and worst values respectively. An extension of the scenario tree is the *probability tree*. Here a conditional probability is attached to each branch on the scenario tree. Each path through the tree represents a feasible scenario whose probability is the product of the conditional probabilities of the branches along that path. A discrete probability distribution for the output, or risk profile, can be obtained by calculating the probability and output value for each scenario.

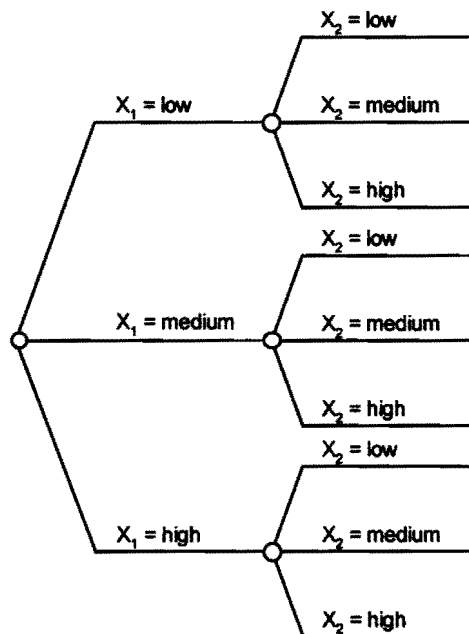


Figure A-2 Scenario tree with three levels

The construction of a probability tree requires that the uncertainty in the input parameters be expressed as discrete probability distributions. However, uncertain quantities are often continuous rather than discrete. It is mathematically too complex for all but the simplest cases to obtain an analytical solution for the probability distribution of a function of a set of continuous random variables. Where there are only a few uncertain variables, it is possible to approximate the continuous distributions by discrete ones and use the probability tree approach. An alternative approach is *Monte Carlo simulation*. In this method, all the combinatorial scenarios are considered by selecting a random sample of scenarios for evaluation. Each scenario is generated by selecting each branch at a node according to its assigned probability, and as the computational effort depends on sample size and not on the number of possible values for each parameter, the branch values may be generated directly from the underlying continuous distribution, avoiding the need to discretise.

The resulting output distribution is inevitably only an approximation of the actual distribution, but it does avoid the approximation due to discretising the continuous distributions. Also the accuracy of Monte Carlo simulation can simply be increased by increasing the sample size, and unlike the probability tree approach, its accuracy can easily be estimated using standard statistical techniques. The appropriate sample size depends on the accuracy required, and thus the application of the model. For models with a large number of uncertain variables, Monte Carlo methods are generally preferable to probability tree methods on the grounds of the computational effort required. In addition, Monte Carlo methods provide a simple measure of uncertainty importance, that of a correlation analysis (explained in section A.1.2 below).

A.1.1. Monte Carlo and Other Sampling Methods

Monte Carlo simulation is the best known and simplest method for sampling from the uncertain input domain. In this method, a value is drawn at random from the distribution for each uncertain input, producing a set of random values. This set, containing one value for each input, defines a scenario used as input to the model, from which the corresponding output is computed. The process is repeated m times, producing m independent scenarios and their output values. The m output values constitute a random sample from the probability distribution over the output induced by the probability distributions over the inputs. Standard statistical techniques can then be used to estimate the precision of the output distribution derived from this random sample. The accuracy of the estimates of the output distribution's parameters depends on the sample size m , and not on the number of uncertain inputs n . This is because the output sample consists of independent random values from the output distribution, and how representative the sample is of the output distribution is irrespective of the number of uncertain inputs. The number of runs required, m , depends on the relative accuracy required of the output distribution. For a given degree of uncertainty, m is thus independent of n , the number of uncertain inputs, although the computational effort to run the model is typically proportional to n .

A.1.1.a Selecting a Sample Size

A compromise needs to be found between the time for each model run and the precision of the results, and will usually be dictated by the application of the analysis. Using standard statistical techniques an estimate of the number of samples required to meet the desired precision can be calculated, either from stating the acceptable uncertainty about the mean, or from stating required confidence intervals for the fractiles. Where the random sample of m output variables is given by $(y_1, y_2, y_3, \dots, y_m)$, the mean and standard deviation of y are estimated by:

$$\bar{y} = \sum_{i=1}^m \frac{y_i}{m}$$

$$s^2 = \sum_{i=1}^m \frac{(y_i - \bar{y})^2}{(m-1)}$$

The required confidence interval, with confidence α , is then given by:

$$\left(\bar{y} - c \frac{s}{\sqrt{m}}, \bar{y} + c \frac{s}{\sqrt{m}} \right)$$

where c is the deviation for the unit normal enclosing probability α . To obtain an estimate of the mean of y with a confidence interval α smaller than w units wide, the width of the interval must be less than w , i.e.

$$2c \frac{s}{\sqrt{m}} < w$$

tain the required number of samples, m , this equation can simply be rearranged to yield:

$$m > \left(\frac{2cs}{w} \right)^2$$

Use this equation, a small number of runs (around ten) first need to be done, so that an initial estimate of the variance, s , can be obtained. The deviation, c , enclosing the specified probability, α , can be obtained from statistical tables, and the value substituted into the above equation, together with the specified interval width w and the estimate of variance, to obtain the number of sample runs.

Alternatively, the number of samples required can be determined by specifying the required precision of the estimate of the median or of the other fractiles. Assuming the m sample values of y are relabelled to be in increasing order, i.e. $y_1 \leq y_2 \leq \dots \leq y_m$, sample value y_i is an estimate of fractile Y_p where $p = i/m$. The confidence interval, α , for a pair of sample values can be shown to be given by (y_i, y_k) , where:

$$i = mp - c\sqrt{mp(1-p)}$$

$$k = mp + c\sqrt{mp(1-p)}$$

the values of i and k are rounded down and up respectively. Supposing confidence, α , of the fractile Y_p is specified as being between the sample value estimates of the $p - \Delta p$ th and $p + \Delta p$ th fractile, i.e. $i = m(p - \Delta p)$ and $k = m(p + \Delta p)$. These expressions for i and k can be combined with the above equations for i and k to yield:

$$m = p(1-p) \cdot \left(\frac{c}{\Delta p} \right)^2$$

With this estimate of m , no previous sample runs need to be done and the number of samples can be obtained directly from the specified precision. For example, a 95% confidence interval for the 50th percentile to be plus or minus one estimated percentile, gives $p=0.5$, $\Delta p=0.01$ and c , the deviation enclosing 95% of the probability of a unit normal, approximately equal to 2 (read from statistical tables). To achieve this very high precision, approximately 10000 runs are required. A reasonable estimate of the desired precision is thus necessary if the number of sample runs is to be kept to a manageable level. Where the empirical uncertainty from the input parameters is high, a very high degree of precision in the propagation of these uncertainties is probably worthless. In this case, the approximation uncertainty due to the number of runs will most probably be dominated by the empirical uncertainty from the input parameters, and a few hundred runs will probably suffice.

1.1.b Selecting a Sampling Method

Monte Carlo sampling is the simplest sampling method, in which each of the sample points (m) for each uncertainty quantity (X) is generated at random from X , with probability proportional to the probability density function $f(x)$. Using the inverse cumulative method, m uniform random variables (u_i for $i=1,2,\dots,m$) between 0 and 1 are generated. The inverse of the cumulative probability distribution is then used to compute the corresponding values of X , i.e.

$$X_i \text{ where } P(x_i \leq X) \text{ for } i=1,2,\dots,m.$$

In this simple Monte Carlo method, each value of every random variable X , including those calculated from other random variables, is a sample of m independent random values from the true probability distribution for X .

The value of the Monte Carlo methods lie primarily in their providing a uniform distribution of points in the parameter space, and not in the randomness of the sampling per se. Stratified sampling methods are able to enhance this primary objective by ensuring a more uniform sample over the input domain. In these methods the sample space for an input parameter is divided up into strata, and input values are obtained by sampling from within each stratum instead of from the whole distribution. One such method is *Latin hypercube sampling* (LHS), in which each input distribution is divided into m equiprobable intervals. A scenario is generated by selecting, without replacement, one value at random from the m samples for each of the inputs, resulting in m nominally independent scenarios. With Median Latin Hypercube sampling the sample points are the medians of the m intervals, that is, the fractiles:

$$X_i \text{ where } \mathbf{P}(x_i \leq X_i) = (i-0.5)/m \text{ for } i=1,2,\dots,m.$$

To avoid non-random correlations among different quantities, the points are then randomly shuffled so that they are no longer in ascending order. Random Latin Hypercube sampling is similar to median Latin Hypercube sampling, except that the sample points are random samples taken from each of the m equiprobable intervals, instead of the medians of the intervals. Thus in random Latin Hypercube sampling each sample is a true independent sample from the distribution, although the samples are not totally independent.

The sample of m values obtained using LHS are for each distribution more uniform than would be obtained by purely random sampling (i.e. Monte Carlo sampling), with median LHS able to yield a still more even spread of sample scenarios over the input domain. LHS is thus able to represent the parameters of the output distribution more accurately than random Monte Carlo sampling, and with median LHS, the mean and variance are often almost exact. Stratified sampling is almost always better than truly random sampling, although for highly non-linear models and those with a large number of uncertain inputs, the improvement may be slight. The sample scenarios, and hence the outputs, are not completely independent when using LHS. Standard statistical techniques are therefore not directly applicable, and more complex measures are required to compute the true precision of the output distribution. The equations presented in the preceding section are thus not able to accurately predict the number of required samples derived from LHS. However, as these equations typically underestimate the true precision, they are able to provide a useful minimum level of precision or maximum number of samples required. In general, median LHS is considerably better than standard LHS, although it is not able to accurately sample functions displaying high frequency periodicity. As such functions are not at all likely to occur in LCA models, median LHS is the sampling method of choice in the uncertainty analysis of empirical quantities in LCA models.

A random sampling technique to generate the random samples, also needs to be chosen. A number of techniques exist, of which three common methods are:

- the *Minimal Standard* random number generator, an implementation of Park and Miller's Minimal Standard, based on a multiplicative congruential method, with a Bays-Durham shuffle;
- the *L'Ecuyer* random number generator, based on L'Ecuyer's algorithm. Also based on a multiplicative congruential method, it gives a series of random numbers with a much longer period but is slightly slower than the Minimal Standard method.
- *Knuth's algorithm*, based on a subtractive rather than a multiplicative congruential method, and is slightly faster than the Minimal Standard method.

The differences between the methods were found to be slight, so the simpler minimal standard method is used.

A.1.2. Measures of Uncertainty Importance

A *sensitivity analysis* is the simplest and most well established method for determining the influence of an uncertain parameter on the overall output. The model is run with only one uncertain input, the other parameters remaining fixed at their most likely value. The analysis is repeated for each uncertain parameter and the resulting sensitivities compared. However, sensitivities calculated in this way do not take into account any interference or correlation between the various uncertainties. They can thus only be taken as an indication of relative sensitivity and not each parameter's actual contribution to the overall uncertainty (Meier, 1997).

An analytical measure of uncertainty importance, able to identify the key parameters contributing to the uncertainty in the results, is demonstrated by Heijungs (1996). Derived from Gaussian approximation, but using absolute error and not standard deviations to characterise the uncertainty, this method rests on the fact that the uncertainty of the results can be expressed as a sum of the individual contributions from the input data. Thus, for the function $y = f(x_1, x_2, \dots)$:

$$\frac{\Delta y}{|y|} = \left| \frac{\partial f}{\partial x_1} \right| \frac{\Delta x_1}{|y|} + \left| \frac{\partial f}{\partial x_2} \right| \frac{\Delta x_2}{|y|} + \dots$$

The above equation is disaggregated, and the individual contributions arranged in decreasing order, so as to give a ranking of the parameters which have the highest influence on the uncertainty of the result (Heijungs, 1996). The method is subject to the same limitations as given for Gaussian Approximation in section A.1.

A powerful global measure of uncertainty importance, applicable to the methods of uncertainty analysis using simulation techniques, is the correlation of the sample output values with the corresponding sample of values for each input variable:

$$U_\rho(x, y) = \rho(x, y)$$

This estimates the effect of uncertainty in x on uncertainty in y , averaged over all possible combinations of values of the other inputs, weighted by their probabilities. Thus for m samples from the output, y_k , and a particular input, x_k , for $k = 1$ to m the sample correlation is given as:

$$U_\rho(x, y) = \frac{\sum_{k=1}^m (x_k - \bar{x})(y_k - \bar{y})}{\sqrt{\sum_{k=1}^m (x_k - \bar{x})^2 \times \sum_{k=1}^m (y_k - \bar{y})^2}}$$

This is truly a global measure of uncertainty because the effect of each input is averaged over the joint probability distribution for all other inputs. The correlation estimates the linear contribution of each input to the output uncertainty, and does not necessarily provide a good measure of non-linear relationships. Where the input or output distributions are far from normal, i.e. if they have long tails, distortion from the effect of outliers is likely to be a problem. To avoid this, rank-order correlations can be used, where the sample values for each input and for the output are rank ordered (i.e. placed in order of increasing magnitude) before computing the correlations. Related measures of uncertainty have been developed that extend the correlation approach, e.g. partial correlation coefficients and principal components, or employ regression coefficients. Scatter plots are also a useful and simple way of visualising the relationship between model inputs and outputs.

A.2. PROBABILITY DISTRIBUTIONS

The following common probability distributions are applicable to characterising the uncertain data inputs of life cycle inventories. Sources referenced for this section are Ott (1999) and Morgan and Henrion (1990).

The Normal Distribution

The normal or Gaussian distribution is commonly employed to represent uncertainty resulting from unbiased measurement errors. It results when a number of unrelated, continuous random variables are added together. However, it is inappropriate for many quantities because negative values can occur (the CV must typically be less than 20% to ensure only positive results). The parameters are estimated from the sample mean and standard deviation.

Probability density function:

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x-\mu}{\sigma}\right)^2}; \quad -\infty < x < \infty$$


Expected value:

$$E[X] = \mu$$


Variance:

$$\text{Var}(X) = \sigma^2$$

PDF: Normal(μ , σ)



CDF: Normal(μ , σ)



The Lognormal Distribution

The lognormal distribution results when the logarithm of the random variable is described by a normal distribution, and is often found to be a good representation for quantities that are non-negative and positively skewed. It results from the product of many independent random variables multiplied together, and is typical of many natural processes, e.g. pollutant concentrations, accident events. It is particularly appropriate for representing large uncertainties that are expressed on a multiplicative or order-of-magnitude basis (e.g. "within a factor of two"). The parameters of the distribution are equivalent to the mean and standard deviation of $Y = \ln X$ (the median and the geometric standard deviation).

Probability density function:

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\ln x - \mu}{\sigma}\right)^2}; \quad -\infty < \mu < \infty$$


Expected value:

$$E[X] = e^{\mu + \frac{\sigma^2}{2}}$$


Variance:

$$\text{Var}(X) = e^{2\mu + \sigma^2} (e^{\sigma^2} - 1)$$

PDF: Lognormal(median, σ_g)



CDF: Lognormal(median, σ_g)



The Triangular Distribution

When a central tendency is expected in the parameter range, or when a certain value is more likely to occur than values near either extreme, a triangular distribution is a simple and convenient distribution to use. It is characterised by specifying the minimum and maximum values, and the mode (the most likely value). The triangular distribution tends to overemphasise the tails and under-emphasise the shoulders of the parameter range, however, the apparently arbitrary shape and sharp corners can be a convenient way to convey that the details of the distribution are not known, and thus help to prevent a false sense of confidence in the model results.

Probability density function (symmetric triangular distribution):

$$f(x) = \frac{b - |x - a|}{b^2}; \quad a - b \leq x \leq a + b$$

Expected value (symmetric triangular distribution):



$$E[X] = a$$

Variance (symmetric triangular distribution):

$$Var[X] = \frac{b^2}{6}$$

PDF: Triangular(min, mode, max)

CDF: Triangular(min, mode, max)



The Gamma Distribution

The Gamma distribution describes the time required for the occurrence of η events in a random Poisson process, where η is the shape parameter of the distribution and λ the scale parameter (corresponding to the rate of the Poisson process). The shape of the distribution runs from highly right skewed to symmetrical and bell-shaped as η increases (equivalent to an exponential distribution when $\eta=1$, and approximates a normal distribution when $\eta>10$). It is applicable to many physical quantities, since, like the lognormal distribution, it returns only non-negative values (e.g. pollutant concentrations). However, it is less positively skewed and less "tail-heavy" than the lognormal distribution, and thus generally prescribes a lower probability to the extremes of the distribution than the lognormal distribution.

Probability density function:

$$f(x) = \frac{\lambda^\eta}{\Gamma(\eta)} x^{\eta-1} e^{-\lambda x}; \quad x \geq 0, \eta > 0, \lambda > 0$$

$$\text{where } \Gamma(\eta) = \int_0^\infty u^{\eta-1} e^{-u} du \text{ or, for integers, } \Gamma(\eta) = (\eta-1)!$$


Expected value:

$$E(X) = \frac{\eta}{\lambda}$$


Variance:

$$\text{Var}(X) = \frac{\eta}{\lambda^2}$$

PDF: Gamma(η, λ)



CDF: Gamma(η, λ)



The Beta Distribution

The Beta distribution provides a flexible means of representing variability over a fixed range, and by simply specifying the first two parameters can be made to reflect both positively skewed or negatively skewed quantities. The second two parameters specify the range endpoints, and if not specified, the range defaults to zero to one.

Probability density function:

$$f(x) = \frac{1}{B(c,d)} x^{c-1} (1-x)^{d-1}; \quad 0 \leq x \leq 1$$

where $B(c,d) = \frac{\Gamma(c)\Gamma(d)}{\Gamma(c+d)}$

Expected value:


$$E(X) = \frac{c}{c+d}$$

Variance:

$$Var(X) = \frac{cd}{(c+d)^2(c+d+1)}$$

PDF: Beta(c, d, min, max)


CDF: Beta(c, d, min, max)



Custom fit Distribution

If a variable does not fit any of the above distribution shapes (e.g. a bimodal distribution), the distribution shape can be custom fit to the data. This is most simply done by defining points on the cumulative probability density curve, e.g. the first column in the table below gives the probability that the outcome will be less than or equal to the corresponding element in the second column.

Probability	Outcome
0%	0
34%	1
48%	2
64%	3
84%	4
92%	5
96%	6
97%	7
99%	8
100%	9



Conditional Probability Distribution

A discrete conditional probability distribution is used where the outcome of one variable depends on the outcome of another uncertain variable. A discrete probability table is first drawn up for the independent variable i.e. an index variable covering the range of possible outcomes of the variable is defined, together with the associated discrete probabilities of that outcome occurring (the probabilities must sum to one). A table of outcomes for the dependent variable is then drawn up, corresponding to the same index variable of the independent variable used to define the probabilities. For example, in the tables below, an outcome of a calorific value of 15.4%, will return ash and carbon values of 45 and 43% respectively.

Index variable	
Calorific Value	Probability
12.6	10%
14.4	12%
15.4	33%
16.3	30%
19.1	15%

Index variable	Dependent variables	
Calorific Value	Ash (%)	C (%)
12.6	59%	32%
14.4	48%	41%
15.4	45%	43%
16.3	43%	45%
19.1	35%	53%

A.3. PRINCIPAL COMPONENT THEORY

PCA is one of the most straightforward geometric techniques, and is widely employed, which makes it the best geometric technique to start with (Murtagh and Heck, 1987). PCA seeks the best, followed by successively less good, summarisation of the data. A graphical representation is typically required to investigate and interpret data. However, in high-dimension spaces we are limited to being able to visualise one-dimensional and two-dimensional representations (axes and planes), or possibly three-dimensional representations at the outside. For example, consider the array of 4 objects by 5 attributes:

$$\begin{pmatrix} 7 & 3 & 4 & 1 & 6 \\ 3 & 4 & 7 & 2 & 0 \\ 1 & 7 & 3 & -1 & 4 \\ 2 & 0 & -6 & 4 & 1 \end{pmatrix}$$

the projection of the 4 objects onto the plane constituted by axes 1 and 3 is simply:

$$\begin{pmatrix} 7 & 4 \\ 3 & 7 \\ 1 & 3 \\ 2 & -6 \end{pmatrix}$$

The projection of points onto axes or planes is thus a trivial operation. However, PCA, first obtains better axes. This it does by seeking the axis to which the cloud of points are closest (see Figure A-3), where the Euclidean distance is usually used to define the closeness to the axes (see section A.4 for a definition of Euclidean distance). This criterion of closeness to the axis can be shown to be identical to a second criterion, which is that the projections of points on the new axis be as elongated as possible, i.e. that the variance of the projections be as great as possible. A second best-fitting axis, orthogonal to the first axis found, will constitute the best-fitting plane. Then a third best-fitting axis, orthogonal to the two already obtained, will together constitute a best-fitting three-dimensional subspace, and so on. Central to the results of a PCA are the co-ordinates of the points on the derived axes (i.e. their projections). The positive and negative orientations of the axes are arbitrary, as they do not influence the objectives of the PCA. The axes are listed in decreasing order of importance, or best-fit. Planar representations of the data can also be generated, i.e. the projections of points on the plane formed by the first and second new axes, then the plane formed by the first and third new axes, and so on (Murtagh and Heck, 1987).




Figure A-3 Two simple examples of centred clouds of points (Murtagh and Heck, 1987).

The objectives of PCA, as given by Murtagh and Heck (1987), are:

1. Reducing the dimensionality of the parameter space to its inherent dimensionality, thereby eliminating “noise”, “cleaning” the data, and lessening the volume of data.
2. Determining the most important linear combinations of the parameters.
3. Determining the important parameters present (feature selection).
4. Determining the underlying (“latent”) variables present.
5. Visualising the data by the selection of the most important planar views of it.
6. Identification of groups of objects or of outliers.

The third and fifth points are particularly useful to an interpretation of LCA results (see section A.3.2). Feature selection investigates opportunities to simplify the task of characterising each object by a set of attributes or variables. Highly correlated variables (i.e. closely located variables in the transformed space) allow some variables to be removed from consideration, whilst the proximity of the variables to the new axes indicate the more relevant and important variables. Whilst PCA provides a convenient planar representation of multidimensional data, an important consideration is the adequacy of the representation, i.e. the percentage of variance explained by the pair of axes defining the plane. Murtagh and Heck (1987) recommend that the new axes should account for approximately 75% of the variance, although they stress that there is no set threshold.

A.3.1. Calculating Principal Components

The basic method of calculating principal components is given here, followed by a simple two-dimensional example. Whilst it is possible to solve the equations analytically for two (or maybe three) variables, iterative solutions are required for higher numbers of variables. This section is merely to give an understanding of the underlying theory of PCA, so these procedures are not presented here, but can be found in the following references on PCA (Jackson, 1991; Jolliffe, 1986; Murtagh and Heck, 1987). The following overview of the method is taken from Murtagh and Heck (1987).

Basic Method

Consider a set of n objects measured on each of m attributes or variables. The $n \times m$ matrix of values will be denoted by $\mathbf{X} = \{x_{ij}\}$ where i is a member of the set of objects, and j a member of the attribute set. The objects may be regarded as row vectors in \mathfrak{R}^m and the variables as column vectors in \mathfrak{R}^n . In \mathfrak{R}^m (the space of objects), PCA searches for the best-fitting set of orthogonal axes to replace the initially-given set of m axes in this space. An analogous procedure is simultaneously carried out for the dual space, \mathfrak{R}^n .

First, the axis that best fits the points in \mathfrak{R}^m is determined. If \mathbf{u} is this vector, and it is of unit length, then the product \mathbf{Xu} of the $n \times m$ matrix by the $m \times 1$ vector gives the projections of the n objects onto this axis. The criterion of “goodness of fit” of this axis to the cloud of points requires that the variance of the points to be maximised when projected onto the new axis, i.e. to minimise the sum of distances between points and axis or equivalently, maximise the sum of squared projections onto the axis. The squared projections of points on the new axis are given by:

$$(\mathbf{Xu})'(\mathbf{Xu})$$

where \mathbf{u} is chosen (arbitrarily) to be of unit length. The maximum of $\mathbf{u}'\mathbf{S}\mathbf{u}$ is sought, where $\mathbf{S} = \mathbf{X}'\mathbf{X}$, and the solution is subject to the constraint that $\mathbf{u}'\mathbf{u} = 1$. This is done by setting the derivative of the Lagrangian equal to zero, i.e. obtaining the derivative of:

$$\mathbf{u}'\mathbf{S}\mathbf{u} - \lambda(\mathbf{u}'\mathbf{u} - 1)$$

where λ is a Lagrange multiplier. This gives:

$$2\mathbf{S}\mathbf{u} - 2\lambda\mathbf{u} = 0$$

The optimal value of \mathbf{u} (i.e. \mathbf{u}_1) is thus the solution of:

$$\mathbf{S}\mathbf{u} = \lambda\mathbf{u}$$

which is a well known occurrence: \mathbf{u} is the eigenvector associated with the eigenvalue λ of matrix \mathbf{S} . Therefore the eigenvector of $\mathbf{X}'\mathbf{X}$, \mathbf{u}_1 , is the axis sought, and the corresponding eigenvalue, λ_1 , indicates the amount of variance explained by the axis.

The second axis is to be orthogonal to the first, i.e. $\mathbf{u}'\mathbf{u}_1 = 0$, and satisfies the equation:

$$\mathbf{u}'\mathbf{X}'\mathbf{X}\mathbf{u} - \lambda_2(\mathbf{u}'\mathbf{u} - 1) - \mu_2(\mathbf{u}'\mathbf{u}_1)$$

where λ_2 and μ_2 are Lagrange multipliers. Differentiating gives:

$$2\mathbf{S}\mathbf{u} - 2\lambda_2\mathbf{u} - \mu_2\mathbf{u}_1.$$

Setting this equal to zero, and noting that multiplying across by \mathbf{u}'_1 implies μ_2 must equal zero, the optimal value of \mathbf{u} (i.e. \mathbf{u}_2), arises as another solution of $\mathbf{S}\mathbf{u} = \lambda\mathbf{u}$. Thus λ_2 and \mathbf{u}_2 , are the second largest eigenvalue and associated eigenvector of \mathbf{S} . The eigenvectors of $\mathbf{S} = \mathbf{X}'\mathbf{X}$, arranged in decreasing order of corresponding eigenvalues, give the line of best fit, the plane of best fit, the three-dimensional hyperplane of best fit, and so on for higher-dimensional subspaces.

In the above derivation, \mathbf{S} is the matrix of the sums of squares and cross products. However, the variables under analysis are often very different (some "shout louder" than others), and it is therefore rare to base the PCA on the sums of squares and cross products (i.e. on the original data). Some transformation of the original data is usually necessary. Murtagh and Heck (1987) recommend basing the PCA on a correlation matrix, which works because the distance between the variables is directly proportional to the correlation between them (see Murtagh and Heck (1987) for proof). In standardising the variables by calculating their correlation coefficients, the row vectors are centred and reduced, i.e. have zero mean and unit standard deviation. Where only centring is appropriate, the PCA can be based on the covariance matrix. Jackson (1991) suggests that the covariance matrix should be used whenever possible, as PCs calculated from it have some useful properties (e.g. the PCs are in the same units of the original data). However, where the data is in different units, as will typically be the case with LCA applications, only the correlation matrix is appropriate (and is consequently the method used in the case studies).

Simple Example

This hypothetical example is taken from Jackson (1990). Although many of the properties of PCA may seem superfluous for this small two-variable example, its size allows one to easily understand these properties and the operations underlying the calculation of PCs.

The concentration of a chemical component in a solution is measured by two different methods. Method 1 is the standard procedure, and Method 2 a proposed alternative. To check that the two methods are interchangeable, the data in Table A-1 is obtained. A plot of this data suggests the use of regression to determine to what extent it is possible to predict the results of one method from the other. The least-squares equation for predicting Method 1 from Method 2 minimises the variability in Method 1 given a specific level of variability in Method 2, and vice versa. However, a single prediction equation is required that could be used in either direction. The line that will perform this role is called the orthogonal regression line, which minimises the deviations perpendicular to the line itself. This line is obtained by the method of principal components.

Table A-1

Observation No.	Method 1	Method 2
1	10.0	10.7
2	10.4	9.8
3	9.7	10.0
4	9.7	10.1
5	11.7	11.5
6	11.0	10.8
7	8.7	8.8
8	9.5	9.3
9	10.1	9.4
10	9.6	9.6
11	10.5	10.4
12	9.2	9.0
13	11.3	11.6
14	10.1	9.8
15	8.5	9.2

In this example, the PCA is based on the covariance matrix, so first the sample means, variances and the covariance between the two methods must be obtained.

Let x_{1k} be the test result for Method 1 for the k th run and the corresponding result for Method 2 be denoted by x_{2k} . The vector of sample means is:

$$\bar{\mathbf{x}} = \begin{bmatrix} \bar{x}_1 \\ \bar{x}_2 \end{bmatrix} = \begin{bmatrix} 10 \\ 10 \end{bmatrix}$$

and the sample covariance matrix is:

$$\mathbf{S} = \begin{bmatrix} s_1^2 & s_{12} \\ s_{12} & s_2^2 \end{bmatrix} = \begin{bmatrix} 0.7986 & 0.6793 \\ 0.6793 & 0.7343 \end{bmatrix}$$

where the covariance is given by:

$$s_{ij} = \frac{n \sum x_{ik} x_{jk} - \sum x_{ik} \sum x_{jk}}{[n(n-1)]}$$

The above derivation shows that to determine the principal components, it is necessary to calculate the eigenvalues (or characteristic roots) of S . This can be done by solving what is termed the characteristic equation:

$$|S - \lambda I| = 0$$

where λ is the matrix of eigenvalues and I the identity matrix.

For the example under consideration, this gives:

$$|S - \lambda I| = \begin{vmatrix} 0.7986 - \lambda & 0.6793 \\ 0.6793 & 0.7343 - \lambda \end{vmatrix} = 0.125 - 1.53\lambda + \lambda^2 = 0$$

$$\text{and } \lambda_1 = 1.45 \text{ and } \lambda_2 = 0.086$$

The eigenvectors can be obtained by the solution of:

$$[S - \lambda I]t_i = 0 \text{ where } u_i = \frac{t_i}{\sqrt{t_i' t_i}}$$

For the example, for $i=1$, this gives:

$$[S - \lambda_1 I]t_1 = \begin{bmatrix} 0.7986 - 1.4465 & 0.6793 \\ 0.6793 & 0.7343 - 1.4465 \end{bmatrix} \begin{bmatrix} t_{11} \\ t_{21} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

To solve the above equation, let $t_{11} = 1$, and solve just the one equation, i.e.:

$$-0.6478 + 0.6793t_{21} = 0$$

which gives $t_{21} = 0.9538$. Substituting this into the above equation for u_i yields:

$$u_1 = \begin{bmatrix} 0.7236 \\ 0.6902 \end{bmatrix}$$

Similarly, for $\lambda_2 = 0.0864$ and letting $t_{22} = 1$:

$$u_2 = \begin{bmatrix} -0.6902 \\ 0.7236 \end{bmatrix}$$

The above procedure is nothing more than a principal axis rotation of the original co-ordinate axes about their means. This principal axis transformation can now be used to transform m correlated variables x_1, x_2, \dots, x_m into m new uncorrelated variables z_1, z_2, \dots, z_m . The co-ordinate axes of these new variables are described by the eigenvectors u_i . The transformation is therefore given by:

$$z = u' [x - \bar{x}]$$

where x and \bar{x} are $m \times 1$ vectors of observations on the original variables and their means. The transformed observations are referred to as the principal component scores to differentiate them from the transformed variables, i.e. the principal components.

For the example, the first observation in Table A-1 gives:

$$\mathbf{z} = \begin{bmatrix} 0.7236 & 0.6902 \\ -0.6902 & 0.7236 \end{bmatrix} \begin{bmatrix} 10.0 - 10.0 \\ 10.7 - 10.0 \end{bmatrix} = \begin{bmatrix} 0.48 \\ 0.51 \end{bmatrix}$$

i.e. the principal component scores for the first observation are $z_1 = 0.48$ and $z_2 = 0.51$.

Interpreting the principal components for the two variable example is quite straightforward. The coefficients of the first eigenvector, 0.72 and 0.69 (often referred to as the PC loadings), are nearly equal and both positive. This indicates that the first PC is a weighted average of both variables. This is related to the variability that x_1 and x_2 have in common, which in the absence of correlated errors of measurement, can be assumed to represent process variability. The coefficients for the second eigenvector, -0.69 and 0.72, are also nearly equal, except for the sign. The second PC can therefore be interpreted as representing differences in the measurements for the two methods, i.e. testing and measurement variability. The variance of the first PC is equal to $\lambda_1 = 1.45$, whilst that of the second PC is equal to $\lambda_2 = 0.086$. The fact that the sum of the original variances is equal to the sum of the characteristic roots enables each PC's contribution to the total variance to be estimated:

$$s_1^2 + s_2^2 = \lambda_1 + \lambda_2 = 0.7986 + 0.7343 = 1.5329$$

and the % variance accounted for by the first PC given by $\frac{1.45}{1.53} = 0.94$

Thus, roughly 94% of the total variability of the data is accounted for or "explained by" the variability of the process, and 6% is due to the variability related to testing and measurement.

3.2. Interpreting Principal Component Plots

The PC representation used in this thesis is after the method of Le Teno (1999), in which the principal component scores (the transformed variables) and PC loadings (the coefficients of the eigenvectors) are merged into a single plot. In two-dimensions (i.e. for a plot of the top two PCs), the plot of the PC loadings gives the plane of best fit (determined by the criteria PCA, see above derivation). The arrows on the PC plots thus give the distance from the origin to the plane of best fit, and indicate each criteria's contribution to the particular PC (the longer the arrow, the greater the contribution). The first PC accounts for the greatest share of the overall variance in the system, the second for the next largest, and so on (as explained above), thus the criterion contributing the greatest loading to the first PC is the criterion responsible for the biggest differences between the systems, and so on. In addition to the strength of the criterion, the PC loadings provide information about the dependence of the criterion. Criteria that plot close together signify strong correlations in the data between those criteria, i.e. the options change to a similar extent and in a consistent direction for each of the criteria. Thus those criteria not capturing any significantly different information about the system not already demonstrated by another criterion can be identified and eliminated.

Merged with the plot of the PC loadings is a plot of the transformed data (the PC scores). These are each of the data points transformed onto the PC plane under consideration (usually that of the first two PCs, as this maximises the variance able to be displayed). For a mid-point analysis, each of the system alternatives plot as single points on the PC plane, whilst for probabilistic output samples, the system alternatives plot as "clouds" of points. These "clouds" can be interpreted as "zones of confidence" and the overlap between the clouds for each option identifies the significance of the differences between the options. In addition, a large spread identifies high uncertainty in the system, and thus identifies which options require better data characterisation.

The merged plot of loadings and scores is required to provide information as to the cause of the particular orientation of the points or "clouds" in the PC space (i.e. the directional "pull" each criterion has on the "clouds"). The PC loadings are represented by arrows on the plot, as only their relative length and direction from the origin of the PC plane are of relevance (i.e. they do not correspond to the values on the axes). The arrows yield significant information about the criteria, with the length of the arrows representing the strength of the criteria, and the clustering of the arrows representing the independence of the criteria. The direction of the arrows on the plot represent the performance of the options relative to the particular criteria, with those points "clouds" plotting in the direction of the arrow performing poorly on that criterion, and those away from the arrow performing well.

As the determination of the PCs is based purely on maximising the variance between the original data and the PC plane, the absolute positioning of the loads and scores is not of significance (and tends to change between plots as the optimum representation is sought). Only their orientation relative to each other is of relevance. In addition, the units in which the output samples are specified for each criteria are not of significance, as the PCA is based on the correlations between the sample points.

APPENDIX B

MEASURES TO ADDRESS DATA UNCERTAINTY IN LCA

B.1. QUALITATIVE AND SEMI-QUANTITATIVE DATA QUALITY ASSESSMENT METHODS

An overview of the major data quality guidelines and frameworks developed for the inclusion of a qualitative or semi-quantitative assessment of data quality into the LCA process.

B.1.1. SETAC Data Quality Workshop

As part of an effort to develop a uniform and consensus-based system for assessing and documenting LCA data quality, SETAC sponsored a workshop in October 1992 on the role of data quality in LCA. The outcome was the development of a conceptual framework for data quality assessment. As well as identifying the general principles for data quality assessment, specific data quality issues were identified and discussed for the specific areas of energy, materials, environmental emissions, ecological health and exposure and human health and exposure (Fava et al., 1994). Whilst "LCA data" was defined as a collective term for all data and information used in performing LCAs, ranging from input data to final results, the emphasis of the workshop was on input data. The reliability of LCA final results depends on the quality of the input data and the way in which they are processed into results using the LCA methodology. A data quality assessment (DQA) was thus defined by Fava et al. (1994) as a systematic approach to identifying and applying measurements of the suitability of LCA data to meet the intended purpose of the study.

The use of data quality goals (DQGs) was identified as an applicable DQA technique for LCA. The DQG process was defined to consist of three activities:

1. *Identification of decision types* focuses on rational and transparent definition of goals, eligibility of potential data sources, and identification of conceptual models for the LCA.
2. *Identification of data uses and needs* involves the identification of data types, data quality requirements and data quality indicators (DQIs) commensurate with the DQGs.
3. *Design of data collection programs* addresses data collection, interpretation and documentation, data use and processing into final results, and assessment relative to goals.

The DQG process is iterative during the course of the study, and is ultimately designed to support and defend the decision-making process related to the findings of the LCA. It is purported to enhance communication, provide a structure for augmenting existing data sets, provide a focussed set of data requirements, and define the resulting uncertainty of the study results. The level of quality achieved is dependent on the level of effort that is allocated to the study, to each subsystem and to each variable, as decided by the study purpose, budget, time constraints, data availability and the need to maintain scientific integrity (Fava et al., 1994).

The workshop identified the need for a DQA framework because the diverse nature of LCA data required a logical, formalised and repeatable method for its analysis. The framework, shown in Figure B-1, was developed to allow for flexible data documentation using quantitative or qualitative measures. The process begins with the establishment of preliminary DQGs during the goal definition and scoping phase. The DQGs are based on the purpose and scope of the study and with reference to a checklist of DQIs. The set of generic DQIs developed by the workshop are listed in Table 3.3. Data collection is begun once the DQGs have been defined, but may be

refined or additional DQGs added as data collection progresses. If the quality of the data collected meets the requirements, the study can proceed, if not, measures must be taken to combat the poor quality data. These include obtaining better data, clearly defining the limitations of the data collected, performing sensitivity analysis to test the uncertainty, modifying the purpose of the study, or abandoning the study. Factors used to decide the inclusion or rejection of data include the quality of the data source, the analyst's degree of knowledge of the product or process, the assumptions made and the calculations used to generate the results (Fava et al., 1994).

Two approaches to assessing data quality were identified during the workshop: qualitative evaluation using a matrix approach, and quantitative evaluation using value trees. The matrix approach uses a qualitative representation of both qualitative and quantitative DQGs to highlight the data believed to reduce confidence in the conclusions of the study. Data quality concerns are identified and preserved at the time they arise, and can be later reviewed for significance. The level of detail of the matrix depends on the nature of the study, but generally the matrix framework avoids difficulties with reporting the data quality for each data source. Whilst comprehensive and reasonably transparent, it is acknowledged by Fava et al. (1994) that it can be biased. In the value tree approach quantitative quality scores are assigned to each DQG in the value tree. The DQGs are then aggregated into an overall index by applying weighting factors to each DQG according to their importance to the goals of the study. These weighting factors are highly subjective and will vary across studies, possibly resulting in a final numerical score of questionable significance and a false implication of precision (Fava et al., 1994). The use of sensitivity analyses during data quality assessment is recommended by Fava et al. (1994) to direct the expenditure of time and money to those areas most likely to improve the overall quality of the study.

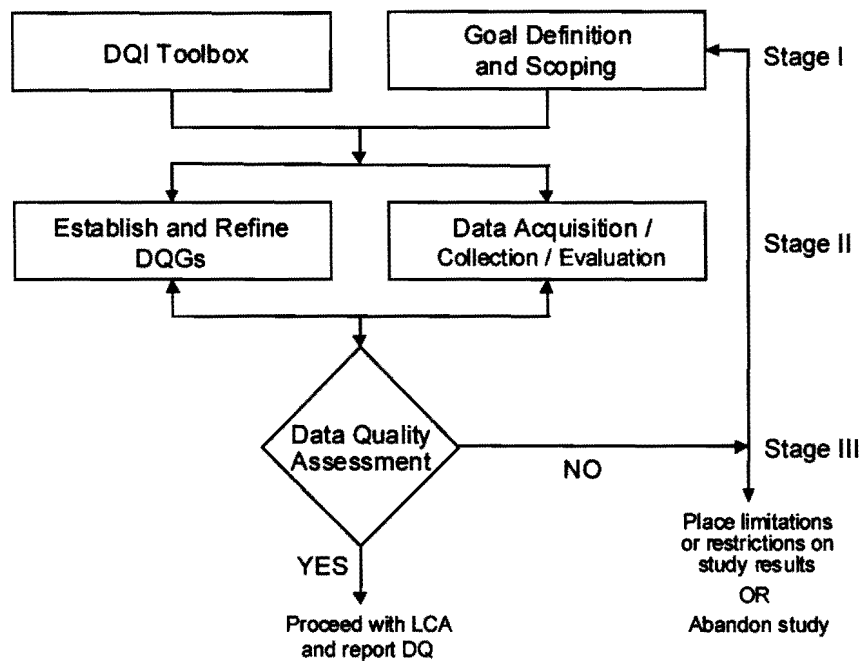


Figure B-1 Conceptual framework for assessing LCA Data Quality (Fava et al., 1994)

B.1.2. USEPA Data Quality Worksheets

The purpose of the USEPA data quality assessment guidelines is to provide guidance for assessing the quality of key LCI data sources. They do not suggest conducting detailed assessments of every data source used, nor do they lay out each step required for data collection and quality assessment. Like the SETAC guidelines, the USEPA approach is based on DQGs and DQIs, although relies on a more contained set of DQIs (see Table 3-3). The guidelines address where and how data quality can be considered and provide a general, systematic approach, which recognises that an iterative process is often required. The approach is based on a series of data quality worksheets which document data-quality information on key data sources. The worksheets include a thorough documentation of the data source, the DQGs and the data assessed, and a qualitative rating of the relevance of the DQIs and the data quality according to each DQI. The rating of “high”, “medium” or “low” is qualified with a brief explanation (USEPA, 1995a).

B.1.3. ISO 14041 Data Quality Guidelines

The ISO standard on goal and scope definition and inventory analysis states that data quality should be characterised by both quantitative and qualitative aspects, as well as by the methods used to collect and integrate those data (ISO, 1998). Although no clear procedure is suggested by the standard, a similar approach to the SETAC conceptual framework can be inferred, i.e. that of setting DQGs, assessed by a number of DQIs. The DQIs mentioned in the standard, although not particularly defined as such, are summarised Table 3-6. The first four relate to data collection, i.e. the desired age of the data and the minimum length of time over which the data should be collected, the geographical area from which the data should be collected, and the technology mix for which the data is collected. The standard states that data quality requirements (or DQGs) should be set for these three parameters. The fourth, the nature of the data relates to the type and source of the data, i.e. whether it is collected from specific sites, or whether it is from published sources, and also whether it is measured, calculated or estimated.

Two quantitative DQIs are stipulated, the precision, as measured by the variability of the data values, and the completeness, as measured by the percentage of locations reporting data from the potential number in existence. The remaining three DQIs included are qualitative in nature. The first, representativeness, is largely redundant, as this has already been covered under the time-related, geographical and technology coverage descriptors, and is an assessment of the degree to which the data set reflects the true population of interest. Consistency relates to the methodology of the study rather than the data, and refers to how uniformly the methodology is applied to the various components of the analysis. Reproducibility relates to the transparency of methodology and the reporting structures, and refers to what extent an independent practitioner could reproduce the results. The ISO standard suggests no procedure for data quality assessment, although it does call for “a data quality assessment and sensitivity analyses on significant inputs, outputs and methodological choices” as part of the interpretation of LCI results.

B.1.4. Data Pedigree

Instead of reporting the usual single value, this approach reports a summary of the data quality characteristics in a numeric code. The basic concept is consistent with that of the methods discussed above, i.e. the definition of DQGs and their assessment against various relevant DQIs, although the pedigree approach provides a structured format aimed at achieving reproducible data quality statements. The NUSAP (Numerical, Unit, Spread, Assessment, and Pedigree scheme) states the spread of values associated with a certain data value, and uses a numerical code to describe various qualities related to the data and the data acquisition methods (Funtowicz and Ravetz, 1990). The code proceeds from more quantitative to more qualitative aspects of the information as it is read from left to right. The spread conveys an indication on the inexactness of the information in the numeral and unit places, and is usually reported as a confidence interval, calculated using standard statistical techniques (e.g. " $\pm n$ ", "with variance σ " etc.). The assessment expresses a judgement of the reliability associated with the quantitative information conveyed in the previous categories. It can be represented through "confidence limits" or "significance levels" of classical or Bayesian statistics, or can be indicated verbally on a qualitative scale (e.g. "total", "high", "medium" or "low" degree of confidence). The assessment category incorporates such concepts as the data's information value based on its spread, its application, its generalisability to other applications and its robustness over time.

The pedigree category is the most qualitative and complex of all the NUSAP categories. Its aims to represent uncertainties operating at a deeper level than those of the other categories, and "conveys an evaluative account of the production process of the quantitative information (e.g. whether the number is based on an exhaustive and detailed measurement process, or upon a snap judgement from someone over the telephone)" (Funtowicz and Ravetz, 1990). The pedigree category thus captures the "border with ignorance" sort of uncertainty, and is represented in a matrix form by assigning a numerical score to a number of elements including: theoretical basis, data input, estimation methods, estimation metric, theoretical structures, level of peer acceptance and colleague consensus (Funtowicz and Ravetz, 1990).

Weidema and Wesnæs (1996) simplified and adapted the NUSAP scheme to produce the "pedigree matrix" reproduced in Table 3-7. Their "pedigree matrix" incorporates the last two qualitative components of the NUSAP scheme, particularly tailored to incorporate relevant metrics of LCA data quality. The matrix gives clear guidelines to obtain a semi-quantitative indication of the reliability of data (including an assessment of the data sources, sampling methods and verification procedures), the completeness of the data (including the statistical representativeness of the data, the number of measurements in the sample and the time period of data collection), and the temporal, geographical and further technological correlations between the data and the data quality goals. Like the NUSAP code, Weidema and Wesnæs (1996) suggest that the five "scores" from the matrix be quoted in brackets behind the reported data quantity.

The reliability indicator and the completeness indicator both relate to the properties of the sample itself, and are independent of the DQGs of the study (i.e. the score given to the indicator would remain the same if the data were used in another study). Whilst the temporal indicator, which represents the time correlation between the year of study and the year of the obtained data, is closely related to the DQGs of the study, as are the geographical and technological indicators. The geographical indicator represents the geographical correlation between the defined area of the study and the obtained data, whilst the technological indicator is concerned with all aspects of correlation other than that covered by the temporal and geographical considerations.

A semi-quantitative score from 1-5 is assigned to each indicator. The scores serve as identification numbers only and cannot be regarded as representing an “amount” of data quality. A score of 2 for one indicator is not necessarily of the same importance of a score of 2 for another indicator, nor is a score of 4 twice as problematic as a score of 2 in any one indicator. The scores can therefore not be compared across indicators, and any attempt to aggregate the numbers is therefore inappropriate (Weidema and Wesnæs, 1996).

The five indicators characterise independent aspects of data quality, and are argued by Weidema and Wesnæs to be “necessary and sufficient” to describe LCI data quality. A number of authors have suggested modifications to this pedigree matrix, which include adjusting the indicators to be more sensitive, adding additional indicators or splitting the existing indicators into sub-indicators, and aggregating across the indicators (Weidema, 1998b). Weidema argues against these adjustments for the following reasons. For communication purposes, the scores should not be adjusted for different studies, and different sensitivities can usually be ascertained from the context of the study. As stated above, the scores do not represent an “amount” of uncertainty, so aggregation is strongly discouraged by Weidema and Wesnæs (1996). An expression of reliability at the system level can be more meaningfully obtained by interpreting the data quality in terms of quantitative estimates of the variance the score introduces, which can then be combined with the statistically measurable uncertainty of the data sample (termed the “basic” uncertainty by Weidema and Wesnæs), and the overall uncertainty in the output computed by simulation techniques (e.g. Monte Carlo analysis). The additional indicators suggested by authors can also be better incorporated by this extended method. Weidema stresses the need to keep the number of indicators as low as possible for operational purposes, and also the need to keep interrelated issues together to ensure the individual indicators are independent (Weidema, 1998b).

The efficacy of the pedigree matrix has been demonstrated by a multi-user test in which seven persons independently scored the same 10 data sets with minimal disagreements (Weidema, 1998b). In addition scoring per data set took on average less than 10 minutes. Deviations in scores were in most cases among neighbouring scores, and less than 10% of the deviations between participants affected the overall assessment of the data quality of the data set. A small number of deviations in scores is inevitable, as there remains an irreducible amount of subjectivity in any data quality assessment, despite any improvements and explanations that can be made to the data quality matrix. Thus Weidema emphasises that the scores should not be seen as objective, but rather as a representation of a subjective judgement of the data quality. This does not compromise the usefulness of the pedigree matrix for internal data quality management or its ability to communicate the subjective judgement of data quality in a comprehensive way (Weidema, 1998b).

B.2. QUANTITATIVE MEASURES TO ADDRESS DATA UNCERTAINTY DEMONSTRATED IN LCA

Brief summaries of the studies demonstrating a quantitative assessment of data uncertainty found in the LCA literature follow below.

B.2.1. Simulation Techniques

Simulation methods for the assessment of the uncertainty of empirical quantities are well established, thus the challenge of using these methods for LCA lies not with the method, but in defining relevant input distributions that realistically characterise the uncertainty in the input data. The LCA studies that have demonstrated the use of this technique are therefore discussed according to their method of defining the input probability distributions.

B.2.1.a Uncertainty Factors

To illustrate parameter uncertainty and uncertainty due to choices, Huijbregts (1998b) compares two types of roof gutters with respect to their potential contribution to global warming and acidification using simplified inventory data. The stochastic model is based on sampling simple triangular probability distributions. Standard sensitivity ranges or “uncertainty factors” are defined for each input quantity, which, together with the mean or “most likely” values, are used to determine the triangular probability distributions. The minimum value is simply found by dividing the most likely value by the uncertainty factor, and the maximum by multiplying the most likely value by the uncertainty factor. Huijbregts (1998b) suggests the use of this method when only a rough first-pass solution is required.

B.2.1.b Subjective Judgements Using DQIs

Kennedy et al. (1996) apply Monte Carlo simulation to a set of LCI models for a beverage company. The deterministic LCI models are converted to stochastic models by seeking expert opinion (i.e. the original developers of the LCI models) to determine a semi-quantitative estimate of the data quality of each input data element. The LCA practitioner is required to assign a value of 1 to 5 on a sliding scale according to the perceived quality of the data, where 1 is the best possible quality data and 5 is the worst, as well as likely minimum and/or maximum values. Each DQI corresponds to a pre-established set of beta probability distribution shape parameters and range endpoints, which, in the absence of specific information, are based on a set percentage deviation from the original single data element value. Kennedy et al. (1996) claim the beta probability distribution is the most practical to use where nothing is known about the actual probability distribution. By simply specifying four parameters, a great variety of distribution shapes can be achieved, thereby allowing for greater modelling flexibility, whilst requiring no prior knowledge of the distribution shape (i.e. the shape can be “built up” using expert judgement). In addition, range endpoints can be specified, thus allowing minimum and maximum values to be incorporated.

The procedure was applied to near complete LCI models, and Kennedy et al. (1996) acknowledge that ideally the data quality should be evaluated as it is being researched, collected and aggregated. They suggest that the stochastic modelling approach should be used at lower tiers of the modelling process, which, together with sufficiently large sample sizes, would allow more representative probability distributions to be determined.

B.2.1.c Combination of DQIs and Measured Uncertainty

Weidema and Wesnæs (1996) develop a method to translate data quality scores into a quantitative estimate of variance. They distinguish between “additional” uncertainty and “basic” uncertainty. The former is related to the data not being of optimal quality (i.e. the uncertainty captured by their pedigree matrix), whilst the latter arises from measurement errors and normal fluctuations of the measured variable (Weidema and Wesnæs, 1996). If the sample of measured data is sufficiently large, the “basic” uncertainty can be calculated, and is combined with an estimate of the “additional” uncertainty to yield an estimate of the total variance of the variable. To estimate the “additional” uncertainty, quantitative estimates of the uncertainty introduced by each poor pedigree matrix score are made by assigning them each a coefficient of variance (CV). Unlike the individual DQI scores, it is possible to aggregate the estimated CVs to get an overall estimate of the variance, because the influences of the individual DQIs are additive as long as the DQIs are determined independently of each other (Weidema and Wesnæs, 1996). In addition to increasing the uncertainty, a low score on a DQI may require an adjustment of the mean value. The CVs are either calculated, where sufficient data on the quantity, or a related quantity exists, or they are approximate from the technical knowledge of “experts” (Weidema and Wesnæs, 1996).

Meier (1997) undertakes a detailed quantitative analysis of both inventory and impact assessment uncertainty to determine the significance of the rankings obtained for four gas purification systems. He develops a method very similar to that of Weidema and Wesnæs (1996), but which is generalised and does not differentiate between “additional” and “basic” uncertainty.

Meier defines default matrices for the following types of data uncertainty identified in his study:

1. Uncertainty in measuring process data (foreground data)
2. Uncertainty due to age of data/temporal variation (foreground data)
3. Uncertainty of the emission measurements (background data)
4. Uncertainty due to the averaged data over space and time (background data)

Those uncertainties characterised by a normal distribution (1-3) are assigned a default CV, whilst the fourth uncertainty is characterised by a lognormal distribution and assigned a standard deviation of the corresponding normal distribution (STD_n). This is because the first three uncertainties are assumed to be symmetrically distributed around the mean, whilst the possible presence of outliers (dirty processes) means that the asymmetric lognormal distribution better characterises the fourth type of uncertainty. Whilst recognising the beta distribution’s flexibility, Meier (1997) states that a trade-off needs to be made between a flexible definition of the function’s shape on one hand and reproducibility and practicability of its application on the other. He therefore advocates the use of normal and lognormal distributions because they may both be conveniently characterised by a CV or as a factor of uncertainties respectively, and both are well-known and easily definable statistical functions. Whilst the assumption of any distribution shape introduces a fair amount of subjective uncertainty, as it pretends the actual distribution is known, the clear and reproducible nature of the normal and lognormal distribution is expected to reduce the necessary expert judgements and subjective assumptions as much as possible (Meier, 1997).

Each type of data uncertainty has a different number of score options (ranging from 4 to 7), which correspond to a pre-defined measure of uncertainty (CV or STD_n). The measures of uncertainty are based on a combination of actual foreground and background data, literature and expert judgement. As with “scoring” the data elements on

Weidema and Wesnæs' pedigree matrix, guidelines are given to facilitate reproducible scoring of the data elements. Once the DQI score has been assigned its corresponding measure of uncertainty is read off the table. The corresponding uncertainty in the results is calculated by simulation using Latin Hypercube sampling (Meier, 1997; Meier and Hungerbühler, 2000).

B.2.1.d Actual Data Ranges

Coulon et al. (1997) use a combination of actual data and expert judgement to incorporate Monte Carlo analysis into an LCA study comparing Polyvinyl Chloride and Polycarbonate/Acrylonitrile Butadiene Styrene. They state that actual data provide the best source of information for characterising the distribution of the input, but suggest the use of "expert judgement" to estimate the data uncertainty where actual data are not available. Coulon et al. (1997) use lognormal and triangular distributions (symmetric or asymmetric) when there is strong reason to expect a central tendency to the data, whilst uniform distributions are used when strong central tendencies are not expected (Coulon et al., 1997).

Maurice et al. (2000) have developed an iterative procedure to streamline quantitative uncertainty assessments in LCI, which they demonstrate with a case study on the production of electricity in French coal power plants. Their method identifies the data most likely to contribute to the uncertainty using qualitative means, probability distributions are then assigned to these data elements according to the amount of information known about them, and Monte Carlo simulation is used to propagate these uncertainties to the LCI results. The contribution to the uncertainty by the data elements not selected for detailed uncertainty analysis is then roughly assessed, and if not found to be negligible, the process is repeated with a more detailed analysis of the non-selected data.

A similar approach to Coulon et al. (1997) is taken by Maurice et al. (2000) to define the uncertain input parameters, although they present a hierarchy of desired approaches, depending on the type and amount of data available. They distinguish between those parameters that are extensively measured, and those that are based on little information. Parameters for which more than 30 measurements exist fall into the first category. This is usually the case for product or service flows, and increasingly so for flows of environmental consequence (e.g. SO₂, particulates etc). Standard statistical techniques are used to fit the data measurements to the relevant probability distribution. The probability distribution is only roughly estimated, since defining it with high precision would be irrelevant given the nature of the information on which it is based (Maurice et al., 2000).

For parameters that are rarely measured or merely estimated, Maurice et al. (2000) recommend basing the estimate of uncertainty on a comparison of different literature values of the parameter, on the uncertainty of a related parameter giving rise to the parameter of interest, or on estimates of uncertainty of similar processes. When several literature values are in existence, and the reasons for their differing from each other is not known, a uniform distribution can be used between the minimum and maximum value. If one value is thought to be of better quality than the others, a triangular distribution may be used. Where no uncertainty information is known about the parameter itself, but the main underlying parameter causing the variability can be identified and its uncertainty is known, then its uncertainty can be used as a proxy for the dependent variable. The other suggestion of Maurice et al. when faced with no uncertainty information about a particular data element, is to apply the confidence interval of a similar process. For a first pass uncertainty assessment (i.e. when no uncertainty information is known about the parameters and a quick estimate is desired), Maurice et al. use large intervals based on the variation seen in LCI databases. Lower and upper limits have been determined by Finnveden (1998) for a number of LCI databases according to a rough breakdown of the type of flow.

Thus, in agreement with Coulon et al. (1997), Maurice et al. (2000) suggest the use of a range of distribution types of different levels of sophistication. However, they choose beta distributions (or a version of it, called the Pert distribution) for precise definitions, as they find these easier to define than normal or lognormal distributions (i.e. lower, upper and most probable value is all that is required instead of the less intuitive standard deviation). They use uniform or triangular distributions where little information is known about the quantity, although they highlight some limitations of these distributions. A problem with the uniform distribution is that the probability abruptly changes to zero at the upper and lower limit. To remove this problem, they suggest the lower and upper limit be specified as the limits of a confidence interval of a normal distribution (e.g. 90% or 99% confidence limits). The triangular distribution, whilst flexible, simple and intuitively easy to specify, tends to overemphasise the tails of the distribution and under-emphasise the shoulders. To achieve less of an emphasis on the minimum and maximum values they suggest using the beta distribution (Maurice et al., 2000).

B.2.2. Analytical Techniques

The use of analytical uncertainty propagation techniques require that the LCI be presented as a function, e.g. Gaussian approximation requires that partial derivatives be taken of the function. This can be achieved using the matrix-based input-output method of inventory calculation, which provides linear algebraic expressions for the quantification and occurrence of each process (Heijungs, 1994).

B.2.2.a Gaussian / First-Order Approximation

Heijungs demonstrates the use of Gaussian approximation to propagate uncertainties through an LCI of sandwich packaging, where the LCI is expressed as a series of matrices (Heijungs, 1996). A limitation of Gaussian or first-order approximation is that it is a local approach and consequently valid only for smooth functions and small uncertainties. Thus the accuracy of the uncertainty assessment is a consideration, although it is purported to give a “fairly good approximation” for larger uncertainties (Heijungs, 1996). The Gaussian approximation method, as applied by Heijungs (1996) is slightly adapted in that absolute error rather than standard deviations are used in the determination of the overall uncertainty. The reason for this is that the method developed by Heijungs is more concerned with identifying the key data elements contributing to the overall uncertainty, than it is with accurately determining the reliability of the LCI results. Linearity was identified as important to finding the key contributing elements to the uncertainty because of the need to be able to represent the total uncertainty as a sum of the ranked individual contributors, although Heijungs (1996) acknowledges that the use of absolute error is less statistically sound.

As with the simulation methods, a greater obstacle to the method than the mathematical procedure, is the definition of the input data uncertainty. This is slightly simplified in this method, as only the margin of error, and not its likely distribution is required. Nevertheless determining margins of error for the input data is acknowledged to be problematic by Heijungs (1996), who suggests that identical margins of error, without reference to actual margins of error, can be used as part of an initial sensitivity analysis to determine key issues.

B.2.2.b Interval Modelling

Chevalier and Le T no (1996) recommend the use of interval modelling to express ill-defined LCA data because of its generality and ease of use. Whilst they define LCA data as fuzzy, the increased effort of defining possibility distributions and the increased computational complexity, led to the choice of intervals over fuzzy sets to characterise the uncertain data. As with the previous analytical method, the matrix method of calculating the inventory, as defined by Heijungs (1994), is used as a basis for the method. A twofold strategy is developed to solve the required matrix calculations with intervals instead of single values. Where the system can be described by an M-matrix (a special type of inverse positive matrix), an exact solution can be calculated using interval arithmetic, whilst the interval equation is converted into a system of inequalities and solved by linear programming in the general case. The exact solution can be obtained provided there are no loops and multiple-output processes in the production tree. A complete mathematical description of the interval calculation procedure can be found in Chevalier and Le T no (1996), where the method is demonstrated on a simple inventory for the manufacture of wallboard products.

APPENDIX E

INVENTORY MODELS

A detailed description of the models underlying the LCI are presented here. Special attention is afforded to the many assumptions made, and methodological decisions taken.

E.1. GENERAL MODEL STRUCTURE

The inventory model is built up in a modular format. It is broken down into the three main life cycle stages (i.e. coal mining, coal stockpiling, and coal combustion), and the two major supporting processes (i.e. ash and effluent disposal, and water treatment). Each of these five modules are comprised of foreground and background processes. The foreground processes are further broken down into major process steps, where applicable. The inventory is calculated at this disaggregated level to aid transparency and allow for easier model checking and/or improvement assessment. The models are based on mass-balance principles, and essentially consist of process data-derived splitter functions that partition inputs into the relevant output streams. Outputs that cannot be directly related to an input via a mass balance, are specified with process factors (sometimes called “technosphere coefficients” e.g. Maurice et al., 2000), which relate the emission to a relevant process metric, e.g. time, area, flow etc.

E.1.1. Functional Unit

The balances are calculated for a specified time interval (with a default of one year). The functional unit is therefore essentially a time period of operation. This is chosen because not all environmental interventions can be directly related to a product-related functional unit, such as the power generated or coal burnt, e.g. seepage from effluent containment dams, and leachate from waste dumps, occur regardless of whether power or coal is produced in that particular time interval. However, for comparative purposes, and to present the data in a useable basis (LCI energy data is usually presented on a per MJ or per MWh basis), a product-related functional unit is required. The LCI calculated for the specified time period is therefore normalised to the net power output. The power output is related to the time of operation by the capacity of the generating units and by specifying an average load factor (i.e. the fraction of the total time available for which the power station was operating, specified as the equivalent percentage of time that the power station was operating at full load). The time period for which the load factor is specified must correspond to the time period chosen for the functional unit.

E.1.2. Time Dependent Emissions

Area dependent emissions will change from year to year, e.g. changing mine areas and waste dump/dam areas (virgin ground decreases and rehabilitated ground increases over the life of the mine and power station). These time-dependent emissions are related back to the time period of interest by calculating the total potential emissions over the entire time span for which they will be generated, and then averaging them over the life of the mine / power station. Area dependent emissions (e.g. dust, and water emissions from leachate and run-off) are given by:

$$m_j = \frac{\sum_{i=1}^L (E_j \times A_{ij})}{L}$$

where m_j is the time dependent emission for area j , L is the expected operating life of the mine or power station, E_j is the emission factor for area j in kg/m^2 and A_{ij} is area j transformed for time period i in m^2 . The total time dependent emission is given by:

$$M = m_1 + \dots + m_j + \dots + m_n$$

where the areas $j = 1$ to n contribute to that emission. For example, in the mining module, rain water run-off arises from the virgin ground, the unrehabilitated spoils, the rehabilitated spoils and the discard dump. Each area has a different emission factor (i.e. different value of E_j), as the percentage of rain water forming surface run-off depends on the vegetation covering the ground and the compaction of the ground, and each is changing over time (i.e. has a different value of A_{ij} for each time period i). The total surface run-off is determined by summing that from the different land areas. A time period of a year is used for the value of i (i.e. L is the total number of years of operation), as a smaller time interval is felt not to be warranted at the level of data accuracy and availability of the study. The average yearly emission is subsequently converted to the time period of interest, as given by the functional unit.

E.1.3. Steady state vs. Time Dependent

An inconsistency arises in the model because the modules predict steady state emissions, other than the waste management modules, which predict time-dependent emissions, i.e. the solid waste flows are calculated on a annual steady state basis, but because they are cumulative, the emissions they give rise to are time-dependent. These emissions are related back to a pseudo steady state emission, consistent with the direct emissions (aerial emissions from stacks, direct effluent discharges etc.), by averaging them over the operating life of the system (see above section). However, the mix of steady state and time dependent modules create some problems within the models. Feedback loops from waste management (e.g. effluent returned from the containment dams to the water plant) have to be based on the value averaged over the life of the operation (e.g. cannot be related to the actual volume of water in the dam). This is particularly problematic for the early years of operation, where the relationship between the waste management and process modules is likely to far from steady state (e.g. in the initial years of operation, the waste volume will be small, contaminated surface volumes will be low, and the effluent return stream will be small, if at all. This will gradually increase, as the waste dump area increases and consequent run-off effluent volumes increase).

The initial years of operation will be at low loads (e.g. not all units commissioned yet, more downtime as initial problems are sorted out etc.). Modelling these early flows at the kind of average loads obtainable later in the plant's life are therefore inappropriate. The waste dump areas are able to be adjusted to some degree for these initial low load years, by allowing the cumulative mass or area to be specified at a particular year. Steady state additions of mass are assumed only after that year. However, whilst allowing for more accurate dump profiles, this creates inconsistencies with the process modules. The steady state waste volumes are not the volumes reflected in the dumps for these early years, which causes problems where waste flows and process flows interact (e.g. in the co-disposal of effluent with solid waste, the effluent volume is inconsistent with the mass of solid waste, and artificially high excess effluent volumes are calculated). The obvious solution to this problem is to also calculate the process modules (mine and power station processes) on a time dependent basis. However, this would considerably increase the size and complexity of the models. It increases the data requirements,

because many data inputs now have to be specified per year of operation, instead of a single average value. This type of data is typically not easily available. Whilst current operating data can be monitored, past process conditions will most likely have to be estimated, and future operating data predicted.

The platform in which the inventory model is built (Analytica[®]) is easily able to handle a time-dependent inventory. The model can simply be adjusted from a single data input to an array, indexed by the time spanned by the operation. However, the size of the models increase considerably, as calculations are performed not once, but for every year of operation. Whilst this is not too significant for a mid-point inventory, if a consideration of uncertainty is included and the data entered as probability distributions, the computing time and power required to calculate the output is greatly increased. Also, iterative calculations are required for modelling the burdens associated with the provision of electricity to mining, and to model the cyclic nature of certain water plant processes. Modelling time-dependent iterative processes would require considerable restructuring of the model, due to the way in Analytica[®] handles iterative calculations. Within the current structure of the LCI model, time dependent emissions can thus only be handled where the simplified water model is used, and mining power burdens are assumed supplied by an average grid inventory.

E.1.4. Processes and Environmental Interventions Considered

Lack of transparency is a significant barrier to the usefulness of LCI results. With many LCIs, what is not considered can almost be more significant than what is included in the inventory. Data and modelling constraints inevitably mean that not all environmental interventions will be reflected in the inventory, whilst simplifications may result in not all processes giving rise to a certain intervention being included. The following sections, and the presentation of the model and data inputs in appendix F, aim to present a clear picture of what is included in the inventory model used to calculate the LCI presented in chapter 6, and to generate the results for the case studies in chapters 6 and 7.

Within the processes considered, there is still the problem of data gaps. These can be seen in the listing of data inputs presented in the spreadsheets in appendix F. A problem with the Analytica[®] model is that a zero value is required for a blank input field, thus a value of zero in the input field does not necessarily mean that the value is zero, but that no value is known for that input. It is therefore important that the data entries be adequately annotated, so that data gaps are not hidden. Data gaps are a problem with the water quality analyses input into the model, in that the power stations and mines do not always monitor for the same set of elements. Thus the few elements that are able to be included are not always available for both inventories (or all processes), and the combined inventory underestimates the total emissions of such elements. This problem arises because it was decided that as complete an intervention list as possible should be drawn up for the SA electricity LCI. This LCI is calculated using single point data entries, and in the relatively rare instances when no information can be found to support the choice of input value, it is set to zero. The inventories used to support the case studies in chapters 7 and 8, however, are based on each input value's possible range. This allows data gaps to be addressed by assigning an approximate value and large uncertainty bounds to entries based on poor or very little information, and assessing the effect of the data input on the overall inventory uncertainty.

Three different databases were used to provide generic LCI information for the background processes. The ESU database consists of an extremely detailed intervention list, whilst the APME and IDEA inventories are considerably less detailed (see appendix F for a listing of background data used). The different databases

contain some incompatible inventory categories due to the higher degree of aggregation of some of the APME categories. These were added to the ETH intervention list to create the background list given in Table E-1. A substantially less detailed inventory was able to be generated for the foreground processes, owing to data constraints (see Table E-1). In the inventories generated for the case studies in chapters 7 and 8, the intervention list is shortened to include only those environmental interventions that contribute to a defined impact category (i.e. are assigned an equivalency factor in the EI 99 impact assessment method). This also substantially decreases the number of environmental interventions considered, particularly around emissions to water (see Table E-1). The lack of an indicator of emissions to water (no correspondence between the columns in Table E-1) prompted the inclusion of waterborne sulphate emissions into the intervention list of chapters 7 and 8 (see section 5.1.3).

Table E-1 Environmental intervention list

Background List ¹	Also in foreground	Has equiv. factor	Background List ¹	Also in foreground	Has equiv. factor
Resources			Emissions to Air cntd.		
process energy (biofuel)			Butene		✓
process energy (hard coal)	✓		Ethane		✓
process energy (hydro power)			Ethanol		✓
process energy (lignite)			Ethene		✓
process energy (natural gas)			Ethylbenzene		✓
process energy (nuclear)			Ethylene dichloride		✓
process energy (oil)	✓		Ethyne		✓
land transformation (II-III)	✓	✓	Formaldehyde		✓
land transformation (II-IV)	✓	✓	Halogenated hydrocarbons		✓
land transformation (III-IV)		✓	Halons		✓
land use (IV-IV)		✓	Heptane		✓
gas reserves		✓	Hexafluoroethane		✓
hard coal reserves	✓	✓	Hexane		✓
lignite reserves		✓	Hydrocarbons (unspecified) ²		✓
oil reserves		✓	Methane	✓	✓
biotic products used (dry weight)			Methanol		✓
Barite			methyl tert-butyl ether		✓
Bauxite		✓	non methane VOC	✓	✓
Bentonite			Polyaromatic hydrocarbons	✓	✓
Chromium reserves		✓	Pentane		✓
copper reserves		✓	Phenol		✓
iron reserves		✓	Propane		✓
lead reserves		✓	Propene		✓
Limestone		✓	Propionaldehyde		✓
Manganese reserves		✓	Propionic acid		✓
nickel reserves		✓	Tetrachloride-dibenzo-dioxin		✓
silver reserves			Tetrafluoromethane		✓
sodium chloride			Toluene		✓
Sulphur reserves			Xylene		✓
Tin reserves		✓	Al		✓
Uranium reserves			As	✓	✓
Zeolite			Ba	✓	
Zinc reserves		✓	Ca		
Water	✓	✓	Cd		✓
Net air			Co		
Other inputs			Cr	✓	✓
Emissions to Air			Cu	✓	✓
Br ₂			Fe		
Cl ₂ ²			Hg		✓
CO	✓	✓	K		
CO ₂	✓	✓	La		
F ₂ ²			Mg		
HCl	✓	✓	Mn	✓	✓
He			Mo		
HF	✓	✓	Na		
H ₂ S			Ni	✓	✓
Hydrogen			Pb	✓	✓
N ₂			Pt		
NH ₃		✓	Sn		
N ₂ O	✓	✓	Sr		
NO _x	✓	✓	Th		
SO ₂	✓	✓	Ti		
Acetaldehyde		✓	U	✓	
acetic acid		✓	V	✓	
Acetone		✓	Zn	✓	✓
Acrolein		✓	Zr	✓	
Aldehydes		✓	metals (unspecified) ²		✓
Alkanes		✓	B	✓	
Alkenes		✓	P	✓	
Aromatics		✓	Se		
Benzaldehyde		✓	Si		
Benzene		✓	Total suspended particulates	✓	✓
benzo(a)pyrene		✓	steam/water vapour	✓	
Butane		✓			

Background List ¹	Also in foreground	Has equiv. factor	Background List ¹	Also in foreground	Has equiv. factor
<i>Emissions to Water</i> waste water	✓		<i>Emissions to Water contd.</i> Halogenated hydrocarbons (waterborne)		
Al (waterborne)	✓		Hydrocarbons (unspecified) (waterborne)		
As (waterborne)		✓	methyl tert-butyl ether (waterborne)		
Ba (waterborne)			organic compounds (unspecified) (waterborne)		
Be (waterborne)			polyaromatic hydrocarbons (waterborne)		✓
Ca (waterborne)			oils & greases (waterborne)		
Cd (waterborne)		✓	Phenols (unspecified) (waterborne)		
cobalt (waterborne)		✓	Sn (tributyl) (waterborne)		
Cr-III (waterborne)		✓	Toluene (waterborne)		✓
Cr-VI (waterborne)		✓	Trichloroethylene (waterborne)		
Cu (waterborne)		✓	Xylene (waterborne)		
Fe (waterborne)	✓		Salts (unspecified) (waterborne)		
Hg (waterborne)		✓	Unsolved substances (waterborne)		
K (waterborne)			Acid as H ⁺ (waterborne)		
Mg (waterborne)			TDS	✓	
Mn (waterborne)	✓		TSS		
Mo (waterborne)			<i>Wastes</i>		
Na (waterborne)	✓		Waste oil		
Ni (waterborne)		✓	Solid waste ²	✓	
Pb (waterborne)		✓	Hazardous waste ²	✓	
Rb (waterborne)					
Se (waterborne)					
Sn (waterborne)					
Sr (waterborne)					
Ti (waterborne)					
V (waterborne)					
Zn (waterborne)		✓			
metals (unspecified) (waterborne) ²					
Ammonia as N (waterborne)					
borates / borides as B (waterborne)					
barite (waterborne)					
Chlorides (waterborne)	✓				
Cyanides (waterborne)	✓				
Fluorides (waterborne)					
H ₂ S (waterborne)					
Hypochlorous acid (waterborne)					
iodides (waterborne)					
nitrates (waterborne)					
Nitrogen (organic bonded) (waterborne)					
Nitrogenous compounds (unspecified) (waterborne)					
P (waterborne)					
Phosphates (waterborne)					
Sulphates (waterborne)	✓				
Sulphides (waterborne)					
sulphite (waterborne)					
alkanes (waterborne)					
alkenes (waterborne)					
Aromatics (waterborne)					
Benzene (waterborne)		✓			
Chlorobenzene (waterborne)		✓			
Ethylbenzene (waterborne)					
fatty acids (waterborne)					
Formaldehyde (waterborne)		✓			
glutaric aldehyde (waterborne)					

1. Environmental interventions included in ESU database. The list was slightly reduced by omitting those flows that contribute less than 10⁻⁶ to the functional unit, as well as less than 10⁻⁶ to the impact category score (the "CML factors" (Heijungs et al., 1992) were used to perform this calculation).
2. Environmental interventions added from the APME database.

APPENDIX G

SUPPORTING INFORMATION FOR SA COAL-FIRED ELECTRICITY LCI

G.2. ANALYSIS OF THE SA ELECTRICITY INVENTORIES

This section highlights some key features of the inventories presented in section 6.3.

G.2.1. Foreground / Background Contributions

Table G-1 presents the percentage contributions from the foreground and background systems to the overall inventory for some key environmental interventions. The huge mass of coal combusted during power generation, and the ensuing large volume of flue gas, results in power generation, and thus the foreground system, dominating practically all air emissions. The only exception is non methane VOCs, because of the high contribution from diesel and fuel oil production to these emissions. Methane also shows a high contribution from these processes, although the total methane is dominated by that produced during mining. As would be expected, the power station and mining processes (foreground) completely dominate land transformation, solid waste mass and water use, because of the very large scale on which these processes are carried out.

The ESU database used to provide generic LCI data for some of the background processes has more detailed environmental interventions list than it was possible to calculate for the foreground processes, and compared to those of the other databases used (see Table E-1 for a full listing of the environmental interventions considered in the foreground and background). These inventory items obviously show 100% contribution from the background processes, although the numerical values in these categories are mostly very small. Exceptions are in the process energy categories (those other than coal), the mineral reserves categories (also in all those other than coal), and a few of the emissions to water. The background system's dominance in the energy categories other than coal is to be expected, as only coal-derived energy is included in the foreground system. This simplification is justified by the relatively small contribution of nuclear, gas and oil to the overall generating mix in South Africa, and the location of the processes under consideration. Thus any nuclear energy reflected in the inventory is due to the use of generic LCI data for background processes, in which the energy mix differs markedly from that of South Africa. The procurement of all fuels and materials used in the foreground processes, other than coal (i.e. fuel oil, diesel, gas and chemicals), are consigned to the background system. Thus all mineral and fuel reserves show 100% contribution from the background system.

A few of the water emission categories show relatively large contributions from the background system. In some cases this is because of difficulties in reconciling the foreground system to the background system (e.g. sulphates and nitrates are recorded in the foreground system data, whilst sulphites and nitrogen (organic bonded) are recorded in the background inventories). Other water components with fairly high values in the background system, but not recorded in the foreground system, include Sr, barite, and oils & greases. Significant contributions from the background system relative to the foreground system are present for Na and chlorides, arising from the high values in the oil and diesel LCIs. This high relative background contribution is an overestimate, as for some foreground effluent streams, notably leachate from ash dumps, coal stockpiles and discard dumps, only sulphates and TDS emissions were able to be estimated.

Table G-1 % Contributions from the foreground and background processes to the overall inventory.

	Wet / Wet		Dry / Dry		Wet / Dry	
	Fore	Back	Fore	Back	Fore	Back
Process energy (hard coal)	99.4%	0.6%	99.8%	0.2%	99.0%	1.0%
Process energy (natural gas)	0.0%	100.0%	0.0%	100.0%	0.0%	100.0%
Process energy (nuclear)	0.0%	100.0%	0.0%	100.0%	0.0%	100.0%
Process energy (oil)	90.3%	9.7%	90.1%	9.9%	90.4%	9.6%
Land transformation (II-III)	99.7%	0.3%	97.1%	2.9%	99.5%	0.5%
Land transformation (III-IV)	99.9%	0.1%	99.9%	0.1%	99.9%	0.1%
Water	99.4%	0.6%	96.8%	3.2%	98.9%	1.1%
CO	99.6%	0.4%	99.1%	0.9%	99.4%	0.6%
CO ₂	99.9%	0.1%	100.0%	0.0%	99.9%	0.1%
HCl	98.8%	1.2%	0.0%	100.0%	0.0%	100.0%
HF	99.8%	0.2%	0.0%	100.0%	0.0%	100.0%
N ₂ O	97.8%	2.2%	97.0%	3.0%	97.8%	2.2%
NO _x	99.8%	0.2%	99.9%	0.1%	99.8%	0.2%
SO ₂	99.9%	0.1%	100.0%	0.0%	99.9%	0.1%
Methane	88.5%	11.5%	97.8%	2.2%	75.6%	24.4%
Non methane VOC	43.0%	57.0%	35.8%	64.2%	43.8%	56.2%
Total suspended particulates	99.6%	0.4%	99.9%	0.1%	97.0%	3.0%
Waste water	97.2%	2.8%	98.1%	1.9%	84.2%	15.8%
Na (waterborne)	50.3%	49.7%	14.8%	85.2%	59.7%	40.3%
Chlorides (waterborne)	46.9%	53.1%	5.6%	94.4%	11.1%	88.9%
Fluorides (waterborne)	99.9%	0.1%	57.5%	42.5%	98.9%	1.1%
Sulphates (waterborne)	100.0%	0.0%	100.0%	0.0%	99.9%	0.1%
Sulphite (waterborne)	0.0%	100.0%	0.0%	100.0%	0.0%	100.0%
TDS	99.9%	0.1%	100.0%	0.0%	99.8%	0.2%
Solid waste	100.0%	0.0%	100.0%	0.0%	100.0%	0.0%

Unfortunately, the background inventories do not reliably report sulphates and TDS (the TDS value reported is very much smaller than that of the individual salts), so the actual contribution of the background processes to the total water emissions can not be assessed. The contribution of the background processes to total waste water volume is fairly small, at just under 3%. However, this is also an unreliable metric, as the definition of “waste water” is open to interpretation. Surface run-off is included in the foreground, although it is likely that only process water is included in the background inventories.

G.2.2. Relative Contributions of the Life Cycle Stages

The relative contributions of the individual stages to the overall inventory are shown in Figure G-1 for selected environmental interventions.

G.2.2.a Process Energy

The foreground electricity requirements (as reflected by “process energy (hard coal)”) are shared more or less equally between mining and power generation. Only overall power station energy use data was available, so energy use for the boiler plant also incorporates the auxiliary power requirements of the water plant and the ash plant. Thus for “process energy (hard coal)”, Figure G-1 should be read as “power plant” and not “boiler plant” shared equally with mining. Similarly electricity use by the stockpile is included either with mining or the power

plant, depending on whether it is managed by the power station or the mine. All electricity used in the foreground is assumed to be 100% coal based. The “process energy (nuclear)” category therefore only reflects background power use. The background coal-based energy contribution is small compared to that in the foreground, so the foreground dominates the relative contributions of the life cycle stages seen for “process energy (hard coal)”. The various stages’ contributions to background electricity use follow the same trends as that of liquid fuel use (as reflected in “process energy (oil)”), except with a slightly higher contribution from the water plant, reflecting the additional energy used in the production of water treatment chemicals. The life cycle stages which are high consumers of heavy fuel oil and diesel, and to a lesser extent ancillary materials, therefore also show high background electricity requirements.

Liquid fuel consumption is dominated by mining, with opencast mining methods showing a larger contribution than underground mining, due to their higher diesel use. The boiler plant is the next highest contributor, primarily as a result of its use of heavy fuel oil during boiler start-ups and during low load conditions. A small contribution from the water plant is apparent, mostly due to the diesel requirements for transport of ancillary materials to the plant. The Wet/Dry plant has the highest treatment chemical requirements, and thus shows a higher water plant contribution to oil-derived process energy than the other two stations.

G.2.2.b Land Use

The first land use category (II-III) incorporates land that has been transformed from its original state, but is not built or paved over, and includes grassed over areas, rehabilitated dumps and spoils, and the areas covered by containment dams and pans. Mining and ash disposal are the main contributors, with the very large areas disturbed and rehabilitated during opencast mining dominating the stations supplied by opencast mines. Very much smaller areas are disturbed during underground mining, and the area of the rehabilitated ash dump is the largest contributor to this category for stations supplied by an underground mine. The second land use category (II-IV) incorporates land that has been severely degraded, and includes built and paved over areas (power station terrace, mine preparation plant etc.), coal stockpiles, mining pit area, and unrehabilitated ash dumps and spoils. The very large areas affected by opencast mining methods results in mining once again dominating the land use category for the stations supplied by opencast mines. In these stations, the next largest contributor is the unrehabilitated ash dump, with far smaller contributions by the water plant area, boiler plant area and coal stockpile. Underground mining disturbs a far lower area than the opencast methods, with only the plant area and discard dump area contributing to this category. The area occupied by the unrehabilitated ash dump, power station (“boiler plant” reflects the combined areas of the boiler plant and water plant for this station) and mine contribute about equally to land degradation in stations supplied by underground mines, with a slightly smaller contribution from the stockpile area.




Figure G-1 Percentage contributions from individual life cycle stages to overall inventory.

G.2.2.c Resource Use

The huge volumes of water lost in the cooling towers of the wet cooled stations result in the water plant dominating the water use category for wet cooling stations. A small contribution from mining is also evident, especially the opencast mining methods, in which large volumes of water collect in the pit and require containment. Stormwater and seepage collected and used or evaporated on site are included in the water use category, as this water is assumed to have been effectively removed from the water catchment of the area, where it could potentially have been used by another process. The dry cooling station uses very much less water, and thus the water plant and mine show approximately equal contributions to water use. Stormwater run-off from the ash dump and coal stockpile contribute to a small degree.

Fuel and mineral reserve depletion (other than coal) are all included in the background and are not shown here, as their distribution between the stages is self explanatory, i.e. wherever they are consumed (e.g. sulphur reserves are dominated by the water plant, where sulphuric acid is consumed). Oil reserves mirror the “process energy (oil)” category, with the life cycle stages with high fuel oil and diesel requirements dominating (i.e. mining and boiler plant).

G.2.2.d Air Emissions

The very large mass of coal burnt, and consequently huge volume of flue gas, results in the boiler plant dominating practically all the air emission categories. Contributions from the combustion of diesel fuel during mining are only evident for the lower volume power station flue gases (e.g. CO and NO_x). A notable exception is methane, for which mining dominates comprehensively. Emissions of non-methane VOCs and nitrous oxide were not estimated for the boiler flue gas, and are present only from fuel oil and diesel combustion and production, and the stage consuming the most liquid fuel therefore dominates (mining). Particulate emissions are primarily contributed by the boiler plant, as the large volume of ash produced means that large volumes of particulates are emitted even at high particulate removal efficiencies. Opencast mining is associated with a fair degree of dust, and thus shows significant contributions to the overall particulate loading of the stations they supply. Underground mining shows a far smaller contribution, allowing the much smaller contributions from ash disposal and the coal stockpile to be just evident for the Dry/Dry station. Dry ash disposal poses more of a dust problem than wet ash disposal, with wet ash disposal showing a negligible contribution to overall particulate volumes. Background processes generally do not contribute much to the overall particulate volume, other than for the Wet/Dry station, for which the contribution from particulates associated with chemical manufacture (predominantly due to limestone quarrying) are just evident.

G.2.2.e Water Emissions

The waste water category is assigned only effluent that is actually emitted, and as the power stations all work on a system of ZLED, the waste water volumes are dominated by the mines. The mines also work on the principle of containing as much effluent on site as possible, and although dominating the category, the actual volumes of water emitted are fairly small. Data gaps exist for many of the components considered under emissions to water, since the mine and power station do not always monitor for the same components, or even for the same components in all streams. Sulphates and TDS are the most uniformly monitored and are thus the most reliable categories for an indication of contaminants emitted from the foreground processes to water (to both surface and underground water bodies, calculated according to the simplified methodology explained in section 6.2.6). Both are dominated by mining to a very large degree. For the Dry/Dry station, the slightly lower mining contribution (underground mining has lower TDS and sulphate emissions than opencast mining) allows for a small

contribution from the coal stockpile to be evident. The Wet/Dry station shows a small contribution from ashing and a slightly lower mining contribution than the Wet/Wet station, also supplied by an opencast mine.

G.2.2.f Solid Waste

For most of the power stations, the solid waste volume generated by the system is dominated by the ash produced in the boilers. The stations burn high ash coals, and the consequent volume of ash produced is considerable. The stations are predominantly supplied by mines employing very limited coal preparation, and consequently little or no discard coal is produced. In addition, only discard that is dumped is recorded as solid waste in the inventory, and the crushing discard used to make ramps or as backfill on opencast mines is not included. However, the mine supplying the Wet/Wet station is a dual product mine, and produces a washing discard. This discard contributes just under half of the total waste produced by this system, and is unusually high for most mine/power station systems. Only a very small portion of the ash produced is sold for use in cement manufacture (approximately 5%), so this has a negligible effect on the dominance of the boiler plant in solid waste production.

G.3. UNCERTAINTY CONSIDERATIONS

This section expands on the information presented in section 6.3.2. It provides details on the estimation of the future coal-fired electricity demand, and additional explanation of the uncertainty analysis results.

G.3.1. Uncertainty in the Future Coal-Fired Electricity Demand

Although clearly linked to the economic growth rate of the country, predictions in the growth in electricity demand will always be highly uncertain. A relatively large range in growth rate was thus assumed for the next 10 years (30 to 75%), based on the prediction of a 50% growth rate between 2000 and 2015 by (Lennon, 1997). This figure is assumed to include any demand-side management measures implemented to slow this growth. A simple triangular distribution was chosen to model the future base-load demand.

Only the coal-fired component of this energy is relevant to the study, and thus the percentage share of coal-fired electricity to the total electricity produced needs to be estimated. Nuclear power makes up 4.8% of Eskom's total installed capacity, although it contributed 7.4% of the total power generated in 1998 (Eskom, 1999a). Other sources (gas turbines, pumped storage etc.) are purely for peaking supply (see Table 6-1), so are not considered. In addition, Eskom purchased 1.5% of the total power it sold in 1998. Possible non-coal energy expansion could be increased nuclear capacity, and/or increased purchases from the SAPP, which has considerable hydro power potential. A possibility of gas-fired stations also exists. A uniform range of 5-20% non-coal supply in 10 years time was estimated (i.e. an increase of about 3% from current levels). Coupled with the uncertainty in supply options is the uncertainty in transmission losses. As the inventory is calculated from the actual coal-fired power sent out, an estimate of transmission losses needs to be included. A simple uniform range of 4-8% was chosen. The median value (5.6%) lies between that of 1997 and 1998 (6.4 and 4.5% respectively) (Eskom, 1999a), and reflect the possibility of higher values should power be sourced from further afield.

A simulation using Latin Hypercube sampling of the uncertainty estimates given above, and a current demand of 175 TWh, predicted the 2010 coal-fired electricity demand to lie in a 90% confidence interval of 200-240TWh. The aim of the analysis is to determine whether additional capacity is required, and not the electricity demand per se, since the inventory is presented per MWhSO. Figure G-2 compares the demand required with the supply from installed capacity, which does not include plants currently in storage, but assumes plants on order have

been completed. An overall load factor of 70% was assumed for the current operating stations, with a maximum possible load of 80%, whilst for new plants these were assumed to be 78% and 85% respectively. Minimum load factors of 65% and 70% were assumed for the current operating stations and new stations respectively. Considering the range of load factors (varied uniformly) and the variable electricity demand, there is only a very low probability that the current installed capacity will cover the demand. Figure G-2 shows that most probably around 20 000GWh additional power will be required.

However, Eskom has considerable excess capacity in storage (see Table 6-2). Figure G-3 shows the probability whether new capacity will be required as a function of the percentage of plants in storage brought back into operation. Considering their age and reliability, a fairly low load factor of 60% was assumed for the plants in storage (varied uniformly between 50 and 70%). Figure G-3 shows that if more than 60% of the capacity in storage is brought back into operation, there is a fairly high probability (greater than 60%) that no new capacity will be required.




Figure G-2 Cumulative probability chart of coal fired electricity demand less that supplied by the currently operating capacity. Negative values mean that supply exceeds demand, whilst positive values mean that plants in storage need to be re-commissioned, or new plants built.

been completed. An overall load factor of 70% was assumed for the current operating stations, with a maximum possible load of 80%, whilst for new plants these were assumed to be 78% and 85% respectively. Minimum load factors of 65% and 70% were assumed for the current operating stations and new stations respectively. Considering the range of load factors (varied uniformly) and the variable electricity demand, there is only a very low probability that the current installed capacity will cover the demand. Figure G-2 shows that most probably around 20 000GWh additional power will be required.

However, Eskom has considerable excess capacity in storage (see Table 6-2). Figure G-3 shows the probability whether new capacity will be required as a function of the percentage of plants in storage brought back into operation. Considering their age and reliability, a fairly low load factor of 60% was assumed for the plants in storage (varied uniformly between 50 and 70%). Figure G-3 shows that if more than 60% of the capacity in storage is brought back into operation, there is a fairly high probability (greater than 60%) that no new capacity will be required.




Figure G-2 Cumulative probability chart of coal fired electricity demand less that supplied by the currently operating capacity. Negative values mean that supply exceeds demand, whilst positive values mean that plants in storage need to be re-commissioned, or new plants built.

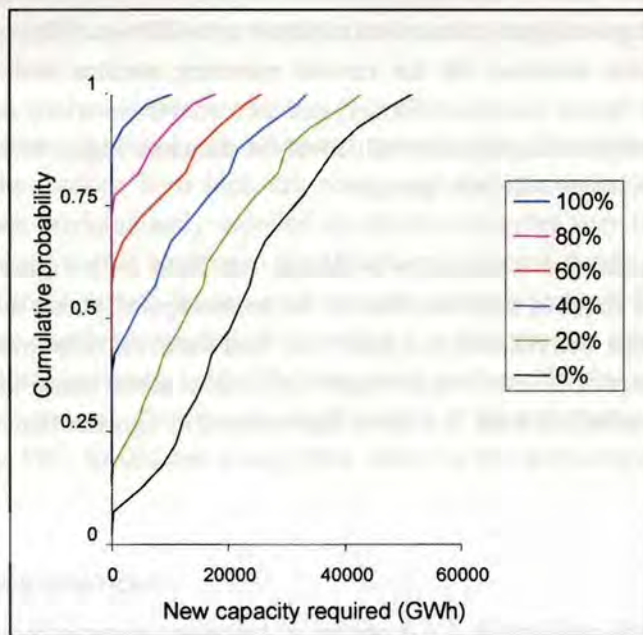


Figure G-3 Cumulative probability of supply required from new capacity, depending on the percentage of plants in storage brought back into operation.

G.3.2. Uncertainty in Generating Mix

The following two sections provide more detailed interpretations of the results of the uncertainty analysis presented in section 6.3.2.a.

G.3.2.a % Plants in Storage Re-commissioned (Figure 6-3)

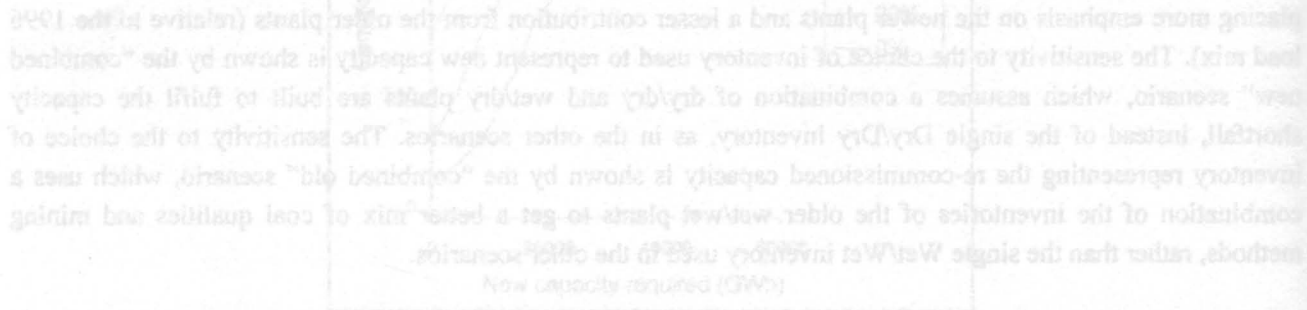
The majority of inventory items was surprisingly insensitive to the percentage of plants in storage re-commissioned, with only those environmental interventions that differ very markedly between the old and new plants significantly increasing the uncertainty range over that due to the uncertainty in future energy demand. These include water use, because of the very large difference in the water requirements of an old wet cooled plant and a modern dry cooled plant, particulate emissions, due to the superior performance of fabric filters to old electrostatic precipitators, and emissions to water (TDS and waterborne sulphate), as a result of different ash disposal and mining methods between the old and new plants.

The Dry/Dry station used to model the new capacity, burns a relatively poor quality ROM coal from an underground mine, whilst the in storage capacity, represented by the Wet/Wet station, is supplied by an opencast mine producing a relatively low-sulphur coal. This explains the slight decrease in sulphur dioxide emissions, and increase in mining-related impacts (land use and emissions to water). These trends are thus more a function of the choice of inventory than the percentage of stations re-commissioned, and are investigated in the following section. A slight increase in carbon dioxide emissions reflects the lower thermal efficiencies of the older stations.

The analysis presented here exaggerates the difference in performance between the re-commissioned stations and a possible new station, as the older plants are unlikely to be brought back into service without significant improvements in their environmental performance. This will have the effect of decreasing the sensitivity to the percentage brought out of storage to an even greater degree.

intervention categories than the assumptions governing the installed capacity inventory. A notable exception is particulates, as the older plants considered have notably poorer ESP efficiencies.

The choices governing the compilation of the future generating mix inventory are shown to introduce more uncertainty than that due to the uncertainty in energy demand, i.e. the range spanned across the scenarios in Figure 6-4 is greater than the range spanned by the individual scenarios (as represented by the “whiskers”). This is because the analysis looked at predicting the inventory of a relatively near-term generating mix. The uncertainty in the energy demand, and the choice of new plant inventory, are likely to play larger roles in a prediction further into the future.



The inventory of the future generating mix is most sensitive to the choices of new plant inventory, which is shown to be the largest contributor to the uncertainty in the inventory. The choices of new plant inventory are shown to be the largest contributor to the uncertainty in the inventory, which is shown to be the largest contributor to the uncertainty in the inventory. The choices of new plant inventory are shown to be the largest contributor to the uncertainty in the inventory, which is shown to be the largest contributor to the uncertainty in the inventory.

shows that the additional level of modeling detail is worth the extra effort and in fact the model results are significantly better than the results from the simpler model. The model results are significantly better than the results from the simpler model. The model results are significantly better than the results from the simpler model. The model results are significantly better than the results from the simpler model.

The Oxy-Dry plant is shown to have a significant impact on the inventory. The Oxy-Dry plant is shown to have a significant impact on the inventory. The Oxy-Dry plant is shown to have a significant impact on the inventory. The Oxy-Dry plant is shown to have a significant impact on the inventory. The Oxy-Dry plant is shown to have a significant impact on the inventory.

APPENDIX H

SUPPORTING INFORMATION FOR TECHNOLOGY CHOICE CASE STUDY (CHAPTER 7)

H.2. ADDITIONAL RESULTS AND EXPLANATIONS

This section presents some additional discussion and results of the case study in chapter 7. The PC tables corresponding to the PC plots in chapter 7 and in the appendix can be found in section H.2.3.

H.2.1. Selection of Impact-Level Criteria

Table H-11 shows the significance of each impact category, according to its contribution to the damage category to which it belongs. The table is only intended to give a rough indication, as the normalisation factors used have been developed for European concerns and conditions. As would be expected, given the huge volumes of CO₂ emitted, climate change shows the dominant contribution to human health damage. Summer smog shows a very low contribution, as only methane and non-methane VOCs are determined for the foreground processes. Even with its low significance, summer smog is retained as a criteria, since the human health damage and normalisation factors are extremely uncertain, and its low importance in Table H-11 may be an overestimate.

The “mineral resources” category was not retained, because in addition to a PC analysis on the output sample indicating a high degree of correlation between fossil fuels and mineral resource extraction, the significance of its contribution to resource depletion is very low (see Table H-11). This is because the most significant mineral resources being consumed by the systems are not reflected in the category. Damage factors are only provided for metallic ores, thus the significant consumption of limestone, both as a desulphurisation sorbent and for water treatment, is not reflected. Neither construction of the process, nor maintenance procedures, are included in the system boundary, so the consumption of mineral resources over the operating life of the system is small. The category is therefore dominated by the production of magnetite for the coal washing process, and burdens from the background system (predominantly the provision of diesel fuel). The high contribution of magnetite for coal washing accounts for the strong correlation between mineral resource extraction and fossil fuel consumption, as the loss of magnetite from the washing circuit is directly related to the mass of coal produced.

The effect of model uncertainty in the models underpinning the criteria to assess water-related impacts is investigated by presenting the results according to two different criteria with very different levels of model complexity. Initially the analyses were conducted with the inventory-level environmental intervention, waterborne sulphates, to reflect the water related impacts of the systems. This was subsequently replaced by the “impacted land footprint”, which relates the waste type to its potential for water contamination far more rigorously than the simple calculation of waterborne sulphates (see section 5.4.2 for details on determining the impacted land footprint). The land occupation criterion was simultaneously removed, to avoid double-counting with the impacted land footprint. The impacted land footprint only includes the area occupied by the waste dumps, but in this study all other types of land transformations are common to the systems (as they each consider a different refurbishment option of an existing plant), so the land transformation indicator can be removed. A comparison of Figure H-7 with Figure H-8 clearly shows the better separation between the systems with the footprint indicator, than with the waterborne sulphates and land indicators (the system “clouds” move further apart). This better representation of the systems is achieved because the footprint better captures the benefits of

removing the discard dumps. The decrease in scatter reflects the lower uncertainty associated with the more sophisticated impacted land footprint indicator than with the rough calculation of waterborne sulphates.

Table H-1 Relative importance of each impact's contribution to its damage category, according to the normalisation factors given in the EI 99 method (Hierarchist perspective). The results shown are for the "most likely" scenario of each system.

	PF	FBC
% Contribution to Damage to Human Health (DALYs)		
Carcinogenic effects on humans	10%	36%
Respiratory effects on humans caused by organic substances (summer smog)	0.06%	0.04%
Respiratory effects on humans caused by inorganic substances (winter smog)	24%	21%
Climate change	67%	43%
% Contribution to Ecosystem Quality (PDF.m².yr)		
Ecotoxicity	33%	36%
Combined effect of acidification and eutrophication	24%	31%
Land occupation	42%	33%
% Contribution to Resource Depletion (MJ surplus energy)		
Extraction of minerals	0.2%	0.2%
Extraction of fossil fuels	99.8%	99.8%




Figure H-1 Output samples (% difference between scenarios and PF “most likely” scenario) for EI 99 damage categories plus waterborne sulphates and water use (see Table H-16 for PC loadings).




Figure H-2 Output samples (% difference between scenarios and PF “most likely” scenario) with land occupation and waterborne sulphates replaced by impacted land footprint (see Table H-17 for PC loadings).

H.2.2. Additional Principal Component Plots




Figure H-3 “Best” and “most likely” operating “spaces”, derived from ratios between the output samples and the PF “most likely” scenario transformed on the PC plane (see Table H-18 for PC loadings).





Figure H-4 “Best” and “most likely” operating “spaces”, according to selected environmental interventions and the impacted land footprint (in original units), transformed on the PC plane (see Table H-19 for PC loadings).

Change in % SO₂ Removal and Sorbent Type



Change in Discard Quality and Allocation of Mining Burdens to Discard




Figure H-5 FBC options, considering sorbent and discard source, degree of desulfurisation and cooling technology. Derived from the % difference between the scenarios and the PF scenario transformed on the 1st and 3rd principal component plane (see Table H-15 for corresponding PC loadings).

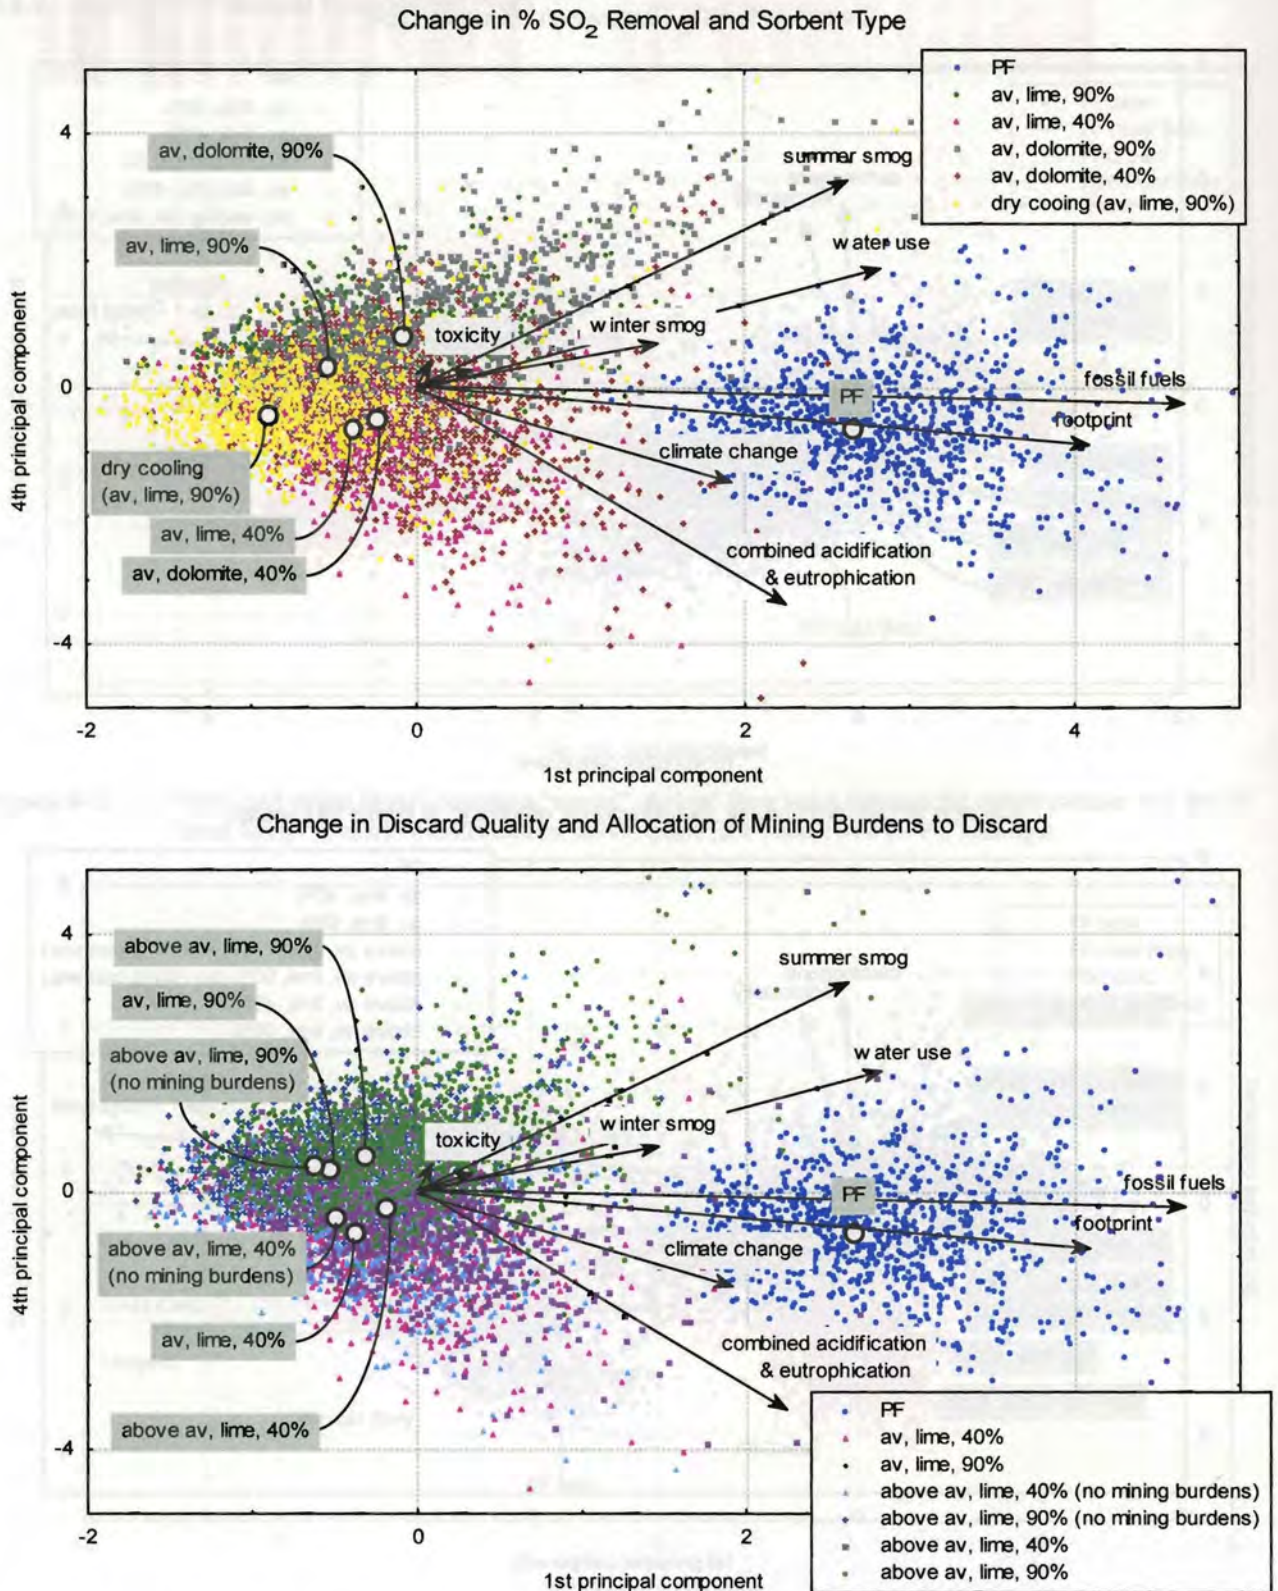


Figure H-6 FBC options, considering sorbent and discard source, degree of desulphurisation and cooling technology. Derived from the % difference between the scenarios and the PF scenario transformed on the 1st and 4th principal component plane (see Table H-15 for corresponding PC loadings).

H.2.3. Principal Component Tables

Table H-2 Principal component loadings corresponding to Figure 7-3.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9
Carcinogenic effects	0.13	0.22	0.74	-0.60	0.13	0.11	0.00	-0.02	0.00
Summer smog	0.67	0.01	-0.12	0.02	-0.06	0.45	-0.57	0.08	0.02
Winter smog	0.02	0.79	-0.17	0.12	0.22	0.32	0.29	0.31	0.00
Climate change	0.53	0.44	-0.12	-0.16	-0.21	-0.57	-0.18	0.28	0.01
Ecotoxic effects	0.20	0.24	0.70	0.63	-0.11	-0.08	-0.06	-0.04	0.00
Acidification and eutrophication	0.57	0.52	-0.22	-0.05	-0.14	0.00	0.10	-0.57	0.02
Fossil fuel consumption	0.89	-0.30	0.02	0.00	-0.12	0.06	0.23	0.07	-0.20
Raw water use	0.55	-0.11	-0.04	0.10	0.78	-0.22	-0.08	-0.08	0.01
Impacted land footprint	0.83	-0.37	0.05	0.00	-0.12	0.07	0.32	0.12	0.18
Explained variance	2.90	1.44	1.14	0.81	0.78	0.70	0.62	0.53	0.07
% of total variance	32%	16%	13%	9%	9%	8%	7%	6%	1%

Table H-3 Principal component loadings corresponding to Figure 7-4.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9
Extraction of fossil fuels	-0.95	0.08	-0.01	0.00	0.00	-0.04	-0.18	0.18	-0.17
Water use	-0.54	-0.09	-0.11	-0.06	-0.74	0.27	0.25	-0.04	0.01
NO _x	-0.84	-0.09	-0.10	-0.08	0.03	-0.20	-0.28	-0.39	0.02
SO ₂	0.35	-0.58	0.48	-0.34	-0.25	-0.02	-0.37	0.08	0.01
Pb	-0.53	-0.22	0.30	-0.17	0.41	0.60	0.14	-0.06	0.00
CO ₂	-0.37	-0.66	-0.23	-0.25	0.17	-0.38	0.36	0.08	0.01
TSP	0.36	-0.34	-0.75	-0.02	0.05	0.33	-0.27	0.04	0.00
Cr	-0.13	-0.50	0.16	0.84	-0.02	0.01	-0.02	-0.01	0.00
Impacted land footprint	-0.92	0.16	-0.04	0.05	0.03	-0.01	-0.18	0.25	0.15
Explained variance	3.43	1.24	0.98	0.92	0.81	0.73	0.57	0.26	0.05
% of total variance	38%	14%	11%	10%	9%	8%	6%	3%	1%

Table H-4 Principal component loadings corresponding to Figure 7-5.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9
Carcinogenic effects	0.12	0.74	0.22	0.05	-0.62	0.00	0.06	-0.03	0.00
Summer smog	0.67	-0.11	0.05	0.15	0.04	-0.04	0.69	0.18	0.05
Winter smog	0.32	0.04	-0.56	-0.70	-0.15	0.23	0.06	0.09	0.00
Climate change	0.57	0.21	-0.44	0.39	0.06	-0.08	-0.30	0.43	0.02
Ecotoxic effects	0.19	0.69	0.23	-0.22	0.62	0.05	0.02	0.00	0.00
Acidification and eutrophication	0.63	0.13	-0.47	0.24	0.07	-0.09	0.03	-0.55	0.02
Fossil fuel consumption	0.88	-0.16	0.26	-0.12	-0.04	-0.18	-0.11	0.00	-0.26
Raw water use	0.53	-0.14	0.23	0.18	-0.01	0.77	-0.11	-0.05	0.02
Impacted land footprint	0.77	-0.19	0.35	-0.25	-0.07	-0.26	-0.24	-0.03	0.22
Explained variance	2.97	1.18	1.07	0.89	0.81	0.77	0.66	0.53	0.12
% of total variance	33%	13%	12%	10%	9%	9%	7%	6%	1%

Table H-5 Principal component loadings corresponding to Figure 7-10, and Figures H-11 and 12.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9
Carcinogenic effects	0.02	0.53	0.65	0.07	0.13	0.07	-0.02	-0.52	0.00
Summer smog	0.48	0.23	-0.20	0.59	-0.43	-0.04	-0.37	-0.04	0.06
Winter smog	0.27	0.47	-0.45	0.13	0.17	0.63	0.25	0.02	-0.01
Climate change	0.35	0.53	-0.21	-0.27	-0.39	-0.37	0.43	-0.06	0.01
Ecotoxic effects	0.06	0.59	0.59	0.05	0.03	0.01	-0.03	0.54	0.00
Acidification and eutrophication	0.40	0.37	-0.27	-0.60	0.18	-0.02	-0.48	-0.02	0.01
Fossil fuel consumption	0.84	-0.36	0.23	-0.05	-0.10	0.09	0.03	0.00	-0.29
Raw water use	0.51	0.04	-0.17	0.34	0.64	-0.42	0.10	0.01	0.01
Impacted land footprint	0.75	-0.45	0.31	-0.16	-0.03	0.16	0.12	0.01	0.27
Explained variance	2.13	1.65	1.31	0.96	0.84	0.75	0.64	0.57	0.16
% of total variance	24%	18%	15%	11%	9%	8%	7%	6%	2%

Table H-6 Principal component loadings corresponding to Figure H-7.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9	Factor 10
Carcinogenic effects	0.10	0.31	-0.71	-0.06	-0.60	0.18	0.03	-0.02	-0.01	-0.01
Summer smog	0.73	-0.07	0.05	0.00	0.03	0.13	0.57	0.33	-0.08	0.01
Winter smog	0.08	0.71	0.27	-0.34	0.15	0.46	-0.12	0.12	0.19	-0.04
Climate change	0.50	0.48	0.18	0.17	-0.20	-0.50	-0.23	0.35	-0.04	0.02
Ecotoxic effects	0.19	0.31	-0.66	-0.07	0.62	-0.20	0.01	0.00	-0.04	0.00
Acidification and eutrophication	0.59	0.47	0.27	-0.01	-0.06	-0.14	0.20	-0.52	-0.14	0.03
Land occupation	0.81	-0.31	-0.05	-0.25	0.00	0.05	-0.16	-0.04	0.14	0.37
Fossil fuel consumption	0.81	-0.20	-0.07	0.21	0.00	-0.08	-0.02	-0.11	0.43	-0.24
Waterborne sulphates	0.68	-0.31	0.00	-0.48	-0.01	0.07	-0.25	0.01	-0.32	-0.21
Raw water use	0.46	0.03	-0.02	0.72	0.13	0.42	-0.21	-0.01	-0.19	0.03
Explained variance	3.16	1.39	1.12	1.00	0.82	0.76	0.56	0.53	0.41	0.24
% of total variance	32%	14%	11%	10%	8%	8%	6%	5%	4%	2%

Table H-7 Principal component loadings corresponding to Figure H-8.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9
Carcinogenic effects	0.11	-0.25	-0.73	-0.60	0.12	0.12	-0.01	-0.03	0.00
Summer smog	0.68	-0.01	0.11	0.03	0.01	0.44	0.56	0.13	-0.02
Winter smog	0.03	-0.78	0.18	0.15	0.27	0.30	-0.32	0.25	0.00
Climate change	0.53	-0.44	0.15	-0.20	-0.28	-0.52	0.13	0.32	-0.01
Ecotoxic effects	0.20	-0.25	-0.68	0.63	-0.13	-0.11	0.08	-0.03	0.00
Acidification and eutrophication	0.58	-0.50	0.24	-0.05	-0.15	0.01	-0.03	-0.57	-0.02
Fossil fuel consumption	0.89	0.30	-0.03	0.00	-0.10	0.08	-0.24	0.05	0.20
Raw water use	0.55	0.11	0.02	0.07	0.75	-0.32	0.08	-0.07	-0.01
Impacted land footprint	0.83	0.38	-0.07	0.01	-0.11	0.09	-0.32	0.09	-0.18
Explained variance	2.93	1.43	1.13	0.82	0.79	0.71	0.60	0.52	0.07
% of total variance	33%	16%	13%	9%	9%	8%	7%	6%	1%

Table 8 Principal component loadings corresponding to Figure H-9.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9
Carcinogenic effects	0.09	0.29	-0.71	-0.54	0.31	0.10	0.03	0.01	0.00
Summer smog	0.52	0.03	0.11	0.27	0.71	-0.36	0.04	0.05	0.01
Winter smog	-0.04	0.75	0.23	0.22	0.15	0.47	0.25	-0.17	0.00
Climate change	0.45	0.56	0.24	-0.28	-0.19	-0.39	-0.30	-0.28	0.01
Ecotoxic effects	0.16	0.34	-0.65	0.58	-0.26	-0.16	-0.10	0.02	0.00
Acidification and eutrophication	0.69	0.35	0.18	-0.14	-0.23	-0.05	0.22	0.49	0.02
Fossil fuel consumption	0.90	-0.26	-0.07	0.00	-0.08	0.06	0.20	-0.16	-0.20
Raw water use	0.63	-0.07	0.05	0.06	0.10	0.43	-0.61	0.12	0.00
Impacted land footprint	0.86	-0.32	-0.10	0.00	-0.11	0.12	0.22	-0.21	0.18
Explained variance	2.94	1.37	1.10	0.86	0.81	0.74	0.68	0.43	0.07
% of total variance	33%	15%	12%	10%	9%	8%	8%	5%	1%

Table H-9 Principal component loadings corresponding to Figure H-10.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9
Extraction of fossil fuels	-0.95	-0.11	0.01	-0.06	0.06	-0.02	-0.15	0.19	0.16
Water use	-0.68	0.04	0.06	-0.34	0.39	-0.07	0.49	-0.11	-0.01
NO _x	-0.84	0.08	0.06	-0.15	-0.10	0.08	-0.34	-0.38	-0.01
SO ₂	-0.07	0.73	0.19	0.35	0.45	0.30	-0.10	0.03	0.00
Pb	-0.56	0.17	0.06	0.57	-0.16	-0.53	0.14	-0.06	-0.01
CO ₂	-0.31	0.67	0.10	-0.22	-0.57	0.18	0.18	0.10	0.00
TSP	0.48	0.47	0.14	-0.41	0.14	-0.53	-0.23	0.05	0.00
Cr	-0.11	0.31	-0.94	-0.01	0.06	-0.03	-0.01	-0.01	0.00
Impacted land footprint	-0.92	-0.17	-0.01	-0.07	0.07	-0.03	-0.16	0.26	-0.15
Explained variance	3.57	1.37	0.96	0.82	0.75	0.70	0.51	0.27	0.05
% of total variance	40%	15%	11%	9%	8%	8%	6%	3%	1%

H.3. INVENTORY VERSUS IMPACT UNCERTAINTY

This section presents cumulative probability graphs of the PF and FBC scenarios. Plots are presented for both the impact categories, and the dominant environmental interventions contributing to each category. The plots are in the actual units of the category or intervention, so that the increase in uncertainty between the impact and inventory results can be seen by the change in shape of the cumulative probability curve. The tables give an indication of how big a share in the impact category each environmental intervention contributes. This is shown by its numerical contribution, i.e. its impact assessment “score” (its inventory value multiplied by its equivalency factor), as well as its uncertainty contribution, represented by the magnitude of its rank-order correlation coefficient.

The plots presented here are for the scenarios defined according to a different basis. To avoid the coal and discard allocation problems, an alternative to the “avoided burden” approach was explored, that of a dual product system (i.e. a combined coal and electricity product). Here the mine and discard dump is included in the product system, and the benefits of removing or avoiding the dumps are therefore reflected in the inventory. The system also captures the benefits of producing an electricity product in addition to the coal product, without increasing mining production (i.e. using the otherwise wasted discard energy). The “per TJ” of the results below is thus the “effective energy” of the combined energy products of the system (i.e. the net power sent out and the higher heating value of the coal). Although avoiding the need to allocate the mining burdens, a problem arises because of the non-equivalent energy function of the coal and electricity products, making the definition of a meaningful functional unit difficult. The “avoided burden” approach is also preferred because it focuses attention on the actual decision system (i.e. the power plant and not the mine as well). The dual-product approach was therefore not taken further. The different basis makes no difference to the shape of the cumulative probability curves, or to the order of the curves (i.e. the relative magnitude of the systems), but does have an effect on the degree of separation between the curves (i.e. the certainty with which differences between the systems can be discerned).

H.3.1. Carcinogenic effects on humans


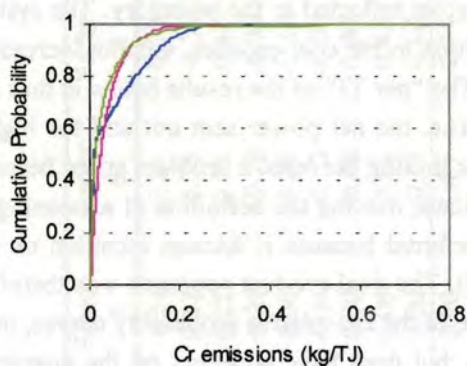



Table H-10 Magnitude of contribution and uncertainty importance of top 10 environmental interventions contributing to carcinogenic effects on humans ("score" in DALYs/TJ)

Carcinogenic effects components	PF System		FBC System	
	Correlation coefficients	Impact "scores"	Correlation coefficients	Impact "scores"
Cr	0.94	1302	0.96	8326
Ni	0.14	46	0.08	90
Ni (waterborne)	0.02	11	0.14	13
As (waterborne)	0.04	8.1	0.00	9.9
Cd	0.05	3.6	0.03	4.1
Cd (waterborne)	0.03	2.8	0.01	3.1
As	0.02	1.8	0.01	8.8
polyaromatic hydrocarbons (waterborne)	0.03	1.1	0.04	1.2
metals (unspecified)	0.02	0.2	0.04	0.1
Cr-VI (waterborne)	0.05	0.04	0.02	0.06

H.3.2. Respiratory effects on humans caused by organic substances (summer smog)




Table H-11 Magnitude of contribution and uncertainty importance of top 10 environmental interventions contributing to summer smog ("score" in DALYs/TJ)

Summer smog components	PF System		FBC System	
	Correlation coefficients	Impact "scores"	Correlation coefficients	Impact "scores"
non methane VOC	0.87	6.72	0.89	7.45
methane	0.33	1.13	0.40	1.13
pentane	0.16	0.05	0.14	0.06
butane	0.11	0.04	0.15	0.04
hexane	0.16	0.02	0.13	0.03
propane	0.10	0.02	0.16	0.02
heptane	0.10	0.01	0.16	0.01
xylene	0.13	0.01	0.09	0.01
toluene	0.07	0.01	0.14	0.01
alkanes	0.09	0.01	0.13	0.01

H.3.3. Respiratory effects on humans caused by inorganic substances (winter smog)

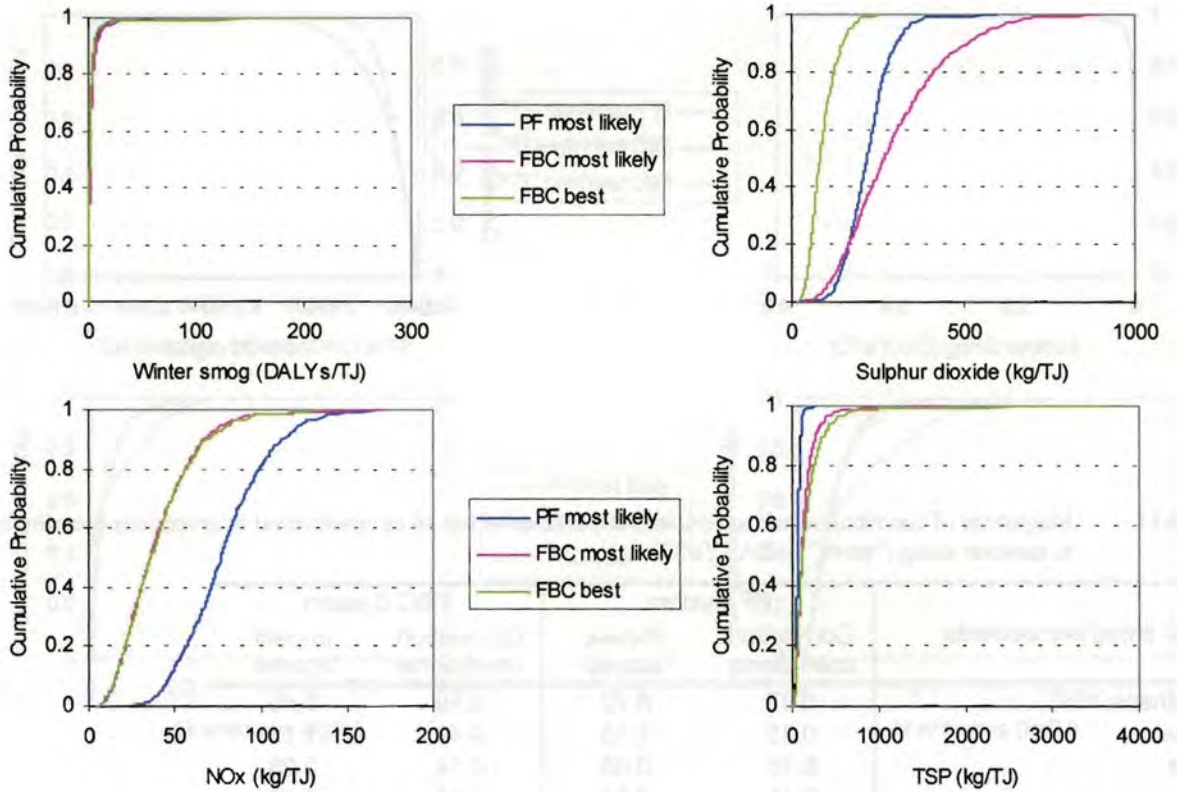


Table H-12 Magnitude of contribution and uncertainty importance of environmental interventions contributing to winter smog ("score" in DALYs/TJ)

Winter smog components	PF System		FBC System	
	Correlation coefficients	Impact "scores"	Correlation coefficients	Impact "scores"
SO ₂	0.60	1596	0.58	2570
NO _x	0.47	927	0.28	511
Total suspended particulates	0.42	875	0.54	1934
NH ₃	0.05	0.03	0.07	0.03

H.3.4. Combined effect of acidification and eutrophication




Table H-13 Magnitude of contribution and uncertainty importance of environmental interventions contributing to combined effect of acidification and eutrophication ("score" in PDF x m² x yr per TJ)

Combined effect of acidification and eutrophication components	PF System		FBC System	
	Correlation coefficients	Impact "scores"	Correlation coefficients	Impact "scores"
NO _x	0.89	6.0E+07	0.74	3.3E+07
SO ₂	0.53	3.0E+07	0.73	4.9E+07
NH ₃	0.14	5.4E+03	0.36	5.3E+03

H.3.5. Climate change




Table H-14 Magnitude of contribution and uncertainty importance of environmental interventions contributing to climate change ("score" in DALYs/TJ)

Climate change components	PF System		FBC System	
	Correlation coefficients	Impact "scores"	Correlation coefficients	Impact "scores"
CO ₂	0.99	9204	0.99	9731
methane	0.05	388	0.28	389
N ₂ O	0.07	23	0.13	1

H.3.6. Land occupation

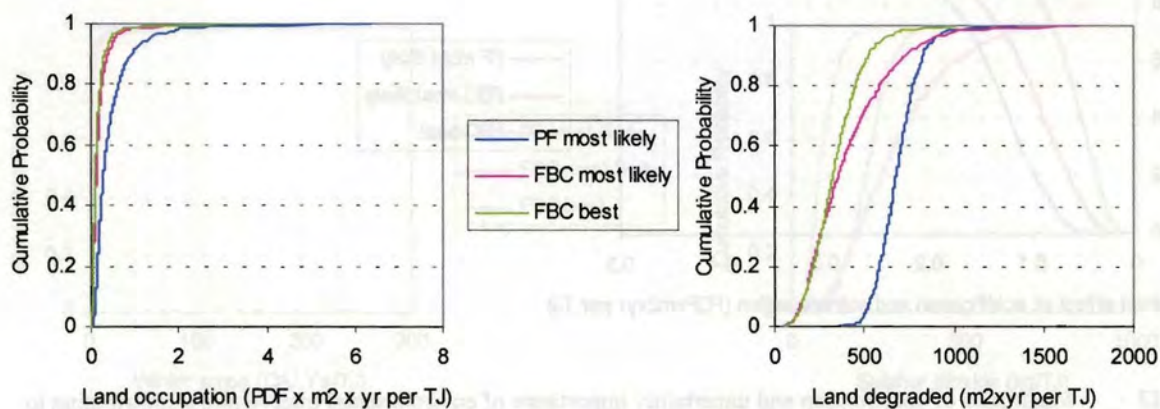


Table H-15 Magnitude of contribution and uncertainty importance of environmental interventions contributing to land occupation ("score" in PDF x m² x yr per TJ)

Land occupation components	PF System		FBC System	
	Correlation coefficients	Impact "scores"	Correlation coefficients	Impact "scores"
land transformation (semi-natural to built)	0.68	8.9E+07	0.83	6.3E+07
land transformation (semi-natural to cultivated)	0.68	6.6E+07	0.46	2.5E+07
land transformation (cultivated to built)	0.01	2.9E+05	0.08	3.2E+05
land use (built to built)	0.01	515	0.01	2161

H.4. ANALYSIS OF VARIANCE: EXAMPLE OF CARCINOGENIC EFFECTS

This section presents an overview of the process summarised in Figure 7-6, followed by an analysis of the variance in the carcinogenic effects impact category for a comparison between the most likely PF and FBC systems. This analysis is built on that presented in section H.3, so is again for the systems defined according to the dual product basis (see section H.3 for a brief description of the different basis). The principles of the analysis are the same regardless of the basis chosen.

Firstly, a distribution is fitted to the result samples, so that their variances can be examined, and the shape and range of the curves investigated to see what reduction in uncertainty is required to enable differences between the systems to be distinguished with a fair degree of confidence, or indeed whether a high degree of confidence can ever be reached. A gamma distribution is chosen, as it approximates the shape of the output probability distribution curves well, and is easily specified with the mean and variance from the result samples. At very high levels of uncertainty (CVs > 1), the gamma distribution over-emphasises the extremes, creating artificially long tails in the relative difference graphs. However, the point of intersection of the axis (i.e. the degree of confidence) still corresponds well to that on plots drawn directly from the data samples.

The variance is examined either in one system or both, until a high degree of confidence in their relative difference is obtained. In many cases, for the particular degree of separation of the sample mean values, a high degree of confidence is never able to be obtained, i.e. decreasing the variance to low levels in both curves, never results in high (or low) y-axis intercepts on the relative difference curve. For a particular level of variance in the impact results (corresponding to a desired level of confidence in the relative difference), the variance in the environmental interventions contributing to the impact category can be determined for a fixed level of variance introduced by impact assessment, e.g. using Gaussian approximation, with fixed mean values for each intervention, and fixed variance in the equivalency factors. Assuming the uncertainty in the equivalency factors is irreducible within the scope of the study sets a limit to the level at which the variance in the impact results is able to be reduced. This is because a certain degree of uncertainty is always introduced by the equivalency factor, even if that contributed by the inventory results is very low. This is only relevant for systems calculated from unrelated uncertainty samples, i.e. the uncertainty arising from impact assessment is not common to both systems and does not cancel out in taking the normalised ratio. If this is not the case, there is no need to incorporate the equivalency factor variance into the analysis, and the target variance for the inventory results is the same as that for the impact results.

A simplification made in the analysis is that the variance in the impact results is attributed to a single intervention, i.e. the top contributor to the variance in the impact category, determined from a rank-order correlation analysis of the contributing environmental interventions. This simplifies the analysis considerably, but is only valid if a single environmental intervention dominates the category. For many of the impact categories, the estimate of uncertainty is the same or similar across all equivalency factors for that category, thus the effect of the simplification is reduced. For categories with more than one intervention with high uncertainty importance, the analysis is repeated considering each intervention singly. The variance in the impact category is checked after each intervention has been investigated and the uncertainty of its key input parameters reduced as far as the scope of the study allows. Once the target variance in the impact category is reached, no further lower order interventions require analysis.

In most impact categories, the uncertainty introduced by the equivalency factors is so extensive that it limits the degree to which the variance in the impact results can be reduced. In this case, instead of the inventory variance solved for a particular level of impact variance, the inventory variance is set to a feasible minimum value, and the impact variance calculated, i.e. the combined minimum variance in the inventory results and the variance in the equivalency factors determines the degree of confidence that can be held in the relative difference between the systems. From the experiences of this study, CVs of less than 0.2 are unlikely to be achievable for most environmental interventions. If the equivalency factor CV is not limiting, the analysis gives a rough indication of the level of variance in the intervention corresponding to the desired level of variance in the impact results. This, or the minimum feasible variance assumed in the limiting case, sets the target for reducing the variance in the intervention's output sample. To direct this reduction in variance, a rank-order correlation analysis of the input parameters contributing to the intervention is conducted. Thus the parameter contributing the greatest to the variance of the intervention is identified, and its uncertainty addressed first. For impact categories with more than one intervention contributing significantly to the impact "score", the analysis is repeated with each of these interventions, until the desired level of variance in the impact category is obtained.

Example

Extremely high uncertainty is present in the carcinogenic effects on humans category (see figure in section H.3.1). This is partly due to the high uncertainty in the toxicity factors, and partly due to the high uncertainty in estimating the emissions of heavy metals from the furnaces. At such high uncertainty, even with fairly significant differences in the means of the systems (80%), no significant differences can be discerned between the systems. This is shown by plotting the cumulative probability of their relative difference (% change between the systems), i.e.

$$\frac{\text{Carcinogens}_{\text{FBC}} - \text{Carcinogens}_{\text{PF}}}{\text{Carcinogens}_{\text{FBC}}}$$

The probability that the FBC has a higher potential for carcinogenic effects is given by the point at which the cumulative probability curve crosses the y axis. Figure H-13 shows that there is only a 50% probability that the FBC system performs worse than the PF system with regard to carcinogenic effects.




Figure H-7 Percentage change in carcinogenic effects between FBC most likely and PF most likely systems.

Cr emissions are seen to be by far the greatest contributor to carcinogenic effects for both the PF and FBC systems, both in terms of uncertainty, and numerical value. Ni emissions are the next most significant contributor, although substantially less so than Cr (see Table H-20). High output variance is calculated for Ni and Cr emissions, although they show less uncertainty than when aggregated as carcinogenic effects. This can be seen by the shape of their probability curves, which can at least be discerned, unlike that of the carcinogenic effects curves (see figures in section H.3.1). The curves between the systems are not stochastically dominant (i.e. their cumulative probability curves intersect each other, representing that at certain probabilities the FBC system appears better than the PF system, whilst at others, the order is reversed). Thus, although significant differences are apparent in their median values (40-80%), the FBC system show only between 40-60% probability of emitting more heavy metals than the PF system (see Figure H-14).

A Gamma distribution was fitted to the carcinogenic effects curves for the two systems so that the effect of decreasing their variance could be investigated. Even at significantly decreased uncertainties (CVs reduced to 1.7) the confidence with which it can be stated that carcinogenic effects are higher for the FBC system is only marginally increased to 60% (see Figure H-16, compared with Figure H-13). The fact that Cr emissions are responsible for such a high proportion of the total carcinogenic effect (see Table H-20), enables the use of the variance in its equivalency factor to relate the variance in carcinogenic effects back to that of the Cr emissions. The equivalency factor relating Cr emissions to carcinogenic effects has extremely high uncertainty (approximated by a CV of 5), thus even should the variance in Cr emissions be reduced to a very low level, the variance in carcinogenic effects can never be decreased to below a CV of 5 (considering only Cr emissions). However, the other environmental interventions contributing to carcinogenic effects have similarly high levels of uncertainty (minimum CV of 1.3), thus even with the effect of the other environmental interventions, the variance in carcinogenic effects is unlikely to be reduced. Thus carcinogenic effects is not a good criteria for comparison between the systems, because at the degree of separation predicted for the systems, it can never be stated with confidence that the one system is better or worse than the other. The comparison should therefore be based on inventory data.




Figure H-8 Percentage change in Cr and Ni emissions between the FBC and PF most likely systems.




Figure H-9 Cumulative probability density of FBC and PF most likely systems, with the CVs of the fitted gamma distributions set to 1.7.




Figure H-10 Percentage change in carcinogenic effects between FBC most likely and PF most likely systems, for the limit case, where the CVs of the fitted gamma distributions have been reduced to 1.7.

A Gamma distribution was fitted to the Cr emissions to investigate the reduction in uncertainty that would be required to be able to determine significant differences between the systems. If the variance in the PF system is kept the same and that of the FBC system is reduced, stochastic dominance of one system over the other can never be obtained. Thus even if the FBC system's CV could be reduced to 0.2, it could be stated with only 60% certainty that the FBC system has higher Cr emissions. For stochastic dominance to be obtained, the CV of the PF system needs to be reduced to at least 0.3. Although unlikely to be achievable, at CVs for both systems of 0.2, it can be stated with 85% confidence that the PF system emits more Cr emissions than the FBC system (see Figure H-17). Similarly high degrees of confidence can be achieved for such levels of uncertainty for Ni emissions (see Figure H-17). The high degree of uncertainty in determining heavy metal emissions, and the variability of heavy metal concentrations in coal, mean that CVs of 0.2 are unrealistic for this system. At the more achievable level of 0.5, a degree of confidence of 65% and 70% can be achieved for Cr and Ni emissions respectively.

The uncertainty in predicting heavy metal emissions is largely irreducible within the current scope of the study, partly because of the high variability in the feedstocks (coal and sorbent), and partly because of the wide variability in furnace partitioning coefficients found in the literature. The above analysis shows that even should significant reductions in uncertainty be achieved, the dominance of one system over the other can not be determined with a high degree of confidence. Cr and Ni emissions can thus be regarded, at best, as weak criteria for the comparison between the PF and FBC systems.










Figure H-11 Percentage change in Cr and Ni emissions between the FBC and PF most likely systems, showing the degree of confidence attainable should the variance in the output be able to be reduced to a CV of 0.2 for both systems.

H.5. FBC OPERATING SPACE PARAMETER ANALYSIS

The following cumulative probability plots show the relative performance of the PF and FBC system for different levels of SO₂ removal and type of sorbent used. The graphs plot the relative difference between the FBC systems and the PF system, i.e:






$$\frac{PF - FBC}{PF}$$

Where *PF* and *FBC* denotes an impact category “score” or emission flow for the PF and FBC system respectively. With the ratio defined in this way, a positive outcome shows the FBC system to be preferred to the PF system (where less is always better). The greater the degree to which the cumulative probability curve falls into the positive region of the graph, the greater the degree of confidence that can be held that the FBC system is better than the PF system. The y-intercept gives the probability that the FBC system’s “score” or emission is always less or equal to that of the PF system, which can be interpreted as the degree of confidence that can be held that the PF system is always worse than the FBC system. For example, in the plot of the relative difference in summer smog potential, the FBC system using lime to achieve a 40% SO₂ removal has a y-intercept of 15%. This indicates that it can be stated with 85% confidence that this FBC system decreases the summer smog potential relative to the PF system.



APPENDIX I

SUPPORTING INFORMATION FOR OPERATIONAL CASE STUDY (CHAPTER 8)



APPENDIX I

SUPPORTING INFORMATION FOR OPERATIONAL CASE STUDY (CHAPTER 8)

I.1. ANALYSIS OF THE PRINCIPAL COMPONENT TABLES

For the analysis with impact criteria (Table I-1 and Figure 8-9), the first PC incorporates the variance arising from the additional power consumption of the system, and the second the mine water related effects. The waterborne sulphate emissions and raw water use are strongly correlated (i.e. fall under the same PC), as they are both primarily related to the volume of mine water used. Also related to the volume of mine water taken are the emissions resulting from increased transport and manufacture of water treatment chemicals, as shown by the contributions of winter and summer smog to the PC.

Similarly, for the analysis based on environmental interventions (Table I-2 and Figure 8-10), the first PC captures the energy-related effects of using mine water (i.e. the increased auxiliary power use), and the second the other mine water related effects (i.e. the reduction in sulphate emissions and increased chemical usage). Here the TSP emissions are separated out from winter smog, where they were somewhat masked by the influence of the SO₂ emissions, and thus show a much stronger correlation to the amount of mine water taken. This correlation reflects the increase in the mass of lime used with the increase in the volume of water treated in the clarifier, because of the very high value for dust from limestone mining in the background inventory used for lime production, and the fact that lime consumption is simply related to the throughput of the clarifier. Coupled to the increase in the mass of chemicals consumed, is their transport to the power station, which is also picked up by the second PC. This is also more clearly shown in the analysis based on interventions rather than impacts, through the high contribution of methane and non-methane VOCs to the second PC. Here the relative increase in the transport-related emissions is not masked by the far greater volumes of coal-related emissions when aggregated into impact data.

The fact that the uncertainty sample for the impacted land footprint is generated independently from the other criteria (see section 8.2.5) results in the footprint falling into its own PC. The impacted land footprint sample is calculated from a different random sample to the other criteria, so no co-variance with the other sample points exists. The PCA thus recognises the footprint variance as a separate source of variance and places it in a separate PC. This, and the fact that relatively small differences between the impacted land footprints are evident, accounts for the footprint only showing up in the 3rd and 4th PC for the impact and inventory criteria, respectively, and the fact that it is only a weak criterion in Figures 8-9 and 10.

Although a fairly strong contributor to the 2nd PC, raw water use is split between this component and a lower order component (the 4th and 3rd PC for the impact and inventory criteria, respectively). This reflects the opposing trends in raw water consumption, i.e. for the use of compartment water it is correlated with the waterborne sulphate emissions (it increases as additional compartment water is taken), whilst for the pit water it shows the opposite trend (it decreases when high volumes of pit water are used). The fact that the PCA is based on the percentage difference between the mine water systems and the base case (no mine water system) results in the energy-related impacts being over emphasised (see section 8.3.3). Although of greater relevance to the systems as a whole, the magnitude of the percentage changes in water use, impacted land footprint and waterborne sulphates are smaller than those arising from increased auxiliary power use.

Table I-1 PC loadings for impact categories, raw water use, waterborne sulphates and the impacted land footprint.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9	Factor 10	Factor 11
Carcinogenic effects	0.94	0.09	-0.02	0.05	0.03	0.03	0.32	-0.02	-0.00	-0.00	-0.00
Summer smog	0.86	-0.24	0.14	-0.27	-0.11	-0.30	-0.02	-0.11	-0.05	-0.00	-0.00
Winter smog	0.61	-0.43	0.15	-0.41	-0.42	0.25	-0.02	0.01	-0.00	-0.00	-0.00
Climate change	0.94	0.23	-0.09	0.16	0.08	0.11	-0.06	-0.02	-0.04	-0.02	0.03
Ecotoxic effects	0.96	0.02	0.03	-0.05	-0.02	-0.13	-0.02	0.25	-0.00	-0.00	-0.00
Acidification and eutrophication	0.94	0.22	-0.09	0.16	0.07	0.10	-0.07	-0.02	-0.03	-0.03	-0.03
Land occupation	0.96	0.20	-0.07	0.13	0.06	0.07	-0.07	-0.03	0.01	0.07	-0.00
Fossil fuel consumption	0.99	0.05	0.01	-0.01	-0.00	-0.08	-0.06	-0.07	0.10	-0.02	-0.00
Raw water use	-0.11	0.59	0.44	-0.60	0.29	0.05	-0.00	-0.00	-0.00	-0.00	-0.00
Waterborne sulphates	-0.31	0.74	-0.01	0.08	-0.59	-0.06	0.01	-0.01	-0.00	-0.00	-0.00
Impacted land footprint	0.06	-0.11	0.88	0.46	-0.04	0.01	-0.00	-0.00	-0.00	-0.00	-0.00
Explained variance	6.69	1.31	1.03	0.90	0.64	0.21	0.12	0.08	0.02	0.01	0.00
Percentage of total variance	61%	12%	9%	8%	6%	2%	1%	1%	0%	0%	0%

Table I-2 Principal component loadings for selected key environmental interventions and the impacted land footprint.

	Factor 1	Factor 2	Factor 3	Factor 4	Factor 5	Factor 6	Factor 7	Factor 8	Factor 9	Factor 10	Factor 11
Raw water use	-0.11	0.31	0.78	-0.37	-0.37	0.07	-0.00	-0.00	-0.00	-0.00	-0.00
Carbon monoxide	0.98	0.12	-0.01	-0.00	-0.00	-0.01	-0.07	0.12	-0.00	-0.00	-0.00
Carbon dioxide	0.96	0.26	-0.06	0.05	-0.01	0.08	0.01	-0.02	-0.02	0.02	0.01
Nitrous oxide	0.95	0.28	-0.07	0.05	-0.01	0.09	0.01	-0.03	-0.02	-0.02	-0.00
Nitrogen oxides	0.97	0.23	-0.05	0.04	-0.01	0.06	-0.00	-0.03	0.06	-0.00	-0.00
Sulphur dioxide	0.95	0.28	-0.07	0.05	-0.01	0.09	0.01	-0.02	-0.02	0.01	-0.01
Methane	0.88	-0.34	0.16	-0.12	0.04	-0.24	0.15	0.02	-0.00	-0.00	-0.00
Non-methane VOCs	0.84	-0.39	0.18	-0.14	0.04	-0.27	-0.12	-0.05	-0.00	-0.00	-0.00
Total suspended particulates	0.38	-0.74	0.21	-0.24	0.33	0.32	-0.00	-0.00	-0.00	-0.00	-0.00
Waterborne sulphates	-0.29	0.63	0.26	-0.15	0.65	-0.09	-0.00	-0.00	-0.00	-0.00	-0.00
Impacted land footprint	0.06	-0.13	0.52	0.84	0.06	0.01	-0.00	-0.00	-0.00	-0.00	-0.00
Explained variance	6.35	1.61	1.06	0.97	0.67	0.28	0.04	0.02	0.00	0.00	0.00
Percentage of total variance	58%	15%	10%	9%	6%	3%	0%	0%	0%	0%	0%

3.3.3. Model Parameter Uncertainty and Uncertainty about Model Form

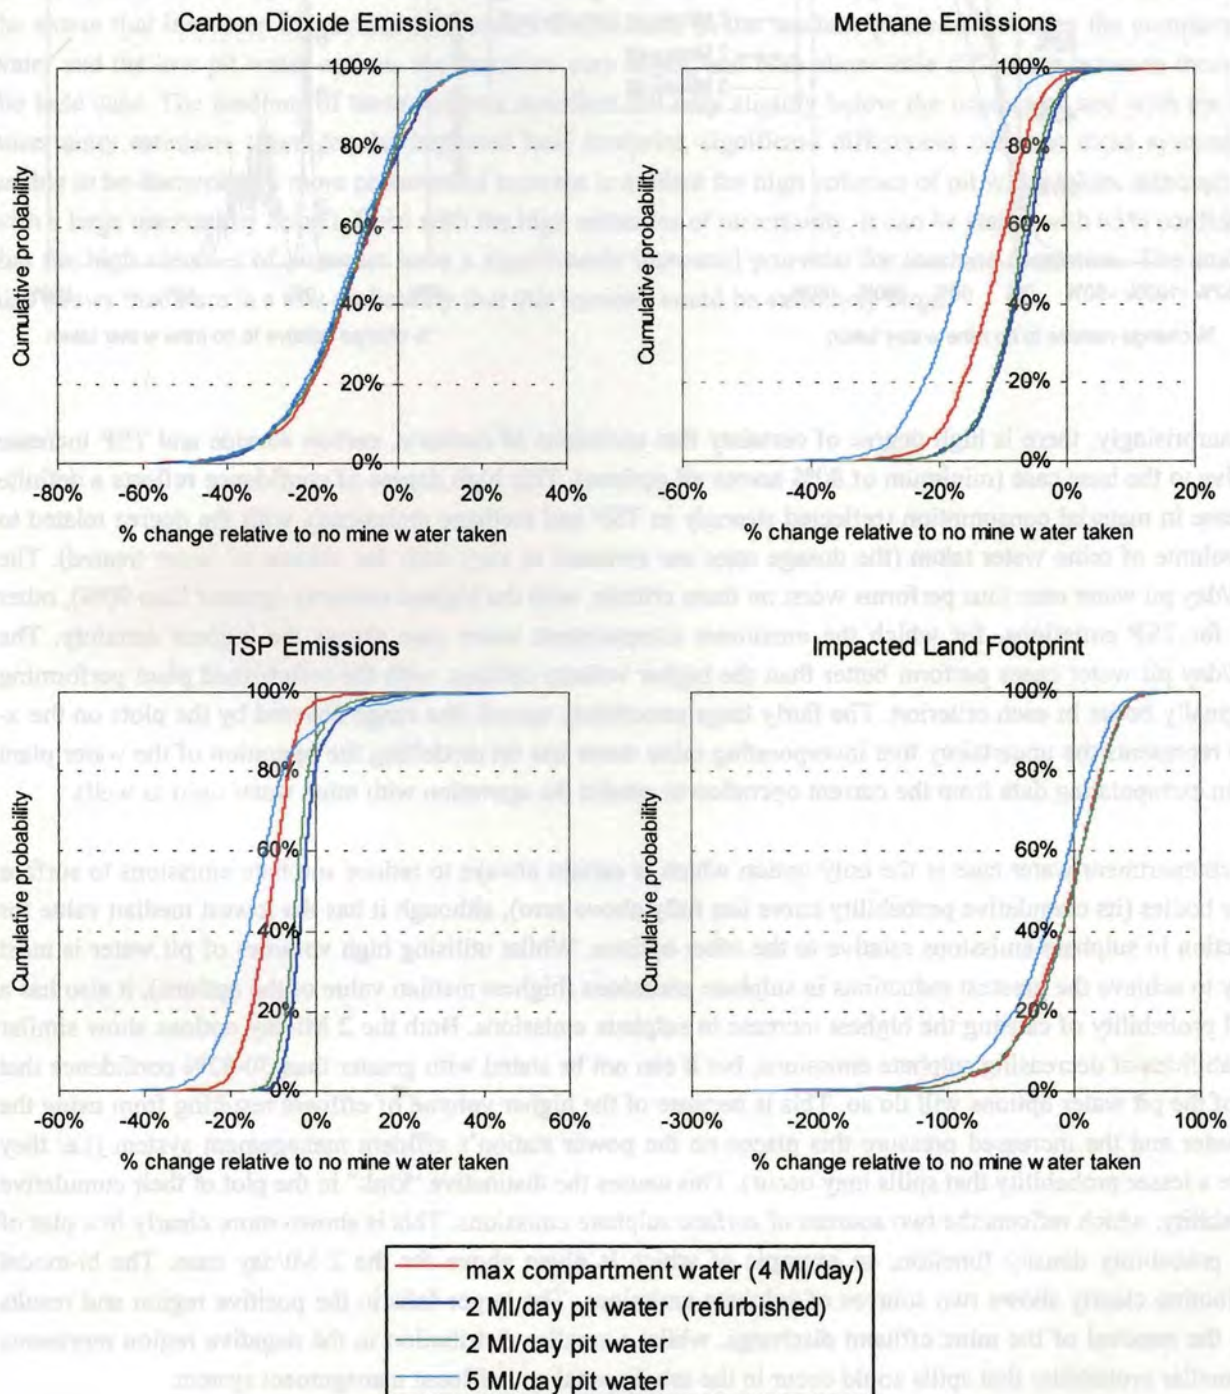
The majority of uncertainty assessments demonstrated in LCA studies have been around empirical uncertainty. However, the use of sensitivity analyses to determine the effect of the many choices in LCA studies is fairly well established, and there is a general consensus on the use of parametric or sensitivity analyses for the assessment of model uncertainties (Huijbregts, 1998b; Maurice et al., 2000; Meier, 1997). Studies assessing uncertainties arising from model choices have predominantly been around the choice of allocation method (Azapagic and Clift, 1999b; Azapagic and Clift, 2000; Huijbregts, 1998b), major model parameters (e.g. system lifetime) (Huijbregts, 1998b; Maurice et al., 2000; Meier, 1997), and the choice of impact assessment method (Baumann and Rydberg, 1994; Meier, 1997; Tolle et al., 2001; Tukker, 1998). These studies demonstrate the extremely significant consequences the choice of LCA method can have on the results.

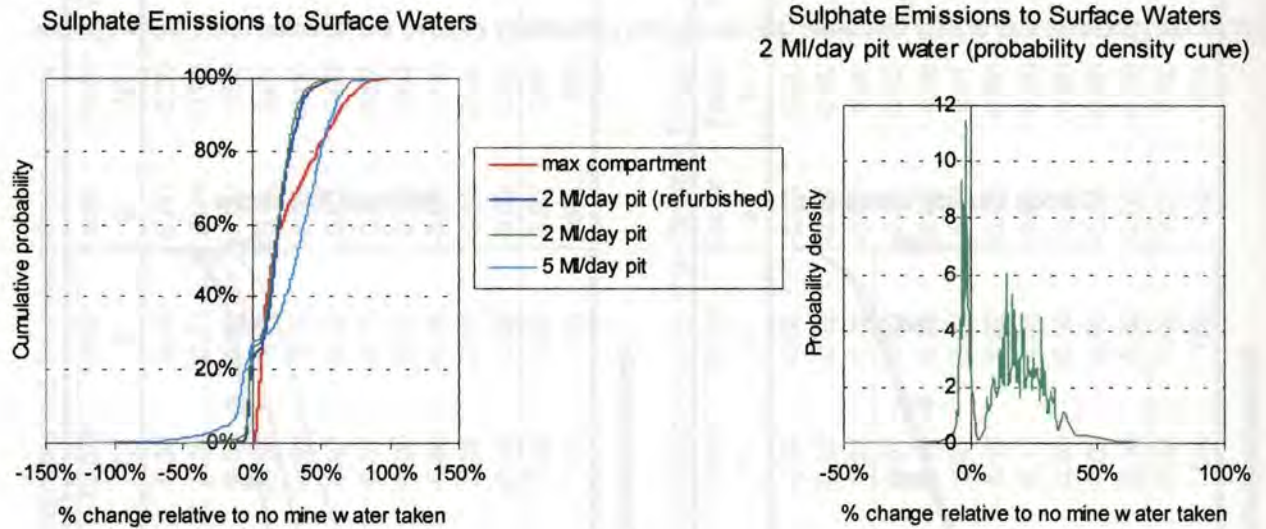
For the assessment of model parameters, Steen (1997) suggests that a probability number could be assigned to each alternative choice, although Morgan and Henrion (1990) argue strongly against this (see section 3.1.1. for reasons). Huijbregts (1998b) identifies that sensitivity and parametric analyses are cumbersome to perform and report, and that it is not feasible to analyse the effect of all possible combinations of choices, and to compute the model output for all ranges of model parameters. Huijbregts (1998b) therefore suggests finding the two extreme options for every choice, and then constructing two extreme combinations of options, and computing the effect of these on the LCA results.

Meier (1997) recognises the need for scenario modelling of discrete options where uncertainties are caused by preferences within a decision-making process. For example, in an assessment of gas purification systems, Meier (1997) repeats the probabilistic assessment with a different background electricity production model. For one metric of ecological indicators, this causes the rankings of the different technologies to be nearly reversed, thereby demonstrating the potential significance of the choice of model parameters. Meier (1997) further distinguishes model uncertainties whose importance can only be qualitatively addressed, i.e. no clear scenarios can be defined to test the model choices by sensitivity analysis. These are model uncertainties based on inherent model assumptions and definitions (Meier, 1997), e.g. spatial and temporal limitations. This is in agreement with Huijbregts (1998b), who recognises that certain model uncertainties (such as the inherent inability of the current LCA method to

I.2. PROBABILITY PLOTS FOR MINE WATER OPTIONS

This section presents and briefly discusses the cumulative probability plots of the selected mine water options.





Not surprisingly, there is high degree of certainty that emissions of methane, carbon dioxide and TSP increase relative to the base case (minimum of 80% across all options). This high degree of confidence reflects a definite increase in material consumption (reflected strongly in TSP and methane emissions), with the degree related to the volume of mine water taken (the dosage rates are assumed to vary with the volume of water treated). The 5 MI/day pit water case thus performs worst on these criteria, with the highest certainty (greater than 90%), other than for TSP emissions, for which the maximum compartment water case shows the highest certainty. The 2 MI/day pit water cases perform better than the higher volume options, with the refurbished plant performing marginally better in each criterion. The fairly large uncertainty spread (the range covered by the plots on the x-axis) represents the uncertainty that incorporating mine water has on modelling the operation of the water plant (i.e. in extrapolating data from the current operation to predict the operation with mine water used as well).

The compartment water case is the only option which is certain always to reduce sulphate emissions to surface water bodies (its cumulative probability curve lies fully above zero), although it has the lowest median value for reduction in sulphate emissions relative to the other options. Whilst utilising high volumes of pit water is most likely to achieve the greatest reductions in sulphate emissions (highest median value of the options), it also has a small probability of causing the highest increase in sulphate emissions. Both the 2 MI/day options show similar probabilities of decreasing sulphate emissions, but it can not be stated with greater than 70-82% confidence that any of the pit water options will do so. This is because of the higher volume of effluent resulting from using the pit water and the increased pressure this places on the power station's effluent management system (i.e. they create a lesser probability that spills may occur). This causes the distinctive "kink" in the plot of their cumulative probability, which reflects the two sources of surface sulphate emissions. This is shown more clearly in a plot of their probability density function, an example of which is given above for the 2 MI/day case. The bi-modal distribution clearly shows two sources of sulphate emissions. The larger falls in the positive region and results from the removal of the mine effluent discharge, whilst a smaller distribution in the negative region represents the smaller probability that spills could occur in the ash disposal and effluent management system.

The uncertainty in the impacted land footprint reflects the empirical uncertainty in the data from which it is calculated. This is less rigorously calculated than for the other criteria, so the uncertainty ranges reflected in the plots are only an estimate, and may be conservatively large (see section 8.2.5). Nonetheless the impacted land footprint analysis shows that the mine water options always increase the potential for leachate formation relative to the no mine water case. Incorporating either compartment water or low volumes of pit water does not significantly increase the volume or concentration of the water plant effluent (i.e. the increase in effluent volume is able to be absorbed by the effluent management system without increasing the moisture content of the ash to the extent that increases in leachate will occur). Differences in the leachate potential between the compartment water and the low pit water options are therefore very slight, and both show little difference between them and the base case. The medians of these systems therefore fall only slightly below the base case, and with the high uncertainty estimates taken for the impacted land footprint, significant differences between these systems are unable to be discerned. A more pronounced increase is evident for high volumes of pit water taken, although still with a large uncertainty bound. Even with the high estimates of uncertainty, it can be stated with 65% confidence that the high volumes of pit water have a significantly increased potential for leachate formation. The analysis also shows that there is a low probability that this increase could be extremely large.

1.3. BOX AND WHISKER PLOTS OF MINE WATER EXTREME OPTIONS

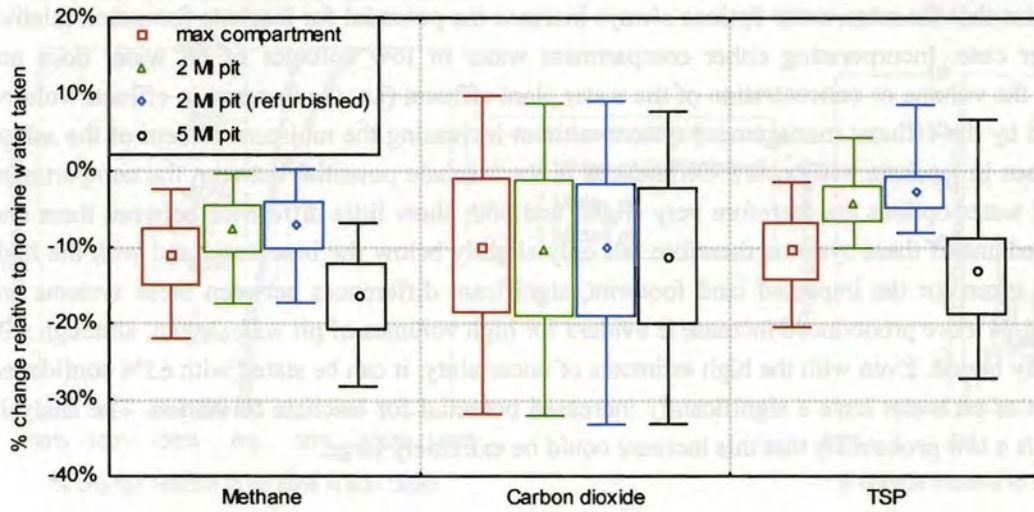


Figure I-1 Emissions to air for extreme mine water options considered.




Figure I-2 Waterborne sulphates emissions to surface waters for extreme mine options considered.




Figure I-3 Impacted land footprint for extreme mine options considered

I.4. MODEL PARAMETER ANALYSIS

This section presents an investigation into the effect of varying the model parameters. The graphs show the mid-points of the result samples calculated at each feasible combination of the model parameters, at the considered resolution of the model parameters. A fairly coarse step interval is chosen, as the greater accuracy of a smaller interval was considered unnecessary given the high empirical parameter uncertainty. The points are joined as the parameters are in fact continuous variables, and any point along a given line can be considered a feasible operating state. Similarly, the volume of water taken is a continuous variable. As it is required that the analysis be represented on a two dimensional plot, one parameter is held constant, whilst the other two are varied. The value chosen to be held constant is selected so as to maximise the operating ranges of the other two parameters.

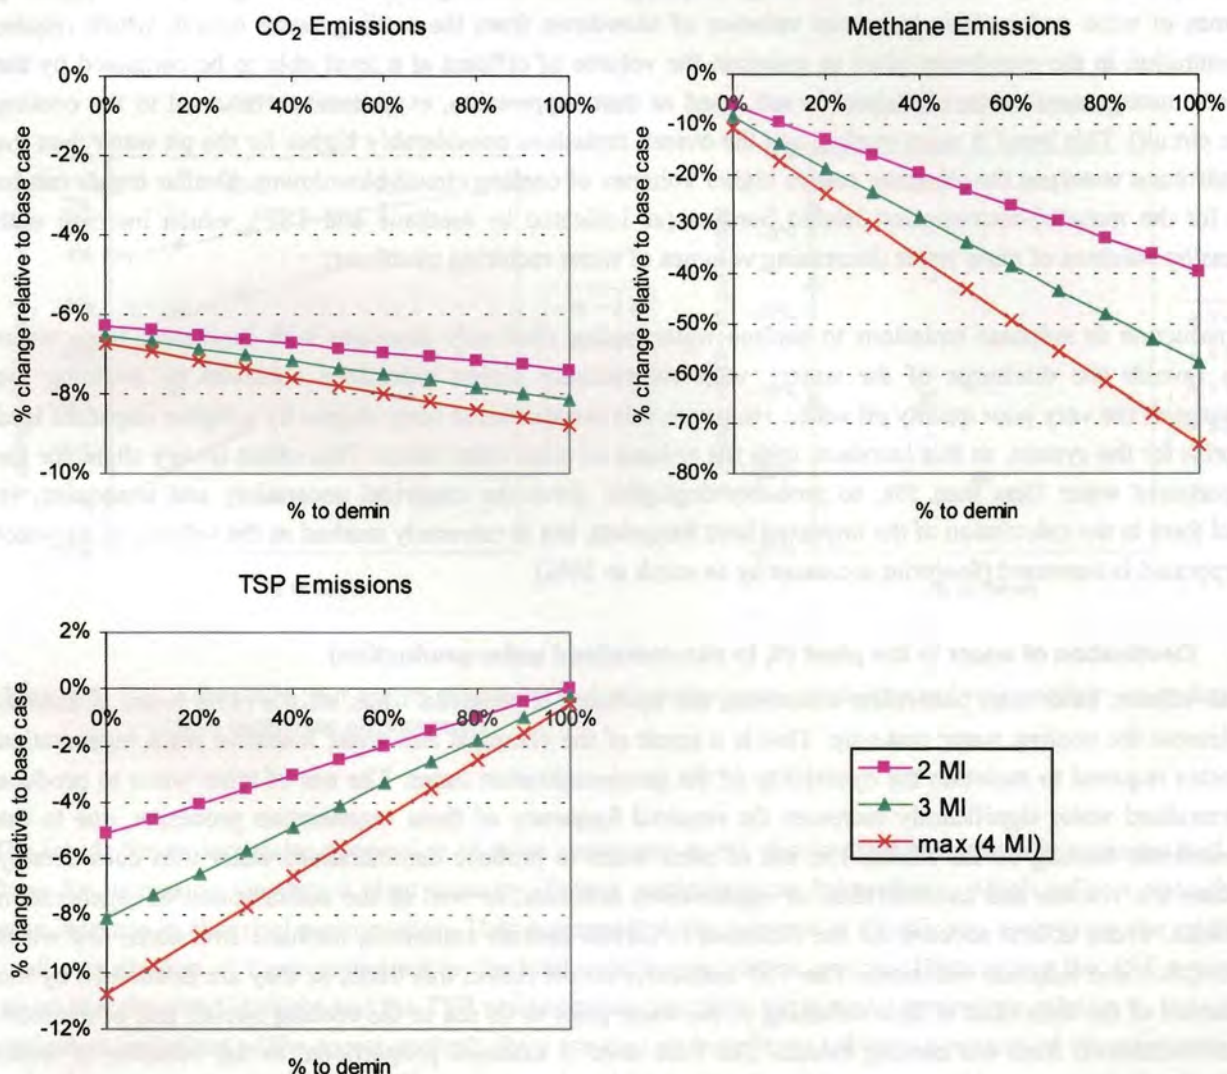


Figure I-4 Effect of % mid-quality mine water to demineralised water production and volume of compartment water taken, at 40% membrane plant capacity.

1.4.1. Volume and quality of mine water

For the compartment water case, increasing the volume of water taken decreases the volume of raw water make-up required. Contrastingly, as the volume of pit water is increased, raw water consumption increases with additional water taken, as the cooling water circuit is forced into decreasing its cycles of concentration. The high sodium concentration in the mine water shifts the operation of the cooling circuit from sulphate limiting to sodium limiting as additional volumes of mine water are added, i.e. increased blowdowns from the cooling circuit are required to meet the sodium ion concentration limit in the circulating cooling water. Significantly greater reductions in raw water consumption are achieved with the mid quality mine effluent, with high volumes of the poor quality pit effluent actually increasing the raw water consumption relative to the no mine water case.

Energy related emissions (as indicated by carbon dioxide) increase with the volume of mine water taken. This is reflecting the increased membrane plant usage required with increasing mine water usage, i.e. the increasing volumes of mine water cause increased volumes of blowdown from the cooling water circuit, which require concentration in the membrane plant to maintain the volume of effluent at a level able to be contained by the effluent management system (retained in ash, used as dust suppression, evaporated or returned to the cooling water circuit). This trend is more marked and the overall emissions considerably higher for the pit water than the compartment water, as the pit water causes higher volumes of cooling circuit blowdowns. Similar trends can be seen for the material-consumption related burdens (as indicated by methane and TSP), which increase with increasing volumes of mine water (increasing volumes of water requiring treatment).

The reduction in sulphate emissions to surface water bodies obviously increases with increasing mine water taken (avoids the discharge of the water), with substantially higher reductions achieved by avoiding the discharge of the very poor quality pit water. However, this is balanced to some degree by a higher impacted land footprint for the system, as this increases with the volume of mine water taken. This effect is very slight for the compartment water (less than 5%, so probably negligible given the empirical uncertainty and uncertainty in model form in the calculation of the impacted land footprint), but is extremely marked as the volume of pit water incorporated is increased (footprint increases by as much as 20%).

1.4.2. Destination of water in the plant (% to demineralised water production)

For all effects, other than particulate emissions, the optimum is achieved when all the mine water is used to supplement the cooling water make-up. This is a result of the chemical and water intensive resin regeneration processes required to maintain the operability of the demineralisation trains. The use of mine water to produce demineralised water significantly increases the required frequency of these regeneration processes, due to the increased salt loading on the resins. The use of mine water to produce demineralised water thus considerably increases the volume and concentration of regeneration effluents, as well as the consumption of regeneration chemicals. These effects account for the increases in carbon dioxide emissions, methane emissions, raw water consumption and sulphate emissions. The TSP emissions do not reflect this trend, as they are dominated by the production of the lime used in lime softening of the water prior to its use in the cooling circuit, and of the side-stream withdrawn from the cooling circuit. The lime used is assumed proportional to the volumes of water treated, both of which decrease with decreasing volumes of mine water placed in the cooling circuit.

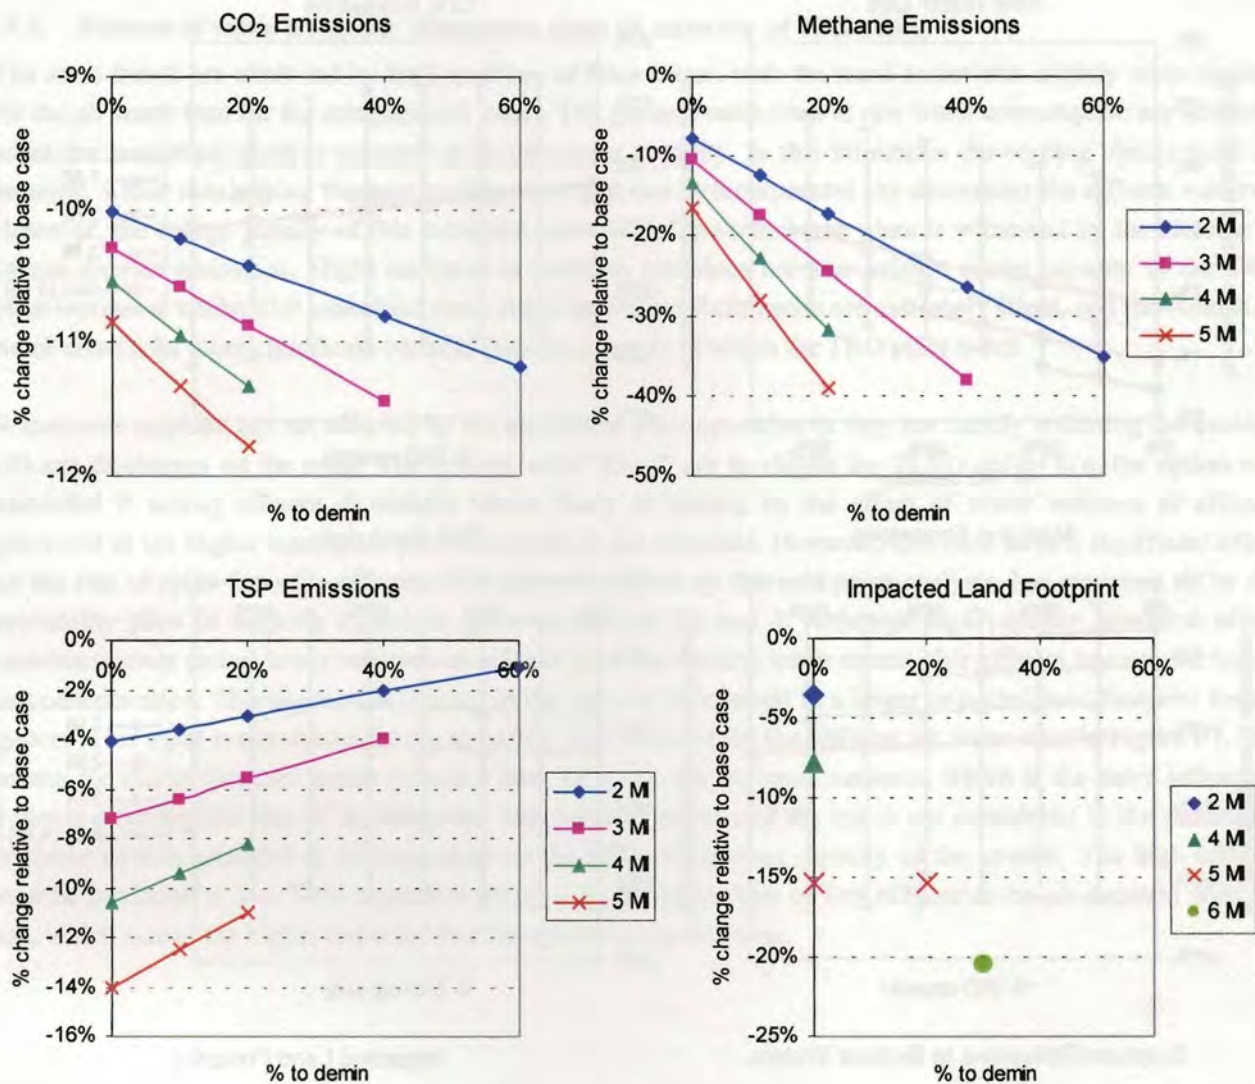


Figure I-5 Effect of % poor-quality mine water to demineralised water production and volume of pit water taken, at 80% membrane plant capacity.

The trends for an increasing proportion of mine water sent to the demin plant are generally more marked than those for increasing membrane plant capacity. This is particularly so for methane, which reflects strongly the large increase in chemical consumption. This suggests that this increase is significant, even given the relatively small contribution of these emissions to the total mine/power station system. Discounting the TSP emissions (given that the trend is slight and the TSP emissions account for a fairly small percentage relative to the overall particulate emissions of the power station), there are thus no benefits to splitting a portion of the mine water off to demineralised water production. The only time in which it may be attempted would be to operate the system at the extreme mine water volume possible (i.e. relax the constraint that raw water consumption must not be more than that of the system not incorporating mine water, to achieve the maximum reduction in sulphate emissions). However, this is at such significant costs to all other impacts, it is unlikely to be attempted.

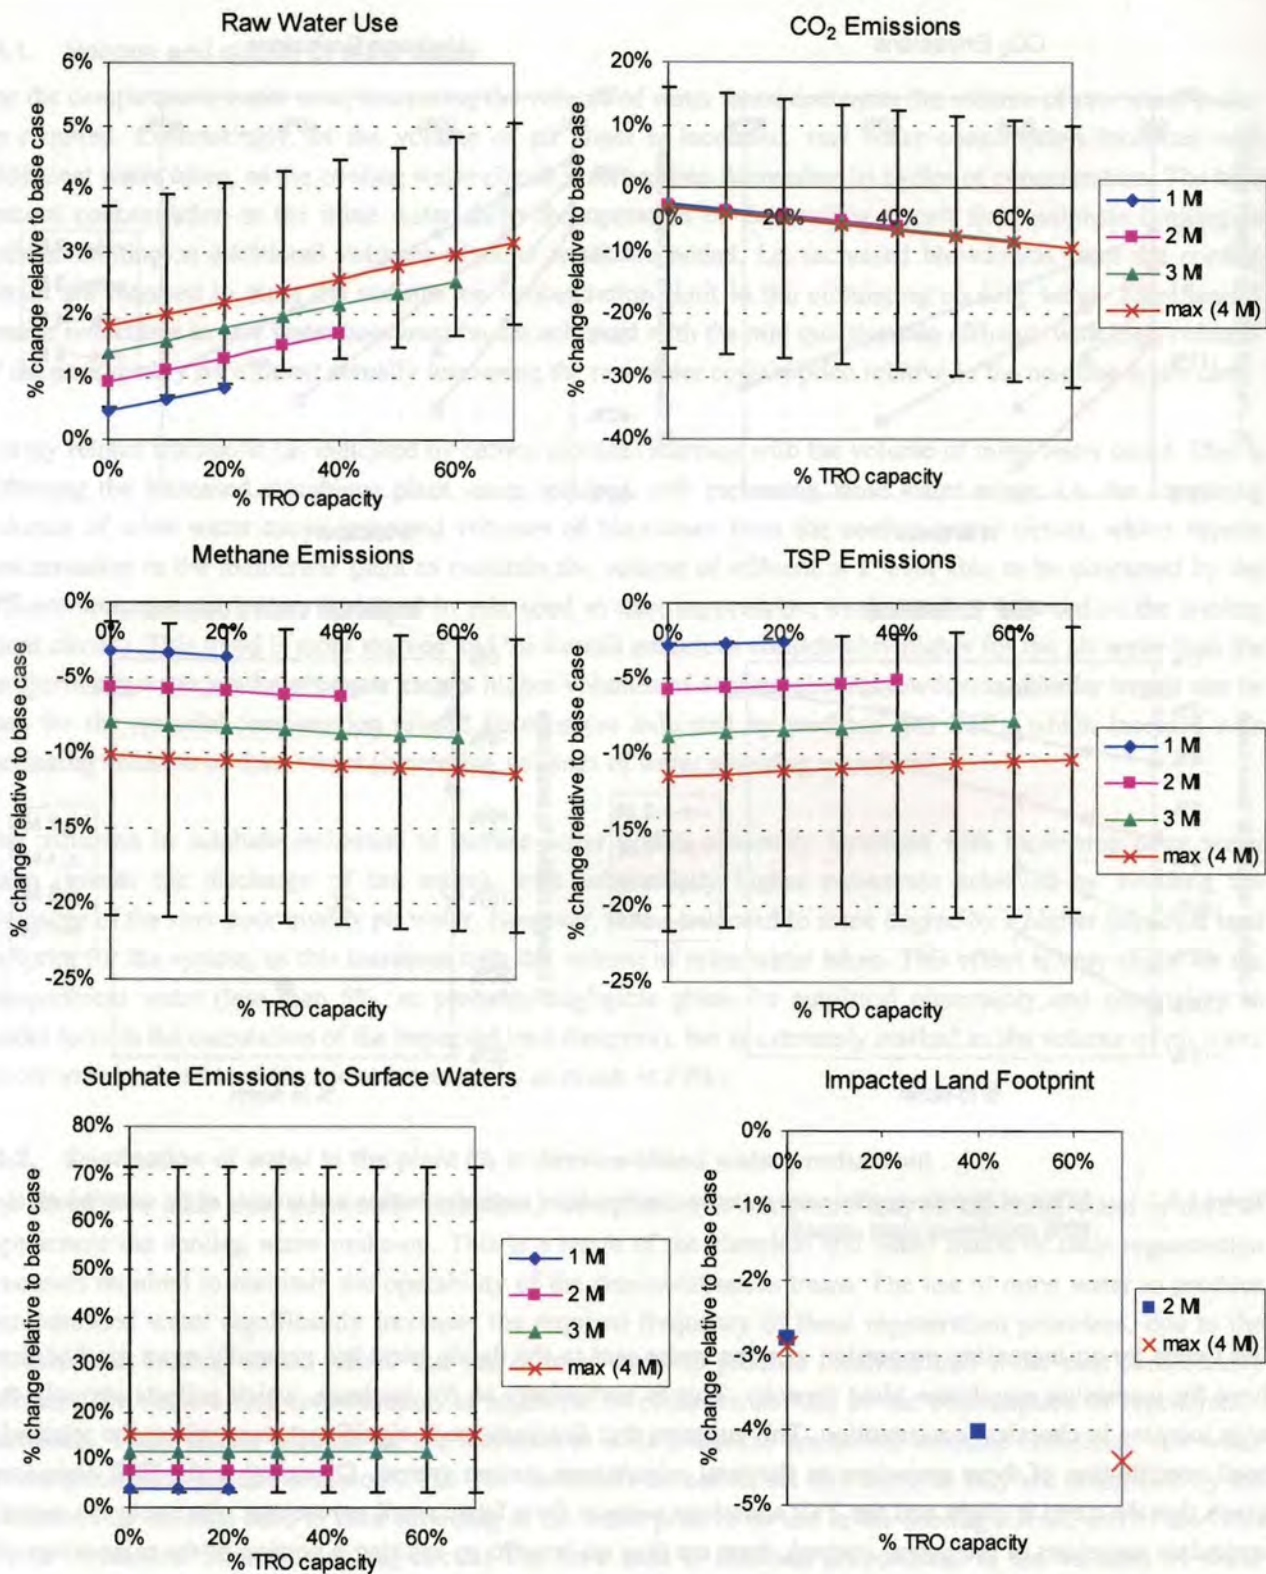


Figure I-6 Effect of volume of water treated in membrane plant and volume of compartment water taken, for all mid-quality mine water used as cooling water make-up. Error bars giving the 90% confidence range are given for the maximum compartment water case (representing empirical parameter uncertainty).

I.4.3. Volume of water treated in membrane plant (% capacity of TRO plant)

The same trends are exhibited by both qualities of mine water, with the trend sometimes slightly more marked for the pit water than for the compartment water. The greatest reductions in raw water consumption are achieved when the membrane plant is operated at its maximum capacity, as this minimises the cooling circuit make-up required, whilst maximising the poor quality water that can be incorporated (by decreasing the effluent volume). However, the energy penalty of this increased operation of the membrane plant is witnessed by the increase in carbon dioxide emissions. Slight increases in methane emissions are also evident as the capacity of the TRO plant increases, whilst TSP emissions show slight decreases. Both trends are extremely slight, and the volume of water taken a far more significant variable than the capacity at which the TRO plant is run.

Waterborne sulphate are not affected by the membrane plant operation as they are merely reflecting the avoided effluent discharges on the mine. The options were “fixed” not to violate the ZLED policy (i.e. the option was discarded if ashing effluent discharges were likely to occur), so the effect of lower volumes of effluent generated at the higher membrane plant capacities is not reflected. However, this does have a significant effect on the risk of spills from the effluent. This effect is missed by this mid-point analysis, but is picked up by the probability plots of sulphate emissions given in sections I-2 and 3. Although high capacity operation of the membrane plant causes lower volumes of effluent from the cooling water circuit, this effluent has a much higher salt concentration. This higher salt loading in the effluent is reflected in a larger impacted land footprint for the system. This trend is not shown for the impacted land footprint for the 4Ml/day pit water case in Figure I-7. The reason for this is that two points reflected have different ash moisture contents, which is the more influential variable affecting the size of the footprint. The moisture content of the ash is not considered in the parametric analysis, as it is assumed to be dependent on the effluent carrying capacity of the system. The high effluent volume produced at low TRO capacities forces a higher proportion of this effluent to be co-disposed with the ash, which causes the higher impacted land footprint seen in this case.

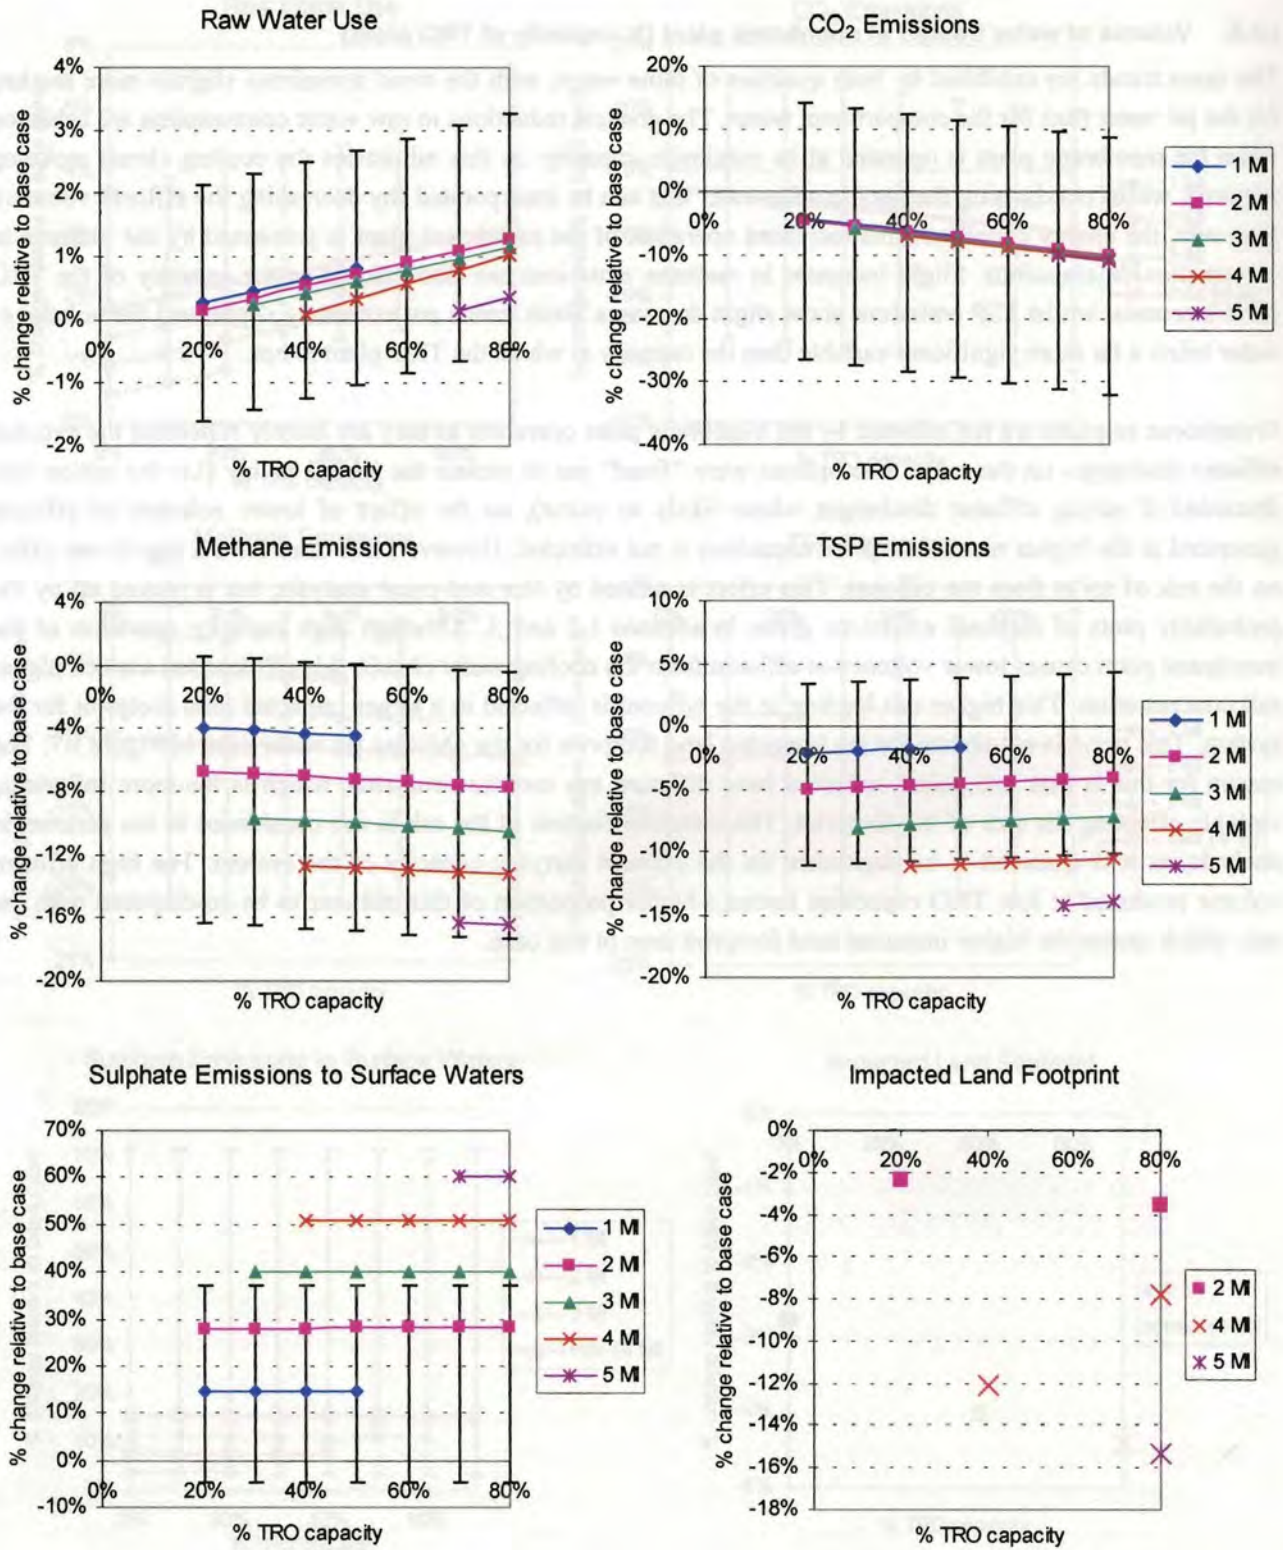


Figure I-7 Effect of volume of water treated in membrane plant and volume of pit water taken, for all the poor-quality mine water used as cooling water make-up. Error bars giving the 90% confidence range are given for the 2Ml/day case (representing empirical parameter uncertainty).

I.4.4. Effect of Membrane Plant Operation

The operation of the membrane plant is central to the inclusion of poor quality water into the cooling water circuit. The salt rejection obtained in the plant (i.e. the proportion of salts rejected to the brine stream) governs the necessary volume of cooling circuit blowdowns, whilst the water recovery affects the quantity of make-up water required. The current membrane plant is operated fairly far from its design specifications, which place stringent requirements on the inlet concentration to the plant and the maximum permeate recovery. The water recovery and salt rejection obtained therefore exhibit a fair degree of variability, which is incorporated into the modelling of the system (the “mid” scenario in Figure I-8). Although this variability was slightly increased for the mine water cases to reflect the uncertainty in incorporating a different quality feed stream and higher operation rates, the magnitude of the potential effect on membrane plant performance for the incorporation of a very poor quality water stream is not known.

An extreme scenario is therefore investigated for sensitivity, where the inclusion of pit water is assumed to significantly decrease the salt rejection obtained in the membrane plant (the “poor” scenario in Figure I-8). A significant consequence of the “poor” salt rejection is that it decreases the maximum volume of pit water that can be incorporated into the cooling circuit before it causes a net increase in raw water consumption (4Ml/day). Also, for the same volumes of mine water incorporated, obtaining a poorer salt rejection in the membrane plant forces it to be operated at higher capacities, and thus with higher energy penalties.

A further possibility is that the membrane plant be refurbished and operated within its design specifications. In this case extremely high salt rejections can be obtained over a small band of variability. This is the “best” scenario in Figure I-8, which exhibits some large differences from the other two scenarios. The reason for these differences are a result of the changes required to the operation of the cooling circuit to ensure the inlet concentration to the membrane plant does not exceed the maximum specified for guaranteed design salt rejections. To keep the inlet concentrations to the membrane plant below their specified limits, the maximum allowable concentration limits for the circulating cooling water have to be decreased. These lower limits decrease the number of cycles of concentration that can be achieved and significantly increase the volume of cooling water blowdowns. To ensure high salt rejections, the permeate recovery is also restricted to a fairly low (70%) and narrow operating band.

These three membrane plant operating states are compared in Figure I-8. The refurbished plant allows higher raw water savings to be achieved, whilst a poor operation of the membrane plant decreases the raw water savings. The higher salt rejections of the refurbished membrane plant causes greater increases in raw water savings with increasing membrane plant capacities, relative to the other two scenarios (this causes the crossovers in the curves for raw water in Figure I-8). The refurbishment was not assumed to affect the power consumption of the plant, so no energy related differences between the systems are seen. The change in methane and TSP emissions is attributed to the change in blowdown volumes, and thus in the volume of water treated, although these effects are slight. As above, there is no change in sulphate emissions due to the “fixed” ZLED policy. The “poor” membrane plant scenario shows an increased impacted land footprint relative to the “mid” scenario, due to its higher effluent volumes.

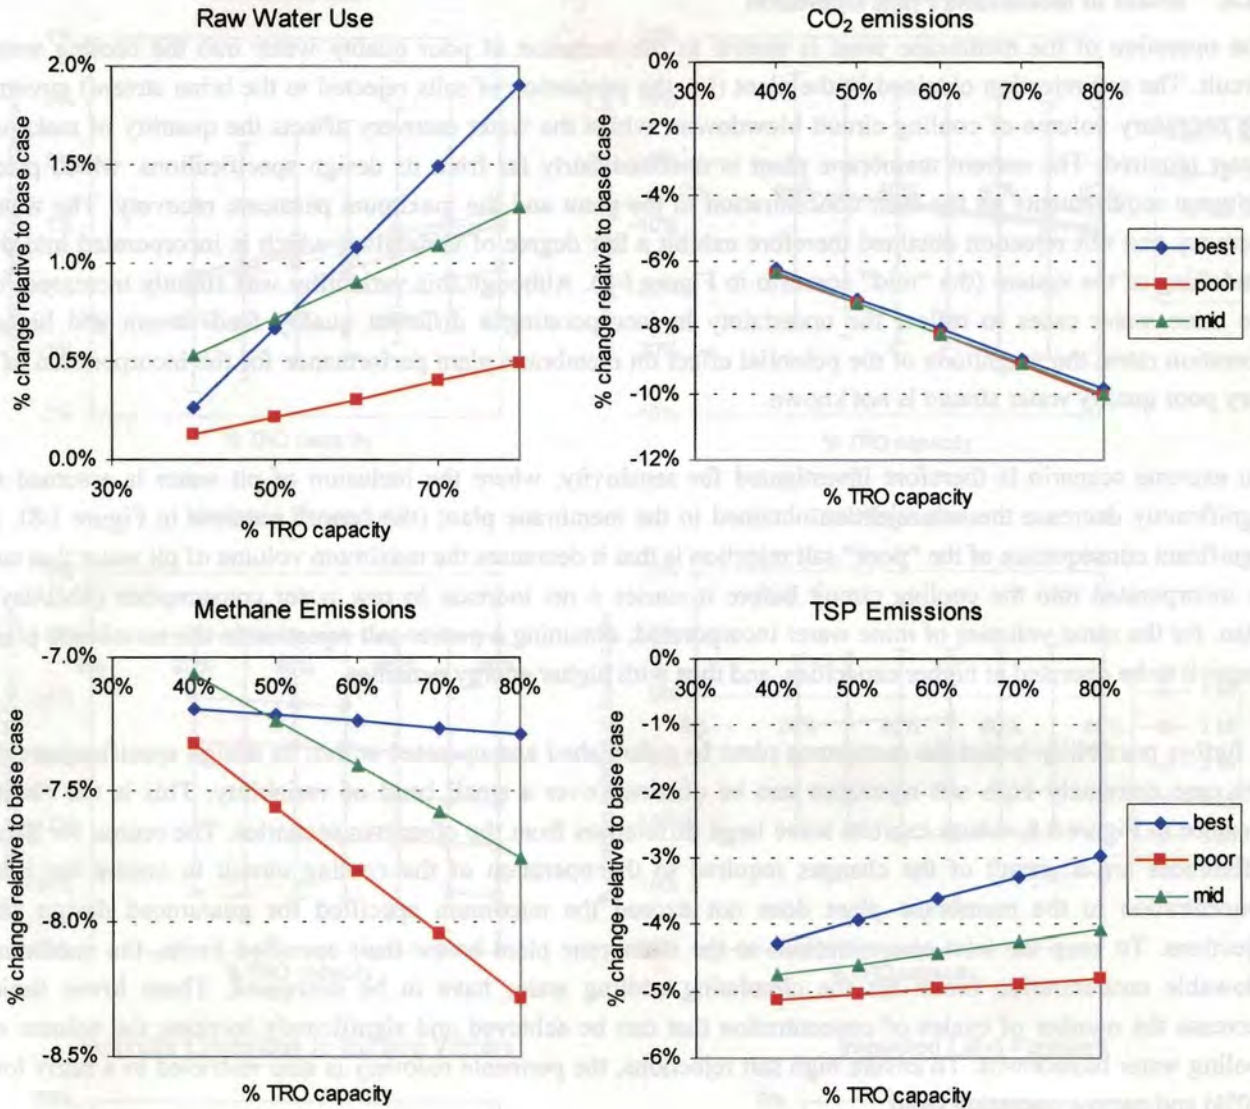


Figure I-8 Effect of membrane plant operation for 2Ml/day pit water incorporated in the cooling water circuit. “Mid” assumes the current operation will be maintained, “poor” assumes the pit water will have a strong effect in decreasing the performance of the membrane plant, whilst “best” assumes the plant has been fully refurbished and is operating at its design specifications.

The differences between the “mid” and “poor” scenarios are generally small and well within the empirical and model parameter uncertainty estimated for the system, i.e. it does not significantly increase the overall uncertainty range. However the fact that the “poor” membrane plant causes a slightly different operating space to be defined (lower maximum volumes and higher membrane plant capacities) is significant, but more from an operability than an uncertainty perspective. The refurbished membrane plant considers a significantly different way of running the water plant, and demonstrates some significantly different trends, particularly for water use.

averaged out in the inventory. Inventory data can therefore not be matched with the real time emissions that are required for an accurate assessment of environmental impact (Huijbregts, 1998a). As with spatial differentiation, presenting a time-dependent inventory will dramatically increase the modelling detail and data required. Calculating an inventory incorporating the daily or hourly variability required by many damage models is unlikely to be feasible, although an inventory with annual temporal differentiation may add a valuable dimension.

The need and/or possibility to include site and/or time dependent information depends on the decision context. In certain studies the decision to be supported affects processes and products whose location is fixed and known, whilst in others the location of the processes are not fixed and remain unknown. The latter case often involves decisions that affect future systems (long-term forecasting), where it may be impossible or even undesirable to include site-specific information. In this way, time and space considerations are interrelated, in that the time horizon to some degree affects the availability and the need for site specific information, with the need for site dependent information decreasing as the time horizon increases (Wenzel, 1999).

2.1.2. System Boundary Definition

A necessary first step in an LCA is the drawing up of the system flowsheet. This identifies all material and energy flows involved in the provision of the product or service, tracing them back up or down the flowsheet to flows either directly from or to the environment. In most studies, there is not sufficient time, data or resources to include all processes identified in the system. Decisions to omit life cycle stages, processes or inputs and outputs are dictated by the application of the study, i.e. each simplification made to the system inevitably comes with an associated increase in uncertainty, which must not exceed that acceptable to the goal of the study. The system boundaries define the exact placement of the system in the overall system environment. Lindfors et al. (1995) identify three different boundaries requiring definition:

- Geographical boundaries
- Life cycle boundaries
- Boundaries between the technosphere and biosphere

Heijungs et al. (1992) add a fourth: The boundary between the product system under consideration and other product systems, which involves the application of allocation procedures.

The geographical boundary specifies the area for which the study is representative. The boundary can be political, e.g. national boundaries, or natural, e.g. ecosystems, or climatically or geologically defined regions. The geographical boundary has implications for scenario development, choice of impact categories and data collection/uncertainty. Certain impact categories may not be relevant on a local scale, and similarly, certain site-specific impacts, are not able to be meaningfully assessed on a global scale. Data collected from an area outside the specified geographical area will introduce an element of uncertainty, depending on the similarity of the areas, and thus the applicability of the data to the system of interest. The spatial positioning also has a large effect on the quantity of data required, particularly around impact assessment, with the level of information required increasing as study moves from a global to a local scale.

To be in accordance with LCA principles, the life cycle boundary should include all stages of the life cycle from primary resource extraction through to final waste disposal, including use, re-use and recycling. However, the products from primary industries spread out to many diverse uses, and a study could not feasibly follow all of these. The system boundary for these systems is therefore “cradle-to-gate” rather than “cradle-to-grave”, as it ends where the material spreads out to its many uses. The amended life cycle analysis tracks the flow of material from its extraction and preparation through to primary processing, and includes all relevant associated processes, including the provision of ancillary materials and services.

The following processes, identified by Lindfors et. al (1995) as usually not quantified as part of an LCA study, are excluded from the life cycle boundary considered in the case studies in this thesis:

- capital equipment and infrastructure,
- accidental spills,
- impacts caused by personnel, and
- human resources.

For many systems the share of releases due to capital goods and infrastructure are less than 5-10% (some notable exceptions are in the transportation sector, renewable energy and agricultural systems) (Frischknecht, 1997). Similarly, personnel related impacts, e.g. canteens or sanitary facilities, are usually not considered because the differences between the scenarios are likely to be small (Lindfors et al., 1995). Accidental spills are not included where process modelling is based on normal or design operating conditions, and even where they are incorporated, e.g. in total annual or monthly emissions, they are usually not considered explicitly. The use of risk assessment has been suggested to address the

and Petrie (1996) suggest the following heuristics to define the appropriate level of flowsheet breakdown (in the context of minerals processing):

- Common function:* If the function of units is common they are integrated into one unit.
- Mass flow rate:* Units which have a high through-put, high reagent addition or high waste generation must be kept separate so that their effect on the total waste stream is not diluted.
- Hazardous waste:* If a unit gives rise to a hazardous emission it must be kept separate so that its point of emission can be pinpointed.
- Energy Intensity:* Units with a high energy consumption must be identified.
- Waste generation:* A unit which generates a specific waste is not always the point at which that waste crosses the system boundary, so it is important to couple waste generation with the process unit where it arises, and not only where it crosses the system boundary.

Another reason for a disaggregated inventory is to facilitate the integration of LCA with tools capable of predicting site-specific impacts. In its current format, LCIA is limited to predicting potential environmental impacts, most often of global significance (see section 5.3). Site-specific time-dependant impacts fall into the gambit of ecological risk assessment (ERA), whose methodology integrates the results of exposure and effects assessments to obtain a probabilistic expression of the effects of environmental change (van Leeuwen and Hermens, 1993). The fate and transport models of ERA require information on the point source of pollution (the type and quantity of emission) as well as spatial information (background information). Information on the point source of pollution is consistent with the type of information generated by a disaggregated LCI study, so to facilitate this integration of LCI with ERA (or other site-specific environmental assessments), the inventory must be structured so that the information is available at a disaggregated level.

However, modelling at a detailed level may not be possible or even desirable in all circumstances. It has been suggested that only the portion of the system of direct interest, i.e. the group of processes over which influence can be exerted by the decision makers, should be modelled in detail (Clift et al., 1998; Trinius and Le Téno, 1999). Clift et al. (1998) distinguish between a foreground and a background system, where the background system is defined as the set of processes whose operation is not directly affected by decisions based on the study, other than the quantity of material (or magnitude of the function) input into the foreground system. Ancillary or secondary material inputs usually fall into this

directly affected by the studied change, i.e. the so-called marginal technologies. This is somewhat analogous to placing certain processes in the foreground system (Weidema et al., 1999).

2.1.4. Inventory Model Structure

The underlying concept of an inventory model is the conservation of mass and energy. The inventory table records all material flows crossing the system boundary. Ideally, mass flows entering should equal mass flows leaving, but usually this is not possible, as there will be some degree of simplification. Also, the vast differences in the magnitudes of certain environmental interventions (often several orders of magnitude) mean that an exact balance is infeasible. However, the mass balance is an indispensable check to see if all the major flows have been accounted for. Similarly, energy entering the system should equal the energy leaving. However, the energy balance is a whole lot more complex to specify, as it requires data on the energy contents and chemical energies of the various inputs and outputs (Heijungs et al., 1992). The lack of standardisation on the definition of energy content, as well as on reference temperatures and pressures, also hampers the drawing up of an energy balance. It is therefore common practice to limit the assessment to a global/approximate energy balance, in which only the major energy flows (usually only electricity, fuels and steam) are incorporated (Heijungs et al., 1992; ISO, 1998).

For practical reasons, the inventory is usually built up from a number of smaller material balances around the individual processes or unit operations included in the system, which are then summed to yield the overall inventory table. The level of breakdown of these sub-balances is essential, as this is the lowest level at which inventory information is available (see section 2.1.3). The sub-balances are calculated on some relevant basis, e.g. 100kg product, which needs to be related to the actual amount required by the system. This is done through the specified reference flow, termed the functional unit, which is the flow required to fulfil the function of the system. Two calculation procedures have been identified to scale the individual balances to the overall inventory; the commonly used sequential method and the mathematically elegant matrix method (Heijungs et al., 1992).

In the sequential method, references are made to earlier processes in a certain ratio, based on the product, material and energy needs of the process which supplies the functional unit, and the quantity of waste processing services required, based on the waste produced by the functional unit. In turn, the requirements of these processes initiate the next round of calculations, following up the entire process

Table 3-3 Data quality indicators/descriptors identified for use in LCA.

	SETAC	USEPA	ISO 14041	Weidema
1	<i>Completeness^a</i> <i>Precision</i> Representativeness Stability <i>Homogeneity</i> <i>Data distribution</i> Applicability / Suitability / Compatibility <i>Bias</i> <i>Accuracy</i> Identification of Anomalies	Precision Representative Data collection method Bias	<i>Completeness</i> <i>Precision</i> Representativeness Time-related coverage Geographical coverage Technology coverage Nature of the data	Completeness Temporal correlation Geographical correlation Technological correlation Reliability of source
2	Consistency Derived models <i>Uncertainty</i>		Consistency	
3	Comparability Accessibility Transparency Peer review	Comparability Referenced Acceptability	Reproducibility	

1. DQIs applicable to empirical quantities

2. DQIs to assess model uncertainties

3. DQIs to assess practical aspects of the study as a whole.

a. Quantitative DQIs are given in italics. All DQIs are treated semi-quantitatively in the pedigree matrix approach (Weidema, 1998b) and qualitatively in the USEPA worksheet approach (USEPA, 1995a).

Whilst the DQI methods predominantly address data quality and not model uncertainties, the ISO 14041 and SETAC list of DQIs include consistency, which refers more to model uncertainty, since it relates to how uniformly the methodology is applied across the various components of the study or across comparative studies. The additional DQIs listed in category 3 in Table 3-3 relate more to the reporting structures of the study, and are thus not relevant to this discussion.

In a distinction similar to that made by Morgan and Henrion (1990) between empirical variables and decision variables, Weidema and Wesnæs (1996) distinguish between environmental data, relating to data on the investigated process, system data, relating to the flow of materials, energy and products through the investigated processes, and performance data, related to the definition of the functional unit. The factors determining data quality are given by Weidema and Wesnæs (1996) as:

- its uncertainty (spread and pattern of distribution),
- its reliability (measurement error, calculations, assumptions and quality control),
- its completeness (sample size and time period, sample representativeness),
- its age,
- the geographical area for which the data is representative, and
- the process technology or technological level for which the data is representative.

the uncertainty assessment into being an “add-on” analysis to a completed mid-point inventory, but it requires arbitrary “cut-off” rules to be set, to determine at which point data can be judged to be unimportant. The qualitative uncertainty assessment has a similar problem in setting the “cut-off” DQI score, in addition to the reservations regarding weighting and aggregating DQI scores, i.e. the fact that ordinal scores can not meaningfully be aggregated, and that weighting assigns the same importance to the DQIs for all types of elementary flows, although their relative importance is likely to change according to the type of flow.

In addition to setting the “cut-off” criteria, subjective judgement is required in setting the criteria at which the “non-selected” data is judged to have no appreciable affect on the results. Maurice et al. identify different criteria according to the type of LCI study. For a relative comparative assessment of products, the aim is to identify the best option rather than to have the best estimation of the confidence interval of the results. The two simulation cases, that including the roughly characterised non-selected data, and that not including the non-selected data, are compared, and if the change in the mean and standard deviation is less than a pre-defined maximum tolerable difference (e.g. 10%), then the influence of the non-selected data can be considered negligible. In an evaluation of cumulative results for a product the aim is to determine the most probable confidence interval of the cumulative results. In this case, it is checked that the simulation including the roughly characterised non-selected data does not change the standard deviation by more than 10% or cause the mean to increase by more than 5%. The percentage change tolerated will depend on the scope of the study, and the values of 10% and 5% are the result of personal experience of Maurice et al. (2000).

The framework developed in this thesis dispenses with either calculating an initial mid-point inventory or a qualitative ranking of uncertainty. It starts from the premise of using broad uncertainty estimates to identify those parameters contributing substantially to the overall uncertainty. An uncertainty importance analysis identifies the influential parameters, so that their definitions can be refined, and an improved estimate of the output uncertainty obtained. This process is shown by the steps in the central shaded area in Figure 4-1. Data ranges are incorporated as they arise during data collection, thereby removing the need to isolate an often unrealistically defined “most likely value”. The model output is computed by simulation, and rank-order correlation coefficients are calculated from the simulated inputs and outputs. The rank-order correlation coefficients provide a relative measure of each input parameter’s contribution to the output uncertainty, which, in turn, is a function both of the magnitude of

The emphasis of the correlation analysis is different to that employed in the framework of Maurice et al (2000). If the uncertainty analysis is incorporated from the start of a study, it is used to shape data collection, e.g. generic estimates can be updated with specific process data if the parameter is returned as important. If the process is applied to a study in which data collection is completed, the emphasis is on identifying those parameters whose distributions require better definition (as in Maurice et al.), and on identifying those parameters limiting the certainty in the results, i.e. those parameters that have high uncertainty importance even once they have been defined as accurately as the current level of data allows.

As in Maurice et al's framework, "stopping" criteria are required, i.e. to judge when a sufficient number of iterations have been completed. This point is thought to be more meaningfully related to the degree of confidence in the results, than a percentage change between the iterations. The parameters can be successively refined until an acceptable variance in the results is obtained, where the acceptable variance will depend on the goal and scope of the study. The acceptable variance is most meaningfully defined in a comparative study, where it can be related to a required degree of confidence in the differences between the options, e.g. iterative model refinements can be made until the confidence intervals show 90% probability of being non-overlapping.

The acceptable confidence limit will depend on the decision context (i.e. the level of risk able to be tolerated by the decision makers), and on the degree of separation between the options. The level of acceptable variance is likely to be revisited as the study progresses, and may differ between selection criteria (i.e. according to the particular impact category or environmental intervention under consideration). Data availability is often the limiting factor in determining the acceptable variance in the results. Additional data collection may be outside the scope of the study, or better data may simply not be available. In this case, the acceptable or limit variance is reached when successive iterations yield little or no reduction in the variance of the results. In this case, a percentage change between the iterations can be taken as the "stopping" criterion, e.g. the 10% suggested by Maurice et al. (2000).

For the first iteration, the uncertainty of each parameter is conservatively estimated (a simple distribution with an overestimate of the range). The rough definitions of the parameters returned with high uncertainty importance are then amended by combining their measurable uncertainty (i.e. the statistically measurable variance of the data sample) with an estimate of the increase in variance due to

Two important parameters governing the accuracy of the probabilistic simulation are the choice of sampling method, and the number of samples to be taken. An overview on the various sampling routines are given in appendix A.1.1. The advantages of a stratified sampling routine outweigh their slight disadvantage (that of not being able to simply calculate the precision of the output), so Median Latin Hypercube sampling is the recommended sampling method. The use of stratified sampling means that the equations for estimating the required sample size given in appendix A.1.1. are no longer strictly applicable (as these are based on Monte Carlo sampling). However, since stratified sampling always returns a precision equal to, or better than, Monte Carlo sampling, these can still be used to calculate the maximum number of samples required for a desired precision, i.e. for stratified sampling they will give an over-estimate of the required sample size (Morgan and Henrion, 1990). For Monte Carlo sampling, 1000 samples give a 95% confidence interval that the 50th percentile (as the least accurate percentile) will be plus or minus 3.5 estimated percentiles. Given that Median Latin Hypercube sampling should be more accurate and the high model parameter uncertainty, sample sizes in excess of this would appear to be unnecessarily precise for most LCI studies.

Whilst the precision of the simulated output can be simply estimated from the sample size, this does not convey the overall accuracy of the probabilistic analysis, which contains aspects that can not be quantitatively determined, e.g. the applicability or comprehensiveness of the assigned input distributions. The accuracy of the model parameter and model form analyses is even less easy to evaluate. In order to streamline the parametric analysis, "best" and "worst" combinations of variables are grouped together, which could result in possible combinatorial effects being overlooked. A trade-off therefore needs to be found between a manageable number of runs and the potential loss of accuracy through the grouping of variables. The accuracy of the model parameter analysis will also depend on the definition of meaningful parameter ranges and intervals at which the parameters are evaluated. The accuracy of an assessment of different model forms can not usually be determined because in most cases the full range of potential model forms is not known, and are continually changing as advancements are made.

The level of accuracy aimed at for each level of the assessment (i.e. each layer in Figure 4-2) will depend on the accuracy of the other levels. For example, there is little point in conducting a highly accurate empirical parameter assessment, with a large sample size and precisely defined probability distributions, in a study where model parameter uncertainty is high, and completely overshadows the

empirical parameter uncertainty. Similarly, if results from different model forms are shown to be widely divergent, a parametric analysis of every possible combination of model parameters would be redundant. The relative importance of the various levels depends on the decision context (see Figure 3-2).

4.1.3.b Limiting Variance / Sensitivity

In the framework outlined above an acceptable variance in the probabilistic output needs to be defined in order to determine whether further iterations are required. The value chosen will be closely related to the decision context of the study, e.g. a decision resulting in large capital expenditure will require far greater certainty than one resulting in a simple process modification. A confidence level in excess of 95% is likely to be required for the former, whilst 60 or 70% may be sufficient for the latter. For comparative systems, the variance does not need to be less than that showing significant non-overlap of the options, i.e. if the options differ widely, a high variance in the output of each option can be tolerated.

The chosen variance limit is likely to be revisited during the course of the study, e.g. if targets have been set that are infeasible within the scope or data availability of the study. In such cases the minimum variance achievable with the data currently available to the study will be reached, and either the acceptable variance must be relaxed or the scope of the study must be amended (e.g. significantly increased modelling detail or data collection undertaken). In some cases, the inherently uncertain nature of the system means that the output variance will never be able to be reduced, even given limitless resources, e.g. strategic modelling of systems for implementation well into the future. A revision of the acceptable variance must be done in accordance with the goals of the study, as increasing the level of uncertainty able to be tolerated in the results may mean the goals of the study can no longer be attained, and will have to be simultaneously revised.

The sensitivity at which the model parameters are considered significant depends on similar considerations, i.e. how robust the results need to be. In a comparative study, it will depend on the degree of overlap between the options, with the choice of a parameter obviously not significant if it does not cause the options to overlap. If the options do overlap, the combined "uncertainty range" from the variance arising from the empirical parameters and the range arising from the model parameters can be set according to the required degree of certainty (e.g. less than 5% overlap between the options). As

4.2. INVENTORY MODELLING TO INCLUDE UNCERTAINTY

This thesis aims to promote uncertainty analysis as an integral part of the LCI model, and not as an add-on to a completed inventory study. The implications that this has on inventory modelling are covered in this section.

4.2.1. Modelling with Variable and Correlated Inputs

In a probabilistic uncertainty analysis a parameter is assigned a range of values, instead of a single value. The parameter is then assigned a value from this range at random, in accordance with the probability that this value is likely to occur. A different value of each uncertain parameter is taken each sample, and the output calculated for each set of values. This random sampling can lead to nonsensical results for correlated variables, i.e. it results in combinations of parameter that are not feasible. The problem of correlations between the input data elements has been identified by Kennedy et al. (1996) and Maurice et al. (2000). Both sets of authors recognise that independently varying certain data elements will result in a violation of the mass balance. Kennedy et al. (1996) recommend that any correlations between the input data elements be specified and incorporated in the model, although they recognise that this would considerably increase the modelling complexity. Similarly, Maurice et al. (2000) recommend that advanced models should be able to represent the interaction between the parameters. They suggest that correlation analyses be used to highlight dependencies between the data and bring to light possible modelling errors.

A probabilistic uncertainty analysis is therefore not applicable to a "black box" modelling approach. In such models nothing is known about the interactions between the variables. Hidden correlations abound, and nonsense is likely to result by varying all variables independently of each other. Deconstructing the model is the best way to avoid correlations. At sufficiently detailed levels, apparent correlations can be broken down into the individual variables and the relationships between them. However, a balance needs to be found between removing the correlations and increasing the number of variables and complexity of the model. From this study it has been found that, at a minimum, emissions need to be related directly to the flow responsible for their generation. Deconstructing certain correlations may be beyond the scope or understanding of the modeller. In such cases conditional probability distributions can be used to specify probabilistic dependencies between variables. These

material resources, fossil fuels and emissions resulting directly from their use, and all other emissions, are adapted from Finnveden and Lindfors' (1998) "rules of thumb" for expected variations in LCI data. An uncertainty factor of 10, 2 and 100 is used for each class respectively. Each inventory item is then specified as a uniform distribution, with the minimum value taken as the mean divided by the uncertainty factor, and the maximum value as the mean multiplied by the uncertainty factor. The specification of such a rough distribution and increase in mean value, is aimed at inflating the uncertainty importance of the background data, as there is considerable uncertainty in applying European data to South African processes and conditions.

Even with these large uncertainty ranges, the background data is found not to dominate the uncertainty in the case studies. The limited information available means that the background empirical parameters are not able to be iteratively refined in subsequent model runs, other than replacing the uniform distribution with a lognormal distribution where this is seen to be warranted. The background processes therefore introduce a certain amount of irreducible uncertainty into the assessment. Apart from obtaining more detailed information on the database, the only way in which the uncertainty could be reduced further would be to bring the process into the foreground and model it in detail.

4.2.2.b Model Parameters

The assessment of model parameters is separated from that of empirical parameters, as they can not meaningfully be analysed probabilistically. Model parameters are broken down into decision variables, model domain parameters and value parameters. Value parameters and decision variables are considered together here, as the same principles apply for their assessment. The treatment of uncertainty in both model domain parameters and decision variables is essentially the same, as both are analysed in a parametric assessment. Their differentiation is thus primarily a function of their different roles within the model.

Decision Variables

Decision variables are those parameters, specifying the scenario, which fall directly under the control of the decision maker. Their most appropriate choice is often not known, and a range of possible values therefore requires investigation. The feasible operating range of the parameter is first identified. The number of samples to be taken over this range then needs to be decided. A very influential parameter would require the range to be sampled at frequent intervals (e.g. 10% intervals), whilst minimum,

Calculating rank-order correlations is extremely computationally intensive, so to decrease the run time it was found better to do duplicate runs with smaller sample sizes, than a single run with a high sample size (a sample size of 700 is found to be sufficient in the case studies). As highlighting the top few variables is all that is required by the analysis, and not the value of the correlation coefficients as such, a smaller sample size does not introduce a loss in accuracy. However, the duplicate run is required because a problem with the analysis, particularly at lower sample sizes, is that co-incidental variance between the parameters is returned as important, i.e. purely random and not deterministic co-variance is picked up between a parameter and the output. By repeating the analysis with a different set of random samples (i.e. re-initialising the random seed of the random number generator between runs), the deterministically co-varying variables can be identified as those returned by both analyses.

The acceptable variance is chosen through a consideration of the decision context and the consequent goals and scope of the study (i.e. the time and resources that can be spent on model development and data collection, and the certainty required in the results). In practice this is likely to be decided upon by first reaching the minimum variance achievable with the data in hand (and thus the confidence with which the decision can be supported), and the decision then taken whether additional data collection is warranted. The uncertainty importance analysis is able to direct this additional data collection effort by ranking the data inputs according to their contribution to the uncertainty, and thus prioritising the data collection.

Ultimately the irreducible variance is reached, where the variance cannot be reduced further without significant changes to the model form, e.g. if an approximate factor or "black box" element is returned with high uncertainty importance, the variance in the output can not be reduced without modelling that element in greater detail. For the background processes, this is analogous to moving the identified parameter into the foreground system, and a disaggregated inventory modelled for the process. Breaking down approximate factors into the variables and relationships underlying them introduces additional uncertain parameters and increases the size and complexity of the inventory models. Increased modelling detail may therefore not always achieve the desired reduction in output variance, or be feasible within the scope of a study. In such cases, the goals of the study will have to be revised instead, to allow for a higher level of variance to be tolerated in the results.

etc.), whilst data quality and availability, and the lack of data quality analysis, are also barriers to the extension of LCA to RA.

Potting et al. (1999) assert that “less is better” and “only above threshold” do not have to be incompatible approaches, and conclude that the only difference between the approaches is the way in which threshold values are dealt with. It is thus possible to combine the methods that predict above or below threshold exposures with the fate and transport models developed for LCA (e.g. by defining sub-categories for exposure situations above and below threshold, and weighting these sub-categories to indicate their relative importance). To unify these approaches additional factors are required, including an identification of source type (e.g. high/low point sources, indoor/outdoor sources), and for each source type, an estimation of the area of increased concentration, the exposure situation (e.g. above or below threshold on the basis of typical background concentrations), and the number of exposed people (Potting et al., 1999). However, this does not overcome the inability of LCA to predict absolute emission quantities. Owens (1997a) therefore disagrees that the “less is better” and “only above threshold” approaches can be compatible, and concludes that the inability to predict absolute emissions, as well as the practice of system-wide data aggregation (that Potting et al. (1999) do address), largely removes spatial and temporal considerations. Thus thresholds for toxic effects do not exist under any circumstances, and all dose-response functions relate linearly to mass loadings. Thus, the inventory accounting procedures of LCA constrain LCIA to providing a relative indicator, able to provide a simplified directional perspective on environmental topics, and the capability of which varies considerably between impact categories (see Table 5-1) (Owens, 1997a).

An interesting study by Tolle et al. (2001) demonstrates conducting LCIA at increasing levels of detail and site-specificity. The study compares two equivalency factor approaches (“less is better” approaches) with simplified RA (“only above threshold” approach) for assessing the toxicity impact potential of emissions to air and water. The two equivalency factor methods, the PBT (Persistence, Bioaccumulation and Toxicity) method and the MFM (Multimedia Fate Model) method, incorporate toxicity benchmarks and data on potential exposure into a single equivalency factor. However, the MFM method incorporates the fate of the emission using Mackay Level III models (see appendix C.4 for an overview of including fate and exposure models in LCIA). The simplified RA approach uses dispersion modelling with site-specific environmental data to determine the environmental concentrations of emissions and compare them with threshold values considered safe for humans and

To quantify the land area affected by the pollution plume, tools to predict leachate generation and pollution plume migration are required, as well as a methodology to define the boundary of the affected land footprint (Hansen et al., 2001b). ERA is used by the LCIA toxicity characterisation methods to compare predicted environmental concentrations with levels believed to cause effects, and thus enable an indication of risk (Mackay and Seth, 1999). It is therefore proposed that ERA can similarly be used to define the boundary between unacceptable and acceptable risk to the environment, i.e. to enable the boundary for the impacted land footprint to be drawn (Hansen et al., 2001b). The differentiation between acceptable and unacceptable risk is based on the environmental value to be protected, and can be translated into acceptable concentration limits of contaminants by a consideration of dose-response and other toxicological data. Alternatively such criteria as water quality guidelines can be used. Basing the indicator on the area of land affected rather than the actual toxicological effects is consistent with the practice of basing the assessment on easily measured midpoints, rather than the endpoint (the actual environmental damage). The use of water quality guidelines is consistent with basing the assessment on midpoints, as these have been worked back from the endpoints (i.e. to ensure protection of the environment) (Hansen et al., 2001b).

The extremely complex nature of the reactions taking place and the many species involved have limited the footprint approach to a consideration of salts. Salts are more readily modelled as the chemical reactions and mechanisms are better understood, whilst trace metal concentrations in leachates and soils are difficult to predict (their mobility is determined by aqueous and heterogeneous chemical reactions, particle coagulation and flocculation mechanisms, and are affected by pH, salinity, redox state and available ligands). Their inclusion would thus increase the modelling complexity significantly (Hansen et al., 2001b). Modelling the salts only is consistent with the level of inventory information available, so this is not seen as a particular limitation of the footprint approach. On the contrary, the ability of the method to characterise the solid waste on this reduced information set, is considered a significant strength of the method. The interrelation between salinity, metals mobility and bio-availability is complicated. With increased salinity, metals mobility increases due to inorganic salts competing with metals for adsorption sites and dissolved complexation with salt anions, whilst metals bio-availability decreases as there are presumably less metals in free ion form. The toxicity of trace metals can thus be said to be inversely related to salinity (Hansen, 2001).

Since salinity and metals mobility are related, it follows that trace metals will follow a salinity pollution plume. In the footprint model, salinity is measured as the ionic strength of the solution rather

Major environmental impacts stem from the high ash content of Southern Hemisphere coals, which require beneficiation to produce coals of acceptable quality for the world markets. Of the 279 Mt of run-of-mine (ROM) production in 1997, only 219 Mt was of saleable quality, resulting in 60 Mt of discard requiring disposal (Prevost, 1998). The resultant discard dumps are responsible for some of the most serious environmental effects of coal mining, including land sterilisation and groundwater contamination. On exposure to air and water the pyrites oxidise to form sulphuric acid, and iron oxides and hydroxides, which cause the pH to drop. The acid produced reacts with bases in the rock to form salts and consequently mobilises any heavy metals present. The resultant acid mine drainage (AMD) contains elevated levels of salts (mainly calcium and magnesium sulphates) and metals (predominantly iron, manganese and aluminium). The pyrite-rich discard is also susceptible to low temperature oxidation (so-called spontaneous combustion), and subsequent release of toxic air pollutants. AMD and spontaneous combustion can be minimised by preventing water and air getting to the pyrite and other *sulphidic minerals*. *The power stations mostly burn ROM coal, so are not responsible for discard production, although some power stations are supplied by dual product mines. These mines maximise their coal production by producing a high-quality coal for export, and a medium quality power station coal, which contains a portion of the washing discard blended in with the ROM coal.*

The location of South Africa's coalfields is significant. The Mpumalanga / Eastern Gauteng / Northern Free State region, where 65% of the reserves are to be found, has been extensively farmed, with little natural environment remaining. Coal mining therefore has little residual impact on natural ecosystems and land rehabilitation is usually able to restore the land to an acceptable state (Wells et al., 1992). However, from a water quality perspective, the coalfields occur in the worst possible location, since most mines are situated in the vulnerable upper reaches of Southern Africa's major river systems (Wells et al., 1992).

In response to the environmentally damaging nature of mining operations, the Minerals Act of 1991 was implemented, which requires any company wishing to mine, or those already doing so, to submit Environmental Management Programmes Reports (EMPRs). These form a legally binding contract between the mine and the Department of Mineral and Energy Affairs (DMEA) to show how they will protect the environment during the working life of the operation, and how they plan to rehabilitate the working areas after closure. An important feature of the legislation is the control over water resources. Mines require authority to abstract water and to discharge wastewater to designated watercourses, and

are responsible for ensuring that the discharged water complies with quality requirements. Contaminated water has to be adequately contained, and rainwater kept free from pollutants as far as possible (Baxter, 1993).

6.1.2. South Africa's Power Generation Technologies and Plant Mix

This study is limited to a consideration of the environmental effects of the coal-fired plants only, with the relatively small contributions from nuclear, gas and hydro not addressed. It is also limited to Eskom power stations, and does not include the small contribution (less than 5%) from old municipal power plants. In addition, the study only addresses generation, i.e. excludes transmission and distribution. An overview of the technologies employed in Eskom coal-fired power plants is given in appendix G.1.2, and covers the three broad aspects of the plants:

- *Boiler Plant* dealing with the "solids section" (milling and combustion of the coal).
- *Water Plant* dealing with the "water circuits" (production of potable, ultra-pure and cooling water).
- *Waste Management* where the "solids section" and "water circuit" overlap (co-disposal of combustion residues with liquid effluent).

South Africa generates electricity on an impressive scale. The most-recently constructed stations were built in a "six-pack" format (6 units of approximately 600 MW each). The generating capacity of the units reach as high as 686 MW, with one station (Kendal), the largest bituminous coal-fired power station in the world. Table 6-2 summarises some main features of Eskom's coal-fired stations.

All the stations employ pulverised fuel firing. The composition of the flue gas depends on the coal quality and conditions in the boiler, with carbon dioxide, water vapour, nitrogen oxides and sulphur dioxide the principal products. Thermodynamic equilibrium may not be reached and other compounds, such as methane, ammonia and carbon monoxide may be formed in small quantities. South African coals are generally low in chlorine and fluorine, although some HCl and HF is emitted, together with other volatile trace elements present in the coal, e.g. Hg and Se. Some of the less volatile elements condense onto the fly ash particles, where they also can be emitted e.g. As, Pb and Cd (Swaine and Goodarzi, 1995).

6.2. THE COAL-ELECTRICITY LIFE CYCLE

Electrical energy is an input into practically all processes, and the burdens associated with its provision have been found to significantly contribute to the overall burdens in many product- and process-life cycle studies. Energy data are thus indispensable in an LCA study, and consequently many life cycle inventories have been developed for energy-generating systems (see summary in Spath et al., 1999). LCA databases typically contain electricity data for many countries, built up by combining life cycle inventories of generating systems in the correct proportions to represent the fuel mix of that country, e.g. the IDEA and ESU databases (Frischknecht et al., 1994; Lubkert et al., 1991). Most LCA studies have tended to focus on the aerial emissions and energy requirements of electricity generation, with a lesser emphasis on solid waste and water emissions. Leachate from ash dumps and coal stockpiles is typically neglected, and water use often not included in the assessment (Spath et al., 1999). This current research study has a particular emphasis on water use and water-related impacts, because of the importance of these in the South African context.

In addition, the deficiencies of LCA with respect to solid waste have been identified (see section 5.2). There is thus a renewed interest in developing methodologies able to assess the implications of waste management in the LCA framework. This study contributes to this process, and suggests some methodological adaptations to ensure the localised impacts arising from coal mining and ash waste disposal are not overlooked (see section 6.2.6).

6.2.1. A Life Cycle Inventory Model of South African Coal-fired Power Generation

The life cycle presented here is that of coal; its extraction, processing, combustion and disposal of its residues. If viewed from the perspective of electricity, the life cycle is one of "cradle-to-gate" and not "cradle-to-grave", as the assessment ends at the generation of electricity. Transmission, distribution and use of the power are not covered. It thus provides an inventory of *undelivered* electricity. In addition, only process-related emissions are assessed. Burdens associated with offices, workshops etc. are not incorporated in the assessment. The primary aim of the study is to provide an inventory of the pollutant and resource flows associated with power generation in South Africa. This information is expected to be of value to other South African LCA studies, where relevant South African data can be used, instead of the European or American data commonly found in LCA databases. According to the application

not always straightforward whether the by-product material should be defined as a waste or as a potential future resource (e.g. coal washing discard). Regardless of its definition, it is the effects of containing the material that are of concern (i.e. the land occupied and the emissions from the dump), and not the mass of material. The dumps/stockpiles are treated as “pseudo” unit operations, in that they generate emissions (dust, leachate and stormwater run-off) which are either emitted or incorporated back into the process (in the case of surface run-off).

All ancillary materials comprise the background system. These include the production of fuels (those used “on site” and in the transport of materials), water treatment chemicals (lime, sulphuric acid etc.), and process ancillaries (filter bags, ion exchange resins etc.). “Cradle-to-gate” product inventories from LCI databases are used to supply the inventory data for these substances, as is consistent with the LCA methodology for processes consigned to the background system. These substances fit the definition of background processes, as only their input quantity is controlled by the foreground processes (which supports the use of aggregated data), and the quantities input into the foreground processes are small (which supports the use of generic data). The ancillary materials are consumed in the foreground system processes, so the emissions from their use are incorporated into the foreground system, although their production falls into the background system.

6.2.3. Flowsheet Construction

The aim of this case study is to provide inventory data on South African power generation (i.e. the Type I system, defined in section 2.3). As proposed in section 2.3.1, systems in this context require a flowsheet of sufficient resolution to allow the generation of high quality inventory data. Thus it should be detailed enough to allow for transparency (i.e. it should be clear what has been included), and for an assessment of uncertainty, either qualitative or quantitative. However, the primary driver of the system resolution is the level of aggregation in the available data. The highest quality data should be used, regardless of its resolution. For example, a metered overall electricity consumption figure may be known, but not the individual machinery requirements. The measured overall value should be used rather than the sum of estimated requirements of pumps, fans etc., as it is of higher quality. The fact that it gives no insight into which unit operations are consuming the most power is not relevant in the context of the study.

net power generated (i.e. MWhSO) in a specified time period, in this case a single year. The period of operation is related to the power sent out by the load factor. This is the fraction of the total time available during which the power station is operating. A simplifying assumption made is that the power station is assumed either not to be operating or operating at full load.

The power sent out needs to be related to the time of operation because certain burdens cannot be directly related to the power sent out (or indirectly the flow of coal through the system), e.g. those emissions related to coal stockpiles, waste dumps and certain aspects of mining. The emissions generated from these processes act over the life of the mine and power station (and in some cases beyond) regardless of whether the station is operating for that specific time period or not. Specifying the functional unit as an operating period of interest, also enables accumulation or depletion of the coal stockpile to be taken into account.

The system boundary ends at generation, so the burdens associated with the transmission, distribution and use of the power are not included. The quantity of power used in an upstream process will not be the same as that sent out at the station, and should be inflated to account for power losses during transmission and distribution, so that the inventory reflect the power sent out by the station, rather than the power consumed by the system of interest. The bulk of the power losses occur during transmission, and are usually estimated per km of power line. The amount by which to inflate the electricity use can therefore be relatively easily estimated, provided the approximate distance between the generating source and the system of use is known. However, identifying the generating source in a power grid system is often not possible, although the fact that the power stations are all located fairly close together in South Africa makes a rough estimate possible.

A small portion of the ash produced ($\pm 5\%$) is used in cement production. Ash is not considered a by-product, but a waste, so a portion of the system burdens are not allocated to its production, nor is the system credited with the burdens avoided by replacing a portion of virgin material with ash in cement manufacture. The system does reflect the benefit of landfilling less ash, although the effect is negligible, since the ash sold makes up such a small percentage of the total ash produced. Eskom's stations are predominantly supplied by dedicated collieries, although three of the stations are supplied by dual product mines. These collieries produce a high quality coal product in addition to the power station coal, and by so doing, also produce a discard coal. A simple mass-based allocation rule has been

collection efficiency of the ESPs, so achieves low particulate emissions even with a high ash coal. Mining particulate emissions affect the overall particulate loading to a lesser degree, with opencast mining generating significantly more dust than underground mining. The ash disposal method also affects particulate emissions, albeit to a much smaller degree, with dry ash disposal creating more of a dust problem than wet ash disposal.

The volume of solid waste generated also does not follow the ash content of the coal exactly, but is influenced by the thermal efficiency, the particulate removal efficiency and any ash sales occurring. The Dry/Dry station, although burning a slightly higher ash coal, produces slightly less ash requiring disposal than the Wet/Wet station, as both it and the Wet/Dry station sell a small portion of their ash. The high mass of solid waste for the Wet/Wet station reflected in Table 6-3 stems from high discard coal volumes, and its ash volume is in fact between that of the Wet/Dry and Dry/Dry station. The Wet/Dry station is burning an extremely high ash coal, and thus has the highest ash volume requiring disposal.

The power station coals are typically supplied without coal preparation (other than crushing), although the mine supplying the Wet/Dry station is mining such poor quality coal that a simple de-stoning wash is required. This has the side-effect of producing a low sulphur coal, which accounts for the lower SO₂ emissions of this station. The discard/shale produced from this wash is not included in the waste volume as it is used to backfill the opencast mine. The mine supplying the Wet/Wet station is a dual product mine and is producing both a discard that is used as backfill (not reflected in the inventory) and one which is dumped, causing the high waste volumes seen in Table 6-3. The allocation of this discard to the power station coal is questionable, and results from the simple mass-based allocation method used to allocate mining burdens between the two coal products (see section 6.2.4).

6.3.2. A Consideration of Uncertainty

The aim of this study is to provide an inventory of South African electricity that can be used in the life cycle design and analysis of South African products and processes. The study uses a combination of recent process-specific data, literature data and generic LCI data to produce life cycle inventories for three representative South African power plants. These inventories are then extrapolated to cover the potential future South African generating mix. This extrapolation is expected to introduce considerable uncertainty, in addition to the uncertainty in the inventory itself, i.e. that due to uncertainty and

storage capacity brought back on line. The inventories used to represent the three capacity “types” are kept the same (the generating mix of 1996 is used to represent currently installed capacity, a slightly amended inventory of the Wet/Wet station to represent the capacity brought out of storage, and the inventory of the Dry/Dry station to represent new coal fired capacity). At a low percentage of the in-storage capacity re-commissioned, new capacity will most likely be required (e.g. if only 20% of the capacity in storage is re-commissioned, 60% of the predicted capacity shortfall will have to be supplied by a new plant). At higher percentages of the in-storage plants re-commissioned, it is likely that no new capacity will be required (e.g. re-commissioning 60% of the plants in storage is sufficient to supply the predicted capacity shortfall (most likely estimate), but taking the uncertainty of this prediction into account, there is a 25% probability that 35% of the capacity will have to be supplied by a new plant. If all the in-storage capacity is re-commissioned, there is only a very low probability (less than 10%) that new capacity will be required at all for the time frame investigated.

The inventories used to represent the three capacity “types” are investigated in Figure 6-4, which varies the inventories chosen to represent the installed, re-commissioned and new plants independently of each other in a sensitivity analysis. The main features of these figures are highlighted below, and are explained in greater detail in appendix G.3.2. Box and Whisker plots are used in Figures 6-3 and 4 as these are able to display the range in the data caused by the uncertainty in the prediction of the capacity required in the mid-term, as well as the most likely value (the median). The “boxes” give the interquartile range (the range spanned by the mid 50% of the data), whilst the “whiskers” give the total range of possible values.

The majority of inventory items is surprisingly insensitive to the percentage of plants in storage re-commissioned, with only those environmental interventions that differ very markedly between the old and new plants significantly increasing the uncertainty range over that due to the future energy demand (as shown by the range spanned by the “whiskers” in Figures 6-3).

Surprisingly, given the high ash content of the Southern African coals, the analysis shows marginally lower total TSP emissions for the South African grid than the UK average inventory. This can probably be attributed the use of limestone sorbent for FGD in the UK power stations, as limestone mining is associated with high particulate emissions. Unfortunately no TSP emissions are reported for the European stations, so the individual coal-fired stations can not be compared across this category. South Africa's high contribution to land transformation is attributed to the high volumes of ash produced, and the consequently large areas required for ash dumps. The increasing number of land-intensive opencast coal mines also contribute to this high land use. The high land use of opencast mines is shown by the marked difference between the two South African stations.

On an individual coal-fired power station level, the potential contribution of South Africa's stations to global warming and acidification is very competitive to that of Europe. The very large discrepancy between the modern and old European station is probably due to the modern station's incorporating FGD. The South African stations do not have FGD units, but the relatively low sulphur Southern African coals results in their releasing only marginally more sulphur dioxide per MWhSO than a plant with FGD. The South African stations show a marginally better performance in global warming potential than the modern European plant. This is attributed to the low methane emissions in South African coal mines, as the South African inventory actually shows marginally higher CO₂ emissions than the modern European plant. Also, the South African system has very low transport emissions, as all of the stations are mine-mouth, further reducing their hydrocarbon emissions. The relatively high contribution to global warming from the "poor" European station is most probably due to the older station's low thermal efficiency.

The South African stations are shown to consume less coal reserves than their European counterparts. This is not because they are more thermally efficient, but is a consequence of their burning near-discard quality coal. In the South African inventory, the mass of coal reserves utilised is calculated by dividing the energy extracted by a reference calorific value of coal reserves (taken as 29 MJ/kg). The comparison is somewhat questionable, as it is not known how the reserves in the ESU database are calculated. The actual mass of material burnt is not a meaningful indicator of reserve depletion, as the South African stations would then show far greater values than the European stations, and not the fact that by burning such poor quality coal they are freeing up higher quality coal for other uses. The higher oil reserve consumption observed for the European stations is attributed to higher coal transport distances, since

the South African power stations are all mine-mouth. Both of the South African power stations are more water-efficient than the two European stations, with the considerable water savings of dry cooling very evident.

The South African inventory shows disproportionately high TDS emissions relative to Europe and the UK. This is most probably due to underestimation of TDS emissions in their inventories, as they show far higher emissions than the South African inventory in other water emission categories, particularly chlorides, sodium and nitrates. Their sulphate emissions are also anomalously low, which may also be due to an error in the database, as the individual coal-fired stations show fairly high sulphate emissions. No TDS values were reported for the individual European stations, although they show higher emissions than the South African stations in most water categories. This can be attributed to the ZLED policy of the South African power stations and many of their tied collieries, which requires that all effluent and storm water run-off be collected and retained on site, where it is either re-used or evaporated in containment dams. Thus very low surface emissions to water occur, with most of the emissions reflected in the inventory due to leachate from the dumps, and seepage from evaporation dams.

6.4.2. Electricity Input into a Zinc Refinery

The high electricity consumption of minerals refining offers some interesting opportunities to explore the application of the SA electricity LCI in Type I decision situations, i.e. as background information into a separate decision system. In particular there is the opportunity to explore the possibility of whether the electricity LCI used could change the outcome of the study. In addition to their high electricity requirements, minerals processing technologies are interesting in that they encounter some classic “problem shifting” situations. Newer technologies often employ electricity rather than burning coal on site, which shift burdens to the site of electricity production, whilst innovative technologies, such as bio-processing, change the nature of the wastes, and can potentially shift burdens from one release medium to another. This section examines a small aspect of a study presented by Stewart (1999). The purpose of the case study was to select the optimal technology for a proposed zinc refinery in the Eastern Cape in South Africa.

The current excess generating capacity in South Africa has led to Eskom providing incentives to encourage local electricity-intensive projects. An example of such a project is the Hillside Aluminium

Smelter in Richards Bay, in which the cost of electricity is linked to the selling price of the metal, thus essentially fixing the profits of the process. The electrowinning of base metals is a similarly electricity intensive process, although the lower throughput of material means that their total electricity requirements are not nearly as high as aluminium purification. Nonetheless, a zinc refinery was proposed, and a detailed and extensive EIA process ensued. Although the project did not go ahead, the technology assessment study showed the kind of considerations that needed to be taken on board in selecting the optimal technology combination for the process.

The initial phase of the technology assessment identified the possible flowsheet options shown in Figure 6-6. Both a single technology route, and a combination of technologies were considered. Goal programming, considering a number of economic, social and environmental criteria, was used to determine the optimum flowsheet. Details of the technologies considered and the criteria used in their selection can be found in Stewart (1999). The analysis found that a combination of technologies was always preferred, with the best option a 44:56% split between pressure leaching and roasting, with all of the roast product treated in the Imperial Smelting Process (ISP). A considerable drop in performance resulted when the single technology route was chosen. When considered singly, the pressure leach/electrowin process was favoured, with the roast/leach/electrowin and roast/ISP processes scoring very similarly for the attributes considered (roast/leach/electrowin was marginally favoured).

The contribution of electricity provision to the total burdens for four key environmental interventions is shown in Table 6-8. The model in Stewart (1999) only considers electricity use in the electrowinning process, so no electricity burdens are reflected for the ISP route. In addition, only the major flows through the zinc refining process are considered (feed extraction and processing, ancillary materials, and transport are not included), so the roaster is the only process, other than electricity production, for which emissions to air are reflected. The contribution to the total emissions from electricity production are surprisingly small. Only water consumption shows a notable contribution, since the zinc refining technologies are not terribly water intensive. The very high emissions of CO₂ and SO₂ from the roasting process completely dominate those from electricity production. It is thus only in the pressure leach process that CO₂ and SO₂ emissions from electricity provision are significant (all are from power generation), although the actual quantities are very small compared to those from the roasting process. The high mass of solid waste produced by zinc refining means that electricity provision is insignificant with respect to the total mass of solid waste generated.

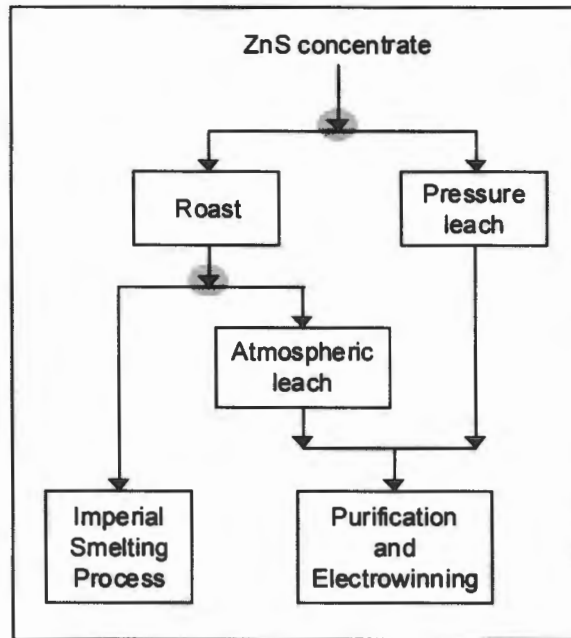


Figure 6-6 Flowsheet options for zinc refining. The shaded circles show where possible stream splits can occur (Stewart, 1999).

The contribution of electricity burdens to each refining process was calculated for the range of values given in Table 6-7, reflecting a range of possible SA electricity LCI combinations, and the uncertainty underlying them. Surprisingly, given that the range in Table 6-7 is fairly extensive, this was found to have no effect on the total values found for the environmental interventions considered. Even for these interventions where the contribution from electricity production is high (water consumption, and CO₂ and SO₂ for the pressure leach process), the quantities are sufficiently low that at the level of reporting (two significant figures), the effect of the range is insignificant. The uncertainty in electricity production is thus not sufficient to change the outcome of the technology selection process. However, the uncertainty in the modelling of the zinc refining process is not investigated. A single value of electricity consumption by the electrowinning process is considered (3.6 kWh/ton zinc produced). If this value were to be higher, uncertainty in electricity production may well be shown to be significant.

Table 6-8 Selected environmental interventions per ton of zinc produced. Also shown is the percentage of the total arising from electricity provision, for the particular intervention considered. Calculated using the models and assumptions of Stewart (1999).

	Pressure Leach, Electrowin		Roast, Imperial Smelter		Roast, Leach, Electrowin		Optimum Technology Combination	
	t / t zinc	% due to electricity	t / t zinc	% due to electricity	t / t zinc	% due to electricity	t / t zinc	% due to electricity
Water consumed	0.05	10-12%	0	0%	0.05	9-12%	0.03	10-13%
Carbon dioxide	0.004	100%	3.4	0%	0.90	0.3-0.4%	1.5	0.1-0.2%
Sulphur dioxide	2.8e-5	100%	0.40	0%	0.39	0.01%	0.18	0.01%
Solid waste	0.98	0.06-0.07%	0.91	0%	0.93	0.07-0.08%	0.95	0.04%

The above brief study shows that incorporating uncertainty information on the background LCI used allows for a greater sense of surety in the results. Where the chosen option can be shown to be robust across the known range of background information, the choice of LCI to use becomes less significant (if a mid-point analysis is to be pursued). However, not only should the sensitivity to the background LCI be checked, but variations within the foreground should be considered as well (i.e. the operating range and the empirical variability within the process(es) under consideration), since these affect the quantity of background product/service input into the foreground. The following two chapters explore case studies where both foreground and background uncertainty are taken into account.

CHAPTER 7

LCA MODELLING FOR DECISION MAKING IN PRIMARY INDUSTRIES: TECHNOLOGY CHOICE CASE STUDY

Strategic and tactical studies are characterised by their lack of precise information, as both require modelling of future systems, where at best, process data from similar systems can be extrapolated to the future system. This constraint is typical of the early phases of a design study, where the first step is to choose among a number of possible technologies and operating conditions. An environmental component in the early phases of design is essential, as it is here, before the design becomes fixed, that the greatest opportunities for affecting environmental improvements exist. However, in the early stages of design the unspecified nature of the system means that full-scale LCA is severely restricted (Azapagic, 1999; Keoleian, 1993). The emphasis of the LCA therefore needs to be on scenario analysis and uncertainty assessment.

This chapter looks at the use of LCA to inform strategic decision making in primary industries, and presents a case study on selecting the optimal technology and operating conditions for re-powering an in-storage power station. The case study is broken into two parts, and represents the progression from a strategic to a tactical decision level (and ultimately down to an operational level) in a design study. The first phase of the study is typical of a strategic technology scanning study, where all relevant technologies are investigated. This is consistent with a screening or first-pass LCA, and requires the evaluation of a large number of options, about which little specific information is known. The inventory is therefore calculated at a relatively low level of resolution, incorporating predominantly generic LCI data. The next phase of the study is typical of the early phases of a design study, where the choice has been narrowed, and the chosen options explored in some detail. Here the system is much more precisely defined, which allows for more detailed data sets to be collected, and a more detailed inventory to be constructed, i.e. it focuses on a particular technology, with the options narrowed to the possible operation of the chosen technology, instead of spanning the "typical" operation of a number of technologies. This phase is typical of that in the early stages of design, where the broad outline has been drawn, although the actual configurations of the system are yet to be decided.

7.1. DECISION CONTEXT

The features of tactical or strategic decisions are discussed in section 2.3.2. These features are consistent with those demonstrated by this case study, as the study is conducted without a specific implementation date or location in mind, and involves elements wider than the primary sphere of influence of the decision maker, i.e. it involves aspects not directly under the control of the decision maker (in this case the power supplier) (e.g. the supply, availability and quality of future discard coal sources). The case study system therefore shares the properties of tactical/strategic decision systems described in section 2.3.2.

Studies to support tactical and strategic decisions are distinct from operational and "historical-type" studies primarily with respect to the quantity and quality of data available to them. Tactical/strategic studies require data to characterise a system different to that from which the data is collected, and one to be implemented some time in the future. The other two study types primarily use data from the actual system of interest to characterise a current system, or one where changes occur in the short term. Tactical and strategic studies thus require data to be predicted for the relevant mid or long term marginal technology, and are associated with a significant amount of inherent uncertainty, especially as strategic decisions covering long time horizons often require that developing technologies be incorporated into the assessment. Coupled with this high empirical uncertainty, is the high uncertainty in model parameters (decision variables) typically encountered in tactical and strategic studies, as a result of the, as yet, loose definition of the system. A feature of tactical studies is thus their emphasis on scenario analysis, and the exploration of possible operating states.

Although the case study is conducted without a specific power station in mind, some site specificity is introduced, as the locations of the power stations currently in storage are known. The older stations that are no longer operating, but have the potential for being reconditioned and brought back into service, all fall into a fairly localised region of the country, which mirrors the localised coal-producing region (see Figure 6-1). Site specific considerations therefore play a more significant role than usually found in tactical studies in the selection of the relevant impacts to be considered, as well as in the importance attached to these impacts.

7.2. PRELIMINARY TECHNOLOGY SCANNING STUDY

The aim of this study is to evaluate technologies able to reduce SO₂ emissions from coal-fired power generation. A preliminary technology screening takes place very early on in a design study, and is used to narrow down the choices to two or three options, which can then be designed in detail, and on which full-scale LCA studies can be conducted. The study is thus required to look broadly at all possible manners of sulphur-removal, relevant to all possible plants and locations.

Technologies developed to reduce SO₂ emissions demonstrate well the dilemma of choosing between disparate technologies offering relative degrees of improvement in environmental performance, and which result in a trade-off between environmental impacts. They range from 'end-of-pipe' control strategies, such as FGD systems, to modifying the combustion conditions, or burning a "cleaner" coal. LCA provides a systematic methodology to organise information about the technologies on a consistent basis, enabling their incorporation into an optimisation strategy, where together with economic and social information, the optimal technology for a particular situation can be identified. The study is thus intended to provide an environmental argument to the usual economic assessment in the preliminary technology-screening phase of a design study.

A full account of the study can be found in appendix H-1, with only a brief overview presented here.

7.2.1. System Definition

The objective of the study is to identify feasible technologies to reduce SO₂ emission from a coal-fired electricity plant, and to evaluate the relative performance of these technologies. A life cycle basis is required to evaluate the technologies because all shift the sulphur burden within the life cycle, i.e. all the technologies result in lower emissions of SO₂ from the power station stack, but at the expense of higher solid and/or liquid waste volumes.

As comprehensive a range of technologies as possible was considered. Three possible intervention "sites" to effect sulphur-removal in the coal-electricity conversion chain were identified, and two representative technologies chosen for each of the three sites, selected according to their applicability to the South African coals and conditions, as well as their level of demonstration. The technologies were required to be equally applicable as a retro-fit on an existing plant, or built with a new plant, i.e.

they had to be compatible with the existing technology employed on the South African power stations.

The following technologies were investigated:

- Before combustion: Conventional dense medium coal washing (at medium and high density), and microbial desulphurisation of the coal.
- During combustion: Combustion in a fluidised bed boiler with in-situ desulphurisation (atmospheric bubbling bed, and circulating fluidised bed).
- After combustion: Flue gas desulphurisation (wet limestone process, and a lime spray dryer process).

The inventories developed are for typical operating conditions, and are intended to give a measure of average operating performance. They are therefore calculated using mid-point values, representing the "most likely" performance of the technologies. The major assumptions and data sources are given in appendix H.1.2.

A significant problem with modelling data deficient systems is the danger that emissions that cannot be quantified are simply ignored. This is a concern when dealing with systems for which the data is not of the same level of completeness, as this leads to the under-specified system appearing better than it is, and the more comprehensive system being prejudiced by its more complete inventory. It is therefore essential that the inventories be specified to the same degree of completeness, i.e. if an environmental intervention is specified for one system, it must be specified for all the systems being compared. A quantitative analysis is therefore inappropriate for interventions which are known for only some of the options, and the comparison should rather be limited to a qualitative assessment of that intervention.

An alternative approach, is to estimate the intervention for the systems where it is not known, and place high uncertainty on this estimate. However, a full quantitative data uncertainty assessment is not feasible at an early screening phase of a design study, so uncertainty management of an estimated data input has to be via a qualitative uncertainty assessment. The simplest of these is placing a "flag" on the uncertain data element (the so-called "post-it-note" approach). However, this approach was not pursued because of problems with aggregating data "flags" through the assessment, and the danger that the estimated data give a false sense of completeness. Environmental interventions that are not known for all options are therefore excluded from the quantitative analysis, and only considered qualitatively. A qualitative data uncertainty matrix, using a pictorial representation developed by Graedel and Allenby (1995), is used to provide an indication of the quality of the data used in the quantitative component of




Figure 7-1 Comparative performance of the most promising technologies to reduce sulphur emissions, relative to an older PF station with no additional environmental controls, and burning ROM coal.

The most significant feature of this combined system is the removal of the discard and the “avoided” burdens which result. This highlights the considerable limitations of reporting only the quantity of waste, and not the potential risks associated with its disposal, as these “avoided” burdens are not reflected at all in the quantitative analysis. Table 7-1 presents a qualitative comparison of the medium density washing system and the combined system. The qualitative assessment is essential, as it reflects the positive effects of converting the polluting discard dumps to the more benign FBC waste, and not merely the benefits of utilising an otherwise wasted energy source in a cleaner combustion technology.

The above analysis shows that quantitative LCA alone can not reliably inform decision making in information deficient systems. At low levels of detail it can potentially give misleading results, especially if an influential aspect of the systems is omitted, or if one system is more complete than the others. In addition, the inability of LCIA to assess the risk of site-specific impacts results in misleading comparisons of technologies in which waste management impacts are important. A combination of a quantitative and qualitative approach is therefore recommended for technology scanning studies of primary industries, where a qualitative assessment is backed up by whatever quantitative information is available.

Table 7-1 Qualitative assessment of the effects of the conventional coal washing system (coal washed at medium density, the coal burnt in an older PF boiler, and the discard stockpiled) and the combined system (coal washing system as before, but the discard burnt in a CFBC system with limestone desulphurisation).

	Medium Density Coal Wash				Combined Washing and FBC System			
	Resource extraction	Feedstock preparation	Power generation	Waste management	Resource extraction	Feedstock preparation	Power generation	Waste management
Resource consumption								
Local air impacts								
Atmospheric impacts								
Water contamination								
Soil contamination								
Land degradation								
Human toxicity								

Explanation of symbols used in impact tables

Concern \ Certainty	Negligible	Minor	Moderate	Significant
25%				
50%				
75%				
100%				

The qualitative assessment matrix highlights the problem areas of the system, and can be used to identify sites in the life cycle system which could possibly “swing” the results of the study. This identifies the sites where qualitative reasoning must be used to determine the relative impacts of the systems, or if the scope of the study allows, where quantitative information needs to be generated. In this study, the need for a quantification of the impacts associated with disposal of the various sulphur-containing residues has been identified.

Although inventories generated for strategic level decision support are highly uncertain, quantitative uncertainty analyses are infeasible because of the large number of options requiring consideration, and the low level of data availability. The qualitative uncertainty matrices used here are therefore more practicable, with the qualitative impact matrix able to provide an indication of the uncertainty of the

In addition to discard allocation for the FBC system, the PF system requires a suitable allocation of the mining burdens between the power station coal and the mine's existing coal product. As with the discard allocation, a marginal or mass-based approach is appropriate as the two products are better regarded as combined production than joint production (i.e. the product volumes are independent of each other, and depend on the relative demand for each product). An alternative approach is to avoid the allocation altogether by extending the system to a consideration of both energy products (i.e. coal and electricity). With this approach, a mining only system is also required, so that the removal or avoidance of the discard dumps can be reflected (i.e. "before" and "after" scenarios are required to reflect the removal/prevention of the discard dump). Whilst such a basis was initially pursued for this case study, problems with the unequal nature of the energy products, and the inclusion of an additional scenario not actually under consideration for the decision, clouded the analysis to such an extent that the problems outweighed the benefits of the simpler allocation (i.e. the normalisation/functional unit uncertainties and valuation uncertainties introduced were greater than the allocation uncertainty removed). The uncertainties arise in the dual product approach because the dual functional unit, based on the total "effective" energy of the combined product, is not able to capture anything of the relative benefits of coal and electrical energy. The system was therefore judged to be better investigated on the basis of a single function (i.e. the electricity produced), and the effect of the coal and discard allocation investigated by way of sensitivity analyses.

7.4.2.d Uncertainty Propagation

The propagation of empirical parameter uncertainty through the inventory model is calculated in a simulation using Latin Hypercube sampling. Latin Hypercube sampling has been shown to be always as good as, and never worse, than the other commonly implemented sampling routines (see section A.1.1.b). A sampling size of 1000 was calculated to be sufficient for a minimum of 95% confidence that the median lies within 3 estimated percentiles. The actual precision achieved may be greater than this estimate, as the equations used to estimate the required sampling size for the specified accuracy only apply for Monte Carlo sampling, and Latin Hypercube sampling is likely to achieve the same or a better precision (see section A.1.1.a). This precision was chosen primarily because of the need to keep the time taken to complete each simulation to a minimum. The precision achieved with 1000 samples is judged sufficient, given the high empirical uncertainty in the system, and the rough distributions used to define the input parameters.

7.4.4.a Empirical Parameters

The study is for implementation in the medium to long term, so data must be predicted to reflect future conditions. Combustion of discard in a fluidised bed boiler has not been demonstrated beyond pilot scale in South Africa, so this data must be extrapolated to represent full-scale implementation. Fairly extensive and recent data are available around power generation for the plants currently in operation, and these data are used to characterise the reconditioned station, although uncertainty is introduced in transferring these data to a different and older power station. As in the South African LCI case study (chapter 6), generic LCI data are used to fill the data gaps.

The uncertainty of the empirical parameters is related to the data source used to define them. The majority are characterised by monitored process data, taken over as long a time period as was available, to allow for variability to be incorporated. The range found for this monitored data incorporates uncertainty from a number of different sources, e.g. variability (process variability, seasonal variability, variability in feedstock quality etc.) and measurement errors. However, because this data must be extrapolated to represent a different and future operating system, a large component of data uncertainty is not quantified in the process data samples. This other component of the uncertainty is addressed by estimating the additional variance (i.e. that not incorporated in the variance of the data sample) according to the set of DQIs proposed by Weidema and Wesnæs (1996). The data is evaluated according to the pedigree matrix (see Table 3-7), and the data quality “scores” extended to estimates of variance (as explained in section 4.2.2.a). Although this method necessarily involves substantial subjective judgement, it is preferable to ignoring the potentially limiting sources of uncertainty.

The subjective estimation of variance is even more important for those data elements for which only a single value or an approximate range is available. These data are obtained primarily from interviews with power station personnel and from the literature (design handbooks etc.), and are assumed to represent “design” or “most likely” values. The matrix scores assigned to these data elements are poorer than those assigned to the monitored data, as typically little is known about how representative the estimate is. Conservative CV estimates are applied to these data types, as they have no quantitative variability information to guide the conversion from qualitative matrix “score” to quantitative variance estimate. Where possible, the range of values available in the literature, or the range for a related quantity, is used to guide this estimate.

An example of the matrix scoring and equivalent CV estimate is given for the two parameters in Figure 7-2. The estimate of variance associated with the assigned score depends on the particular variable, as the variables have widely varying degrees of intrinsic variability. The manner in which the DQIs are used differs slightly from that defined by Weidema and Wesnæs (1996), in that the measured variance of the data sample is included under the completeness indicator, when it is not known to which indicator this observed variance should be attributed, i.e. whether the variance in the data sample is arising from measurement errors, from temporal or geographical variability et cetera. If the parameter gets a low “completeness score”, then this variance is increased by an amount estimated to include the sources of variance not reflected by the data sample, but if it scores well, then the variance incorporated by the data sample is judged sufficient to capture the full variance in the parameter. Appendix F.2 gives the scores and the associated estimates of variance for those variables returned with high uncertainty importance, i.e. those parameters whose conservative initial definitions required re-definition to more accurately reflect the quality of the data. The scores assigned reflect the source and quality of the data, whilst the estimates of variance depend on the particular parameter (i.e. the inherent variability exhibited by the parameter), and two parameters with a particular score do not necessarily have the same estimate of variance.

Generic LCI data are used for processes in the background system. The approximate uncertainty in this data is estimated using the ranges found in recent studies on the quality of LCI data, as described in section 4.2.2.a.

7.4.4.b Model Parameters

Decision variables comprise the majority of model parameters investigated in this study. The choice of these variables is under the direct control of the decision makers, and are predominantly related to the degree of refurbishment of the plant (i.e. the specification of the boiler and water plant), and the choice of mine supplying the coal or discard. To determine the sensitivity of the system to the choice of model parameter, the feasible limits of operation and the most likely value have to be defined. Not all decision variables can be defined as feasible ranges, and instead require a discrete choice that involves the choice of a number of related parameters, e.g. wet or dry ash disposal. In these cases, the choice of one parameter determines the choice of others, e.g. the choice of dry ash disposal requires a certain level of refurbishment on the water plant to limit the effluent volume to that able to be handled by the ashing system. Also, it does not always make sense to vary the parameters independently. In these instances, groups of parameters are rather defined together as operating “states”, e.g. the volume of storm-water

7.5. RESULTS AND INTERPRETATIONS

The uncertainty analysis process is an iterative one, where successive revisions are made until an acceptable resolution is obtained in the results. Figure 7-3 shows the final output from this process. The results presented here are primarily concerned with the outcome of this process, although reference is made to the key factors driving this progression in the sections that follow.

Figure 7-3 clearly shows the benefits of burning discard in a fluidised bed boiler. The figure shows a principal component representation of the output data, where the output samples have been projected onto a best-fit plane, so as to obtain a two-dimensional (planar) representation of the data, maximising the information able to be displayed in two-dimensions. A plot of the PC loadings (the coefficients of the eigenvectors corresponding to the PCs) is overlain on the plot of the transformed output samples. These are plotted as arrows (or “stressor vectors”) which represent the distance from the origin to the best-fit plane, and give information on the strength and independence of the criteria (see section A.3.2 for information on interpreting PC plots).

The strongest differences between the PF and FBC systems are in their fossil fuel use and their impacted land footprint (as seen by the magnitude of their stressor vectors). The length of the footprint vector reflects the considerable benefits of removing the discard dumps, i.e. the ash/gypsum dump has a consistently smaller footprint than the discard dump. The considerable savings in fossil fuel resources reflect the use of a “waste” energy source to generate power. The discard is defined as a waste from the mining system, and as such, is not allocated any mining burdens other than the “avoided” burdens resulting from the removal of the dumps. It therefore does not reflect any fossil fuel resource consumption, as all fossil fuels consumed and extracted during mining are allocated to the coal product. Also caused by the avoidance of mining burdens are the lower contributions to climate change, water use and summer smog of the FBC system relative to the PF system. This is less marked for climate change, because if compared on the basis of the power station alone (i.e. without the effects of mining), the FBC system has a slightly higher climate change burden than the PF system (caused predominantly by the use and transport of limestone in the FBC system).




Figure 7-3 Representation of FBC and PF operating “spaces”, given by ratios between the output samples and the PF “most likely” scenario for the chosen impact categories, transformed on the 1st and 2nd PC plane (see Table H-12 for corresponding PC loadings).

The lower contribution to the combined effect of acidification and eutrophication from the FBC system stems from the far lower NO_x emissions from the FBC boiler. The effect is much more marked for a consideration of NO_x alone, as SO_2 emissions are not always lower in the FBC system (they depend on the degree of desulphurisation). This is shown in Figure 7-4, which presents the same results as Figure 7-3, but on the level of environmental interventions rather than impacts. The high sulphur content of the discard burnt in the FBC system relative to the coal burnt in the PF system, means that only at high levels of desulphurisation does the FBC boiler have lower SO_2 emissions than PF boiler. The SO_2 emissions of the “worst” FBC scenario (with only 30% SO_2 removal) are higher than those of the PF scenarios, whilst those of “best” FBC system are significantly lower. This causes the strongly negative trend in SO_2 emissions in Figure 7-4, where the maximum variance in SO_2 emissions is between the FBC scenarios and not between the PF and FBC systems. This is a consequence of the PC plots’ ability to only show the variance incorporated in the first two PCs. PCA maximises the variance able to be viewed in two dimensions, and thus picks up the greater source of variance in SO_2 emissions

(i.e. between the FBC scenarios), whilst the variance between the two systems is picked up by a lower order component. From Figure 7-4 it appears that FBC systems are always worse than PF systems with respect to SO_2 emissions, although Figure H-10, which considers only the “best” and “most likely” scenarios, shows this is not the case.

Winter smog is shown to be ambivalent between the PF and FBC systems (i.e. neither system plots strongly with or against the winter smog vector). The winter smog vector falls between those of the SO_2 emissions and TSP emissions, with the direction it takes dependent on the relative magnitude of these two emissions (see Figures H-9 and 10 or Figures 7-3 and 4). Although the FBC system always causes an increase in particulate emissions, the volume of SO_2 emitted can “swing” winter smog to being significantly worse or slightly better than the PF system. The toxicity categories also do not show a strong tendency to be better or worse in either of the systems. The high variability in the output of these categories leads to the PCA being unable to distinguish any clear trends between the systems.




Figure 7-4 Representation of FBC and PF operating “spaces”, derived from the results for selected environmental interventions and the impacted land footprint (in original units), transformed on the 1st and 2nd PC plane (see Table H-13 for corresponding PC loadings).

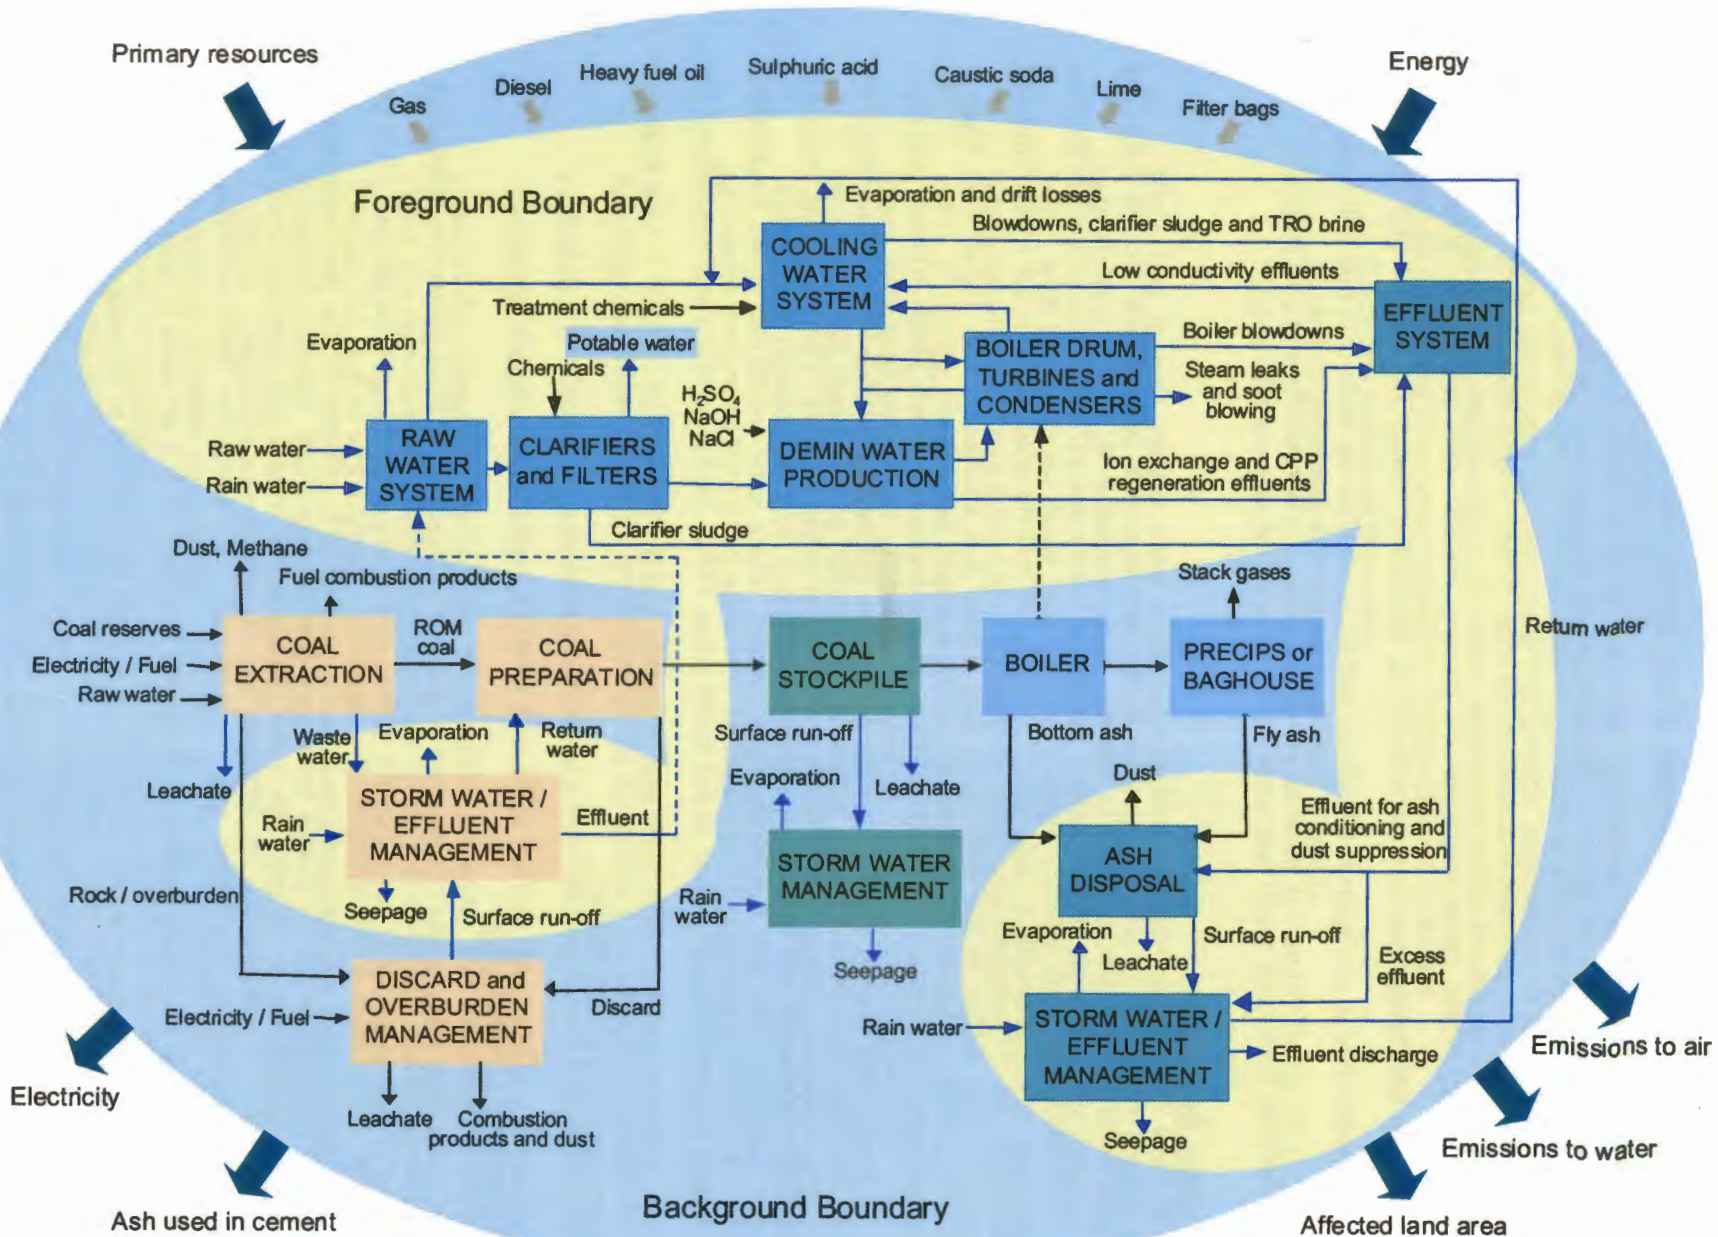


Figure 8-2 Schematic of the coal-electricity life-cycle, showing foreground boundary for mine water case study.

use, the refurbished and non-refurbished 2 Ml/day options plot extremely close, demonstrating that refurbishing the membrane plant introduces surprisingly little model parameter uncertainty. However, the fact that it shifts the raw water use into the 100% confidence range is significant, as it suggests that for full confidence, some degree of refurbishment would be required.

From an analysis of the trends in the model parameter analysis, it can be concluded that the box and whisker plots represent the full uncertainty range (both empirical and model) for the inclusion of mine water into the cooling water circuit (with the exception of TSP emissions, which show a very slight increase over the range predicted for the extreme scenarios). The uncertainty ranges predicted are only for the case where all the mine water is used as cooling water make-up. Utilising a portion of the mine water for demin production causes more extreme effects in all criteria considered, and thus extended ranges would be seen in the box and whisker plots. The significant increases in the negative impacts of the system (power and chemical consumption, footprint), with little or no benefits in the positive (water use and waterborne sulphate reduction) mean that utilising mine water for demineralised water production is unlikely to be considered (see discussion in appendix I.4.2).

8.4. CONCLUSIONS

From the analysis above and those in the appendices, it can be concluded that incorporating mine water does not always have the positive benefits one would expect. This is particularly true for the very poor quality pit water, which if the full range of uncertainty is included, has a low probability of both increasing the raw water consumption and salinity emissions to surface and underground water bodies (the exact opposite of the desired effect). However, the maximum compartment water can be incorporated with confidence, as this was shown to always decrease raw water consumption and salinity emissions to surface waters, whilst only marginally increasing the leachate potential from the ash dump. However, the use of compartment water does increase the energy and chemical requirements of the water plant, which translates to an increase in potential global and regional impacts (see Figure 8-9). A trade-off between these increases and the water and sulphate benefits thus has to be found, which will depend on the significance placed on each of the indicators. The model parameter analysis shows that these benefits will still be achieved for the full range of membrane plant capacities and volumes of water.

For certainty to be achieved that the maximum water-related benefit can be attained for the 2 Ml/day pit water case, some degree of refurbishment of the membrane plant is required. As with the compartment water case this system shows a relatively small increase in leachate potential from the ash, and increases in energy and material related emissions. To obtain certainty that sulphate emissions to surface waters may not actually be increased due to spills from the ashing system, some amendments to the ashing system may be required (e.g. additional spare dam capacity). This could also be achieved by mixing a higher percentage of effluent with the ash, although this then trades the potential for surface emissions with that of leachate emissions. The effect of increased moisture content in the ash has been shown to be potentially extremely significant for the size of the impacted land footprint, so this would not be recommended.

The incorporation of higher volumes of very poor quality mine water would appear not be worth the risks. This carries a substantially higher potential of leachate from the ash dump, as well as increased risks of spills from the power station effluent management system. Thus, although associated with high median sulphate reductions, the high variability in the system could result in these systems actually increasing sulphate emissions through spills and leachate. These systems also show the highest increase in energy and material consumption, which would appear not to be worth the comparatively slim chance of raw water savings or reductions in sulphate emissions.

The results of the study have been significantly enhanced by investigating the options as probabilistic ranges instead of mid-points. A particular benefit of the uncertainty analysis is its ability to recognise the risks of incorporating the high volumes of pit water. These would be missed if only a mid point analysis were conducted, which would only reflect the benefits of high waterborne sulphate removal, and not the possibility of spills from the ashing system. Also key to warning against the use of high volumes of pit water is the inclusion of the impacted land footprint indicator. This picks up the potential of significant increases in leachate from the ash dump, which would otherwise be missed from the assessment.

The enhanced interpretation of the system that can be achieved through PCA is again demonstrated by this case study. In particular, the ability of the PC plot to portray a continuous decision space, incorporating both empirical and model parameter uncertainty, against a number of criteria is notable. However, some limitations of the assessment were also encountered; that of their inability to represent the significance of the changes in the criteria with respect to each other. The three representations of

uncertainty used in the study complement each other, and are each able to enhance different aspects of the uncertainty of the system, i.e. the PC plots clearly present the trade-offs involved and the “spread” of the operating space, whilst probability plots clearly display the level of confidence able to be held in each criteria for each option. Box and whisker plots are a good representation of model parameter uncertainty, able to show the degree of shifting between the options, and the full range over which the systems act.

This case study demonstrates the markedly different approach and methods suitable to the analysis of continuous operational type decisions, as opposed to tactical type decisions. These differences are particularly notable for the type and quantity of the data available to characterise the system models, and the levels of foreground/background detail required. In addition, in operational type decisions, the few model parameters allow a definite and contained operating space to be defined, whereas tactical type decisions require the investigation of a number of possible scenarios spanning a large and undefined operating space. Model parameter uncertainty thus plays a dominant role in tactical studies, whilst for operational studies, although required to define the operating space, model parameters tend to contribute less to the overall uncertainty in the system, i.e. they are generally dominated by empirical parameter uncertainty, even though this empirical uncertainty is typically much lower than that present in tactical decision systems. This is clearly shown in this case study, where the choice of model parameter generally caused a small shift in the operating space of the option relative to the range spanned by its empirical uncertainty.

A *scenario tree*, an example of which is shown in Figure A-2, is a useful way to represent possible combinations of inputs. Each node represents an uncertain quantity or event, and each branch from the node one of its possible outcomes. Each path through the tree represents a sequence of event outcomes determining a specific scenario. The number of scenarios increases exponentially with the number of uncertain inputs, so the computational effort to evaluate every scenario rapidly becomes infeasible with increasing numbers of inputs, as does the ability to display and analyse the results. For this reason, often only a few special interest scenarios are examined, e.g. “most likely”, “best case” and “worst case” scenarios, where the input parameters are set to their nominal, best and worst values respectively. An extension of the scenario tree is the *probability tree*. Here a conditional probability is attached to each branch on the scenario tree. Each path through the tree represents a feasible scenario whose probability is the product of the conditional probabilities of the branches along that path. A discrete probability distribution for the output, or risk profile, can be obtained by calculating the probability and output value for each scenario.




Figure A-2 Scenario tree with three levels

The construction of a probability tree requires that the uncertainty in the input parameters be expressed as discrete probability distributions. However, uncertain quantities are often continuous rather than discrete. It is mathematically too complex for all but the simplest cases to obtain an analytical solution for the probability distribution of a function of a set of continuous random variables. Where there are only a few uncertain variables, it is possible to approximate the continuous distributions by discrete ones and use the probability tree approach. An alternative approach is *Monte Carlo simulation*. In this method, all the combinatorial scenarios are considered by selecting a random sample of scenarios for evaluation. Each scenario is generated by selecting each branch at a node according to its assigned probability, and as the computational effort depends on sample size and not on the number of possible values for each parameter, the branch values may be generated directly from the underlying continuous distribution, avoiding the need to discretise.

To obtain the required number of samples, m , this equation can simply be rearranged to yield:

$$m > \left(\frac{2cs}{w} \right)^2$$

To use this equation, a small number of runs (around ten) first need to be done, so that an initial estimate of the variance, s , can be obtained. The deviation, c , enclosing the specified probability, α , can be obtained from statistical tables, and the value substituted into the above equation, together with the specified interval width and the estimate of variance, to obtain the number of sample runs.

Alternatively, the number of samples required can be determined by specifying the required precision of the estimate of the median or of the other fractiles. Assuming the m sample values of y are relabelled to be in increasing order, i.e. $y_1 \leq y_2 \leq \dots \leq y_m$, sample value y_i is an estimate of fractile Y_p where $p = i/m$. The confidence interval, α , for a pair of sample values can be shown to be given by (y_i, y_k) , where:

$$i = mp - c\sqrt{mp(1-p)}$$

$$k = mp + c\sqrt{mp(1-p)}$$

and the values of i and k are rounded down and up respectively. Supposing confidence, α , of the fractile Y_p is specified as being between the sample value estimates of the $p-\Delta p^{\text{th}}$ and $p+\Delta p^{\text{th}}$ fractile, i.e. $i=m(p-\Delta p)$ and $k=m(p+\Delta p)$. These expressions for i and k can be combined with the above equations for i and k to yield:

$$m = p(1-p) \cdot \left(\frac{c}{\Delta p} \right)^2$$

For this estimate of m , no previous sample runs need to be done and the number of samples can be obtained directly from the specified precision. For example, a 95% confidence interval for the 50th percentile to be plus or minus one estimated percentile, gives $p=0.5$, $\Delta p=0.01$ and c , the deviation enclosing 95% of the probability of the unit normal, approximately equal to 2 (read from statistical tables). To achieve this very high precision, approximately 10000 runs are required. A reasonable estimate of the desired precision is thus necessary if the number of sample runs is to be kept to a manageable level. Where the empirical uncertainty from the input parameters is high, a very high degree of precision in the propagation of these uncertainties is probably worthless. In this case, the approximation uncertainty due to the number of runs will most probably be dominated by the empirical uncertainty from the input parameters, and a few hundred runs will probably suffice.

A.1.1.b Selecting a Sampling Method

Monte Carlo sampling is the simplest sampling method, in which each of the sample points (m) for each uncertainty quantity (X) is generated at random from X , with probability proportional to the probability density for X . Using the inverse cumulative method, m uniform random variables (u_i for $i=1,2,\dots,m$) between 0 and 1 are generated. The inverse of the cumulative probability distribution is then used to compute the corresponding values of X , i.e.

$$X_i \text{ where } \mathbf{P}(x_i \leq X_i) \text{ for } i=1,2,\dots,m.$$

In this simple Monte Carlo method, each value of every random variable X , including those calculated from other random variables, is a sample of m independent random values from the true probability distribution for X .

The Beta Distribution

The Beta distribution provides a flexible means of representing variability over a fixed range, and by simply specifying the first two parameters can be made to reflect both positively skewed or negatively skewed quantities. The second two parameters specify the range endpoints, and if not specified, the range defaults to zero to one.

Probability density function:

$$f(x) = \frac{1}{B(c, d)} x^{c-1} (1-x)^{d-1}; \quad 0 \leq x \leq 1$$

$$\text{where } B(c, d) = \frac{\Gamma(c)\Gamma(d)}{\Gamma(c+d)}$$


Expected value:

$$E(X) = \frac{c}{c+d}$$


Variance:

$$\text{Var}(X) = \frac{cd}{(c+d)^2(c+d+1)}$$

PDF: Beta(c, d, min, max)



CDF: Beta(c, d, min, max)



where \mathbf{u} is chosen (arbitrarily) to be of unit length. The maximum of $\mathbf{u}'\mathbf{S}\mathbf{u}$ is sought, where $\mathbf{S} = \mathbf{X}'\mathbf{X}$, and the solution is subject to the constraint that $\mathbf{u}'\mathbf{u} = 1$. This is done by setting the derivative of the Lagrangian equal to zero, i.e. obtaining the derivative of:

$$\mathbf{u}'\mathbf{S}\mathbf{u} - \lambda(\mathbf{u}'\mathbf{u} - 1)$$

where λ is a Lagrange multiplier. This gives:

$$2\mathbf{S}\mathbf{u} - 2\lambda\mathbf{u} = 0$$

The optimal value of \mathbf{u} (i.e. \mathbf{u}_1) is thus the solution of:

$$\mathbf{S}\mathbf{u} = \lambda\mathbf{u}$$

which is a well known occurrence: \mathbf{u} is the eigenvector associated with the eigenvalue λ of matrix \mathbf{S} . Therefore the eigenvector of $\mathbf{X}'\mathbf{X}$, \mathbf{u}_1 , is the axis sought, and the corresponding eigenvalue, λ_1 , indicates the amount of variance explained by the axis.

The second axis is to be orthogonal to the first, i.e. $\mathbf{u}'\mathbf{u}_1 = 0$, and satisfies the equation:

$$\mathbf{u}'\mathbf{X}'\mathbf{X}\mathbf{u} - \lambda_2(\mathbf{u}'\mathbf{u} - 1) - \mu_2(\mathbf{u}'\mathbf{u}_1)$$

where λ_2 and μ_2 are Lagrange multipliers. Differentiating gives:

$$2\mathbf{S}\mathbf{u} - 2\lambda_2\mathbf{u} - \mu_2\mathbf{u}_1.$$

Setting this equal to zero, and noting that multiplying across by \mathbf{u}_1' implies μ_2 must equal zero, the optimal value of \mathbf{u} (i.e. \mathbf{u}_2), arises as another solution of $\mathbf{S}\mathbf{u} = \lambda\mathbf{u}$. Thus λ_2 and \mathbf{u}_2 , are the second largest eigenvalue and associated eigenvector of \mathbf{S} . The eigenvectors of $\mathbf{S} = \mathbf{X}'\mathbf{X}$, arranged in decreasing order of corresponding eigenvalues, give the line of best fit, the plane of best fit, the three-dimensional hyperplane of best fit, and so on for higher-dimensional subspaces.

In the above derivation, \mathbf{S} is the matrix of the sums of squares and cross products. However, the variables under analysis are often very different (some "shout louder" than others), and it is therefore rare to base the PCA on the sums of squares and cross products (i.e. on the original data). Some transformation of the original data is usually necessary. Murtagh and Heck (1987) recommend basing the PCA on a correlation matrix, which works because the distance between the variables is directly proportional to the correlation between them (see Murtagh and Heck (1987) for proof). In standardising the variables by calculating their correlation coefficients, the row vectors are centred and reduced, i.e. have zero mean and unit standard deviation. Where only centring is appropriate, the PCA can be based on the covariance matrix. Jackson (1991) suggests that the covariance matrix should be used whenever possible, as PCs calculated from it have some useful properties (e.g. the PCs are in the same units of the original data). However, where the data is in different units, as will typically be the case with LCA applications, only the correlation matrix is appropriate (and is consequently the method used in the case studies).

Simple Example

This hypothetical example is taken from Jackson (1990). Although many of the properties of PCA may seem superfluous for this small two-variable example, its size allows one to easily understand these properties and the operations underlying the calculation of PCs.

The concentration of a chemical component in a solution is measured by two different methods. Method 1 is the standard procedure, and Method 2 a proposed alternative. To check that the two methods are interchangeable, the data in Table A-1 is obtained. A plot of this data suggests the use of regression to determine to what extent it is possible to predict the results of one method from the other. The least-squares equation for predicting Method 1 from Method 2 minimises the variability in Method 1 given a specific level of variability in Method 2, and vice versa. However, a single prediction equation is required that could be used in either direction. The line that will perform this role is called the orthogonal regression line, which minimises the deviations perpendicular to the line itself. This line is obtained by the method of principal components.

Table A-1

Observation No.	Method 1	Method 2
1	10.0	10.7
2	10.4	9.8
3	9.7	10.0
4	9.7	10.1
5	11.7	11.5
6	11.0	10.8
7	8.7	8.8
8	9.5	9.3
9	10.1	9.4
10	9.6	9.6
11	10.5	10.4
12	9.2	9.0
13	11.3	11.6
14	10.1	9.8
15	8.5	9.2

In this example, the PCA is based on the covariance matrix, so first the sample means, variances and the covariance between the two methods must be obtained.

Let x_{1k} be the test result for Method 1 for the k th run and the corresponding result for Method 2 be denoted by x_{2k} . The vector of sample means is:

$$\bar{\mathbf{x}} = \begin{bmatrix} \bar{x}_1 \\ \bar{x}_2 \end{bmatrix} = \begin{bmatrix} 10 \\ 10 \end{bmatrix}$$

and the sample covariance matrix is:

$$\mathbf{S} = \begin{bmatrix} s_1^2 & s_{12} \\ s_{12} & s_2^2 \end{bmatrix} = \begin{bmatrix} 0.7986 & 0.6793 \\ 0.6793 & 0.7343 \end{bmatrix}$$

A.3.2. Interpreting Principal Component Plots

The PC representation used in this thesis is after the method of Le T no (1999), in which the principal component scores (the transformed variables) and PC loadings (the coefficients of the eigenvectors) are merged onto a single plot. In two-dimensions (i.e. for a plot of the top two PCs), the plot of the PC loadings gives the plane of best fit (determined by the criteria PCA, see above derivation). The arrows on the PC plots thus give the distance from the origin to the plane of best fit, and indicate each criteria's contribution to the particular PC (the longer the arrow, the greater the contribution). The first PC accounts for the greatest share of the overall variance in the system, the second for the next largest, and so on (as explained above), thus the criterion contributing the greatest loading to the first PC is the criterion responsible for the biggest differences between the systems, and so on. In addition to the strength of the criterion, the PC loadings provide information about the independence of the criterion. Criteria that plot close together signify strong correlations in the data between those criteria, i.e. the options change to a similar extent and in a consistent direction for each of the criteria. Thus those criteria not capturing any significantly different information about the system not already demonstrated by another criterion can be identified and eliminated.

Merged with the plot of the PC loadings is a plot of the transformed data (the PC scores). These are each of the data points transformed onto the PC plane under consideration (usually that of the first two PCs, as this maximises the variance able to be displayed). For a mid-point analysis, each of the system alternatives plot as single points on the PC plane, whilst for probabilistic output samples, the system alternatives plot as "clouds" of points. These "clouds" can be interpreted as "zones of confidence" and the overlap between the clouds for each option identifies the significance of the differences between the options. In addition, a large spread identifies high uncertainty in the system, and thus identifies which options require better data characterisation.

The merged plot of loadings and scores is required to provide information as to the cause of the particular orientation of the points or "clouds" in the PC space (i.e. the directional "pull" each criterion has on the "clouds"). The PC loadings are represented by arrows on the plot, as only their relative length and direction from the origin of the PC plane are of relevance (i.e. they do not correspond to the values on the axes). The arrows yield significant information about the criteria, with the length of the arrows representing the strength of the criteria, and the clustering of the arrows representing the independence of the criteria. The direction of the arrows on the plot represent the performance of the options relative to the particular criteria, with those points or "clouds" plotting in the direction of the arrow performing poorly on that criterion, and those away from the arrow performing well.

As the determination of the PCs is based purely on maximising the variance between the original data and the PC plane, the absolute positioning of the loads and scores is not of significance (and tends to change between plots as the optimum representation is sought). Only their orientation relative to each other is of relevance. In addition, the units in which the output samples are specified for each criteria are not of significance, as the PCA is based on the correlations between the sample points.

contain some incompatible inventory categories due to the higher degree of aggregation of some of the APME categories. These were added to the ETH intervention list to create the background list given in Table E-1. A substantially less detailed inventory was able to be generated for the foreground processes, owing to data constraints (see Table E-1). In the inventories generated for the case studies in chapters 7 and 8, the intervention list is shortened to include only those environmental interventions that contribute to a defined impact category (i.e. are assigned an equivalency factor in the EI 99 impact assessment method). This also substantially decreases the number of environmental interventions considered, particularly around emissions to water (see Table E-1). The lack of an indicator of emissions to water (no correspondence between the columns in Table E-1) prompted the inclusion of waterborne sulphate emissions into the intervention list of chapters 7 and 8 (see section 5.1.3).

Table E-1 Environmental intervention list

Background List ¹	Also in foreground	Has equiv. factor	Background List ¹	Also in foreground	Has equiv. factor
<i>Resources</i>			<i>Emissions to Air cntd.</i>		
process energy (biofuel)			Butene		✓
process energy (hard coal)	✓		Ethane		✓
process energy (hydro power)			Ethanol		✓
process energy (lignite)			Ethene		✓
process energy (natural gas)			Ethylbenzene		✓
process energy (nuclear)			Ethylene dichloride		✓
process energy (oil)	✓		Ethyne		
land transformation (II-III)	✓	✓	Formaldehyde		✓
land transformation (II-IV)	✓	✓	Halogenated hydrocarbons		✓
land transformation (III-IV)		✓	Halons		
land use (IV-IV)		✓	Heptane		✓
gas reserves		✓	Hexafluoroethane		
hard coal reserves	✓	✓	Hexane		✓
lignite reserves		✓	Hydrocarbons (unspecified) ²		
oil reserves		✓	Methane	✓	✓
biotic products used (dry weight)			Methanol		✓
Barite			methyl tert-butyl ether		✓
Bauxite		✓	non methane VOC	✓	✓
Bentonite			Polyaromatic hydrocarbons	✓	✓
Chromium reserves		✓	Pentane		✓
copper reserves		✓	Phenol		
iron reserves		✓	Propane		✓
lead reserves		✓	Propene		✓
Limestone		✓	Propionaldehyde		✓
Manganese reserves		✓	Propionic acid		✓
nickel reserves		✓	Tetrachloride-dibenzo-dioxin		✓
silver reserves			Tetrafluoromethane		
sodium chloride			Toluene		✓
Sulphur reserves			Xylene		✓
Tin reserves		✓	Al		
Uranium reserves			As	✓	✓
Zeolite			Ba	✓	
Zinc reserves		✓	Ca		
Water	✓	✓	Cd		✓
Net air			Co		
Other inputs			Cr	✓	✓
<i>Emissions to Air</i>			Cu	✓	✓
Br ₂			Fe		
Cl ₂ ²			Hg		✓
CO	✓	✓	K		
CO ₂	✓	✓	La		
F ₂ ²			Mg		
HCl	✓	✓	Mn	✓	✓
He			Mo		
HF	✓	✓	Na		
H ₂ S			Ni	✓	✓
Hydrogen			Pb	✓	✓
N ₂			Pt		
NH ₃		✓	Sn		
N ₂ O	✓	✓	Sr		
NO _x	✓	✓	Th		
SO ₂	✓	✓	Ti		
Acetaldehyde		✓	U	✓	
acetic acid		✓	V	✓	
Acetone		✓	Zn	✓	✓
Acrolein			Zr	✓	
Aldehydes		✓	metals (unspecified) ²		✓
Alkanes		✓	B	✓	
Alkenes		✓	P	✓	
Aromatics		✓	Se		
Benzaldehyde			Si		
Benzene		✓	Total suspended particulates	✓	✓
benzo(a)pyrene		✓	steam/water vapour	✓	
Butane		✓			

intervention categories than the assumptions governing the installed capacity inventory. A notable exception is particulates, as the older plants considered have notably poorer ESP efficiencies.

The choices governing the compilation of the future generating mix inventory are shown to introduce more uncertainty than that due to the uncertainty in energy demand, i.e. the range spanned across the scenarios in Figure 6-4 is greater than the range spanned by the individual scenarios (as represented by the “whiskers”). This is because the analysis looked at predicting the inventory of a relatively near-term generating mix. The uncertainty in the energy demand, and the choice of new plant inventory, are likely to play larger roles in a prediction further into the future.

H.3.6. Land occupation

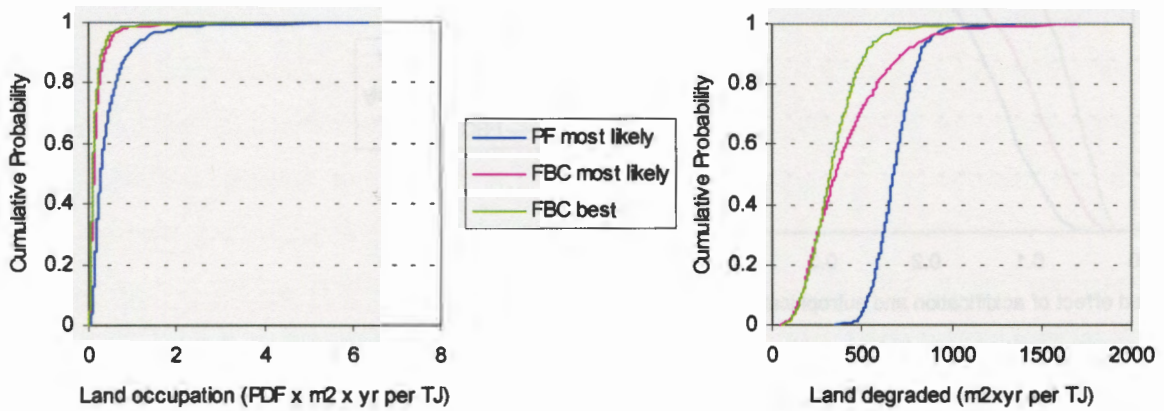


Table H-15 Magnitude of contribution and uncertainty importance of environmental interventions contributing to land occupation (“score” in PDF x m² x yr per TJ)

Land occupation components	PF System		FBC System	
	Correlation coefficients	Impact "scores"	Correlation coefficients	Impact "scores"
land transformation (semi-natural to built)	0.68	8.9E+07	0.83	6.3E+07
land transformation (semi-natural to cultivated)	0.68	6.6E+07	0.46	2.5E+07
land transformation (cultivated to built)	0.01	2.9E+05	0.08	3.2E+05
land use (built to built)	0.01	515	0.01	2161