THE APPLICATION OF BAYESIAN STATISTICS AND MAXIMUM ENTROPY TO ION BEAM ANALYSIS TECHNIQUES

by

Jayanethie Padayachee

A dissertation submitted in fulfilment of the requirements for the degree of

Master of Science

in the

Department of Physics
University of Cape Town

September 1997

University of Cape Town

The University of Cape Town has been given the right to reproduce this thesis in whole or in part. Copyright is held by the author.

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.

University of Cape Town

DST 530 PADA

98/3498

Preface

This dissertation presents work conducted at the Van de Graaff Group of the National Accelerator Centre in Faure, South Africa from January 1996 to July 1997.

Acknowledgements

I would like to take this opportunity to express my appreciation to the following people:

- Victor Prozesky, my supervisor, for introducing me to the world of Maximum Entropy and for his guidance during the course of this work.
- David Britton, my co-supervisor, for interest shown.
- Chris Theron, for reading through the *entire* dissertation and for the many helpful discussions.
- Nieldane Stodart, Karl Springhorn and Colin Doyle for keeping my PC in working order.
- Kevin Meyer, for always lending an ear and a shoulder when I needed either.
- The Stone Family – Ginny, Wimpy, Lauren and Emma – and Judy Blewett for keeping my feet on the ground and for keeping me from completely losing my mind!
- My grandmothers for their love and support.
- My brother, Thavashen for his love, encouragement and understanding and for painstakingly editing the dissertation. Sorry for the headaches, Thavash!
- My mother, for everything she has done to keep my mind, body and spirit healthy. I would have nothing without her.

Thank you.

Abstract

The elimination of some blurring property, such as the detector response function, from spectra has received a considerable amount of attention. The problem is usually complicated by the presence of noise in the data, and in general, there exists an infinite set of possible solutions which are consistent with the data within the bounds imposed by the noise. Such a problem is known, generally, as an ill-defined inverse problem. Many techniques have been developed in an attempt to solve inverse problems, for example the problem of deconvolution, but these techniques employ *ad hoc* modifications to solve different problems.

Bayesian Statistics has been proved to be the only consistent method for solving inverse problems of the type where the information is expressed in terms of probability distributions. This dissertation presents results of applying the Bayesian formalism, together with the concepts of maximum information entropy and multiresolution pixons, to various inverse problems in ion beam analysis.

The results of this method of deconvoluting Rutherford Backscattering Spectrometry (RBS) and Proton Induced X-ray Emission (PIXE) spectra are compared to the results from other deconvolution techniques, namely Fourier Transforms, Jansson's method and maximum entropy (MaxEnt) without pixons. All the deconvolution techniques show an improvement in the resolution of the RBS spectra but only the MaxEnt techniques show a significant improvement in the resolution of the PIXE spectra. The MaxEnt methods also produce physically acceptable results.

The MaxEnt formalism was applied to the extraction of depth profiles from RBS and PIXE spectra and yielded good results. The technique was also used to deconvolute the beam profile from one-dimensional nuclear microprobe scans.

Table of Contents

Bayesian Statistics and Maximum Entropy

Techniques of Ion Beam Analysis

Results

List of Symbols

α	regularisation parameter
A_Z	atomic mass of element of atomic number Z
$\mathcal{A}$	algebra on a space $\mathcal{S}$
B	pre-blurring function
b	width of pre-blurring function
b_Z	branching ratio
$\mathbf{C}_0$	capacity of a channel to transmit a message
C_Z	concentration of element of atomic number Z
d_i	i^{th} data point
d	data set $\{d_i\}$
$D(\omega)$	Fourier Transform of data d
ε	stopping cross-section
ε_i	ionisation energy
ε_Z	detector efficiency
E_0	energy of incident ions
f_i	i^{th} value of the physics f
f	set of $\{f_i\}$
$\tilde{f}$	simulated pseudo-data
$\hat{f}^{(k)}$	k^{th} estimate to $\{f_i\}$
$F(\omega)$	Fourier Transform of f
ϕ_f	power spectrum of f
F	Fano factor
$\lvert G(\omega)\rvert^2$	power spectrum of g
h	hidden image

I	general information pertaining to the problem
$\mathbf{K}$	general operator that modifies the underlying physics
K	kinematic factor
λ	Lagrange Multiplier
m	model function or Lebesgue measure
m_0	mass of incident ion
M	mass of target atom
$M^n(g)$	nth moment of a function g, $M^n(g) = \int x^n \cdot g(x)dx$
$M^0(g)$	zeroth moment of a function g, or area
$M^1(g)$	first moment of a function g, or mean
$M^2(g)$	second moment of a function g, or variance
n	noise
ϕ_n	power spectrum of noise
N_0	Avogadro's number
N	atomic density
N_p	number of protons in beam
N_{data}	number of data points
N	number of pixels in image
P_B	binomial distribution
P_G	Gaussian distribution
P_P	Poisson distribution
r	detector response function
$R(\omega)$	Fourier Transform of detector response function r
$\rho(\hat{f}^{(k)})$	Jansson's relaxation function
σ	scattering cross-section
σ_Z	ionisation cross-section of element of atomic number Z
σ_i	standard deviation of quantity i
σ_i^2	variance of quantity i
S_S	Shannon's entropy, $\mathsf{S}_S = -k\sum_i p_i \ln p_i$
S	generalised entropy, $\mathsf{S} = \sum_i f_i - m_i - f_i \ln \frac{f_i}{m_i}$

$S(E)$	stopping power of a target for an ion at energy E
$\mathcal{S}$	space
τ_i	thickness of layer at a depth x_i
ω_Z	fluorescence yield
y	inverse filter
$Y(\omega)$	Fourier Transform of inverse filter y
Y_w^n	Wiener inverse filter
$Y_0(Z)$	x-ray yield off an element Z from a *thin* target
$Y(Z)$	x-ray yield off an element Z from a *thick* target
Z_L	partition function for likelihood
Z_S	partition function for prior probability distribution

Symbols for §3.3.5.

a_i	number of holes containing i cherries discharged into tin
α'_{ij}	number of holes emptied
α''_{ij}	number of holes not emptied
α_{ij}	number of times the probability p_{ij} occurs
d_i	decision variable associated with hole i
N	number of cherries in tin
n_i	actual contents of hole i
p_{ij}	probability that a hole contains i cherries while the measurement indicates that it contains j
q_i	number of holes measured as containing i cherries
r_i	dielectric measurement of hole i

Notation

Probability Notation

A	a proposition or a well-posed statement
$\overline{A}$	NOT *A* or *A* false
$A \cdot B$	*A* AND *B*
$A+B$	*A* OR *B*
$p(A \mid I)$	probability of *A* being true given that *I* is true
$p(A \cdot B \mid I)$	probability of *A* AND *B* being true, given that *I* is true
$p(A + B \mid I)$	probability of *A* OR *B* being true, given that *I* is true
$P(A \mid I)$	probability density or distribution of *A* based on information *I*

Mathematical Notation

$h(x) * g(x)$	convolution of $h(x)$ and $g(x)$
$g \leftrightarrow G$	*g* transforms to *G*

Abbreviations

EDS	energy dispersive spectroscopy
EMP	electron microprobe
ERDA	Elastic Recoil Detection Analysis
FFT	Fast Fourier Transform
FWHM	full width at half maximum
IBA	Ion Beam Analysis
IBIC	Ion Beam Induced Charge
MaxEnt	maximum entropy
NAC	National Accelerator Centre
NMP	nuclear microprobe
MRI	Magnetic Resonance Imaging
PIXE	Proton Induced X-ray Emission
RBS	Rutherford Backscattering Spectrometry
Si(Li)	Lithium-drifted Silicon
SSB	Silicon Surface Barrier
STIM	Scanning Transmission Ion Microscopy
VDGG	Van de Graaff Group
WDS	wavelength dispersive spectroscopy

Chapter 1

Introduction

1.1. Science, Inverse Problems and Ion Beam Analysis

Man has always been on a quest to understand how the world around him works. His thirst for knowledge has led him from observing Nature to formulating theories about how Nature works and finally to testing these theories by conducting experiments. The study of the natural world based on observation and experimentation is known as Science, and Physics is that branch of Science which is concerned with the properties of matter and energy, and the relationships between them.

The ultimate aim of physicists is to be able to explain and predict how the physical aspects of Nature work within one context, but in order to do this it is necessary first to formulate a theory. Inspiration accounts for some theories, and experiments are then performed to confirm or refute the hypothesis. If, on the other hand, one has experimental results which cannot be explained on the basis of existing knowledge then one would prefer to 'extract' the theory from the experimental data. Such a process of reasoning: drawing a general conclusion from a set of premises, based on experience or experimental evidence, is known as inductive reasoning. The general process involving inductive reasoning is therefore an inverse problem.

The general inverse problem can be described by

$$d = \mathbf{K}f + n \tag{1-1}$$

where d is the data measured in an experiment, n the noise present in a measurement and $\mathbf{K}$ an operator that modifies the underlying physics f of the experiment. The aim of solving the inverse problem described in (1-1) is to extract the physics f from knowledge of the data d and distorting operator $\mathbf{K}$. The problem is, however, complicated by the noise which introduces an element of uncertainty into the measured data. There now exists an infinite set of possible solutions that are consistent with the data within the bounds imposed by the noise. It is no longer possible to obtain a unique answer and the problem becomes ill-defined.

The solution to (1-1) presented here uses the concept of *information entropy* coupled with *Bayesian Statistics* to calculate and rank the possible solutions, and finally to choose the solution most consistent with the data. In this formulation, all known facts about the problem are expressed in terms of probability distributions which are linked together by Bayes' equation, a fundamental relationship of probability theory.

1.2. The Inverse Problem in Energy Dispersive Spectral Analysis

This work focuses on solving ill-defined inverse problems in Ion Beam Analysis (IBA). In IBA, beams of charged particles are focussed on a target resulting in various interactions (e.g. Coulombic interactions, excitations and nuclear reactions) between the atoms in the target and the charged particles in the beam. The radiation that emerges from the interaction (scattered particles in Coulombic interactions, emitted photons from excited atoms and reaction products in nuclear reactions) are detected and their properties, such as energy, are measured yielding information on the composition of the target and distribution of the elements in the target.

The measuring instruments consist of gas-filled detectors, scintillation counters or semiconductor detectors. Each of these use the effects produced by radiation passing through the different media of the detectors as a means of determining the intensities and energies of the radiation impinging on the detector. There are two main techniques of measurement in IBA: wavelength dispersive spectroscopy (WDS) and energy dispersive spectroscopy (EDS). WDS applies Bragg's Law to crystals, to measure the energy of the emerging radiation. It is restricted to the detection of radiation in the x-ray region of the electromagnetic spectrum and has the additional disadvantage that only a single energy can be measured at any one time. In contrast, EDS (using semiconductor detectors) is not restricted to the detection of electromagnetic radiation only, and a range of energies can be measured in a single measurement. One disadvantage of EDS is the fact that only discrete spectra are measured. This means that information is always lost in a measurement. Another disadvantage is that the output of the experiment is always smoothed by the detector resulting in a further loss of information. The following work is based on EDS.

The response of a detector to a beam of monoenergetic radiation is described by the detector response function. The effect of the response function is to smooth the input to the detector, smearing features that are too close and resulting in a loss of information (*Figure 1.1*).

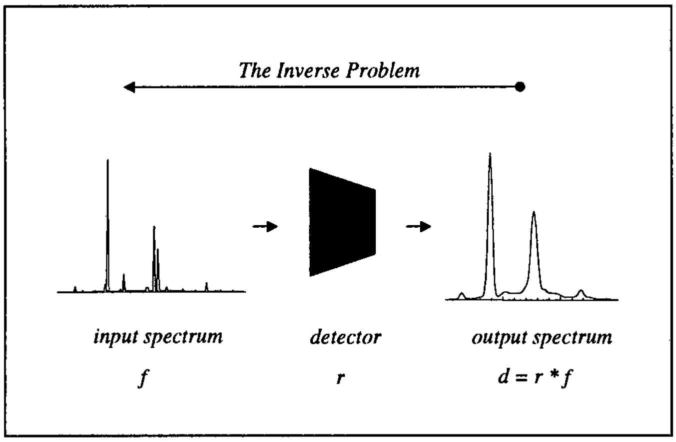

Figure 1.1. The effect of the detector is to smooth the input signal. Consequently, features that are too close together are smeared out and cannot be distinguished. The inverse problem involves extracting the input spectrum f from knowledge of the detector response function r and data d. The problem becomes ill-defined because the presence of noise in the data (not included in figure) means that a unique result cannot be obtained.

Detector response functions are usually Gaussian or Lorentzian in shape and the degree of smoothing or resolution of a detector is then given in terms of the full-width-at-half-maximum (FWHM) of the response function. For example, if a detector has a FWHM of 150 eV then only features separated by an energy greater than 150 eV will be resolved by the detector i.e. will be identified as separate entities. Features in the spectrum with an energy separation of less than 150 eV will be smeared out and will be indistinguishable.

The physical process of smoothing or blurring is described by the mathematics of convolution,

$$d(x) = \int_{-\infty}^{\infty} r(x-y) f(y) dy \qquad (1\text{-}2)$$

where *d* represents the output of the detector or data, *r* the response function and *f* the input to the detector or the physics of the process (i.e. the output of the experiment) as shown in *Figure 1.1.* The measured data is therefore a smoothed version of the actual output of the experiment. Equation (1–2) and *Figure 1.1* represent the ideal experiment. Unfortunately, real world experiments are subject to noise due to counting statistics. The practical problem can be described by

$$d(x) = \int_{-\infty}^{\infty} r(x-y) f(y) dy + n(x) \qquad (1\text{–}3a)$$

or in the discrete case by

$$d_i = \sum_j r_{i-j} f_j + n_i \qquad (1\text{–}3b)$$

where *n* represents the total noise component in the data. Equations (1-3a) and (1-3b) describe a special case of (1-1) where **K** is the convolution operator and the inverse problem is one of deconvolution.

1.3. Bayesian Statistics and the Inverse Problem

When people started using inductive reasoning, one of the first things they noticed was the need for a set of consistency conditions that the result of the reasoning process had to satisfy, including uniqueness and independence of method. The consistency conditions for the method of inductive reasoning, where degrees of plausibility are represented by numbers were proved to be the product and sum rules of probability theory in 1946 by Richard T. Cox [1946]. This signalled the beginning of Bayesian Statistics.

The product rule

$$p(f \cdot d \mid I) = p(f \mid d \cdot I) p(d \mid I) = p(d \mid f \cdot I) p(f \mid I) = p(d \cdot f \mid I) \qquad (1\text{-}4)$$

and sum rule

$$p(f \mid I) + p(\bar{f} \mid I) = 1 \tag{1-5}$$

form the basis of Bayesian Statistics, where $p(A \mid I)$ represents the probability that a proposition A is true given that information I is true. The most important consequence of the product rule is Bayes' equation, which arises from the commutativity of the logical product (1-4):

$$p(f \mid d \cdot I) = \frac{p(f \mid I) p(d \mid f \cdot I)}{p(d \mid I)} \tag{1-6}$$

This equation is named after Thomas Bayes (1702–1752), an 18th century clergyman and amateur mathematician who did some early work in probability theory. Bayes' equation forms the basis of solving inverse problems. The left-hand side of (1-6) represents the probability of the physics f being true given that the data d and other information I is true. I includes, for example, a description of the detector response function. The probability $p(f \mid d \cdot I)$ is then maximised with respect to the physics f to obtain the reconstruction most consistent with the data.

The aim of this work is to use Bayesian Statistics together with maximum entropy (MaxEnt) to solve inverse problems in IBA. Attention will be focussed on the inverse problem described in (1-3a) and (1-3b) but the technique is not restricted to deconvolution problems. MaxEnt will also be used to extract all available information from noisy data in the extraction of depth-concentration information from spectra. Gull & Skilling [1984] give examples of the use of MaxEnt in solving inverse problems in tomography, radio interferometry and magnetic resonance imaging (MRI), just a few examples of the various fields in which MaxEnt is used.

An overview of conventional deconvolution techniques will be given in Chapter 2 while the theory of Bayesian Statistics and Maximum Entropy as applied to spectral analysis will be described in Chapter 3. The experimental/practical part of the investigation will be described in Chapter 4 and results are presented and discussed in Chapter 5. Conclusions and prospects for future work are discussed in Chapter 6.

Chapter 2

A Review of Conventional Deconvolution Techniques

Many techniques have been developed to solve deconvolution problems. Some of these will be briefly described to give an overview of attempts to solve the deconvolution problem and the difficulties experienced. The techniques described in this section are summarised from P.A. Jansson's book, *Deconvolution with Applications in Spectroscopy* [Jansson 1984].

2.1. Convolution

Since the process of deconvolution lies at the heart of the inverse problem described in (1-3a) and (1-3b), it is crucial to understand the convolution process before any attempt at *de*convolution is made. Convolution can be considered to be a moving weighted average of a function with the smoothing function playing the role of the weighting function. The convolution of two functions q and g, written as $q*g$, is defined as

$$h(x) = q(x) * g(x) = \int_{-\infty}^{\infty} q(x-y)g(y)dy \tag{2-1a}$$

for continuous functions and for discrete functions as

$$h_i = q_i * g_i = \sum_j q_{i-j} g_j \tag{2-1b}$$

An example of discrete convolution within limits

As an example of discrete convolution, consider a function $g = \{0, 0, 1, 7, 1, 0, 0, 0, 0\}$ and a smoothing function $q = \{\ldots, 0, 0.1, 0.3, 0.4, 0.2, 0, \ldots\}$. The convolution of g and q can be calculated in terms of matrices as shown below. g is written as a column matrix and q as a square matrix with elements $[q]_{i,i} = 0.1$, $[q]_{i+1,i} = 0.3$, $[q]_{i+2,i} = 0.4$, $[q]_{i+3,i} = 0.2$ and the remaining elements 0.

$$q*g = \begin{bmatrix} 0.1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.3 & 0.1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.4 & 0.3 & 0.1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0.2 & 0.4 & 0.3 & 0.1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0.2 & 0.4 & 0.3 & 0.1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0.2 & 0.4 & 0.3 & 0.1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0.2 & 0.4 & 0.3 & 0.1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0.2 & 0.4 & 0.3 & 0.1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0.2 & 0.4 & 0.3 & 0.1 \end{bmatrix} \bullet \begin{bmatrix} 0 \\ 0 \\ 1 \\ 7 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ 0.1 \\ 1.0 \\ 2.6 \\ 3.3 \\ 1.8 \\ 0.2 \\ 0 \end{bmatrix}$$

q is asymmetric (see inset of *Figure 2.1*) and normalised ($\sum_j q_{i,j} = 1$). *Figure 2.1* shows the smoothing effect of convolution. The convolution of q and g is smoother than g and it has "adopted" the asymmetric nature of the blurring function q. The aim of deconvolution is to remove these effects (smoothing and asymmetry in the case of this example) and extract the original function.

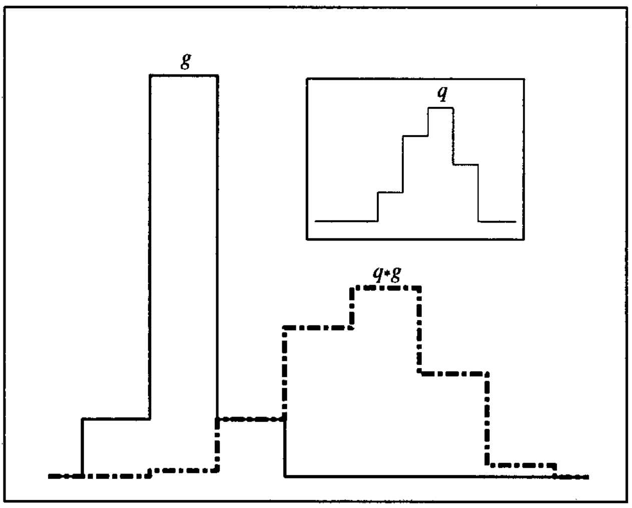

Figure 2.1. The smoothing effect of convolution can be clearly seen. The convolution is a smoothed version of g. ***Inset***: The smoothing function q is asymmetric and the convolution has "adopted" this assymmetry.

Convolution and the Extraction of Information

As described in §1.2, the inverse problem involves extracting the underlying physics of a process from experimental data and knowledge of the distorting function. The amount of information that can be extracted is, however, limited by how much prior information can be supplied. As an example, consider the problem of deconvolution. Consider a function $h(x)$ as defined in (2-1a). The n^{th} moment of h, $M^n(h)$, is then:

$$\begin{aligned}
M^n(h) &= \int_{-\infty}^{\infty} x^n \cdot h(x)dx \\
&= \int_{-\infty}^{\infty} x^n \int_{-\infty}^{\infty} q(x-y)\cdot g(y)dydx \\
&= \int_{-\infty}^{\infty} (y+z)^n \int_{-\infty}^{\infty} q(z)\cdot g(y)dydz \\
&= \int_{-\infty}^{\infty}\int_{-\infty}^{\infty}\sum_{i=0}^{n}\binom{n}{i} y^{n-i}\cdot z^i \cdot q(z)\cdot g(y)dydz \\
&= \sum_{i=0}^{n}\binom{n}{i}\int_{-\infty}^{\infty} y^{n-i}\cdot g(y)dy \int_{-\infty}^{\infty} z^i \cdot q(z)dz
\end{aligned}$$

which can be written in terms of the moments of q and g as

$$M^n(h) = \sum_{i=0}^{n}\binom{n}{i} M^{n-i}(g)\cdot M^i(q) \tag{2-2}$$

The transformation $x = y + z$ has been used in step 3 and the binomial theorem has been used in step 4:

$$(x+z)^n = \sum_{i=0}^{n}\binom{n}{i} x^{n-i} z^i$$

where

$$\binom{n}{i} = \frac{n!}{i!(n-i)!}$$

Consider the zeroth, first and second moments of h. The zeroth moment of a function represents the *area* of the function, the first moment represents the *mean* and the second moment represents the *variance*, and according to (2-2) these are

$$M^0(h) = M^0(q)\cdot M^0(g) \tag{2-3a}$$

$$M^1(h) = M^0(q) \cdot M^1(g) + M^1(q) \cdot M^0(g) \qquad (2\text{-}3b)$$

$$M^2(h) = M^0(q) \cdot M^2(g) + 2M^1(q) \cdot M^1(g) + M^2(q) \cdot M^0(g) \qquad (2\text{-}3c)$$

The functions q and g can always be transformed so that they are normalised to unit area and have zero mean, without any loss of generality. Associate h with the data d, q with the detector response function r and g with the physics f and assume that r and f are normalised with zero mean then (2-3c) becomes:

$$\sigma_d^2 = \sigma_f^2 + \sigma_r^2 \qquad (2\text{-}4)$$

where $\sigma_i{}^2$ represents the variance of the specified quantity. The aim of deconvolution is to improve the variance (or standard deviation σ_i) of the data. For a factor 2 improvement in the resolution[1] of the data i.e.

$$\sigma_f = \tfrac{1}{2}\sigma_d$$

(2-4) yields

$$\sigma_r =\sim 0.87\sigma_d$$

i.e. the standard deviation of the detector response function r must be known to within 13% of the true value. Alternately, the description of r used in the deconvolution procedure must contain at least 87% of the information contained in the real response function. Similarly, for a factor 5 improvement, the standard deviation of the detector response function must be known to within ~2% of the true value. So, to extract more information, a more accurate description of the detector response function is needed [V. Dose, private communication].

2.2. Linear Deconvolution Methods

The aim of linear deconvolution methods is to express the physics f as a linear combination of the data d. Most of the linear techniques described below do not explicitly consider noise and are therefore usually simple to implement and not as computationally demanding as other

---

[1] The resolution of a detector has been defined in terms of the FWHM of the detector response function. For a Gaussian, the FWHM is related to the standard deviation by

$$\mathrm{FWHM}_i = 2\sqrt{2\ln 2}\,\sigma_i \approx 2.35\sigma_i$$

methods. Only discrete deconvolution will be discussed because of its applicability to EDS and computation.

2.2.1. Fourier Transforms

The first linear technique involves the use of Fourier Transforms. The basic definitions of a Fourier transform $G(\omega)$ and its inverse $g(x)$ are

$$G(\omega) = \int_{-\infty}^{\infty} g(x) \cdot e^{i2\pi\omega x} dx \tag{2-5a}$$

$$g(x) = \int_{-\infty}^{\infty} G(\omega) \cdot e^{-i2\pi\omega x} dx \tag{2-5b}$$

An important property of the Fourier Transform is the convolution theorem given below:

$$\begin{aligned} h(x) * g(x) &\leftrightarrow H(\omega)G(\omega) \\ h(x)g(x) &\leftrightarrow H(\omega) * G(\omega) \end{aligned}$$

The convolution theorem allows one to compute Fourier Transforms and multiply instead of shifting, multiplying and adding as for discrete convolution. This is computationally easier because of algorithms such as the Fast Fourier Transform (FFT). The FFT of an N-element one-dimensional function f_k is defined as [Press *et al* 1992]

$$F_s = \frac{1}{N} \sum_{k=0}^{N-1} f_k \exp\left(\frac{-i2\pi sk}{N} \right) \tag{2-6a}$$

and the inverse FFT as

$$f_k = \frac{1}{N} \sum_{s=0}^{N-1} F_s \exp\left(\frac{i2\pi sk}{N} \right) \tag{2-6b}$$

The convolution theorem also provides an attractive option to working with problems involving convolution.

One method of deconvolution is immediately obvious. Compute the Fourier Transform $D(\omega)$ of the data d, divide by the Fourier Transform $R(\omega)$ of the detector response function r and compute the inverse Fourier Transform to get the physics f. The method is simple and quick but the results are not good because of the presence of noise. Noise consists of high frequency Fourier components and in order to eliminate these components, some form of filtering is required. Applying a low-pass filter eliminates the noise frequencies but unfortunately also removes the high frequency components of the data, which results in a reconstruction which does not include these components of the data. The solution also contains negative components because there are no constraints imposed on the reconstruction.

2.2.2. The Direct Linear Method

The direct linear method of deconvolution [Jansson 1984] is a simple, iterative procedure. The only requirement is that the response function r be non-zero over a finite domain, or be transformed so that this condition holds. The standard problem in (1-3b) then becomes

$$d_i = \sum_{j=-L}^{L} r_j f_{i-j}$$

where only values of r in the domain $[-L,L]$ are non-zero. It is also assumed that the detector response function is centred on r_0 and that $f_k = 0$ for $k < 1$. The situation is summarised below for $L = 1$ and $k = 1\ldots8$.

$$\begin{bmatrix} 0 \\ d_0 \\ d_1 \\ d_2 \\ d_3 \\ d_4 \\ d_5 \\ d_6 \\ d_7 \end{bmatrix} = \begin{bmatrix} f_2 & f_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ f_3 & f_2 & f_1 & 0 & 0 & 0 & 0 & 0 & 0 \\ f_4 & f_3 & f_2 & f_1 & 0 & 0 & 0 & 0 & 0 \\ f_5 & f_4 & f_3 & f_2 & f_1 & 0 & 0 & 0 & 0 \\ f_6 & f_5 & f_4 & f_3 & f_2 & f_1 & 0 & 0 & 0 \\ f_7 & f_6 & f_5 & f_4 & f_3 & f_2 & f_1 & 0 & 0 \\ f_8 & f_7 & f_6 & f_5 & f_4 & f_3 & f_2 & f_1 & 0 \\ 0 & f_8 & f_7 & f_6 & f_5 & f_4 & f_3 & f_2 & f_1 \\ 0 & 0 & f_8 & f_7 & f_6 & f_5 & f_4 & f_3 & f_2 \end{bmatrix} \bullet \begin{bmatrix} 0 \\ 0 \\ r_{-1} \\ r_0 \\ r_1 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} 0 \\ f_1 r_{-1} \\ f_2 r_{-1} + f_1 r_0 \\ f_3 r_{-1} + f_2 r_0 + f_1 r_1 \\ f_4 r_{-1} + f_3 r_0 + f_2 r_1 \\ f_5 r_{-1} + f_4 r_0 + f_3 r_1 \\ f_6 r_{-1} + f_5 r_0 + f_4 r_1 \\ f_7 r_{-1} + f_6 r_0 + f_5 r_1 \\ f_8 r_{-1} + f_7 r_0 + f_6 r_1 \end{bmatrix}$$

The system of equations is iteratively solved, with the estimate to the first value $\hat{f}_1$ of the set $\{f_i\}$ calculated from

$$d_0 = r_{-1} f_1$$

$$\hat{f}_1 = \frac{d_0}{r_{-1}} \tag{2-7}$$

The estimate of the second value $\hat{f}_2$ is obtained from:

$$d_1 = r_0 f_1 + r_{-1} f_2$$

$$\hat{f}_2 = \frac{d_1 - f_1 r_0}{r_{-1}} \tag{2-8}$$

This process continues until an estimate $\hat{f}$ of f is obtained. Although this iterative process is not computationally intensive, errors in the estimate of the first value are propagated through to the estimate of the final value. There is also no provision for the treatment of noise. The advantage of using such an iterative process is that a stopping criterion is not required since the estimate $\hat{f}$ is obtained from computing estimates for each element in the set.

2.2.3. Van Cittert's Method

Van Cittert's method [Van Cittert 1931] is an improvement of the direct linear method in the sense that it applies a correction to the k^{th} estimate $\hat{f}^{(k)}$ to obtain a new estimate. The correction is based on the difference $f - \hat{f}^{(k)}$ between the exact solution and the current estimate but this is an unknown quantity. The difference $d - \hat{d}^{(k)}$ is however, related to $f - \hat{f}^{(k)}$, and can therefore be used as the correction factor with

$$\hat{d}^{(k)} = r * \hat{f}^{(k)} \tag{2-9}$$

The new estimate is then calculated from

$$\hat{f}^{(k+1)} = \hat{f}^{(k)} + [d - (r * \hat{f}^{(k)})] \tag{2-10}$$

beginning with $\hat{f}^{(0)} = d$.

Van Cittert's method has the disadvantage that the resulting solutions may be negative. This method, like the direct approach, makes no provision for the treatment of noise. Furthermore, Van Cittert's method is an iterative procedure that requires a stopping criterion.

2.2.4. Inverse Filters

A method of deconvolution similar to Fourier Transforms is the use of an inverse filter, $y(x)$ such that

$$f = y * d \tag{2-11}$$

Taking Fourier Transforms of (2-11) and (1-2) yields

$$F(\omega) = Y(\omega)D(\omega)$$

$$D(\omega) = R(\omega)F(\omega)$$

Therefore,

$$Y = R^{-1} \tag{2-12}$$

$R(\omega)$ is usually small for large ω [Jansson 1984] which means that the reconstruction of a spectrum with a high level of noise will have even more pronounced noise if an inverse filter of this basic type is used.

Norbert Wiener [Goldman 1953] reformulated the problem considering noise explicitly

$$d = r * f + n$$

Under the assumption of a perfect detector response function (i.e. the response function is a δ-function), he derived a smoothing filter to reduce the noise in a spectrum. The smoothing filter $Y_W^n(\omega)$ is

$$Y_W^n(\omega) = \frac{\phi_f(\omega)}{\phi_f(\omega) + \phi_n(\omega)} \tag{2-13}$$

where ϕ_f and ϕ_n are the power spectra of the physics f and the noise n respectively. The power spectrum of a function g is defined by $|G(\omega)|^2$. Wiener derived $Y_W^n(\omega)$ by minimising the mean-square-error E_n between the estimate $\hat{f}_n$ due to noise only and the true reconstruction f where

$$E_n = \int \left| \hat{f}_n - f \right|^2 dx \tag{2-14}$$

and $\hat{f}_n = y * (f + n)$. Therefore, by minimising

$$E = \int \left| \hat{f} - f \right|^2 dx \tag{2-15}$$

with $\hat{f} = y * (r * f + n)$ he was able to derive an inverse filter that accounted for noise. This yields a linear inverse $y(x)$ with its Fourier Transform given by

$$Y(\omega)=\frac{R^*(\omega)\phi_f(\omega)}{\phi_r(\omega)\phi_f(\omega)+\phi_n(\omega)} \tag{2-16}$$

where ϕ_r is the power spectrum of the response function and R^* denotes the complex conjugate of R. The advantage of using a filter is that it is a single step process and thus does not require the stopping criteria of an iterative procedure.

These four methods are the basic linear deconvolution techniques. Jansson [1984] describes a few more linear methods that have been used for specific applications. The problems with linear techniques are the lack of constraints on the reconstruction and the poor treatment of noise, resulting in solutions that are negative with a high noise component.

2.3. Non-linear Deconvolution Methods

Linear deconvolution methods have no provision for constraints on the reconstruction. The non-linear techniques that follow take the positivity constraint into account so that the resulting solutions are positive.

2.3.1. Gold's Ratio Method

Gold's ratio method [Gold 1964] is the non-linear equivalent of Van Cittert's method. It uses multiplicative corrections to eliminate negative components in the reconstruction. The correction factor is given by $\frac{d}{\hat{d}^{(k)}}$ with $\hat{d}^{(k)}$ as in (2-9). A new estimate of the reconstruction can then be made from

$$\hat{f}^{(k+1)}=\hat{f}^{(k)}\cdot\left(\frac{d}{r*\hat{f}^{(k)}}\right) \tag{2-17}$$

with $\hat{f}^{(0)}=d$ as in Van Cittert's method. Gold's method eliminates the negative components in the reconstruction but has a poor treatment of noise.

2.3.2. Jansson's Method

Jansson's method [Jansson 1984] forces the solution to lie within the physical bounds of the problem. His first attempt was to introduce a relaxation parameter ρ into the difference equation of Van Cittert (2-10) where ρ depends on the k^{th} estimate $\hat{f}^{(k)}$. The $(k+1)^{\text{th}}$ estimate is then,

$$\hat{f}^{(k+1)} = \hat{f}^{(k)} + \rho(\hat{f}^{(k)}) \cdot [d - r * \hat{f}^{(k)}] \tag{2-18}$$

with

$$\rho = \begin{cases} \rho_0 & \text{for } \hat{f}^{(k)} \text{ physical} \\ 0 & \text{for } \hat{f}^{(k)} \text{ non-physical} \end{cases} \tag{2-19}$$

The purpose of ρ is to truncate non-physical components of $\hat{f}^{(k)}$. The constant ρ_0 is ideally large so that large corrections are made and the solution converges quickly. However, a too large value may result in over-corrections to the solution and may also cause a solution to diverge. This means that small values have to be used which requires much computation.

Jansson's second attempt was to use a relaxation function that modulated the correction between the physical bounds imposed by the problem. He also had to find a way to force the non-physical components towards the physical bounds. He noticed that the spurious components due to noise grew with successive iterations and he changed the sign of the correction for these non-physical components of the solution. He arrived at a relaxation function of the form:

$$\rho(\hat{f}^{(k)}) = \rho_0[1 - 2\left|\hat{f}^{(k)} - \tfrac{1}{2}\right|] \tag{2-20}$$

illustrated in *Figure 2.2* where the physical bounds of the problem lie between 0 and 1.

The main requirement of ρ is that the correction is not abruptly set to zero when $\hat{f}^{(k)}$ is non-physical but is allowed to be gradually corrected through successive iterations. Although Jansson's method produces physically meaningful results in that there are no negative solutions, there is no treatment of noise.

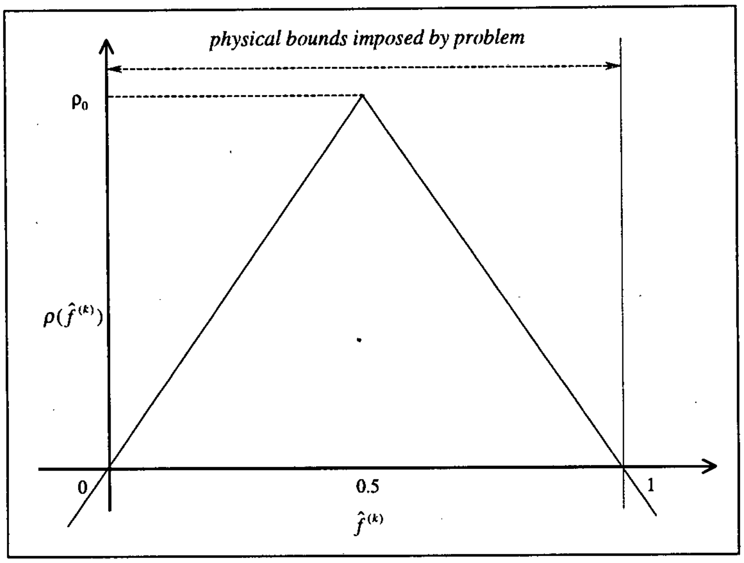

Figure 2.2. Example of Jansson's relaxation function ρ described in (2-20). ρ varies linearly from a maximum value midway through the physical region to zero at the boundaries and negative beyond.

2.3.3. Frieden's Method of Maximum Likelihood

B.R. Frieden [1972] realised that the constraints of positivity and least square error could still yield more than one solution since the hypersurface generated by the error term (e.g. equations (2-14) and (2-15)) could have many local minima. He was the first to incorporate *prior knowledge* into the solution by proposing a prior guess as to what the solution might look like. He also specified a probability of occurrence for each value that may be assumed by the estimate of f. This probability distribution expresses the degree of plausibility of obtaining an estimate $\hat{f}$, with the prior spectrum and probability distribution obtained by statistical analysis of previous results.

Frieden considered the reconstructed image to be divided into J equally sized cells of width Δx and position x_i, with Δx chosen such that it was the required spatial resolution of the reconstructed image. He further assumed that each cell contained $f_i(x_i)$ counts with,

$$f_i = c_i \Delta f$$

and Δf is the resolution in the counts. c_i is a number with

$$\sum_{i=1}^{J} c_i = C \qquad \text{(2-21)}$$

If the filling of cell i is independent of the filling of cell j (for $i \neq j$) then the number of ways that a specific set $\{c_i\}$ can occur is

$$W(c_1, c_2, \dots, c_J) = \frac{C!}{c_1! c_2! \dots c_J!} \tag{2-22}$$

Taking the natural logarithm and using Stirling's Approximation $x! \cong \sqrt{2\pi x}\, x^x \exp(-x + \dots)$ (valid for $x > 10$) yields

$$\ln W(c_1, c_2, \dots, c_J) = -\sum_{i=1}^{J} c_i \ln c_i \tag{2-23}$$

Frieden maximised (2-23) subject to the constraints of the data and obtained the most likely reconstruction consistent with the data. He then showed that for a given $\{c_i\}$ satisfying constraint (2-21), the entropy defined by

$$\mathbf{S}_F = -\sum_{i=1}^{J} c_i \ln c_i \tag{2-24}$$

is a maximum when $c_i = 1/J$ ($\forall$ i). The requirement that the entropy be a maximum, even with other constraints (e.g. data), forces c_i towards a constant value and inhibits large excursions from this constant value if it is not supported by the constraints. The entropy therefore plays the part of a smoothing term and suppresses noise in the reconstruction.

Frieden's method of maximum likelihood selects the $\hat{f}$ that conveys maximum information about the data d. It was the first non-linear method to address the problem of selecting a preferred solution from the large number of possible solutions by incorporating statistical concepts into the deconvolution problem.

2.4. Summary of Deconvolution Techniques

The presence of noise in real data leads to many possible solutions in deconvolution problems. Although linear deconvolution methods are quick and easy to apply, there are no provisions for constraints on the reconstruction and no treatment of noise. As a result, solutions contain negative components and the reconstruction is seldom consistent with the data.

Carter *et al* [1994] used the Fourier technique described in the text to deconvolute spectra, with poor results. Maj *et al* [1996] used Fourier deconvolution in the determination of the beam profile in a proton microprobe. They noted that noise amplification required the Fourier coefficients to be modified by a filter, but had problems choosing an appropriate filter. Schaffer *et al* [1984] used an FFT based method to deconvolute the detector response function from positron annihilation spectra. They eliminated noise by reducing the number of Fourier coefficients retained in the reconstruction, by using the power spectrum of the experimental data.

Non-linear methods allow for constraints such as positivity to be imposed on the reconstruction and therefore produce solutions that are physically meaningful. Imposing constraints reduce the number of possible solutions. Naylor *et al* [1991], Sprenger & Anderson [1991], Evans & Smith [1992], Marchetti & Mignerey [1993] and Coote & Kwan [1995] all used Jansson's method in the deconvolution of spectra. Although Jansson's method produces better results than linear methods [Naylor *et al* 1991, Evans & Smith 1992], there are problems with determining the spectral bounds required for the determination of the relaxation parameter.

Naylor *et al* [1991] compared three methods (Van Cittert, Jansson, maximum entropy) of deconvoluting the spectral transfer function of a Fabry-Perot interferometer from spectra. Their use of Van Cittert's method shows the weakness of linear methods in dealing with noise. The results produced by Jansson's method were better than those of Van Cittert's method, but there was still a considerable amount of noise in the reconstruction. The MaxEnt result was smooth and yielded the best reconstruction of the three techniques compared.

Chapter 3

Bayesian Statistics and Maximum Entropy

The foundation of this chapter lies with the single property that unites all problems that use probabilities or probability distributions: the assignment of an initial probability or initial probability distribution, based on available information. For example, a coin has only two faces so in the absence of additional information, an equal probability of ½ should be assigned to each face. This is similar in the case of rolling a die. A die has six faces and in the absence of information concerning the fairness of the die (i.e. whether it is loaded or not), an equal probability of $1/6$ should be assigned to each face. In the above examples there is a known, finite number of possibilities and it is rather simple to assign initial probabilities. But how does one assign an initial probability when the possibilities are infinite, or worse, unknown? This is a question that had plagued the minds of many workers in probability theory but remained unsolved until 1957 when E.T. Jaynes formulated the Principle of Maximum Entropy [Jaynes 1957].

The fundamental principles of the calculus of probabilities first appeared in Jacob (Jacques) Bernoulli's (1654–1705) *Ars Conjectandi*, published in 1713. The first definition of *probability* appeared as a part of Bernoulli's *Principle of Insufficient Reason*, which is summarised below [Jaynes 1978]:

1. The assignment of a probability is a means of describing a state of knowledge.
2. If we are given two propositions A and B as well as information which gives us no reason to choose one over the other, then a consistent way of describing that state of knowledge will be to assign them equal probabilities.

3. Taking the above into account, define the probability of A given information I, as

$$p(A \mid I) = M/N = \frac{\text{number of cases favourable to } A}{\text{total number of equally possible cases}} \tag{3-1}$$

(3-1) could easily be used to make an initial probability assignment if the total number of equally possible cases were known (as for coins or dice) but Bernoulli noted that the Principle of Insufficient Reason would fail if the total number of possibilities were infinite or unknown.

In an attempt to solve this problem, Bernoulli suggested that if a probability cannot be calculated by application of (3-1) then it may sometimes be possible to obtain an estimate for M/N, a theoretical probability, by observing experimental frequencies in a repeated experiment and then working backwards; an inverse problem. He demonstrated the link between a theoretical probability and an observable frequency by considering a binary experiment. He assigned a constant probability of success $b = M/N$ to each independent trial and determined that the probability of s successes in t trials was described by the binomial distribution[2] [Jaynes 1978]:

$$P_B(s \mid tb) = \binom{t}{s} b^s (1-b)^{t-s} \tag{3-2}$$

He then showed that as $t \to \infty$, the observable frequency $f = s/t \to b$. This was, however, only the first step to the solution of the problem—he now needed an inverse to this distribution.

With the binomial distribution, Bernoulli had determined the probability of s given (M, N, t) but he wanted the inverse probability of M given (s, t, N). An inverse to the binomial distribution was provided by Thomas Bayes in 1763. Given s and t, Bayes showed that the probability that M/N lies in the interval $b < M/N < b + \mathrm{d}b$ is given by

$$P(\mathrm{d}b \mid st) = \frac{(t+1)!}{s!(t-s)!} b^s (1-b)^{t-s} \mathrm{d}b \tag{3-3}$$

---

[2] A small "p" will be used to denote discrete probabilities while a capital "P" will be used to denote probability distributions.

Pierre Simon Laplace (1749–1827) later generalised the idea of inverse probabilities by considering an event E with equally likely causes $C_1, C_2, \ldots, C_t$. Given $P(E \mid C_i)$, for $i = 1 \ldots t$, Laplace noted that the probability that the event was caused by any specific C_i is given by

$$P(C_i \mid E) = \frac{P(E \mid C_i)}{\sum_i P(E \mid C_i)} \tag{3-4}$$

He generalised this result even further by considering causes that were not equally likely but which had a prior probability $P(C_i \mid I)$ based on some prior information I. The terms in (3-4) then have to be weighted according to this prior probability, which yields [Jaynes 1978]

$$P(C_i \mid E \cdot I) = \frac{P(C_i|I) P(E \mid C_i \cdot I)}{\sum_i P(C_i|I) P(E \mid C_i \cdot I)} \tag{3-5}$$

The final result is a general form of Bayes' equation and although it is the only consistent method of solving inverse problems, it does not solve Bernoulli's problem since in order to apply (3-5), a prior probability distribution $P(C_i \mid I)$ is still required. (3-5) is now known to be the entirely correct solution to the inverse problem but Laplace merely stated the results in (3-4) and (3-5) without providing proofs. This led to his work, as well as the concepts of probability theory as put forward by Bernoulli, being ridiculed and rejected. The reason: interpretation of the word *probability*.

3.1. *Probability* — What is it?

The basis of Statistics could be said to rest on the interpretation of one word: probability. The concept of probability was at the heart of a controversy that continued for about 150 years (circa 1800–1946). The problem arose because of two interpretations of the word, namely the Bayesian interpretation and the frequentist interpretation.

The frequentist interpretation views a probability as a proportion or frequency in a repeatable experiment. It is the more common view but it is not entirely correct. The most obvious problem arises when 'one-off' events are considered. It *is* possible to obtain a probability for an 'one-off' event but definitely not by conducting repeatable experiments! For example, there was a finite (albeit small) probability that the Titanic would sink on her maiden voyage which was definitely not determined by conducting repeatable maiden voyages and counting the number of times she sank!

This is where the Bayesian interpretation steps in — the Bayesian formalism interprets a probability as a numerical representation of a *state of knowledge* or alternatively as a *degree of plausibility* (Bernoulli's interpretation). The main difference between the frequentist and Bayesian interpretations comes from the fact that the state of knowledge or degree of plausibility is based on a given set of information. Such an interpretation may seem to be very subjective but the *only* factor which influences a probability assignment is the set of information on which assignment is based. A probability is therefore a number, assigned on the basis of available information, expressing how plausible an outcome is. A probability distribution is then a set of probabilities representing the degrees of plausibility of various outcomes, based on the same set of information.

Even though the Bayesian interpretation of the word probability does not suffer from the problems that plague the frequentist interpretation, there still has to be a way of assigning an initial probability. In general, a probability assignment may be based on an incomplete set of information. For example, if the probability that a group of symptoms is due to a specific disease is required, then this is a probability assignment based on an incomplete set of information. It is impossible to know about every possible disease [Jaynes 1978]. This incompleteness of information leads to an uncertainty in a probability assignment, and therefore in a probability distribution. This uncertainty associated with a probability distribution was quantified by C.E. Shannon in 1948 and he called it *entropy*.

3.2. Entropy

> *You should call it* **entropy** *for two reasons. In the first place, your uncertainty function has been used in statistical mechanics under that name, so it already has a name. In the second place, and more importantly, no one knows what* **entropy** *really is, so in a debate you will always have the advantage.*
>
> — J. von Neumann's advice to C.E. Shannon as to what to call the function $-k\sum_i p_i \ln p_i$ [3]

---

[3] [Tribus & McIrvine 1971]

3.2.1. Information Entropy

The concept of information entropy was born in the mind of C.E. Shannon. In 1948 he published a paper entitled *A Mathematical Theory of Communication* [Shannon 1948] where he described the problem of encoding a message such that the capacity $\mathbf{C}_0$ of a channel to transmit the message at a given rate is a maximum. In order to do this he introduced a function which represented the information content of the message (from the sender's point of view) or, alternatively, represented the uncertainty that the recipient associated with the message before the message had been received.

Shannon derived the form of the measure of uncertainty by considering a set of events with probabilities $p_1, p_2, \ldots, p_t$ of occurring. The measure of uncertainty $\mathbf{S}_S$ must then satisfy the following conditions [Shannon & Weaver 1949]:

1. $\mathbf{S}_S$ must be continuous in the p_i
2. If the p_i are equal ($p_i = 1/t$) then $\mathbf{S}_S$ must be a monotonic increasing function of t.
3. If a choice can be broken down into two successive choices, the original $\mathbf{S}_S$ should be the weighted sum of the individual values of $\mathbf{S}_S$.

The only function that satisfies these conditions has the form:

$$\mathbf{S}_S = -k\sum_i p_i \ln p_i \tag{3-6}$$

where k is a positive constant. $\mathbf{S}_S$ is called the entropy of the set of probabilities $\{p_i\}$ and is a measure of information, uncertainty or choice — information from the point of view of the transmitter and uncertainty or choice from the point of view of the receiver.

Properties of $\mathbf{S}_S$ which make it a reasonable measure of uncertainty/choice

1. $\mathbf{S}_S$ is always positive except if the outcome is certain then $\mathbf{S}_S = 0$. This corresponds to a situation where only one outcome is possible and is known.
2. If $p_i = 1/t$ (for a given t) then $\mathbf{S}_S$ is a maximum and equals ln t. This corresponds to the most uncertain situation.
3. Consider two events, x and y with s and t possibilities respectively. If $p(i{\cdot}j)$ represents the probability of a joint occurrence of i_x and j_y ($i = 1,\ldots,s$ and $j = 1,\ldots,t$) respectively then it can be shown that:

$$\mathbf{S}_S(x,y) \le \mathbf{S}_S(x) + \mathbf{S}_S(y) \tag{3-7}$$

with equality if and only if the events are independent i.e. $p(i{\cdot}j) = p(i){\cdot}p(j)$. This shows that the uncertainty of a joint event is less than the sum of the individual uncertainties. (3-7) also illustrates the fact that correlations between events are not considered by the form of the entropy in (3-6) since there is no 'cross-term' i.e. $\mathbf{S}_S(x){\cdot}\mathbf{S}_S(y)$.

4. Any change in the probabilities that cause them to be more equal results in an increase in $\mathbf{S}_S$ ($\mathbf{S}_S$ is maximal for all probabilities equal)

The value of the entropy ranges from 0 (for the most certain situation) to ln t (for the most uncertain situation). So, for two events x and y, each with t possible outcomes and if $\mathbf{S}_S(x) > \mathbf{S}_S(y)$, then the outcome of x is more uncertain than the outcome of y.

3.2.2. The Principle of Maximum Entropy

Bernoulli's problem of assigning an initial probability based on available information was solved by E.T. Jaynes in 1957. Jaynes presented his ideas in two papers entitled *Information Theory and Statistical Mechanics I & II* [Jaynes 1957], where (using the uncertainty interpretation of Shannon's entropy) he formulated the *Principle of Maximum Entropy* which states [Tribus 1969]

> *The minimally prejudiced probability distribution is that which maximises the entropy subject to the constraints supplied by the given information.*

In the case where, for example, the data is in the form of mean values of some quantities, then the Principle of Maximum Entropy can be stated mathematically as:

Maximise

$$\mathbf{S}_S = -k\sum_i p_i \ln p_i$$

subject to

$$\sum_i p_i = 1 \tag{3-8a}$$

$$\sum_i p_i g_j = \overline{g}_j \tag{3-8b}$$

where (3-8a) forces $\{p_i\}$ to be normalised and (3-8b) is the constraint on the mean values. If there are J mean values then there are $(J+1)$ constraints which the probability distribution must satisfy. The resulting *minimally prejudiced* distribution is

$$p_i = \exp[-\lambda_0 - \lambda_1 g_1(x_i) - \lambda_2 g_2(x_i) - \cdots - \lambda_J g_J(x_i)] \tag{3-9}$$

where the λ_j are the Lagrange multipliers determined by applying the constraints to (3-9). Tribus [1969] shows that applying the constraints (3-8a) and (3-8b) to (3-9) yields:

$$\lambda_0 = \ln\left(\sum_i \exp\left(-\sum_j \lambda_j g_j\right)\right) \tag{3-10a}$$

and

$$\frac{\partial \lambda_0}{\partial \lambda_j} = -\sum_i p_i g_j \tag{3-10b}$$

3.2.3. Application of the Principle of Maximum Entropy

Jaynes' Principle of Maximum Entropy provides a method of determining an initial probability distribution that is based on a set of incomplete information. A few examples of the use of the Principle of Maximum Entropy in deriving some standard probability distributions are given below [Tribus 1969].

3.2.3.1. The Uniform Distribution

The outcome of any event is mutually exclusive if it is logically independent of all other possible outcomes and exhaustive if there are a limited number of possible outcomes, all of which are known. For example, the tossing of a coin is an event with mutually exclusive and exhaustive outcomes since each outcome is logically independent of all others and there are only two possible outcomes.

Consider an event with t mutually exclusive and exhaustive outcomes. All that is known about the event is that one out of the t outcomes must occur. This information is encoded into the maximum entropy distribution by maximising $\mathbf{S}_S$ subject to (3-8a). From (3-9),

$$p_i = \exp[-\lambda_0]$$

and from (3-10a),

$$\lambda_0 = \ln t$$

The probability distribution is therefore,

$$p_i = 1/t \tag{3-11}$$

This result is identical to Bernoulli's Principle of Insufficient Reason, which is a special case of the Principle of Maximum Entropy.

3.2.3.2. The Gaussian Distribution

The Gaussian distribution is obtained if all that is known about a quantity x_i is its mean $\bar{x}$ and variance σ^2 and the fact that it can take on discrete values between -∞ and ∞. The constraints now consist of (3-8a) together with

$$\sum_i p_i x_i = \bar{x}$$

$$\sum_i p_i (x_i - \bar{x})^2 = \sigma^2$$

The MaxEnt distribution is

$$p_i = \exp\left[-\lambda_0 - \lambda_1 x_i - \lambda_2 x_i^2\right] \tag{3-12}$$

Tribus [1969] considers the case where the values of x_i are uniformly spaced. He is then able to solve for the values of the Lagrange multipliers in (3-12) to obtain the Gaussian Distribution.

$$P_G(x \mid X) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left[-\tfrac{1}{2}\left(\frac{x - \bar{x}}{\sigma}\right)^2\right] \tag{3-13}$$

3.2.4. Generalised Entropy in Image Reconstruction

The axioms set out by Shannon stem primarily from information theory. John Skilling [1988] adapted these axioms to derive a measure of uncertainty applicable to the reconstruction of images[4] which differs slightly from Shannon's entropy. Skilling only considers positive, additive quantities in his derivation. A positive, additive quantity q_i is any quantity that satisfies the Kolmogorov axioms of positivity and additivity:

1. $q_i \geq 0$
2. $q_{i \cup j \cup \ldots} = q_i + q_j + \cdots$ (for $i \neq j \neq \ldots$)

where $q_{i \cup j \cup \ldots}$ denotes the quantity formed by combining quantities q_i, q_j, …

The axioms of MaxEnt summarise the properties required of a reconstructed image [Skilling 1988, 1989]. $f\,[I, m]$ represents the reconstructed image f obtained by maximising the entropy subject to information I over a Lebesgue measure m on x. (See **Appendix A** for a summary of Measure Theory).

Axiom 1 Subset Independence

The axiom of subset independence stems from the requirement that separate treatment of separate distributions should give the same assignment as joint treatment of their union.

Let I_1 be information pertaining only to $f(x)$ for $x \in D_1$ and similarly let I_2 pertain only to $f(x)$ for $x \in D_2$. Then, if D_1 and D_2 are disjoint,

$$f[I_1, m] \cup f[I_2, m] = f[I_1 \cup I_2, m]$$

Justification

Information about one domain should not influence the reconstruction in a different domain, providing there is no constraint directly linking the domains.

Axiom 2 Co-ordinate Invariance

Let Γ be a co-ordinate transformation from x to Γx. Then $f\,[I,m]$ transforms as

---

[4] A spectrum is simply a one-dimensional image, so the term *image* will be used to denote the more general case.

$$\Gamma(f[I,m]) = f[\Gamma I, \Gamma m]$$

Justification

The same result should be obtained when the same problem is solved in two different co-ordinate systems with the reconstructions in the two systems related by the co-ordinate transformation.

Axiom 3 System Independence

If a proportion q of a population has a certain property, then the proportion of any sub-population having that property should be assigned as q [Gull & Skilling 1984].

Let $m(x_1,x_2) = 1$ on the unit square $0 \leq x_1 \leq 1$, $0 \leq x_2 \leq 1$. Let the constraints I be values of the marginals $\int f(x_1,x_2)dx_2 = a_1(x_1)$, $\int f(x_1,x_2)dx_1 = a_2(x_2)$ themselves obeying the consistent normalisation condition, $\int a_1(x_1)dx_1 = \int a_2(x_2)dx_2 = 1$. Then the reconstructed image should be the direct product $f(x_1,x_2) = a_1(x_1)a_2(x_2)$.

Justification

$f(x_1,x_2)$ represents a distribution of proportions because it is constrained to satisfy $\iint f(x_1,x_2)dx_1dx_2 = 1$. If all that is known about f are its marginal distributions $a_1(x_1)$ and $a_2(x_2)$, then, in the absence of any contrary bias, the uncorrelated reconstruction $f = a_1 \cdot a_2$ should be obtained. Any other choice of $f(x_1,x_2)$ would imply correlations for which there is evidence neither in the data nor in the measure.

Axiom 4 Scaling

$$f[\varnothing, m] = m$$

where $\varnothing$ represents the absence of additional information.

Justification

In the absence of any additional information, the initial measure should be recovered.

These axioms lead to a generalised entropy of the form:

$$\mathbf{S}=\int\left[f(x)-m(x)-f(x)\ln\frac{f(x)}{m(x)}\right]dx \tag{3-14a}$$

or in the discrete case,

$$\mathbf{S}=\sum_{i}\left[f_i-m_i-f_i\ln\frac{f_i}{m_i}\right] \tag{3-14b}$$

where m is the measure on the space of x and plays the role of a *model* function. Skilling's form of the entropy applies to all positive, additive distributions, normalised or not.

3.2.5. The Model

Information is always measured relative to some initial state. As a result, entropy is a relative quantity [Gull & Skilling 1984]. Although m was introduced as an abstract measure on the space of x [Skilling 1988], a consequence of **Axiom 4** is that it can also play the role of an initial model of the reconstruction. m can be thought of as a "weak" constraint on the reconstruction in the sense that maximising the entropy will try to keep the reconstruction as close to the model as possible ($\mathbf{S}$ is maximal for $f = m$) but the default value of the model can still be balanced by the data [von der Linden 1995]. The formalism enables properties of m, such as smoothness and positivity, to be transmitted to the reconstruction.

The choice of m depends on the prior knowledge of the problem. If there is no prior information then m is a constant. If, however, prior information is available, then it must be incorporated into the model so that reconstructions comply with that prior information. Gull & Skilling [1985] give examples of how prior information can be incorporated into the formula for the entropy via the model, m.

3.3. The Product and Sum Rules

In 1946, R.T. Cox [1946] published the proof to Laplace's result in (3-5). Cox constructed a set of mathematical rules for carrying out inductive reasoning, and found that representing degrees of plausibility by numbers led to consistency conditions in the form of functional equations, whose general solution could then be found. Thus, referring to degree of plausibility as probability, Cox [1961] found the consistency conditions to be the product rule

$$p(A \cdot B \mid I) = p(A \mid B \cdot I)p(B \mid I) = p(B \mid A \cdot I)p(A \mid I) \tag{3-15}$$

and sum rule

$$p(\overline{A} \mid I) + p(A \mid I) = 1 \tag{3-16}$$

of probability theory, thereby proving that any method of inductive reasoning which represents degrees of plausibility by real numbers is necessarily equivalent to Laplace's method or is inconsistent [Jaynes 1978].

There have been various methods of deriving the product and sum rules e.g. Jaynes [1995], Tribus [1969] and Cox [1961]. The basic requirements or desiderata for inductive reasoning are [Jaynes 1995, Tribus 1969]:

1. the representation of degrees of plausibility by real numbers
2. correspondence with common sense — e.g. for propositions A and B that depend on some information I,

   $$p(A \mid I) > p(B \mid I) \Leftrightarrow A \text{ is more plausible than } B$$
3. consistency, in the sense that if a problem can be solved using different methods, then each method must yield the same result
4. universality — the method of reasoning/consistency conditions must hold for *all* problems
5. unambiguity – only unambiguous statements should be used in the description of the problem

These desiderata form the basis of the derivation of the consistency conditions for inductive reasoning.

3.3.1. The Product Rule

The product rule arises from the dependence of $p(A \cdot B \mid I)$ on $p(A \mid I)$, $p(B \mid I)$, $p(A \mid B \cdot I)$ and $p(B \mid A \cdot I)$ under the assumption that I is known to be true. If

$$u = p(A \cdot B \mid I)$$

$$v = p(A|I)$$
$$w = p(A|B \cdot I)$$
$$x = p(B|I)$$
$$y = p(B|A \cdot I)$$

then,

$$u = F(v,w) = F(x,y) \tag{3-17}$$

is the only functional form u can satisfy. Tribus [1969] arrives at (3-17) by using the desiderata as constraints. F must be single valued, continuous and monotonically increasing [Tribus 1969, Jaynes 1995]. If F were not single-valued then u could be multi-valued which would fail to satisfy the uniqueness requirement. If F were not continuous or monotonically increasing then a small change in the arguments could lead to a large change in u.

Applying the requirement of consistency together with the fact that the logical product is commutative as well as associative yields the following functional equation for F,

$$F[F(x,w),z] = F[x,F(w,z)] \tag{3-18}$$

The solution to (3-18) is

$$F(x,y) = p^{-1}(p(x)p(y)) \tag{3-19}$$

where p^{-1} denotes the inverse operator of p, i.e.

$$p^{-1}[p(x)] = x$$

Associativity and commutativity of the logical product leads to the product rule.

$$p(A \cdot B|I) = p(A|B \cdot I)p(B|I) = p(B|A \cdot I)p(A|I) \tag{3-15}$$

3.3.2. The Sum Rule

The sum rule relates the probability of a single proposition A being true to the probability of that same proposition being false. A function of the form

$$w(A|I) = S(w(A|I)) \tag{3-20}$$

is required where w is some monotonic function of the plausibility $(A \mid I)$. Applying the consistency requirement and the product rule leads to a functional equation of the form

$$xS\left[\frac{S(y)}{x}\right] = yS\left[\frac{S(x)}{y}\right] \tag{3-21}$$

where $x = w(A \mid I)$ and $y = w(B \mid I)$. Tribus [1969] derives (3-21) by applying (3-15) and (3-20) to $w(A \cdot B \mid I)$, and imposing the conditions of consistency and universality. The solution to (3-21) is

$$S(x) = \left(1 - x^m\right)^{\frac{1}{m}} \tag{3-22}$$

with $0 \le x \le 1$ and $0 < m < \infty$. This yields

$$w^m\left(\bar{A} \mid I\right) + w^m\left(A \mid I\right) = 1 \tag{3-23}$$

The product rule can be written as

$$w^m(A \cdot B \mid I) = w^m(A \mid B \cdot I)w^m(B \mid I) = w^m(B \mid A \cdot I)w^m(A \mid I) \tag{3-24}$$

The choice of m is therefore arbitrary since a new function $p(A \mid I)$ could be defined where

$$p(A \mid I) = w^m(A \mid I)$$

which yields the sum rule:

$$p(\bar{A} \mid I) + p(A \mid I) = 1 \tag{3-16}$$

3.3.3. A Useful Result

The product and sum rules form the basis of probability theory and all other results can be derived from them. The following is an example of the use of the product and sum rules to derive $p(A + B \mid I)$.

$$\begin{aligned} & p(A+B \mid I) \\ & = p\left(\overline{\overline{A+B}} \mid I\right) \\ & = p\left(\overline{\bar{A} \cdot \bar{B}} \mid I\right) \\ & = 1 - p\left(\bar{A} \cdot \bar{B} \mid I\right) \\ & = 1 - p\left(\bar{A} \mid I\right) p\left(\bar{B} \mid \bar{A} \cdot I\right) \\ & = p(A \mid I) + p\left(B \mid \bar{A} \cdot I\right) - p(A \mid I) p\left(B \mid \bar{A} \cdot I\right) \\ & = p(A \mid I) + p\left(\bar{A} \cdot B \mid I\right) \\ & = p(A \mid I) + p(B \mid I)[1 - p(A \mid B \cdot I)] \\ & = p(A \mid I) + p(B \mid I) - p(A \cdot B \mid I) \end{aligned} \tag{3-25}$$

3.3.4. The Extension Rule or Marginalisation

The extension rule is derived by applying the sum rule to an event with mutually exclusive and exhaustive outcomes A_1, A_2,...,A_t yields

$$\sum_{i=1}^{t} p(A_i \mid B \cdot I) = 1 \tag{3-26}$$

Multiply (3-26) by $p(B \mid I)$:

$$\begin{aligned} & p(B \mid I) \\ & = \sum_{i=1}^{t} p(A_i \mid B \cdot I) p(B \mid I) \\ & = \sum_{i=1}^{t} p(A_i \cdot B \mid I) \end{aligned} \tag{3-27}$$

which is a process is known as marginalisation. Marginalisation can also be expressed in a continuous form:

$$\begin{aligned} & P(B \mid I) \\ & = \int P(A \mid B \cdot I) P(B \mid I) dA \\ & = \int P(A \cdot B \mid I) dA \end{aligned} \tag{3-28}$$

Marginalisation is usually used to eliminate parameters that are of interest to the calculation, but not to the final result, from the probability distribution functions. For example, A would be the parameter that is being eliminated in (3-28) above.

3.3.5. An Industrial Application of the Product Rule

This is an industrial problem concerning the efficient filling of tins with fruit [Garrett *et al* 1997]. Tins are sold as containing nine ½-cherries. To fill the tins, cherries are fed into a drum containing 12 holes arranged in a helix. Air is sucked through the holes, causing either 0, 1 or 2 halves to stick in each hole. Dielectric measurements r_i $(i = 1,\dots,12)$ are then made on each hole. Each measurement can indicate that the hole is either empty, contains one ½-cherry or contains an uncertain number of ½-cherries. The measurements are not completely reliable and a measurement indicating, for example, one ½-cherry may actually contain two because the cherries can stick together. The decisions to empty specific holes are based on the dielectric measurements. The aim is to maximise the number of tins containing nine ½-cherries.

Let N represent the number of ½-cherries in a tin and let n_i represent the actual number of ½-cherries in a hole. Define a decision variable d_i which indicates whether or not the contents of the i^{th} hole are discharged into the tin with $d_i = 1$ if the contents of the i^{th} hole are discharged into the tin and $d_i = 0$ if not. Let X denote all other information related to the problem; in particular that there are 12 holes and that the remaining cherries are removed from the drum between each filling, so that the filling of each tin is logically independent.

The probability of having N cherries in a tin given the measurements and decisions related to each hole i.e.

$$\Pi = p(N \mid r_1 r_2 \dots r_{12} d_1 d_2 \dots d_{12} X) \tag{3-29}$$

must be maximised, where N can be written as:

$$N = d_1 n_1 + d_2 n_2 + \cdots + d_{12} n_{12} \tag{3-30}$$

(3-29) can be expressed in terms of $\{n_i\}$ by marginalising (3-29) over $\{n_i\}$:

$$\Pi = \sum_{\{n_i\}=0,1,2} p(N n_1 n_2 \dots n_{12} \mid r_1 r_2 \dots r_{12} d_1 d_2 \dots d_{12} X) \tag{3-31}$$

Marginalisation has been described as being useful in the elimination of certain parameters but it can also be used to introduce parameters into a calculation as in (3-31). $\{n_i\}$ is important to the calculation but it is not required in the final result in (3-29), so (3-29) is written as (3-31) and *not* vice versa. Using the product rule, (3-31) can be rewritten as:

$$\Pi = \sum_{\{n_i\}=0,1,2} p(N \mid n_1 \ldots n_{12} r_1 \ldots r_{12} d_1 \ldots d_{12} X) p(n_1 \ldots n_{12} \mid r_1 \ldots r_{12} d_1 \ldots d_{12} X) \qquad (3\text{-}32)$$

From (3-30), it follows that

$$p(N \mid n_1 \ldots n_{12} r_1 \ldots r_{12} d_1 \ldots d_{12} X) = \delta(d_1 n_1 + \cdots + d_{12} n_{12}, N) \qquad (3\text{-}33)$$

where δ denotes the Kronecker delta defined by

$$\delta(i,j) = \begin{cases} 1 & \text{if i = j} \\ 0 & \text{if i} \neq \text{j} \end{cases}$$

The holes are logically independent since knowledge about the status of one hole contains no information about the rest. Also, the knowledge that a hole is emptied contains no information about the actual number of cherries in it. Therefore, using the product rule,

$$p(n_1 \ldots n_{12} \mid r_1 \ldots r_{12} d_1 \ldots d_{12} X) = \prod_{i=1}^{12} p(n_i \mid r_i X) \qquad (3\text{-}34)$$

Substituting, (3-33) and (3-34) into (3-32) yields

$$\Pi = \sum_{\{n_i\}=0,1,2} \delta(d_1 n_1 + \cdots + d_{12} n_{12}, N) p(n_1 \mid r_1 X) \ldots p(n_{12} \mid r_{12} X) \qquad (3\text{-}35)$$

which requires knowledge of the sequence of decisions $d_1 d_2 \ldots d_{12}$ of the 2^{12} possibilities that maximises (3-35) for $N = 9$ and a given set of measurements $r_1 r_2 \ldots r_{12}$. Each probability $p(n_i \mid r_i X)$ in the sum in (3-35) can take one of nine values according to the measured number of cherries ($r_i = 0,1,u$) and the actual number of cherries possible in experiments ($n_i = 0,1,2$). Denote this probability by p_{ij} where i denotes the actual number of cherries while j denotes the measured number so, for example, p_{21} represents the probability that a hole contains two cherries given that the measurement indicates one. The numerical values of these 9 probabilities can be obtained by doing calibration runs where the holes are inspected to compare their contents with those indicated by the measurements. Ideal values of p_{ij} would be

$$p_{ij} = \delta(i,j)$$

i.e. when there is no uncertainty between the measured number and the actual number of cherries in a hole.

The expression in the sum in (3-35) can be simplified to take into account the fact that many of the terms are zero (because of the Kronecker delta) while many of the remaining terms have identical values. It is therefore of great use to express (3-35) in terms of permutations of identical terms.

Assume that there are q_0 holes measured as containing 0 ½-cherries, q_1 holes measured as containing 1 ½-cherry and q_u holes measured to hold an uncertain number with

$$q_0 + q_1 + q_u = 12 \tag{3-36}$$

and that a_0 of q_0, a_1 of q_1 and a_u of q_u holes are discharged into the tin. Further assume that the sum in (3-35) contains α_{00} occurrences of the probability p_{00} and so on for the remaining 8 probabilities. The product of the 12 probabilities in (3-35) then reduces to

$$p_{00}^{\alpha_{00}} \, p_{10}^{\alpha_{10}} \, p_{20}^{\alpha_{20}} \, p_{01}^{\alpha_{01}} \, p_{11}^{\alpha_{11}} \, p_{21}^{\alpha_{21}} \, p_{0u}^{\alpha_{0u}} \, p_{1u}^{\alpha_{1u}} \, p_{2u}^{\alpha_{2u}} \tag{3-37}$$

with

$$\alpha_{00} + \alpha_{10} + \alpha_{20} = q_0 \tag{3-38a}$$

$$\alpha_{01} + \alpha_{10} + \alpha_{21} = q_1 \tag{3-38b}$$

$$\alpha_{0u} + \alpha_{1u} + \alpha_{2u} = q_u \tag{3-38c}$$

Partition each of the α parameters into an α' (indicating the holes emptied) and an α'' (indicating the holes not emptied). It then follows that

$$\begin{aligned} \alpha'_{00} + \alpha'_{10} + \alpha'_{20} &= a_0 \\ \alpha''_{00} + \alpha''_{10} + \alpha''_{20} &= q_0 - a_0 \end{aligned} \tag{3-39a}$$

$$\begin{aligned} \alpha'_{01} + \alpha'_{11} + \alpha'_{21} &= a_1 \\ \alpha''_{01} + \alpha''_{11} + \alpha''_{21} &= q_1 - a_1 \end{aligned} \tag{3-39b}$$

$$\begin{aligned} \alpha'_{02} + \alpha'_{12} + \alpha'_{22} &= a_2 \\ \alpha''_{02} + \alpha''_{12} + \alpha''_{22} &= q_2 - a_2 \end{aligned} \tag{3-39c}$$

The Kronecker delta in (3-35) can now be expressed in terms of the α parameters. Only holes which are emptied will make a contribution and the Kronecker delta becomes

$$\delta(\alpha'_{10} + 2\alpha'_{20} + \alpha'_{11} + 2\alpha'_{21} + \alpha'_{1u} + 2\alpha'_{2u}, N) \tag{3-40}$$

With the summation in (3-35) expressed in terms of the α parameters and taking the multiplicity of terms into account, (3-35) now becomes

$$\Pi = p(N \mid q_0 q_1 q_u a_0 a_1 a_u X) = \sum_{\{\alpha'\}\{\alpha''\}} \begin{aligned}[t] & \delta(\alpha'_{10} + 2\alpha'_{20} + \alpha'_{11} + 2\alpha'_{21} + \alpha'_{1u} + 2\alpha'_{2u}, N) \cdot \\ & \left[\frac{a_0!}{\alpha'_{00}!\alpha'_{10}!\alpha'_{20}!} \cdot \frac{(q_0 - a_0)!}{\alpha''_{00}!\alpha''_{10}!\alpha''_{20}!} \right] \cdot p_{00}^{(\alpha'_{00}+\alpha''_{00})} p_{10}^{(\alpha'_{10}+\alpha''_{10})} p_{20}^{(\alpha'_{20}+\alpha''_{20})} \cdot \\ & \left[\frac{a_1!}{\alpha'_{01}!\alpha'_{11}!\alpha'_{21}!} \cdot \frac{(q_1 - a_1)!}{\alpha''_{01}!\alpha''_{11}!\alpha''_{21}!} \right] \cdot p_{01}^{(\alpha'_{01}+\alpha''_{01})} p_{11}^{(\alpha'_{11}+\alpha''_{11})} p_{21}^{(\alpha'_{21}+\alpha''_{21})} \cdot \\ & \left[\frac{a_u!}{\alpha'_{0u}!\alpha'_{1u}!\alpha'_{2u}!} \cdot \frac{(q_u - a_u)!}{\alpha''_{0u}!\alpha''_{1u}!\alpha''_{2u}!} \right] \cdot p_{0u}^{(\alpha'_{0u}+\alpha''_{0u})} p_{1u}^{(\alpha'_{1u}+\alpha''_{1u})} p_{2u}^{(\alpha'_{2u}+\alpha''_{2u})} \end{aligned} \qquad (3\text{-}41)$$

with the sum running over all non-negative α.

Since there are 12 holes there are 91 different combinations of $\{q_0, q_1, q_u\}$ which sum to twelve. All that remains is to determine the number of holes $\{a_0, a_1, a_u\}$ which should be emptied to maximise (3-41) for $N = 9$. These values are then stored in computer memory and looked up when required. (3-41) was maximised using the actual measured values of p_{ij} as given in *Table 3-1*. The resulting look-up table is shown in *Table 3-2*. For example, if $q_0 = 1$, $q_1 = 9$ and $q_u = 2$, then the 9 holes measured as having one ½-cherry should be emptied and none of the others.

Table 3-1. The best values of p_{ij} measured under industrial conditions

p_{ij}	$j = 0$	$j = 1$	$j = u$
$i = 0$	0.990	0.010	0.010
$i = 1$	0.005	0.970	0.012
$i = 2$	0.005	0.020	0.870

To demonstrate the efficiency of the Bayesian solution, the filling of tins was simulated. The Bayesian filling mechanism was also compared to a more *ad hoc* filling scheme. In the *ad hoc* scheme the holes with the highest probability of containing the same number of cherries as indicated by the measurement (other than an empty hole) are emptied first followed by the next precise holes, if necessary. Since the measurement of one ½-cherry is most precise, these holes are emptied first. If the number of holes containing one ½-cherry is less than 9, then some of the holes measuring two are emptied. The results of the simulation are shown in *Figure 3.1.*

Table 3-2 Table generated by maximising (3-41) using values of p_{ij} given in *Table 3-1*.

q_0	q_1	q_u	a_0	a_1	a_u	$p(9\|\ldots)$	q_0	q_1	q_u	a_0	a_1	a_u	$p(9\|\ldots)$
0	0	12	0	0	5	0.3437	4	0	8	0	0	5	0.3437
0	1	11	0	1	4	0.5620	4	1	7	0	1	4	0.5620
0	2	10	0	1	4	0.5620	4	2	6	0	1	4	0.5620
0	3	9	0	3	3	0.6172	4	3	5	0	3	3	0.6172
0	4	8	0	3	3	0.6172	4	4	4	0	3	3	0.6172
0	5	7	0	5	2	0.6714	4	5	3	0	5	2	0.6714
0	6	6	0	5	2	0.6714	4	6	2	0	5	2	0.6714
0	7	5	0	7	1	0.7233	4	7	1	0	7	1	0.7233
0	8	4	0	7	1	0.7233	4	8	0	4	8	0	0.1414
0	9	3	0	9	0	0.7719	5	0	7	0	0	5	0.3437
0	10	2	0	9	0	0.7719	5	1	6	0	1	4	0.5620
0	11	1	0	9	0	0.7719	5	2	5	0	1	4	0.5620
0	12	0	0	9	0	0.7719	5	3	4	0	3	3	0.6172
1	0	11	0	0	5	0.3437	5	4	3	0	3	3	0.6172
1	1	10	0	1	4	0.5620	5	5	2	0	5	2	0.6714
1	2	9	0	1	4	0.5620	5	6	1	5	6	1	0.1074
1	3	8	0	3	3	0.6172	5	7	0	5	7	0	2.9483E-02
1	4	7	0	3	3	0.6172	6	0	6	0	0	5	0.3437
1	5	6	0	5	2	0.6714	6	1	5	0	1	4	0.5620
1	6	5	0	5	2	0.6714	6	2	4	0	1	4	0.5620
1	7	4	0	7	1	0.7233	6	3	3	0	3	3	0.6172
1	8	3	0	7	1	0.7233	6	4	2	6	4	2	7.8359E-02
1	***9***	***2***	***0***	***9***	***0***	***0.7719***	6	5	1	6	5	1	2.7305E-02
1	10	1	0	9	0	0.7719	6	6	0	6	6	0	3.8975E-03
1	11	0	0	9	0	0.7719	7	0	5	0	0	5	0.3437
2	0	10	0	0	5	0.3437	7	1	4	0	1	4	0.5620
2	1	9	0	1	4	0.5620	7	2	3	7	2	3	5.3843E-02
2	2	8	0	1	4	0.5620	7	3	2	7	3	2	2.6005E-02
2	3	7	0	3	3	0.6172	7	4	1	7	4	1	3.0798E-03
2	4	6	0	3	3	0.6172	7	5	0	7	5	0	6.5628E-04
2	5	5	0	5	2	0.6714	8	0	4	8	0	4	3.3485E-02
2	6	4	0	5	2	0.6714	8	1	3	8	1	3	2.5335E-02
2	7	3	0	7	1	0.7233	8	2	2	8	2	2	2.2283E-03
2	8	2	0	7	1	0.7233	8	3	1	8	3	1	6.5126E-04
2	9	1	0	9	0	0.7719	8	4	0	8	4	0	7.1914E-05
2	10	0	0	9	0	0.7719	9	0	3	9	0	3	1.3499E-03
3	0	9	0	0	5	0.3437	9	1	2	9	1	2	6.7354E-04
3	1	8	0	1	4	0.5620	9	2	1	9	2	1	5.5971E-05
3	2	7	0	1	4	0.5620	9	3	0	9	3	0	1.2171E-05
3	3	6	0	3	3	0.6172	10	0	2	10	0	2	3.5212E-05
3	4	5	0	3	3	0.6172	10	1	1	10	1	1	1.3265E-05
3	5	4	0	5	2	0.6714	10	2	0	10	2	0	1.0632E-06
3	6	3	0	5	2	0.6714	11	0	1	11	0	1	7.0563E-07
3	7	2	0	7	1	0.7233	11	1	0	11	1	0	2.1548E-07
3	8	1	0	7	1	0.7233	12	0	0	12	0	0	1.1807E-08
3	9	0	0	9	0	0.7719							

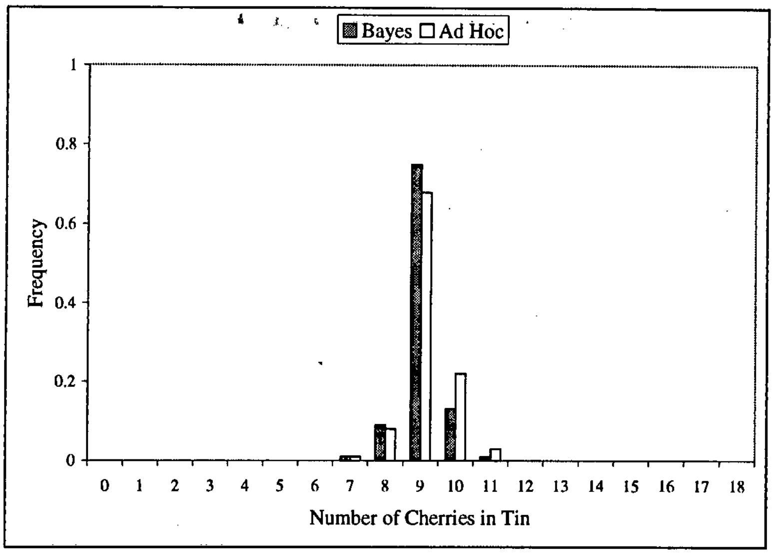

Figure 3.1. Frequency distribution of the number of cherries in each tin obtained using the best values of p_{ij} measured under industrial conditions, as shown in *Table 3-1.* The Bayesian selection process optimises the number of tins containing 9 ½-cherries but *only* 75% of all tins filled contain the desired number of cherries.

The proportion of tins containing 9 ½-cherries increases as p_{ij} becomes more precise. This was demonstrated by performing the simulation for the 'close-to-ideal' values of p_{ij} shown in *Table 3-3.* The results of this simulation appear in *Figure 3.2.*

Table 3-2. Close-to-ideal values of p_{ij}.

p_{ij}	$j = 0$	$j = 1$	$j = u$
$i = 0$	0.995	0.025	0.025
$i = 1$	0.025	0.995	0.025
$i = 2$	0.025	0.025	0.995

The number of tins containing 9 ½-cherries increases to 95% for the 'close-to-ideal' values of p_{ij}. So, in order to increases the number of tins containing 9 ½-cherries, the measurements to determine the contents of each hole in the drum must be more precise.

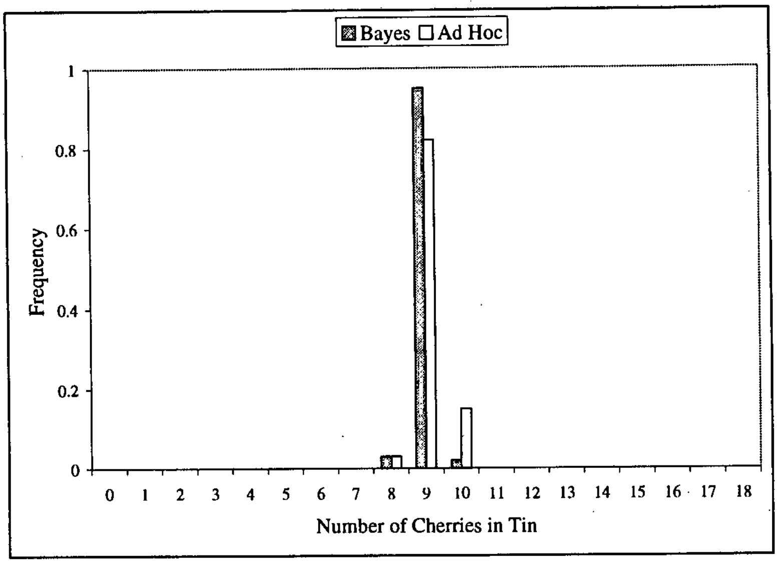

***Figure** 3.2.* Frequency distribution of the number of cherries in each tin obtained using close-to-ideal values of p_{ij} as shown in *Table 3-3.* 95% of all tins filled contain 9 ½-cherries.

3.4. Bayes' Equation

Bayes' equation is a simple but powerful relationship that allows the modification of a known probability to take new information into account. For clarity (1-6) and (3-5) can be re-written as

$$p(B|A\cdot I)=\frac{p(B|I)p(A|B\cdot I)}{p(A|I)} \tag{3-42}$$

If the known probability is $p(B|I)$ and new information A must be considered, then Bayes' equation expresses the probability of B taking both information A and I into account. $p(B|A\cdot I)$ is known as the posterior probability and represents the probability after the new information A has been obtained. $p(A|B\cdot I)$ is known as the likelihood and $p(B|I)$ is known as the prior probability.

Example of the use of Bayes' equation in the updating of probabilities [Tribus 1969]

This example is a simple illustration that Bayes' equation takes *all* necessary information into account.

There are three prisoners **A**, **B** and **C**. One of them is to be executed in the morning. **A** goes to the warden and says: "I know you are not allowed to tell me whether I will be executed but tell me which of the other two will not be executed." The warden, after some thought, replies: "**B** will not be executed."

Assuming the warden is telling the truth, has his reply changed the probability that **A** should attach to his chances to escape execution?

Formulation of Problem

$A \equiv$ "**A** will be executed"
$B \equiv$ "**B** will be executed"
$C \equiv$ "**C** will be executed"
$X \equiv$ "The warden said that **B** will not be executed"
$Z \equiv$ "The remaining facts of the problem" e.g. the fact that one *will* be executed, etc.

$p(A \mid X \cdot Z)$ must be calculated to take the new information into account. The probability that **A** will be executed without considering the new information can be taken as being 1 out of 3 since there are three prisoners and one *will* be executed, i.e. $p(A \mid Z) = \frac{1}{3}$. The probability that the warden said that **B** will not be executed is 1 out of 2 since the warden is not allowed to say anything about **A**'s fate i.e. $p(X \mid Z) = \frac{1}{2}$. The probability that the warden said that **B** will be not be executed given that **A** will be executed is also 1 out of 2, i.e. $p(X \mid A \cdot Z) = \frac{1}{2}$. From Bayes' equation,

$$\begin{aligned} p(A \mid X \cdot Z) &= \frac{p(A \mid Z)\, p(X \mid A \cdot Z)}{p(X \mid Z)} \\ &= \frac{\left(\frac{1}{2}\right)\left(\frac{1}{3}\right)}{\left(\frac{1}{2}\right)} \\ &= \frac{1}{3} \end{aligned}$$

and **A**'s chances of escaping execution remain unchanged.

Now consider the same problem but **A** reads a notice that says: "**B** will not be executed." X now represents the fact that the notice said that "**B** will not be executed." The probabilities $p(X \mid A \cdot Z)$ and $p(A \mid Z)$ remain unchanged but the probability of the notice indicating that **B** will not be executed is now 1 out of 3 since the notice isn't restricted, as the warden was, to not saying anything about **A**'s fate. Bayes' equation now becomes

$$\begin{aligned} p(A \mid X \cdot Z) &= \frac{p(A \mid Z)p(X \mid A \cdot Z)}{p(X \mid Z)} \\ &= \frac{(\frac{1}{2})(\frac{1}{3})}{(\frac{1}{3})} \\ &= \tfrac{1}{2} \end{aligned}$$

and the chances of **A** being executed have increased.

At a first glance it may seem that **A** learns the same fact from the warden and the notice, but whereas the warden is restricted to not telling **A** anything about his fate, the notice is not. It is this subtle difference that leads to a different result.

3.5. Development of MaxEnt in Spectral Analysis

As discussed in §1.3, Bayes' equation can also be used to solve inverse problems. This section traces the development of Bayesian Statistics and Maximum Entropy in spectral analysis and serves as the theoretical basis for applying the MaxEnt formalism to inverse problems in IBA. Bayes' equation in terms of the data d and physics f is

$$P(f \mid d \cdot I) = \frac{P(f \mid I)P(d \mid f \cdot I)}{P(d \mid I)} \tag{3-43}$$

where $P(d \mid f \cdot I)$ is known as the likelihood and $P(f \mid I)$ is known as the prior probability distribution. In order to evaluate the posterior $P(f \mid d \cdot I)$, expressions for $P(d \mid f \cdot I)$ and $P(f \mid I)$ are required and these are described below.

3.5.1. The Likelihood

The likelihood $P(d \mid f \cdot I)$ expresses how well the data would be predicted assuming the physics f were correct. This assumes that the information I contains all necessary information to simulate pseudo-data $\tilde{f}$ given physics f. For the deconvolution problem f must be convoluted with a detector response function as in (1-3a) and (1-3b) to simulate the pseudo-data. For the determination of depth profiles, a spectrum must be simulated from a concentration profile to compare the profile to the data. If the data set is $d \equiv \{d_1 d_2 \ldots d_{N_{data}}\}$ then the product rule (3-15) yields,

$$P(d_1 \ldots d_{N_{data}} \mid f \cdot I) = P(d_1 \mid d_2 \ldots d_{N_{data}} f \cdot I) \ldots P(d_{N_{data}} \mid f \cdot I) \tag{3-44}$$

The data d is assumed to be logically independent so for a given f, knowledge of data d_i contains no information about data d_j ($\forall j \neq i$). (3-44) therefore becomes,

$$P(d_1 \ldots d_{N_{data}} \mid f \cdot I) = \prod_{i=1}^{N_{data}} P(d_i \mid f \cdot I) \tag{3-45}$$

The probability of measuring data d is analogous to the probability of measuring the associated noise n as in (1-1) since if f is correct then the only uncertainty in the measured data would be due to the noise. The noise can therefore be considered to be the deviation between the measured data d and the simulated pseudo-data $\tilde{f}$,

$$n = d - \tilde{f}$$

Assuming that the noise has a Gaussian distribution with zero mean and standard deviation σ_i then

$$P(n_i \mid \sigma_i) \propto \exp\left(-\frac{n_i^2}{2\sigma_i^2}\right) \tag{3-46}$$

Now,

$$P(d_i \mid f \cdot I) \propto P(n_i \mid \sigma_i) \qquad \text{(for } i=1\ldots N_{data}\text{)} \tag{3-47}$$

so the likelihood can be expressed in terms of the parameters that describe the noise.

$$P(d_1 \ldots d_{N_{data}} \mid f \cdot I) \propto \prod_{i=1}^{N_{data}} \exp\left(-\frac{n_i^2}{2\sigma_i^2}\right) \tag{3-48}$$

Substituting for n_i from (3-46) yields,

$$P(d \mid f \cdot I) = \frac{\exp\left(-\frac{1}{2}\chi^2\right)}{Z_L} \tag{3-49}$$

for the expression for the likelihood, with

$$\chi^2 = \sum_{i=1}^{N_{data}} \frac{\left(d_i - \tilde{f}_i\right)^2}{\sigma_i^2} \tag{3-50}$$

$$Z_L = \int \exp\left(-\tfrac{1}{2}\chi^2\right) \mathrm{d}^{N_{data}} d = \prod_{i=1}^{N_{data}} \left(\sqrt{2\pi}\sigma_i\right)^{-1} \tag{3-51}$$

Z_L is a normalisation function. It is assumed that the noise is described by counting statistics so the variance of each data point is equal to number of counts in that point.

$$\sigma_i^2 = d_i \tag{3-52}$$

(See **Appendix B** for details on Counting Statistics).

χ^2 is the usual misfit statistic (or "goodness of fit" indicator) which is used in the fitting of data d using a model $\tilde{f}$. In this case, it is combined with prior information (in Bayes' equation (3-43)) so that the reconstruction is not based on the evidence of the data alone.

The form of the likelihood described above is widely used but it does depend on the noise obeying Gaussian statistics. If this is not the case in counting experiments, then one would have to use the complete description of the Poisson distribution [Robinson 1991] since the Gaussian distribution is an approximation to the Poisson distribution. The Gaussian approximation is valid when there are a large number of counts but if the counts are low then the Poisson distribution would have to be used to describe the distribution of noise in the data. Assuming Poisson statistics, the likelihood now becomes [Robinson 1991]

$$P(d \mid f \cdot I) = \prod_{i=1}^{N_{data}} \exp\left(-\tilde{f}_i\right) \frac{\left(\tilde{f}_i\right)^{d_i}}{d_i!} \tag{3-53}$$

If, however, the experiment does not obey counting statistics, then neither (3-49) nor (3-53) will be valid and some other expression for the likelihood will have to be derived.

3.5.2. The Prior Probability Distribution

The term $P(f \mid I)$ in Bayes' equation is known as the prior probability distribution – prior because it describes how well f can be predicted on the basis of I alone, i.e. prior to knowledge of d. The prior distribution is important because it allows for existing knowledge about a problem to be incorporated into the reconstruction.

Skilling [1989] derived the form of the prior by considering it to be a monotonic function of the entropy. He argued that if a general theory existed then it must apply to special cases. So, if the answer to a special case is known then it can be used to constrain the theory by eliminating all possibilities that lead to an incorrect answer. Skilling then used the example of a team of monkeys throwing balls at cells to derive a form of the prior for this special case and then generalised the special case, to obtain the prior in (3-54) below. Skilling showed that the only consistent prior distribution for positive, additive images is

$$P(f \mid I) = \frac{\exp(\alpha \mathbf{S}(f))}{Z_{\mathrm{s}}(\alpha)} \tag{3-54}$$

where $\mathbf{S}$ is the entropy as defined by (3-9a) and (3-9b) and $Z_{\mathbf{S}}(\alpha)$ is a normalisation function, given by

$$Z_{\mathbf{S}}(\alpha) = \int_0^\infty \frac{\exp(\alpha \mathbf{S}(f))}{\prod_{i=1}^{N} \sqrt{f_i}} \mathrm{d}^N f = \prod_{i=1}^{N} \left(\frac{\alpha}{2\pi f_i} \right)^{\frac{1}{2}} \tag{3-55}$$

α is a scaling factor or regularisation parameter which is introduced as a dimensional constant in the derivation of the prior. α has dimensions of $\mathbf{S}^{-1}$ ($\mathbf{S}$ has the units of f). The factor $\prod_i (\sqrt{f_i})^{-1}$ represents the measure on the space of f. The "existing knowledge" incorporated into (3-54) includes positivity, additivity as well as all information contained in the axioms of MaxEnt as described in §3.2.4. The reconstructed image will therefore satisfy all these criteria.

The posterior distribution can be written as

$$P(f \mid d \cdot I) \propto \exp\left(\alpha \mathbf{S} - \tfrac{1}{2}\chi^2\right) \tag{3-56}$$

for Gaussian statistics.

Maximising the posterior distribution with respect to f yields the reconstruction most consistent with the data. If $\mathbf{S}$ is taken to be a measure of uncertainty then maximising $\mathbf{S}$ alone yields a reconstruction with the maximum uncertainty i.e. the reconstructed image will be featureless. Taking the data into consideration, as in (3-56), yields a reconstruction which contains features if and only if it is supported by the data.

3.5.3. The Regularisation Parameter, α

The regularisation parameter α that was introduced in (3-54) is known generally in statistics as a hyperparameter or, perhaps more appropriately, as a nuisance parameter. These parameters appear in distribution functions but are usually not of any interest and have to be eliminated in some way. In MaxEnt, α can be taken to be a weighting factor which weights the entropy to the data in the posterior in (3-56). For a large value of α, the reconstruction favours the entropy and therefore the model, whereas for a small value of α the reconstruction favours the data. It is therefore important how α is dealt with in the calculation. There are two basic methods of eliminating the hyperparameter, both of which have been used in MaxEnt:

1. setting it to a constant value which is then used in subsequent calculations, or
2. the proper Bayesian technique of marginalisation, which is integrating over the hyperparameter to eliminate it from the calculation.

There are two methods of choosing a constant value for α, the first is known as Classic Maximum Entropy and the second as the Evidence Approximation.

Classic Maximum Entropy

In Classic MaxEnt, α is chosen such that $\chi^2 = N_{data}$, the number of data points – a frequentist assumption. Gull [1989] discusses the use of this choice of α, and says that the $\chi^2 = N_{data}$ criterion is *ad hoc*. Its use leads to an average one standard deviation bias towards the model which defeats the purpose of the entropic term in the posterior, since the main aim of entropy is to deal with noise and missing information, and not be biased to a specific result.

The Evidence Approximation

In the evidence approximation [Strauss *et al* 1993, Wolpert & Strauss 1995], α is chosen on the basis of the data. The probability distribution for the evidence $P(\alpha|d)$ is maximised with respect to α, and α is then set to $\overline{\alpha}$ where $P(\overline{\alpha} \mid d)$ is a maximum. Wolpert & Strauss [1995] discuss the validity of the evidence approximation and give details of how to test its validity. It is evident, however, that this approximation is not a good choice in some cases since $P(\alpha \mid d)$ may not always be a peaked function of α.

α-marginalisation

The fully Bayesian method of dealing with the hyperparameter is to marginalise (§3.3.4) over α by calculating $P(f\alpha \mid I)$ and then integrating over α to obtain $P(f \mid I)$:

$$P(f \mid I) = \int_0^\infty P(f \cdot \alpha \mid I) d\alpha = \int_0^\infty P(f \mid \alpha \cdot I) P(\alpha \mid I) d\alpha \tag{3-57}$$

Marginalisation is a consistent method of eliminating α and does not favour any specific value of α. Although marginalisation is the correct method of eliminating α, the evaluation of the multi-dimensional integrals is computationally intensive. With this in mind, Strauss *et al* [1993] decomposed the multi-dimensional integral into a product of one-dimensional integrals. This simplifies the calculation to a certain extent but the calculation is still computationally demanding. Results using α-marginalisation show an improvement over results using classic MaxEnt and the evidence approximation, but over-fitting in the reconstruction still occurs [von der Linden *et al* 1997a]. A reconstruction "over-fits" the data if it follows the data too closely i.e. it considers features in the data that are due to noise only, to be important. Over-fitted reconstructions therefore tend to be noisy.

3.5.4. Multiresolution Pixons

Reconstructions made on the basis of the theory as described thus far are prone to spurious sources or *ringing* as a result of over-fitting the data i.e. the reconstruction follows the data too closely and the result is not flat in regions of the reconstruction where it should be. This is due primarily to the entropic prior having no local smoothing, a result of the axiom of system independence (**Axiom 3**, §3.2.4). The fact that an image may possess local smoothness may be important prior information which is not included in the standard MaxEnt formalism. The other factor which may lead to over-fitting of the data concerns the grid over which the image is reconstructed. Thus far, it has been ignored in the treatment of the reconstruction problem and each pixel of the grid is treated independently from *all* other pixels. This causes problems if features in the data extend over a few pixels since there will be no correlations between neighbouring pixels in the reconstruction. This means that features extending over a few pixels will not be smooth. Fischer *et al* [1997] incorporate local smoothness or correlations into the reconstruction based on earlier work by Gull [1989], Skilling [1992] and Puetter [1995].

S is maximal for $f = m$ and maximising the entropy tries to keep the reconstruction as close to the model as possible. Therefore, properties of m can be transmitted to the reconstruction. Gull [1989] applied this idea and introduced smoothness into the model m of the entropic prior, by generating it from a blurred image of hidden variables $\tilde{m}$, i.e. $m = B * \tilde{m}$, where B is the pre-blurring function. Each point in $\tilde{m}$ is considered to be logically independent from neighbouring points and by smoothing these values, the resulting model m contains correlations between neighbouring points. The concept of a pre-blurring function is completely separate from that of convolution with the detector response function. Convolution is a special case of the distortion operator **K** as in (1-1) but the pre-blurring function applies to all inverse problems where correlations have to be introduced into the reconstruction. Skilling [1992] introduced smoothness via a hidden image which was blurred with a constant *global* smoothing function.

Puetter [1995] introduced the concept of *pixons* by incorporating the size of the grid over which the image is reconstructed into Bayes' equation. Let M represent the set of pixels and their sizes. Then Bayes' equation becomes

$$P(f \cdot M \mid d \cdot I) \propto P(d \mid f \cdot M \cdot I) P(f \mid M \cdot I) P(M \mid I) \tag{3-58}$$

Consider an image with a total of N pixels and c_i counts in each pixel. Then the probability of a particular reconstruction is given by [Puetter 1995]

$$P(\{c_i\} NC) = \frac{C!}{N^C \prod_i c_i!} = P(f \mid M \cdot I) \tag{3-59}$$

where $C = \sum_i c_i$ is the total counts in the image. In order to obtain the reconstruction most consistent with the data, (3-58) is maximised with respect to f. In (3-58), $P(d \mid f \cdot M \cdot I)$ will not change much as M is changed since the generation of the pseudo-data is independent of the set of pixels and their sizes. $P(M \mid I)$ will also not change much as M is varied so the only term in (3-58) that has a strong influence on the posterior $P(f \cdot M \mid d \cdot I)$ is the prior $P(f \mid M \cdot I)$. (3-59) should therefore be as large as possible. This can be achieved by decreasing N or by increasing the number of counts in each pixel. This leads to a few pixels of variable width, with as many counts in each pixel as possible. Puetter refers to the pixels of variable sizes as *pixons* with the pixon widths determined by the information density present in the data, i.e. small pixons for a high information density and large pixons for a low information density. This is illustrated in *Figure 3.3.*

Puetter proposed the use of a local smoothing function with the width of the smoothing function equal to the pixon width, calculated on the basis of the data. This effectively reduces the number of degrees of freedom available for the fitting of the data.

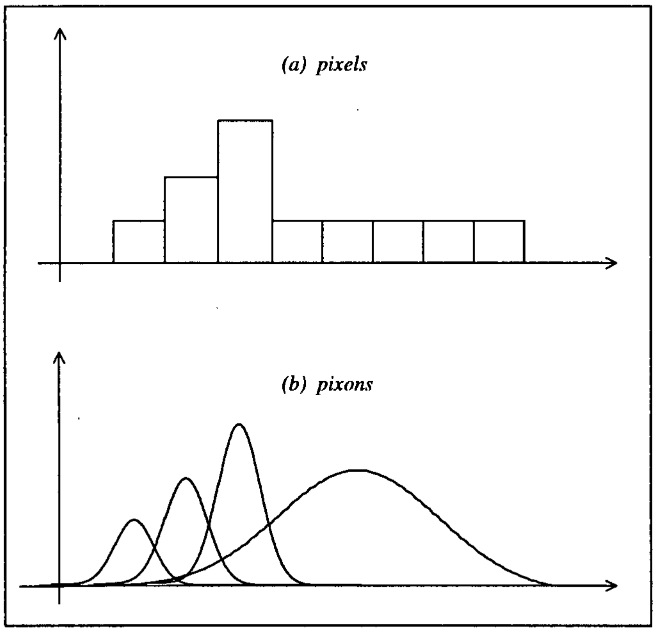

Figure 3.3. Illustration of Pixons. ***(a)*** Conventional image reconstruction with pixels of equal sizes. ***(b)*** Pixons of varying widths. The widths of the pixons are inversely proportional to the information density of the image and the number of pixons controls the degrees of freedom of the fit. In (a) there are 8 degrees of freedom whereas this has been reduced to 4 in (b). Pixons also lead to correlations between neighbouring pixels since the pixons overlap.

Fischer *et al* [1997] incorporated the concept of local smoothness into the Bayesian formulation of MaxEnt. The effective number of degrees of freedom of the reconstruction are reduced with the introduction of local smoothness. The driving force that leads to high correlations and a few degrees of freedom enters into Bayesian Statistics via Occam's Razor[5]. Fischer *et al* [1997] assume that f is a blurred version of a hidden image h with a pre-blurring function B, i.e.

$$f(x) = \int B\left(\frac{x-y}{b(y)}\right) h(y) dy \tag{3-60}$$

where $b(y)$ is a measure of the local width of the pre-blurring function. The hidden image contains no correlations i.e. neighbouring pixels in the hidden image are independent. The form of B is derived by applying the Principle of Maximum Entropy with specified mean and variance, as in §3.2.3.2. This yields a Gaussian for the pre-blurring function B [Fischer *et al* 1997],

$$B\left(\frac{x-y}{b(y)}\right) = \frac{1}{\sqrt{2\pi} b(y)} \exp\left(-\tfrac{1}{2}\left(\frac{x-y}{b(y)}\right)^2\right) \tag{3-61}$$

The pre-blurring width b determines the number of degrees of freedom and cannot be marginalised out of the calculation. Bayes' equation now becomes

$$P(f \cdot b | d \cdot I) = \frac{P(f | b \cdot I) P(d | f \cdot b \cdot I) P(b | I)}{P(d | I)} \tag{3-62}$$

Taking (3-60) into account, (3-50) becomes:

$$\chi^2 = \sum_{i=1}^{N_{data}} \left(\frac{d_i - \tilde{f}_i(B*h)}{\sigma_i}\right)^2 \tag{3-63}$$

where $\tilde{f}(B*h)$ represents simulated pseudo-data generated from a hidden image blurred by a blurring function B. (3-63), combined with (3-49) and (3-51), supplies the likelihood $P(d | f \cdot b \cdot I)$ in (3-62). The prior $P(f | b \cdot I)$ is given by the entropic prior for h since the hidden image does not include correlations.

---

[5] The simplest hypothesis consistent with the data is the best one.

$$P(f \mid b \cdot I) \propto P(h \mid b \cdot I) = \frac{\exp(\alpha \mathbf{S}(h))}{Z_{\mathbf{S}}(\alpha)} \tag{3-64}$$

with

$$Z_{\mathbf{S}}(\alpha) = \int_0^{\infty} \frac{\exp(\alpha \mathbf{S}(h))}{\prod_{i=1}^{N} \sqrt{h_i}} \mathrm{d}^N h = \prod_{i=1}^{N} \left(\frac{\alpha}{2\pi h_i} \right)^{\frac{1}{2}} \tag{3-65}$$

The prior probability $P(b \mid I)$ must contain all available information about the pre-blur widths b. The upper limit b_u of each width is given by the image range and corresponds to a global smoothing. The lower limit b_l corresponds to the width that leads to no smoothing. Fischer *et al* [1997] use a pre-blur width distribution where widths differ by no more than one pixel width from neighbouring pixels. This condition implies that the pre-blur widths are a smooth, slowly varying function of the pixel numbers. Fischer *et al* [1997] also use a Gaussian prior for the derivative of b with a standard deviation of one for $b \in [b_l, b_u]$,

$$P(b \mid I) = \begin{cases} \exp(-\frac{1}{2} \sum_i (b_i - b_{i-1}) & \text{for } b \in [b_l, b_u] \\ 0 & \text{elsewhere} \end{cases} \tag{3-66}$$

Combining equations (3-49), (3-51) and (3-63)-(3-66) into (3-62), yields the posterior probability distribution that incorporates local smoothness into the reconstruction

$$P(f \cdot b \mid d \cdot I) \propto \prod_{i=1}^{N} \left(\frac{\alpha}{2\pi h_i} \right)^{\frac{1}{2}} \prod_{j=1}^{N_{data}} \left(\sqrt{2\pi} \sigma_j \right)^{-1} \cdot \exp\left(\alpha \mathbf{S}(h) - \frac{1}{2} \chi^2 - \frac{1}{2} \sum_k (b_k - b_{k-1})^2 \right) \tag{3-67}$$

where

$$\mathbf{S}(h) = \sum_{i=1}^{N} h_i - m_i - h_i \ln\left(\frac{h_i}{m_i} \right) \tag{3-68}$$

Equation (3-67) is then maximised with respect to both h and b to yield the reconstruction that is most consistent with the data within the bounds imposed by the noise. The prior information incorporated into the reconstruction includes positivity, additivity and local smoothness.

Jansson [1984] suggested that certain questions should be asked about any deconvolution technique. Questions that should be raised include:

- Is noise dealt with in a fair manner?
- Does the method converge?
- Are the reconstructions physically meaningful?
- Is the method applicable to all problems?
- Does the method have a sound mathematical basis?

In this light, most of the deconvolution methods described in Chapter 2 fail to satisfy one or more of the above criteria. Bayesian Statistics, using the concept of multiresolution pixons, produces smooth, physically meaningful results, and the formalism is general and can be applied to all problems.

3.6. Implementation of Theory

The theory, incorporating local smoothness into the reconstruction, was programmed into FORTRAN-77 using the algorithm described below. The program was written by R. Fischer of the Surface Physics Division, Max Planck Institute for Plasma Physics, Garching, Germany [Fischer *et al* 1997]. The original program is a general-purpose program for solving inverse problems of the type described in (1-1) and was modified for the purpose of this investigation. The program maximises the posterior probability distribution (3-67) with respect to the hidden image h as well as the widths b of the pre-blurring function.

Since the dependence of the reconstruction on the prior is reduced by the use of pixons, the value of α was set to a constant in the calculation [Fischer *et al* 1997]. The value of m was also unimportant and was also set to a constant value equal to the average of the data.

The routine `E04UCF` used for the maximisation of the posterior is part of the NAG FORTRAN Library of numerical algorithms [NAG 1993].

```
Program PIXONS
     Initialise variables
     Read in data
     Determine values for $m$, $\alpha$, $\sigma$
```

```
    Maximise $P(f \cdot b | d \cdot I)$ with respect to $h$ and $b$ using NAG routine E04UCF
    Use results of maximisation to determine $f$ by blurring with the pre-blur function $B$
END of Program
```

Chapter 4

Techniques of Ion Beam Analysis

4.1. Introduction

The need to determine the elemental composition of materials began with the discovery that matter is composed of different atoms. Each technique of elemental analysis uses characteristic properties of each element (e.g. mass, charge of nucleus or electromagnetic radiation emitted or absorbed) to determine the composition, concentration and distribution of various elements in materials. The use of high-energy ion beams for elemental analysis is well established, as are the many techniques of IBA. Rutherford Backscattering Spectrometry (RBS), Proton Induced X-ray Emission (PIXE), Nuclear Reaction Analysis (NRA), Elastic Recoil Detection Analysis (ERDA) and Scanning Transmission Ion Microscopy (STIM) are some of the techniques routinely used at the Van de Graaff Group (VDGG) of the National Accelerator Centre (NAC) in Faure, South Africa. These techniques are described briefly below. The techniques of interest to this investigation are RBS and PIXE and are discussed further in §4.3 and §4.4, respectively.

4.1.1. General IBA Techniques

Rutherford Backscattering Spectrometry, as the name may suggest, had its origins in the experiment designed to test Rutherford's theory of the structure of the atom. In RBS a beam of charged particles (usually ^{4}He) is focussed on a target. Incident ions that collide with heavier nuclei are scattered in the backwards direction, and analysis of these backscattered particles yields information about mass (and therefore atomic number), depth distributions and concentrations of elements in the target [Chu *et al* 1978].

In PIXE, a beam of protons is normally used to ionise the atoms in a target. The characteristic x-rays that are emitted when the atom returns to its ground state are detected and analysed to yield trace element information. In comparison to RBS, no depth information is obtained directly from PIXE but it is a multi-elemental analysis technique [Johansson & Campbell 1988].

In NRA, products from nuclear reactions such as charged particles or γ-rays are analysed to obtain information on depth and concentration. NRA is typically a single element analysis technique, but is useful in detecting low-Z elements in high-Z matrices (e.g. N in Fe) for which RBS is not optimal [Watt & Grime 1987].

Elastic scattering also forms the basis of ERDA. When incident ions scatter elastically off lighter atoms in the target, the recoiling target atom can be detected at forward angles. The incident particles in ERDA must be heavier than the target atoms, which means that heavier elements can only be analysed if one has access to a heavy-ion accelerator. ERDA is however very useful in the analysis of hydrogen where a ^{4}He-beam is used [Tirira *et al* 1996].

STIM is based on the principle that if a sample is very thin ($\leq$ 50 μm in the case of protons and ~10 μm in the case of alphas) a high-energy ion beam will only lose part of its energy while passing through it. The analysis of the energy losses yields information on target density and thickness. STIM is popular in tomography where a three dimensional image is reconstructed from one-dimensional scans [Johansson *et al* 1995].

4.1.2. Nuclear Microprobes

The advantage of using ion beams focussed to small lateral sizes lies in the ability to scan over the target obtaining lateral information about elemental distributions using any of the techniques described above. The fact that the lateral resolution could be improved by the use of smaller beams was only realised long after IBA had become established. This led to the development of microprobes using ion beams with sizes of the order of a few microns. Watt & Grime [1987] give a historical account of the development of microprobes from the first use of a collimated beam in 1953 to the use of the focussed beams for trace element analysis.

4.2. The Nuclear Microprobe Facility at NAC

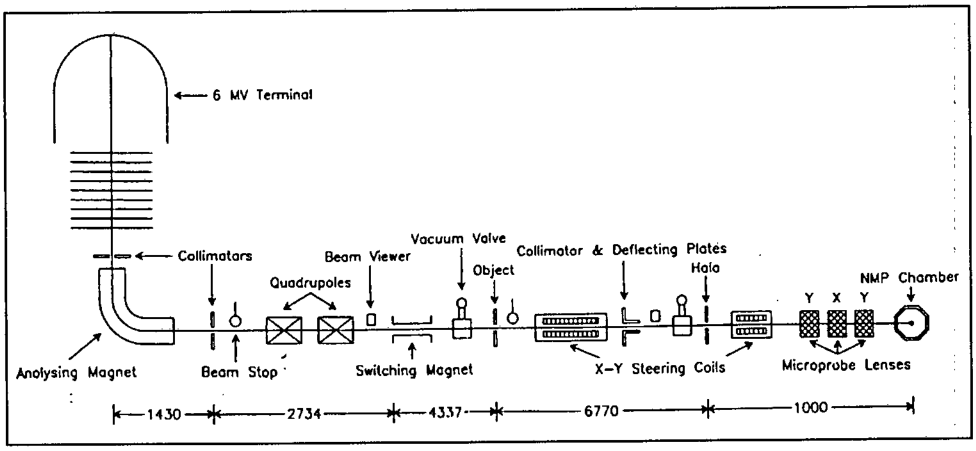

Figure 4.1. A schematic of the nuclear microprobe layout at NAC. The distances are in mm and the layout is not drawn to scale. [Prozesky *et al* 1995]

The nuclear microprobe (NMP) at NAC was installed on the 0° beam line of the single-ended 6 MV Van de Graaff accelerator in 1991 (*Figure 4.1*). The NMP system is now highly automated with most of the NMP being computer-controlled. Since its installation it has been successfully used in the analysis of samples from fields including archaeology, biology, geology, materials science and medicine.

4.2.1. The Van de Graaff Accelerator

The Van de Graaff accelerator at NAC accelerates ions vertically downwards. Energy selection of the accelerated particles is made by a 90° analysing magnet situated at the base of the accelerator (*Figure 4.1*). After passing through the analysing magnet, the beam travels through energy stabilisation slits, situated before the main beam stop. The quadrupole doublet after the beam stop focuses the beam at the object slits. The X-Y steering coils are controlled by a PC, which facilitates easy positioning and scanning of the beam. The focussing of the beam at the target is achieved by means of a quadrupole triplet. The beam line is evacuated to a base pressure of less than 10^{-6} torr.

4.2.2. The Chamber

The microprobe target chamber at NAC (*Figure 4.2*) is a modified version of the standard Oxford NMP chamber from Oxford Microbeams (model OM-70e). The features include a lithium-drifted silicon (Si(Li)) x-ray detector ⑥ (in *Figure 4.2*) for PIXE measurements, situated 35 mm from the target at 135° to the direction of the incoming beam. The x-ray detector is separated from the evacuated chamber by a 25 μm Mylar window ⑦ which protects the 8 μm Be window of the Si(Li) detector during the evacuation of the chamber and also stops backscattered protons from entering the detector. The annular silicon surface barrier (SSB) detector ⑨ is at 176°±1° with respect to the incoming beam and is used for RBS measurements while the SSB detector ⑤ situated at 30° with respect to the beam is used for ERDA as well as for off-axis STIM measurements. The ERDA/STIM detector is equipped with a 7 μm Al foil which prevents scattered ions from reaching the detector during ERDA. The foil is removed for STIM measurements. Other features in the chamber include an optical microscope ② at 45° to the sample normal, a Faraday Cup ④ which collects the charge during analyses of thin targets that do not stop the beam and a secondary electron suppressor ①. A negative voltage on the secondary electron suppressor prevents the emission of secondary electrons from the target surface and ensures accurate charge collection of incident ions.

A custom-made lid has been installed that allows for stepper motor control of the target ladder in the *X*-, *Y*- and *Z*-directions. The stepper motors are controlled by computer. A permanent set of standards has been mounted on the target ladder and the positions of these standards are stored in computer memory to simplify the calibration process. The chamber is evacuated down to a pressure of about 10^{-5} torr, with a diffusion pump backed by a roughing pump, in about five minutes, which is adequate for analytical purposes. A pressure of 10^{-6} torr is reached after about an hour.

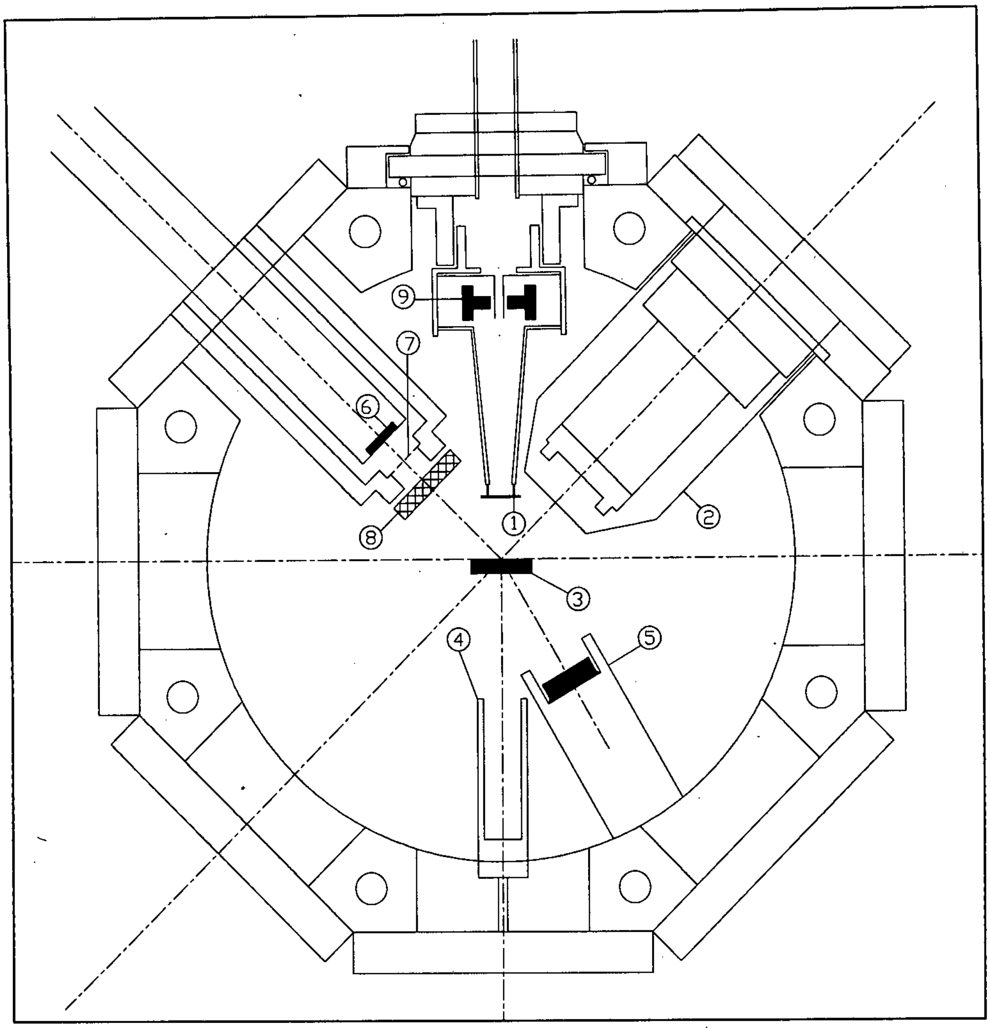

Figure 4.2. NMP target chamber at NAC. ① electron suppressor ② optical microscope ③ target ④ Faraday Cup ⑤ ERDA/STIM detector ⑥ PIXE detector ⑦ Mylar window ⑧ x-ray filter ⑨ RBS detector. The beam enters the chamber from the top of the figure.

4.2.3. Data Acquisition and Analysis

A 4000/VLC VAXstation controls the acquisition and analysis of data from a single terminal at the VDGG. The acquisition of data is handled by the general-purpose data acquisition system, XSYS [1985], which is linked to the VAX computer network via a VME crate. The data acquisition, steering and positioning of the beam is controlled by a local PC interfaced to the steering coils and the VAX using an in-house developed software package [Churms *et al* 1993].

4.3. Rutherford Backscattering Spectrometry

In RBS a beam of energetic charged particles impinges on a target and incident ions interact with nuclei in the target. If the target nuclei are heavier than the incident ions, the incident particles are scattered in the backward direction as a result of the repulsive Coulomb interaction between the incident ion and the target nucleus. The energy of the backscattered particles are analysed to obtain information about mass (and therefore atomic number), depth distributions and concentration of elements in the target.

4.3.1. The Kinematics of RBS

The interaction between the incident particle and the target atom is considered to be an elastic collision, and the mechanics of the system is solved by considering the conservation of energy and momentum. The quantities that are required to analyse an RBS spectrum are the kinematic factor *K*, which is the ratio of the backscattered energy to that of the incident energy, the scattering cross-section σ, which is a measure of the likelihood of an interaction occurring and the stopping cross-section ε, which gives an indication of how the incident particle loses energy as it moves through the target.

4.3.1.1. Kinematic Factor, *K*

Consider an elastic collision between masses m_0 and M (*Figure 4.3*). Mass m_0 has energy E_0 before the collision while mass M is stationary. After the collision, m_0 moves in a direction θ with energy E_1 while M moves in a direction ϕ with energy E_2. The kinematic factor is then [Chu *et al* 1978]

$$K = \frac{E_1}{E_0} = \left[\frac{m_0 \cos\theta + (M^2 - m_0^2 \sin^2\theta)^{\frac{1}{2}}}{m_0 + M} \right]^2 \qquad (4\text{-}1)$$

The concept of the kinematic factor leads to the mass determination capability of RBS since in a standard RBS experiment the only unknown quantity in (4-1) is M.

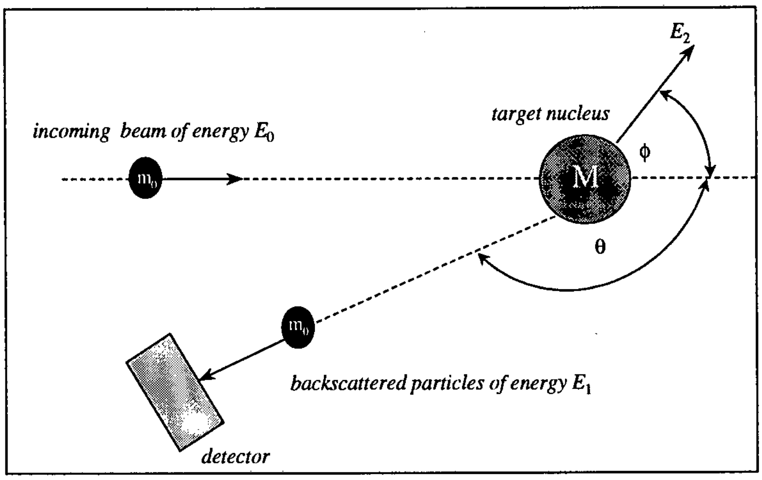

Figure 4.3. Schematic representation of an elastic collision in RBS. The incoming particle m_0 and target nucleus M undergo a Coulombic interaction and the energy of the backscattered particle is measured at an angle θ (relative to the incoming beam direction). m_0 has an energy E_0 before the collision while M is stationary. After the collision m_0 moves in a direction θ with energy E_1 while M moves in the direction ϕ with energy E_2. (Figure adapted from [Chu *et al* 1978])

4.3.1.2. Scattering Cross-Section, σ

The scattering cross-section σ gives an indication of the likelihood of a scattering event occurring while the differential scattering cross-section $d\sigma/d\Omega$ gives an indication of detecting a scattered particle with the detector spanning a solid angle $d\Omega$ situated at an angle θ with respect to the direction of the incident beam (*Figure 4.4*).

For an elastic Coulomb collision, the differential cross-section is given by Rutherford's formula [Chu *et al* 1978]

$$\frac{d\sigma}{d\Omega}=\left(\frac{Z_{m_0}Z_M e^2}{4E}\right)^2\frac{4}{\sin^4\theta}\frac{\left[\left(1-\left(\frac{m_0}{M}\sin\theta\right)^2\right)^{\frac{1}{2}}+\cos\theta\right]^2}{\left[1-\left(\frac{m_0}{M}\sin\theta\right)^2\right]^{\frac{1}{2}}} \quad (4\text{-}2)$$

where Z_i is the atomic number of mass i, E is the energy of the incident atom immediately before scattering and $e^2 = q^2/4\pi\varepsilon_0$ is the electronic charge. (4-2) shows that heavy atoms are more efficient scatterers than light atoms because of the Z_M^2 dependence. This means that RBS is more sensitive to the detection of heavy elements than light elements.

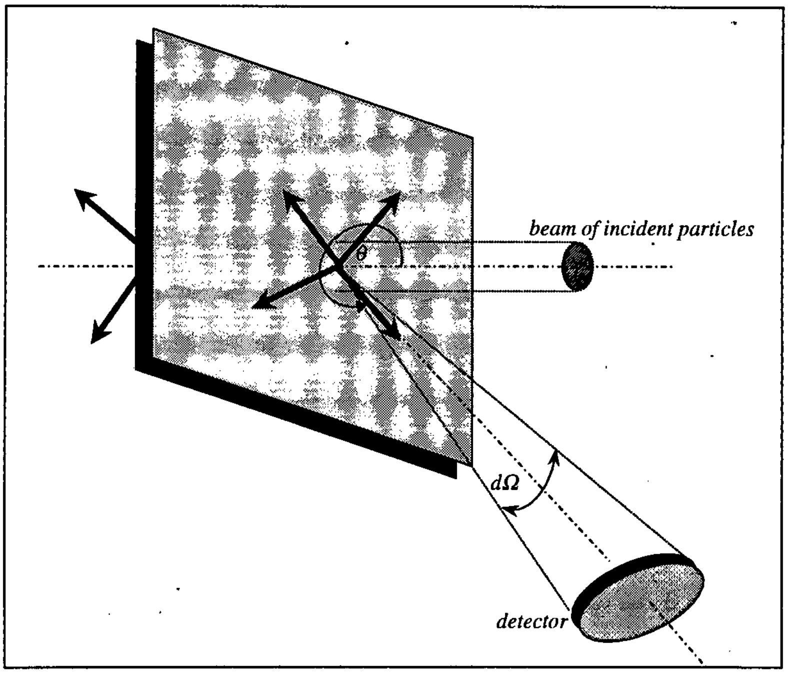

Figure 4.4. Scattering cross-section, σ. A narrow beam of particles strikes a target. The scattering cross-section gives an indication of the likelihood of a scattering event occurring. (Figure adapted from [Chu *et al* 1978])

4.3.1.3. Stopping Cross-Section, ε

When an energetic particle strikes a target, it will most likely penetrate it. As the particle moves through the target, it will lose kinetic energy via collisions and slow down. The amount of energy ΔE lost per distance Δx travelled (*Figure 4.5*) is referred to as the specific energy loss dE/dx and if known, can be used to obtain the energy E of a particle at a depth x in the target.

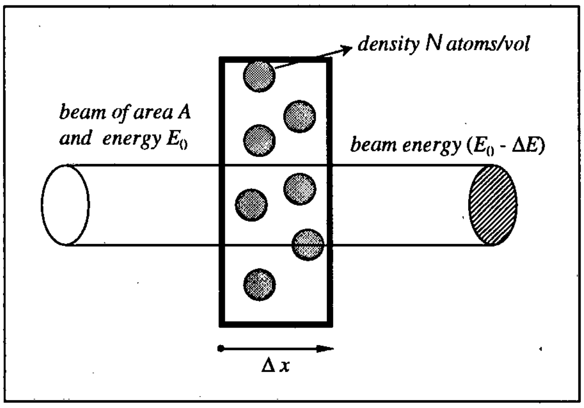

Figure 4.5. Stopping Cross-Section, ε. A beam of energy E_0, travelling a distance Δx through a material containing N atoms/vol. will lose an amount of energy ΔE. (Figure adapted from [Chu *et al* 1978])

The stopping cross-section ε is defined as,

$$\varepsilon \equiv \frac{1}{N}\frac{dE}{dx} \tag{4-3}$$

for a target containing N atoms/vol. When a target contains more than one element, the stopping cross section $\varepsilon^{A_m B_n}$ of a molecule $A_m B_n$ or a mixture with an equivalent composition of $A_m B_n$ is given by Bragg's Rule [Chu *et al* 1978]:

$$\varepsilon^{A_m B_n} = m\varepsilon^{A} + n\varepsilon^{B} \tag{4-4}$$

4.3.2. Analysis of RBS Spectra

4.3.2.1. A Typical RBS Spectrum

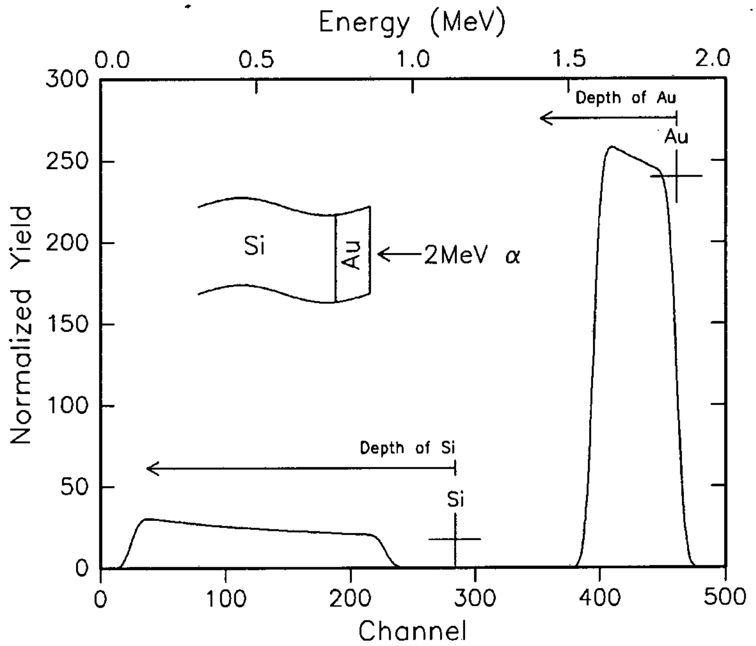

Figure 4.6. Simulated spectrum of Au on Si. The crosses indicate surface positions of the element (vertical bar) and the expected yield (horizontal bar). The co-ordinate systems for the determination of depths are also shown.

Figure 4.6 shows a typical RBS spectrum of a 2000 Å layer of Au on a Si substrate. The spectrum was simulated for a beam of 2 MeV alphas and a detector resolution of 50 keV. The crosses indicate the surface positions (vertical bar) and expected yields (horizontal bar) of Si and Au. There are basically two rules to interpreting an RBS spectrum [Chu *et al* 1978]:

- high atomic numbers give high yields; low atomic numbers give low yields with the scattering cross-section scaling the yields
- heavy elements are at high energies; light elements are at low energies with the incident energy and kinematic factor determining the leading edge of the signal.

Each signal also gives depth information about the element in the sample. As seen in *Figure 4.6* each element has its own co-ordinate system for the determination of depth profiles. Elemental information on Au can be determined in a number of ways from the spectrum:

- the width of the Au signal
- the height of the Au signal
- the offset of the Si signal from its surface position

4.3.2.2. RBS Depth Scale and Spectrum Height

Depths in RBS are calculated by using the specific energy loss, dE/dx. Consider the situation where a beam of particles of energy E_0 impinges on a target at an incident angle θ_1, with respect to the normal to the surface of the target. A scattering event occurs at a depth x with the beam energy reduced to E (*Figure 4.7*). The scattered particle exits the target with an energy E_1 at an angle θ_2. Along the inward path

$$\frac{x}{\cos\theta_1} = -\int_{E_0}^{E}\left(\frac{dE}{dx}\right)^{-1} dE \qquad (4\text{-}5)$$

and along the outward path,

$$\frac{x}{\cos\theta_2} = -\int_{KE}^{E_1}\left(\frac{dE}{dx}\right)^{-1} dE \qquad (4\text{-}6).$$

E_1 is the measured energy of the particle and has to be calculated from (4-5) and (4-6). There are two problems that have to be solved before depths can be calculated: the determination of dE/dx at a depth x and the fact that only E_0 and E_1 are known. These problems can be resolved in the following ways:

1. use tabulated values of dE/dx and integrate numerically, determining x and E iteratively

2. assume dE/dx is constant over each path and eliminate E from the integral (surface approximation)
3. assume some functional form for dE/dx and use the same technique as in 1.

Chu *et al* [1978] discuss the details of solving equations (4-5) and (4-6).

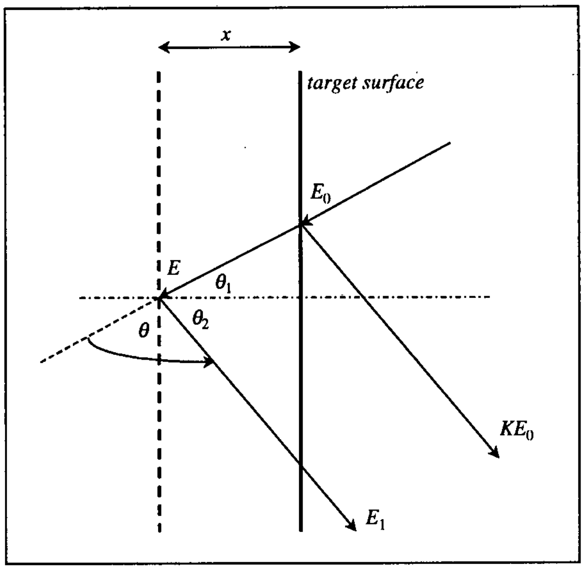

Figure 4.7. RBS depth scale. A beam of particles of energy E_0 impinges on a target at an incident angle θ_1, with respect to the normal to the surface of the target. A scattering event occurs at a depth x with the beam energy reduced to E. The scattered particle exits the target with an energy E_1 at an angle θ_2. (Figure adapted from [Chu *et al* 1978])

The number of counts at a specific energy can be used to determine the concentration of the corresponding element at a specific depth. The counts in channel i can be seen to originate from a layer at depth x_i of thickness τ_i. The number of counts C_i in channel i therefore depends on the thickness τ_i of the layer and the number of atoms contained in this layer. It is therefore possible to determine the number of atoms, $N\tau_i$, contained in the layer (assuming τ_i is known) from the spectrum height. N is the atomic density of the sample element. The number of counts in channel i is given by [Chu *et al* 1978]

$$C_i = \frac{d\sigma}{d\Omega}(E_i)\Omega Q \frac{N\tau_i}{\cos\theta_1} \tag{4-7}$$

where $d\sigma/d\Omega(E_i)$ is the differential cross-section evaluated at energy E_i, Q is the number of particles incident on the sample and θ_1 is the angle of incidence of the beam on the target as in *Figure 4.7.*

4.3.2.3. Spectrum Analysis and RUMP

RBS yields both elemental and depth information and this information is often mixed in the spectrum. The analysis of RBS spectra is made easier with computer programs where the most useful programs are able to simulate spectra from given target information. A popular package is called RUMP [Doolittle 1986] and it includes all above calculations. RUMP was used to simulate the spectrum in *Figure 4.6.*

4.3.3. Silicon Surface Barrier Detectors

SSB detectors are used for RBS measurements. In order to deconvolute the SSB response function from RBS spectra, a very accurate description (see §2.1) of the response function is needed. A model is obtained by examining the processes that lead to the detection and measurement of the energy of a backscattered particle in the semiconductor material of the detector.

4.3.3.1. Basic Principles of Operation of Semiconductor Detectors

The use of semiconductor materials as detectors is based on the principle that charged particles or electromagnetic radiation passing through a semiconductor produces electron-hole pairs. Each detector has an active volume consisting of the depletion region created when a *p*-*n* junction is reverse biased. The radiation or charged particle must be stopped within the active volume for an accurate energy measurement. The charge carriers that are created are swept out of the active volume by an applied electric field and their motion constitutes a signal pulse. The energy of the ionising radiation is then determined from the amplitude of the pulse since this is proportional to the number of electron-hole pairs created in the active volume. The average energy required to produce an electron-hole pair is known as the ionisation energy ε_i. The ionisation energy for Si and Ge are ~ 3eV [Knoll 1979].

The first semiconductor detectors were known as diffused junction detectors because they were fabricated by diffusing impurity atoms into *n*- or *p*-type semiconductors. Problems arose when the donor concentration in the *n*-type material was higher than the acceptor concentration in the *p*-type material. In this case, the depletion region extends far into the *p*-

type semiconductor leaving a layer of *n*-type material outside the depletion region. This surface layer of *n*-type material represents a dead layer through which incident radiation must pass before reaching the depletion region. The presence of a dead layer introduces uncertainties into the energy measurement because the ionising radiation loses energy as it passes through the dead layer. The energy loss is not uniform and it leads to a spread of energies. Since the energy of the ionising radiation is determined by the number of electron-hole pairs created in the active volume alone, energy measurements will be inaccurate if the thickness of the dead layer is not accurately known.

SSB detectors were developed to reduce the effect of the dead layer. These detectors are based on the principle that the role of the *p*-type material in forming the *p-n* junction can be assumed by a high density of electron traps, and are fabricated by evaporating an element with a high affinity for electrons (e.g. Au) onto the *n*-type material resulting in a detector with a dead layer much thinner than diffused junction detectors. The notable disadvantage of SSB detectors is their sensitivity to light which can result in a high noise level at room temperature. They are, however, ideal for the detection of α–particles and other light ions at room temperature [Knoll 1979].

4.3.3.2. The Response Function of SSB Detectors

The response of an SSB detector to a beam of monoenergetic particles is, unfortunately, not a δ-function as one would hope for but instead consists of a peak that has been broadened by noise and statistical fluctuations in the charge creation process. If these were the only sources of peak broadening then the response function of an SSB detector would be a Gaussian but other effects in the detector, such as incomplete charge collection due to trapping and recombination in the surface region, lead to low energy tails on the main Gaussian peak.

Noise

The main sources of noise in SSB detectors are leakage current (generated in the bulk and in the surface regions of the detector) and Johnson noise associated with series resistance or poor electrical contacts to the detector. These sources add approximately in quadrature to give the overall broadening due to noise and is, in general, independent of the ionising radiation:

$$(\Delta E_{noise})^2 = (\Delta E_{bulk})^2 + (\Delta E_{surface})^2 + (\Delta E_{Johnson})^2 \qquad (4\text{-}8)$$

Statistical Fluctuations in the Generation of Charge Carriers

The creation of electron-hole pairs in the detector is a random process and the fluctuation in the number of charge carriers contributes to peak broadening. If all the events along the path of the ionising particle were independent, then the charge creation process would be governed by Poisson statistics and the variance in the total number of electron-hole pairs would equal the total number produced i.e. E/ε_i, where E is the energy of the ionising particle and ε_i is the ionisation energy of the detector semiconductor material. The observed variance differs from that predicted by a Poisson model and the Fano factor F is introduced to relate the observed variance to the Poisson variance [Knoll 1979]

$$F \equiv \frac{\text{observed statistical variance}}{E/\varepsilon_i} \qquad (4\text{-}9)$$

The Fano factor is ~0.1 for both Si and Ge. The contribution to the FWHM of the Gaussian peak by statistical fluctuations $\Delta E_{statistical}$ is given by

$$(\Delta E_{statistical})^2 = 2.35^2\left(F\,E/\varepsilon_i\right)\varepsilon_i^2$$

and the final energy resolution is then [Johansson & Campbell 1988]

$$\Delta E = \left[(\Delta E_{noise})^2 + 2.35^2 F\varepsilon_i E\right]^{\frac{1}{2}} \qquad (4\text{-}10)$$

where ΔE represents the FWHM of the main Gaussian peak. The factor of 2.35 arises from the relationship between the FWHM and standard deviation σ_i of a Gaussian,

$$\text{FWHM}_i = 2\sqrt{2\ln 2}\,\sigma_i \approx 2.35\sigma_i$$

Low Energy Tails

Low energy tails on the main Gaussian peak arise from incomplete charge collection due to trapping and recombination effects.

If electrons and holes recombine before they reach the electrode, charge carriers are lost. Charge carriers also become ‘trapped’ by impurity atoms with potential wells deeper than those of the donor/acceptor impurities. The amount of trapping is proportional to the distance travelled before the charge carrier reaches the electrode and is therefore variable. This leads to low energy tails on peaks which correspond to pulses in which less than the total amount of charge generated by the ionising radiation has been collected. Tailing can also be due to energy losses in the dead layer, where charge collection is inefficient [Knoll 1979].

4.3.3.3. Description of SSB Response Function

Taking the above factors into account, the response function of the SSB detector can be described by a Gaussian to which two low-energy exponential tails have been added:

$$r_{SSB} = \frac{1}{\sqrt{2\pi}\Delta E}\exp\left(-\tfrac{1}{2}\left(\frac{y-x}{\Delta E}\right)^2\right) + E_1(A_1,s_1) + E_2(A_2,s_2) \tag{4-11a}$$

$$E_i = \tfrac{1}{2}A_i\exp\left(\frac{y-x}{s_i}\right)\mathrm{erf}\left(\frac{y-x}{\sqrt{2}\Delta E} + \frac{\Delta E}{\sqrt{2}s_i}\right) \tag{4-11b}$$

for (i = 1,2). A_i represents the amplitude of the exponential tail, s_i represents the slope of the exponential tail and ΔE represents the FWHM of the main Gaussian as described by (4-10). The error function is

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}}\int_0^x \exp(-t^2)dt \tag{4-12}$$

The parameters A_i, s_i and ΔE are determined by doing a least squares fit to a spectrum corresponding to the detector response function, such as the RBS signal from a thin layer of a heavy element on a light substrate.

4.4. Proton Induced X-ray Emission

4.4.1. X-ray Emission Analysis

When atoms are bombarded by high-energy particles, inner shell electrons are either knocked out or elevated to unfilled upper states thereby creating a vacancy in the inner shell. An electron from an upper level then drops to fill the vacancy and the excess energy is emitted either as an x-ray or as an Auger electron. The process continues until the atom has returned to the ground state. X-ray emission analysis is therefore a two step process of ionisation and emission.

X-rays are classified according to the level which has the original vacancy e.g. x-rays arising from transitions from the L, M, N,... levels to the K level are K_α, K_β, K_γ, ... x-rays respectively. L x-rays arise when there is a vacancy in the L level, M x-rays when there is a vacancy in the M level and so on (*Figure 4.8*) [Beiser 1987].

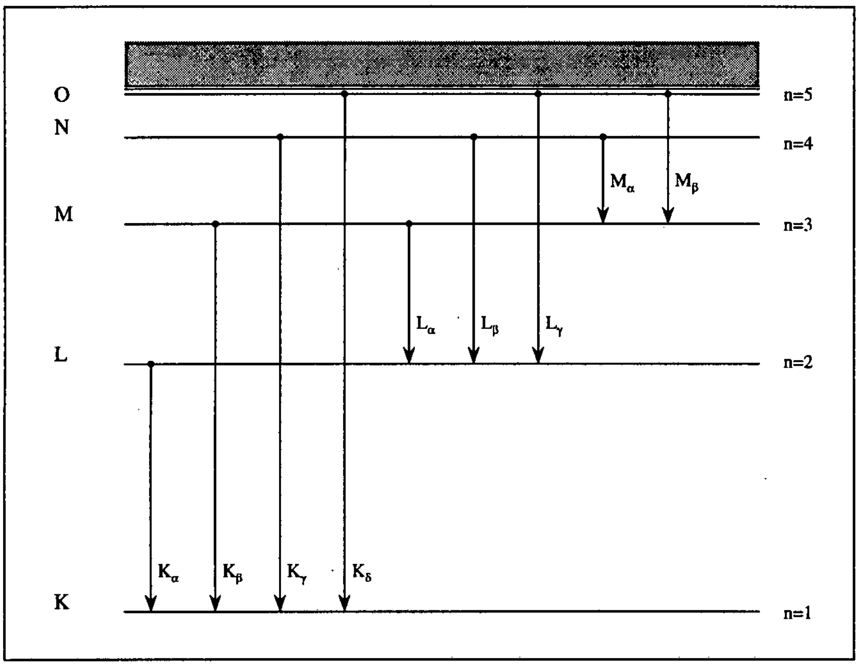

Figure 4.8. The x-rays are classified according to the levels between which the transition occurs [Beiser 1987].

In general, x-ray emission analysis consists of the excitation of atoms by various types of ionising radiation and the detection and analysis of the emitted x-rays. The x-ray technique used at the VDGG is called particle induced x-ray emission or PIXE. Proton beams have become the favoured source of excitation and PIXE usually refers to *proton induced x-ray emission*. PIXE was first used as an analytical technique at the Lund Institute of Technology in Sweden. The first results were published in 1970 [Johansson *et al* 1970] and hailed the beginning of a new era in multi-elemental analysis.

Figure 4.9 shows a typical PIXE spectrum, measured at the VDGG, of a synthetic mineral/glass standard that demonstrates the multi-elemental nature of the technique. The x-ray spectrum produced consists of peaks superimposed on a background due to various bremsstrahlung processes. The area of each peak is related to the concentration of the element in the sample. Analysis of PIXE spectra is made easier by various computer packages (e.g. GeoPIXE [Ryan & Cousens 1992] and GUPIX [Maxwell *et al* 1989]) with analysis times of the order of seconds. PIXE is a highly sensitive technique with maximum sensitivity (~ppm) in two regions of the periodic table ($20 < Z < 35$, $75 < Z < 85$) [Johansson & Campbell 1988]. PIXE, using Si(Li) detectors allows multi-elemental analysis to depths of the order of a few tens of microns. Small samples can also be analysed by using microprobes, and this technique is then referred to as micro-PIXE.

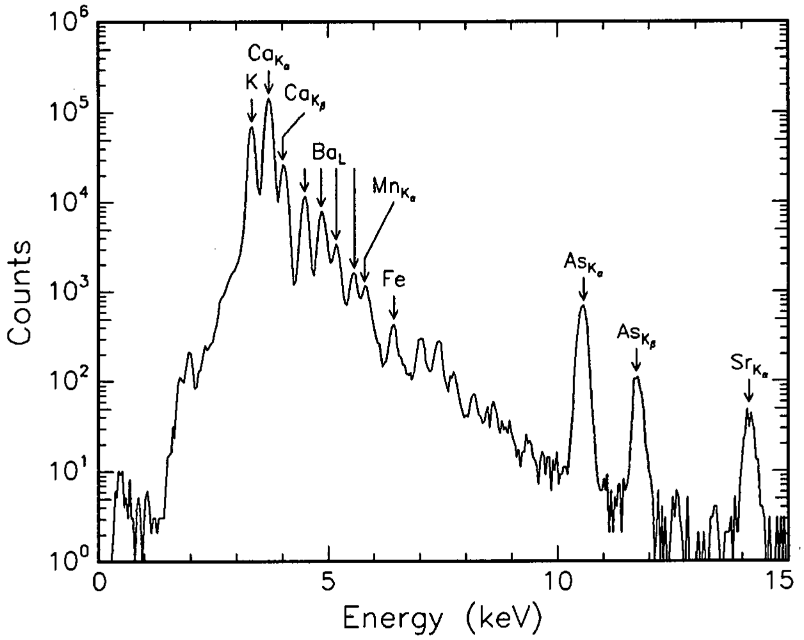

Figure 4.9. A typical PIXE spectrum of a synthetic mineral/glass standard shows the multi-elemental nature of the technique.

4.4.2. The Essentials of PIXE

4.4.2.1. Ionisation Cross-Section

The ionisation cross-section gives an indication of the likelihood that an incident particle will ionise an atom in the target. There are three basic theoretical descriptions of the inner-shell ionisation process. The Binary Encounter (or Impulse) Approximation (BEA) considers the ionisation process to be a binary encounter between two charged particles while the Plane Wave Born Approximation (PWBA) applies a perturbation to a transition mediated by a

Coulomb field. The BEA and PWBA models apply to situations where the projectile energy is much greater than the electron binding energies. In cases where this condition does not hold, a semi-classical approximation (SCA) model has been developed which treats the interaction between the projectile and target nucleus through an impact-parameter formalism [Johansson & Campbell 1988].

The PWBA model was later modified by Brandt & Lipicki [1981] to yield the ECPSSR theory. This model considers the energy loss during a collision (E), deflection and change in velocity of projectile due to Coulomb field of target (C), perturbation of atomic stationary states (PSS) and relativistic effects (R) [Johansson & Campbell 1988]. PIXE software analysis packages like GeoPIXE use tabulated values for the ionisation cross-section based on the ECPSSR theory.

4.4.2.2. X-ray Yield

The x-ray yield depends on the thickness of the sample being analysed. For a thin homogeneous target of mass M, the proton energy will remain essentially constant through the target. The x-ray yield off an element of atomic mass A_Z and atomic number Z in the thin target bombarded with a beam of area A containing N_p protons is given by [Johansson & Campbell 1988]

$$Y_0(Z) = \frac{N_p M}{A} \frac{\sigma_Z(E_0)\omega_Z b_Z \varepsilon_Z N_0}{A_Z} \qquad (4\text{-}13)$$

where $\sigma_Z(E_0)$ is the ionisation cross-section evaluated at the incident energy, b_Z is the branching ratio which is the fraction of x-rays belonging to the ionised shell that are in that particular x-ray line, ε_Z is the detector efficiency, ω_Z is the fluorescence yield, which is the probability of an x-ray being emitted from a particular shell after ionisation, and N_0 is Avogadro's number.

The situation is more complicated for thick targets, which completely stop the beam. These are considered to be a series of thin targets with the total x-ray yield calculated by integrating over the proton energy from the incident energy E_0 to 0. The yield is therefore

$$Y(Z) = \frac{N_p C_Z \omega_Z b_Z \varepsilon_Z N_0}{A_Z} \int_{E_0}^{0} \frac{\sigma_Z(E) T_Z(E)}{S(E)} dE \qquad (4\text{-}14)$$

where C_Z is the concentration of the element of atomic mass A_Z and atomic number Z. $S(E)$ is the stopping power for a target of density ρ given by

$$S(E) = \frac{1}{\rho}\frac{dE}{dx}$$

$T_Z(E)$ describes the transmission of photons through the successive layers of the material on its outward path [Johansson & Campbell 1988].

4.4.2.3. Secondary Fluorescence

Mechanisms other than proton excitation can lead to the generation of x-rays: fluorescence by produced x-rays, ionisation by secondary electrons, fluorescence by bremsstrahlung, and ionisation by proton-induced Auger electrons. Johansson & Campbell [1988] give details of the enhancement of the x-ray yield due to secondary fluorescence.

4.4.2.4. Depth Profiling in PIXE

In NRA, resonances in the reaction cross-section are used to obtain depth profiles. Changing the energy of the incident beam shifts the depth at which the resonance occurs and the reaction products at each depth are detected. The ionisation cross-section in PIXE is, unfortunately, smooth with respect to beam energy, and regular PIXE cannot be used to obtain depth information. It is, however, possible to obtain depth information by changing the angle of incidence or exit of x-rays [Smith *et al* 1995, Livesey & Smith 1994]. A second possibility is to use two detectors at different detection angles (*Figure 4.10*) or to perform two measurements by tilting the target and using a single detector. The advantage of using two detectors in a single measurement is that the beam charge is autonormalised. Depth information can then be extracted by using the information contained in the different spectra.

Figure 4.10 is a schematic diagram of a two-detector set-up that can be used for PIXE depth profiling. The path lengths of the x-rays in the target is different for each detector and one can therefore extract depth information by using both spectra.

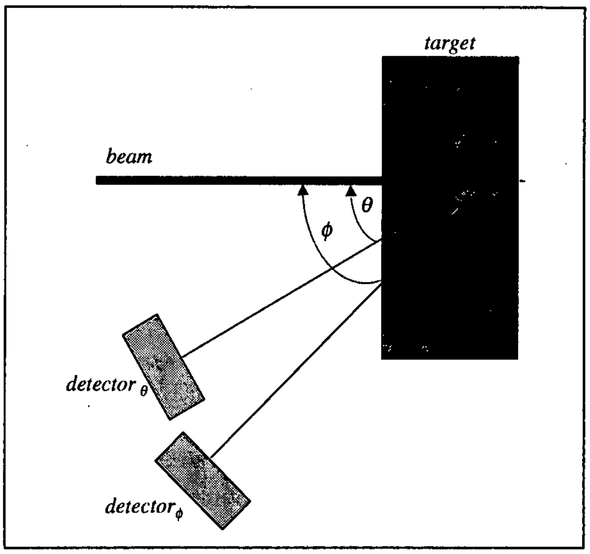

Figure 4.10. Schematic diagram of the two-detector set-up for PIXE depth profiling. The path lengths of the x-rays in the target are different for each detector.

4.4.3. The PIXE Background

Projectile bremsstrahlung and secondary electron bremsstrahlung are the significant contributors to the PIXE background.

Projectile Bremsstrahlung

The deceleration experienced by the proton in Coulombic interactions with bound electrons results in the emission of bremsstrahlung. This is the main contributor to the PIXE background at high energies but it is still significantly lower than that for a beam of electrons as in an electron microprobe (EMP).

Secondary Electron Bremsstrahlung

The incident protons not only cause an ionisation of the target atoms but also cause electrons to be ejected which in turn produce bremsstrahlung. The maximum energy $E_{\max}$ of secondary electron bremsstrahlung is that transferred to a free stationary electron in a head-on elastic collision

$$E_{\max} \approx 4\frac{m}{M}E_0 \tag{4-15}$$

where E_0 is the energy of the incident protons. The background due to secondary electron bremsstrahlung is therefore very intense below E_{max} and decreases rapidly before reaching E_{max}. For 3 MeV protons, E_{max} ~ 6 keV and the background due to secondary electrons is quite significant in low energy regions of the spectrum (*Figure 4.9*).

4.4.4. Lithium-Drifted Silicon Detectors

The active volumes of diffused junction and SSB detectors are too small for more penetrating radiation, like x-rays and γ-rays, to deposit all their energy in the depleted region. Lithium-drifted detectors were developed because the active volumes of junction detectors could not be made larger. Si(Li) detectors are used for the detection of x-rays while lithium drifted germanium (Ge(Li)) and Intrinsic Germanium (IG) detectors are used for high energy x-rays and γ-rays.

4.4.4.1. Fabrication of Si(Li) Detectors

Si(Li) detectors are fabricated though a process known as ion drifting. Ion drifting consists of creating a region of compensated material in which the number of donor impurities is exactly balanced by the number of acceptor impurities. This region then has properties which resemble that of an intrinsic semiconductor. The initial crystal is usually *p*-type and donor atoms (usually Li) are added to create a balance between donor and acceptor atoms. An excess of Li is drifted through one surface of the *p*-crystal creating an *n*-type layer near the surface. The junction is reverse biased and the temperature raised to increase the mobility of the Li ions in the direction of the electric field into the *p*-type region (*Figure 4.11*).

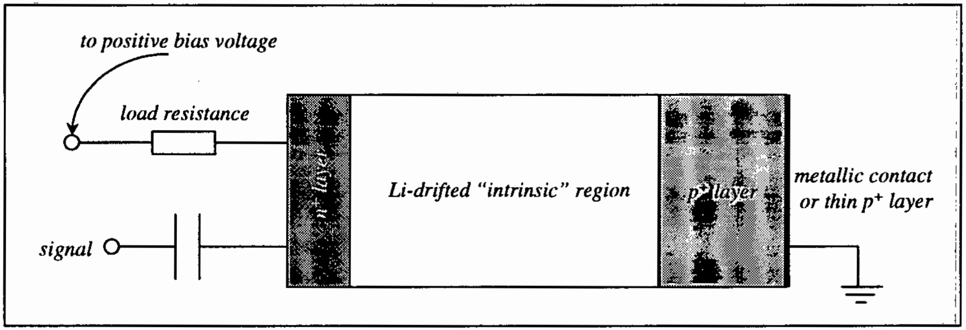

Figure 4.11. Basic configuration of a Li drifted *p-i-n* junction detector [Knoll 1979].

The electric field in the intrinsic region is uniform, because the net charge is ideally zero, but the field drops sharply at the boundaries. The dimensions of the intrinsic region therefore determine the dimensions of the active volume. The signal arises from charge migration to the *p-i* and *i-n* boundaries [Knoll 1979].

4.4.4.2. Description of Si(Li) Response Function

In the absence of detector imperfections and spectral artefacts the response of a Si(Li) detector to a beam of monoenergetic particles would be a Gaussian peak. The FWHM of the peak would be determined by electronic noise and statistical fluctuations in the Si(Li) crystal, as described by equation (4-10). Real peaks, however, have low energy tailing due to surface dead layers and other imperfections that affect the charge collection process (e.g. loss of charge to the weak field regions at the edges of the detector) [Johansson & Campbell 1988]. These effects are also coupled with the intrinsic line-widths of the x-ray lines, which are Lorentzian and is based on the life-time of the excited state.

The ideal method of determining the response function of a Si(Li) detector is to record spectra of monochromatic x-rays (using a crystal monochromator), seek an empirical function that fits the peaks and determine the energy dependence of the parameters. Johansson & Campbell [1988] use two exponentials to describe the low-energy tailing but only one was used in this investigation. The following function describes the response function a Si(Li) detector:

$$r_{Si(Li)} = \frac{1}{\sqrt{2\pi}\Delta E}\exp\left(-\tfrac{1}{2}\left(\frac{y-x}{\Delta E}\right)^2\right) + \tfrac{1}{2}A\exp\left(\frac{y-x}{s}\right)\operatorname{erf}\left(\frac{y-x}{\sqrt{2}\Delta E}+\frac{\Delta E}{\sqrt{2}s}\right) \qquad (4\text{-}16)$$

A represents the amplitude of the exponential tail, s represents the slope of the exponential tail and ΔE represents the FWHM of the main Gaussian as described by (4-10). The parameters A, s and ΔE of the response function are energy-dependent and vary smoothly with energy [Campbell *et al* 1985].

Chapter 5

Results

The theory as described in Chapter 3 was coded into FORTRAN-77 and was applied to the following inverse problems in IBA:

- deconvolution of the detector response function from RBS and PIXE spectra
- extraction of RBS and PIXE depth profiles from spectra
- deconvolution of the beam profile from one-dimensional NMP scans

The deconvolution of the detector response function from spectra, using multiresolution pixons, is also compared to reconstructions obtained using Fourier Transforms, Jansson's Method and the MaxEnt formalism with a pre-blurring function of constant width.

5.1. Response Functions

An accurate description of the detector response function is required for the deconvolution of spectra, for reasons given in §2.1. The empirical formulae for SSB detectors and Si(Li) detectors are described in (4-11) and (4-16) respectively. The determination of the parameters for the detectors used in this investigation is described below.

Channel numbers are directly proportional to energy and it is convenient to use channel numbers in the computations. For this reason, all results will be presented as a function of channel numbers instead of energy.

5.1.1. The SSB Detector

The parameters of the SSB detector response function were determined by collecting an RBS spectrum of a thin layer (~10 Å) of iridium on a silicon substrate (*Figure 5.1*) using the identical experimental parameters as those of the spectrum to be deconvoluted. The measurement was made with a 2 MeV ^{4}He beam and the experimental set-up shown in *Figure 4.2.* The Ir layer was assumed to be thin enough so that the Ir peak would be a δ-function in the absence of detector broadening. The broadening of the Ir peak in the measured spectrum would therefore be due to the response of the detector and could be used to obtain the parameters of the SSB response function in (4-11).

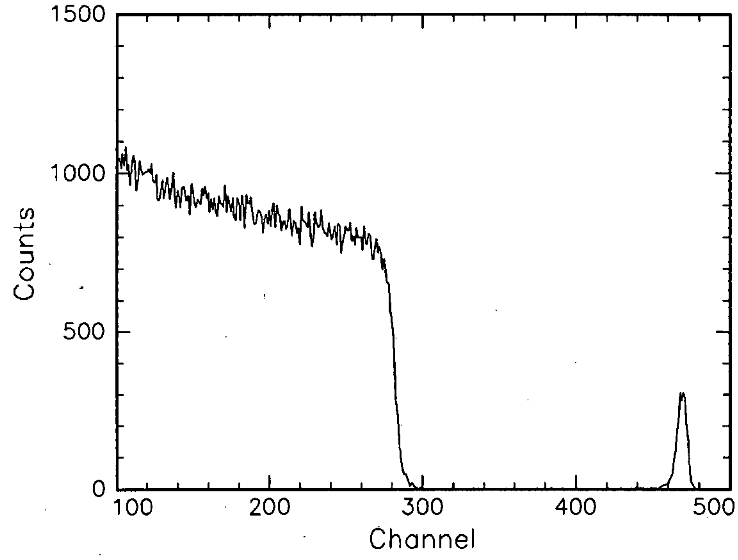

Figure 5.1. Experimental RBS spectrum of iridium on silicon. A thin layer of Ir (~10 Å) was used and the broadening of the Ir peak was assumed to be due entirely to the detector response function. The parameters in the empirical formula of the detector response function for the SSB detector were determined by performing a least-squares fit to the Ir peak.

The values of the coefficients were determined by performing a least-squares fit of (4-11) to the Ir peak. Spectra of Ir on Si were collected for various incident energies to determine the energy dependence of the parameters of the detector response function in (4-11). These values were found to be constant within the measured precision. The values of the parameters used in the deconvolution of RBS spectra are shown in *Table 5-1.*

Table 5-1. Values (in channels) of the parameters of the SSB response function as described in (4-11).

Width of Gaussian ΔE	1.9500
Amplitude of first exponential A_1	0.0309
Slope of first exponential s_1	4.7000
Amplitude of second exponential A_2	0.2090
Slope of second exponential s_2	2.0000

5.1.2. The Si(Li) Detector

The response of a Si(Li) detector is dependent on the incident x-ray energy [Campbell *et al* 1985]. It was therefore necessary to determine the energy dependence of the parameters in (4-16). This was done by doing a least-squares fit of (4-16) to single element peaks covering the energy range of interest. Details of the elements used are given below in *Table 5-2.*

Table 5-2. Elements and energies of peaks used to determine the energy dependence of the parameters in the response function of the Si(Li) detector in (4-16).

Element	K_{α_1} (keV)	K_{α_2} (keV)	L_{α_1} (keV)	L_{α_2} (keV)
V	4.951	4.944	-	-
Ni	7.477	7.460	-	-
Pt	-	-	9.441	9.360
Au	-	-	9.712	9.626
Nb	16.612	16.518	-	-
Mo	17.476	17.371	-	-
Rh	20.213	20.070	-	-
Pd	21.174	21.017	-	-
Sn	25.267	25.040	-	-

Three parameters (width of Gaussian ΔE, amplitude A and slope s of exponential) were determined for each peak. The results of the fits to determine the energy dependence of the parameters are shown in *Figure 5.2.* This description of the response function was then used in the deconvolution of PIXE spectra.

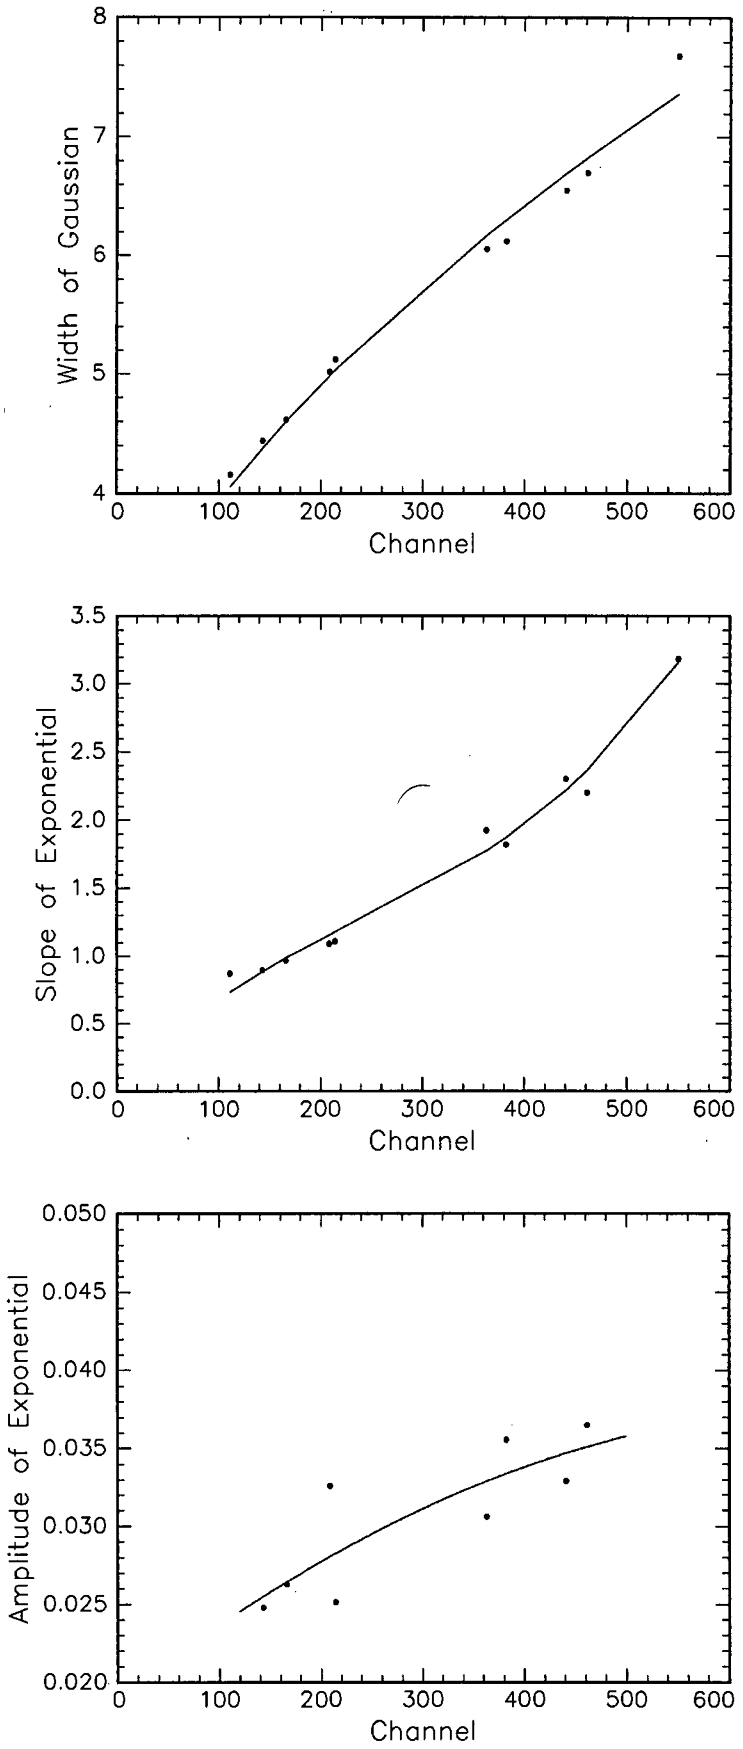

Figure 5.2. Fits to the parameters of the response function of the Si(Li) detector to determine their energy dependence.

5.2. The Spectra

5.2.1. RBS Spectra

A thick sample of GaAs was chosen to demonstrate the deconvolution of RBS spectra. Ga has two stable isotopes of masses 69 (60.1%) and 71 (39.9%) while As has a single stable isotope of mass 75. An ideal spectrum of GaAs (simulated in RUMP) with no detector broadening or noise can be seen in *Figure 5.3.*

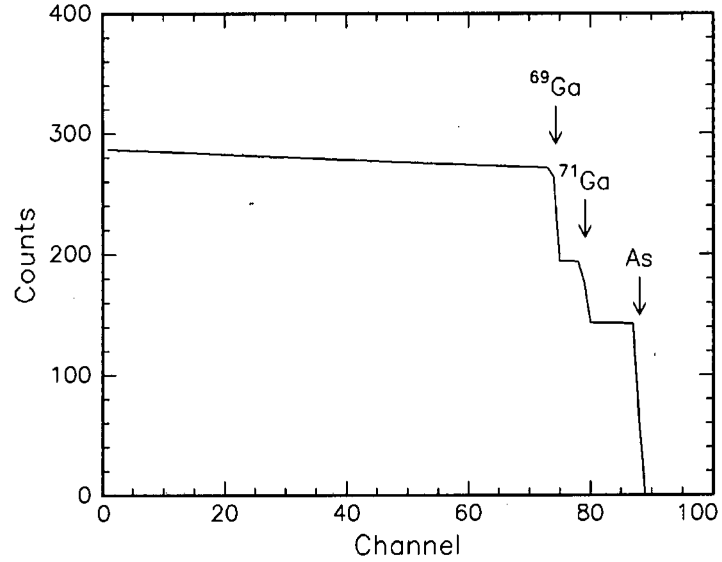

Figure 5.3. Theoretical spectrum of GaAs. The spectrum was simulated in RUMP, with no detector broadening or noise.

A measurement of GaAs was made using a 2 MeV ^{4}He beam and the experimental set-up shown in *Figure 4.2.* The annular SSB detector has an energy resolution of 25 keV at 5.486 MeV. For a 2 MeV incident beam, the energies of the backscattered particles off As and a weighted average of Ga are 1.615 MeV and 1.590 MeV respectively ($\Delta E_{\text{As-Ga}} = 25$ keV). The resolution of the detector is insufficient to resolve these elements as seen in the experimental GaAs spectrum in *Figure 5.4.*

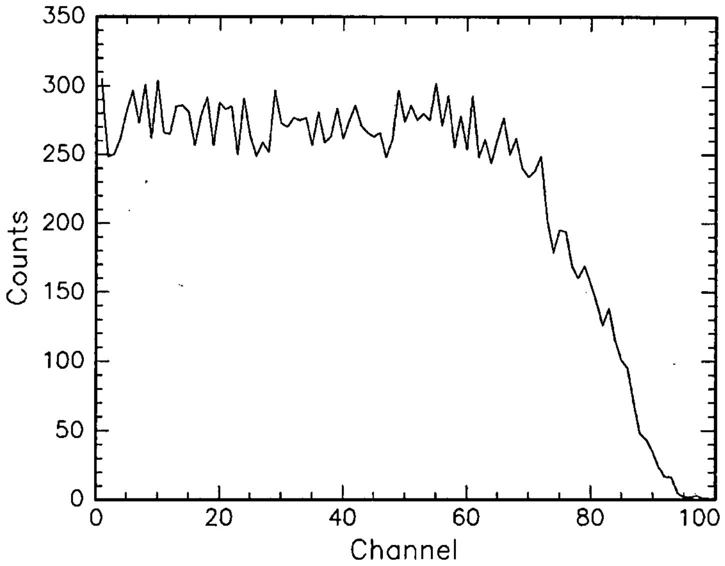

Figure 5.4. Experimental spectrum of GaAs. The measurement was made using a 2 MeV ^{4}He beam with the backscattered particles measured by an annular SSB detector. The SSB detector has an energy resolution of 25 keV at 5.486 MeV, which is insufficient to resolve the Ga and As edges.

5.2.2. PIXE Spectra

Data was simulated to demonstrate the deconvolution of PIXE spectra. The simulated spectrum contained monoenergetic peaks, which were convoluted with the detector response function of the Si (Li) detector described in (4-16), and Gaussian noise was added to simulate counting statistics. Two cases were considered: a spectrum without a background and a spectrum with a background that is quadratic on a logarithmic scale. The form of the background chosen is similar to that described in §4.4.3. The ideal spectra for these two cases can be seen in *Figure 5.5* and *Figure 5.6*, respectively.

PIXE spectra differ from RBS spectra in two main aspects: PIXE spectra have a significant background, and the intensities of the peaks in a PIXE spectrum usually vary over a few orders of magnitude. RBS spectra do not have a background and intensities usually have approximately the same order of magnitude.

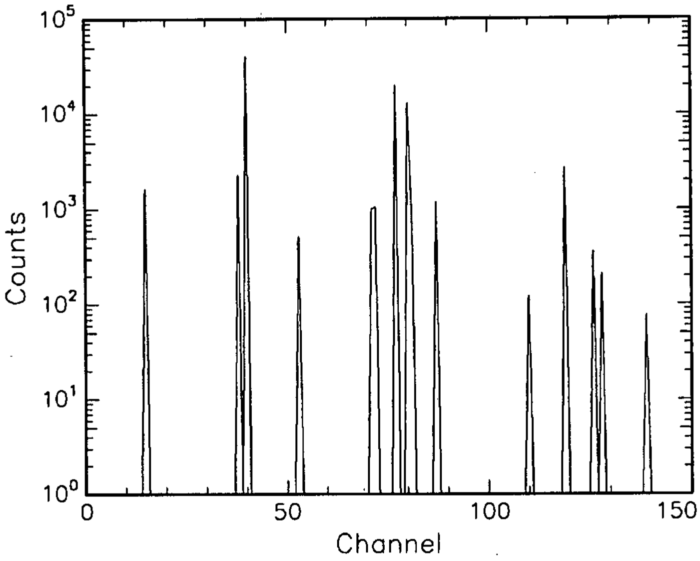

Figure 5.5. Ideal PIXE spectrum without a background.

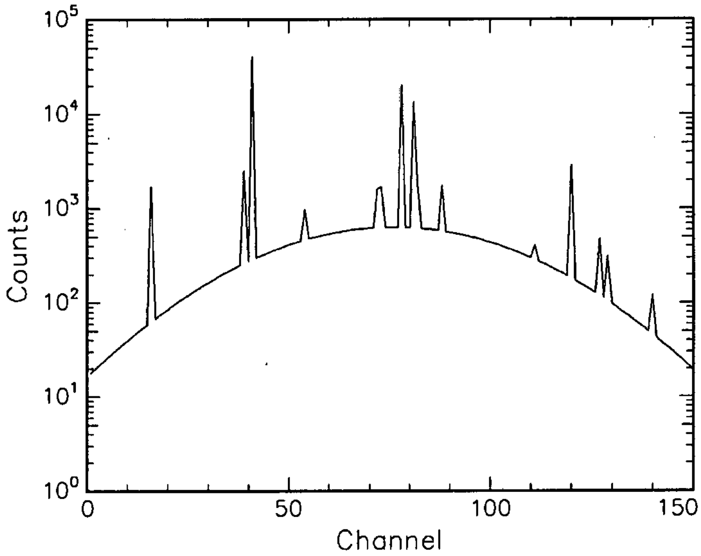

Figure 5.6. Ideal PIXE spectrum with background simulated by adding a quadratic (in the log) background to the spectrum of *Figure 5.5.*

The convoluted spectra with noise added can be seen in *Figure 5.7* and *Figure 5.8.* The peaks have been smeared out in both spectra and some of the smaller peaks have been swamped by the noise. The aim of deconvoluting the PIXE spectra shown in *Figure 5.7* and *Figure 5.8* is to recover the ideal spectra shown in *Figure 5.5* and *Figure 5.6,* respectively.

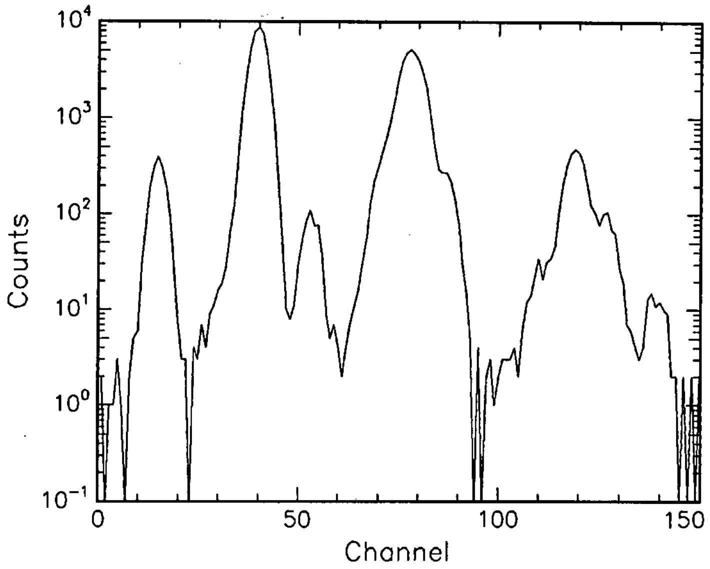

Figure 5.7. Simulated PIXE spectrum without a background. The ideal spectrum of ***Figure*** *5.5* was convoluted with the detector response function of the Si(Li) detector described in (4-16) and Gaussian noise was added to simulate counting statistics. Most peaks have been smeared out and cannot be distinguished as individual entities.

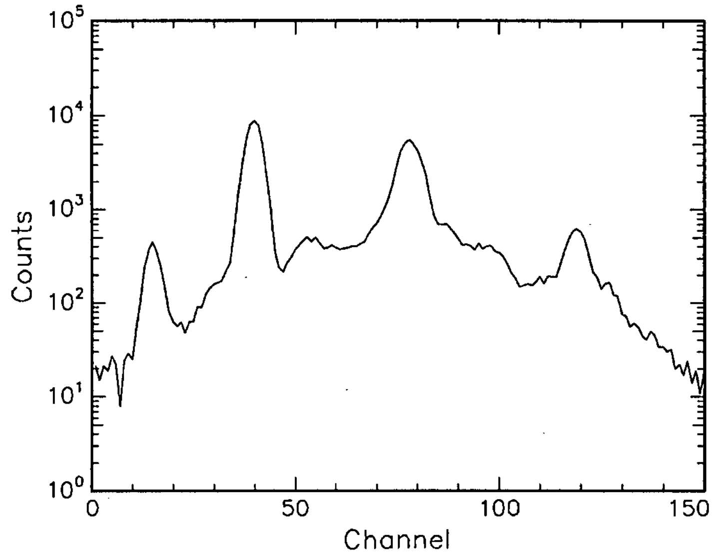

Figure 5.8. Simulated PIXE spectrum with a background. The ideal spectrum in ***Figure*** *5.6* was convoluted with the detector response function of the Si(Li) detector described in (4-16) and Gaussian noise was added to simulate counting statistics. The peaks have been smeared out and some have been swamped by the noise.

5.3. Deconvolution Using Fourier Transforms

Deconvolution of RBS and PIXE spectra using Fourier Transforms was done using the FFT routine in IDL [1993]. The advantage of processing data in IDL is that the results can be displayed immediately and changes can be made accordingly. This was useful with the

Fourier Transforms because the Fourier Transform of a function could be displayed to determine the parameters of the filter, to eliminate noise in the reconstruction.

The spectra were deconvoluted by first computing the FFTs of the data, $D(\omega)$ and the response function $R(\omega)$ and then doing a point-by-point division of the former by the latter. The FFT of the result $F(\omega)$ was then displayed to determine which components had to be filtered. The filter consisted of a function which was zero over the range which was to be eliminated and one elsewhere. After filtering, the inverse FFT was computed to obtain the reconstruction f. Computation time was of the order of a few seconds.

5.3.1. Deconvolution of RBS Spectra

The results of deconvoluting the GaAs spectrum using Fourier Transforms can be seen in *Figure 5.9* and the resolution has been improved in comparison to the original GaAs spectrum in *Figure 5.4*.

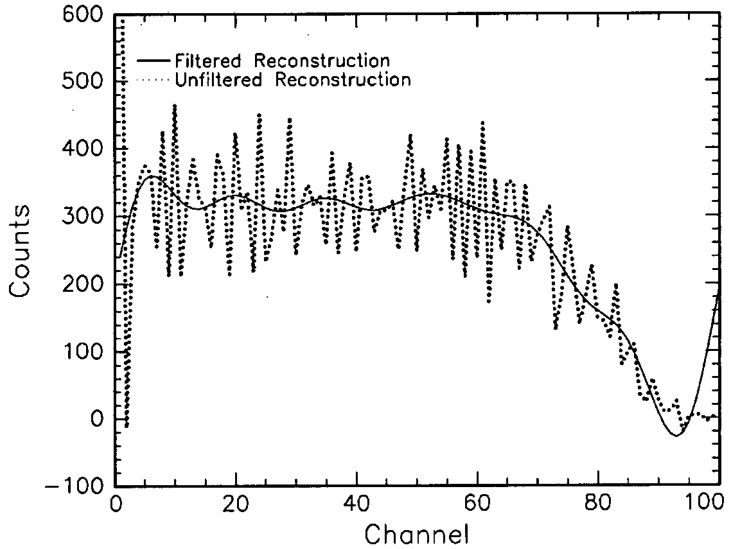

Figure 5.9. Deconvolution of a GaAs spectrum using Fourier Transforms. The unfiltered result (dotted line) shows an amplification of noise, which can, to a large extent, be eliminated by use of a filter. This was done in the filtered solution (solid line) and there is some evidence of a step around channel 80. Both solutions show non-physical negative values.

The solid line in *Figure 5.9* depicts the Fourier Transform reconstruction where the high frequency components have been filtered out. The reconstruction is smoother than the unfiltered result and shows some evidence of a step around channel 80. There is, however, a 'spike' around channel 100, for which there is no evidence in the data. One feature common to both filtered and unfiltered results is the presence of negative components which is definitely non-physical.

Figure 5.10 shows that the noise component of the unfiltered reconstruction is even greater than that of the original data and there is no improvement in resolution. The amplification of the noise is expected since the Fourier Transform of the response function $R(\omega)$ is usually small for high ω [Jansson 1984]. The result of the division of $D(\omega)$ by $R(\omega)$ is therefore large at the high (noise) frequencies and the reconstruction exhibits an amplification of the noise.

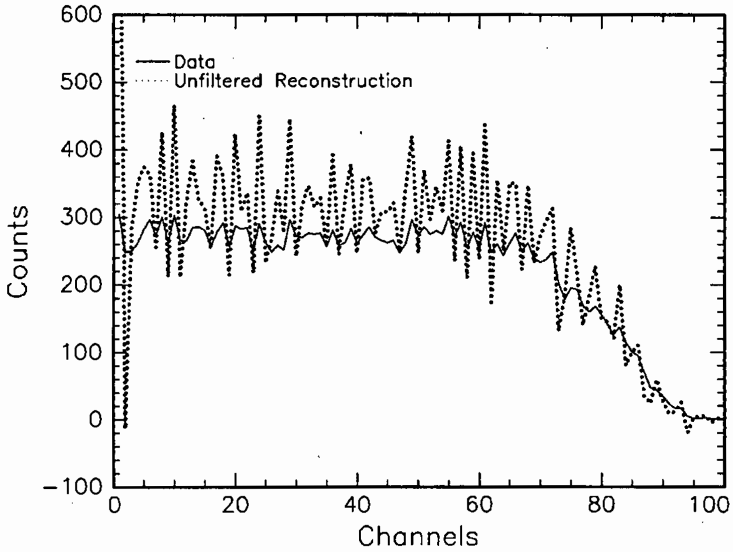

Figure 5.10. Comparison of the unfiltered Fourier Transform reconstruction and original data to show the amplified noise component in the unfiltered reconstruction. The number of counts is not preserved in the reconstruction.

In order to compare reconstructions, a measure of the "goodness of fit" χ^2_{gof} was calculated according to (5-1):

$$\chi^2_{gof} = \frac{1}{N_{data}} \sum_{i=1}^{N_{data}} \left(\frac{d_i - r_i * f_i}{\sigma_i} \right)^2 \tag{5-1}$$

where

$$\sigma_i^2 = d_i \tag{3-52}$$

as for counting statistics.

A χ^2_{gof} of 1 implies that, on average, the reconstruction agrees with the data to within one standard deviation of the noise. A value greater than 1 implies that there is under-fitting of the data, i.e. the reconstruction exhibits features for which there is no evidence in the data, like spurious peaks. Under-fitted reconstructions are, on average, not "close" to the data. Similarly, a χ^2_{gof} of less than 1 implies that there is over-fitting of the data and, on average, the reconstruction follows the noise in the data. An over-fitted reconstruction is usually as noisy as the data.

A value of 225 was obtained for the χ^2_{gof} of the filtered FFT (*Figure 5.11*). The large value arises from the fact that the reconstruction differs significantly from the data at the ends of the spectrum while the reconstruction agrees quite well with the data from around channel 5 to around channel 90.

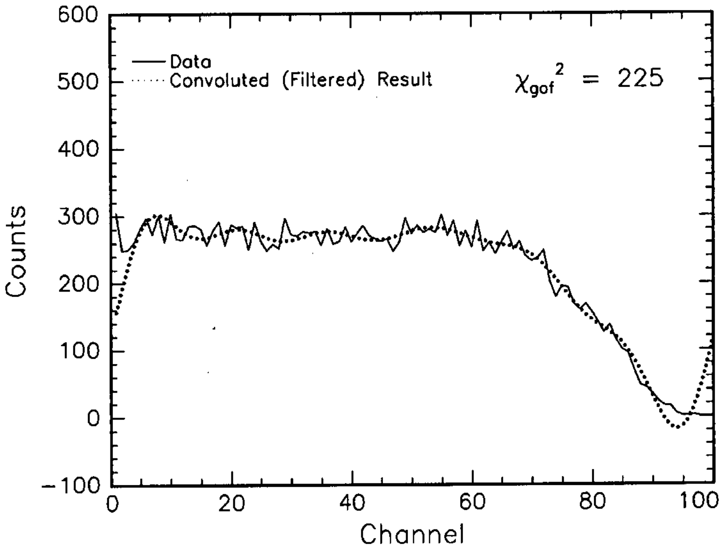

Figure 5.11. Comparison of convoluted result and original GaAs data. The filtered FFT reconstruction in *Figure 5.9* was convoluted with the response function of the SSB detector, and a measure of the "goodness of fit" was calculated. The value of 225 for the χ^2_{gof} confirms the existence of spurious sources in the reconstruction.

5.3.2. Deconvolution of PIXE Spectra

The result of using Fourier Transforms to deconvolute the PIXE spectra without a background can be seen in *Figure 5.12*. The comparison of the data with the convoluted result of *Figure 5.12* can be seen in *Figure 5.13*. *Figure 5.14* shows the Fourier Transform reconstruction of the PIXE spectrum with a background while the comparison of the data with the convoluted result appears in *Figure 5.15.* The unfiltered spectra in *Figure 5.12* and *Figure 5.14* exhibit the same amplification of noise as seen in the unfiltered reconstruction of the GaAs spectrum. Unlike the RBS spectrum, however, there is almost no improvement in resolution and there are many spurious peaks present in the filtered result. The small improvement in resolution in the filtered reconstruction may be because the data contained many high frequency components which were eliminated by the filter. For purposes of plotting, the negative components of the reconstruction were set to 10^{-8} in the program. There are many negative components in the reconstruction.

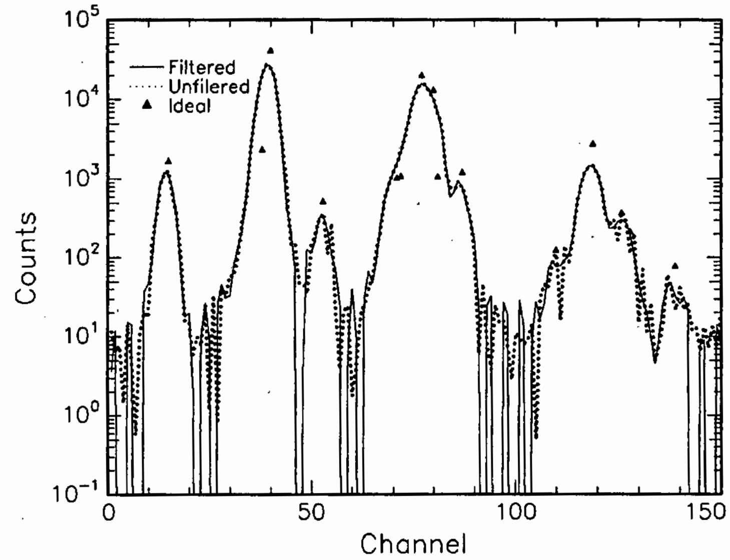

Figure 5.12. Fourier Transform reconstruction of PIXE spectrum without a background. As with the RBS reconstruction, the unfiltered result shows a high noise component. The filtered result shows no improvement in resolution and is still noisy. For purposes of plotting, the negative components were set to 10^{-8} in the program, and the filtered result shows many negative components. The ▲ indicates the position and amplitude of the original peak.

The χ^2_{gof} for the PIXE spectrum without a background was calculated to be 6.9 (*Figure 5.13*). This implies that, on average, there is some under-fitting of the data and that the reconstruction differs slightly from the data in some parts of the spectrum. The FFT reconstruction of the PIXE spectrum without a background agrees quite well to the data in regions of peaks with high counts while the reconstruction is not so good for lower counts.

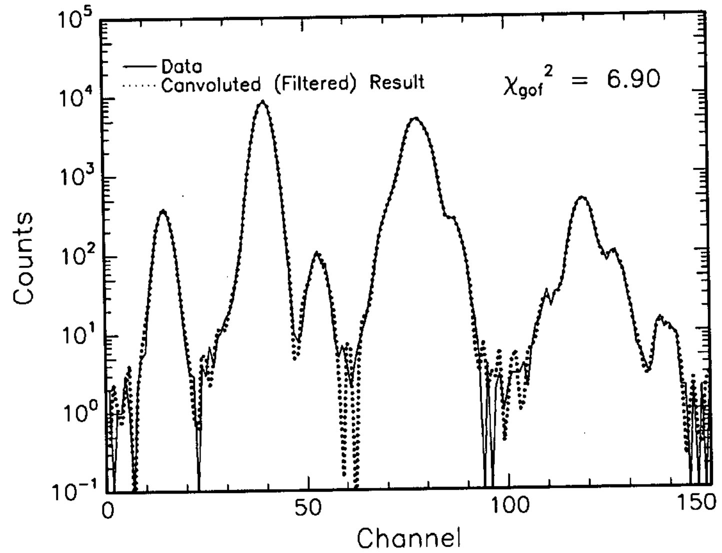

Figure 5.13. Comparison between convoluted result and original data of the PIXE spectrum without a background. The value of χ_{gof}^2 implies that there is some under-fitting of the data and the final reconstruction does not follow the data in some parts of the spectrum.

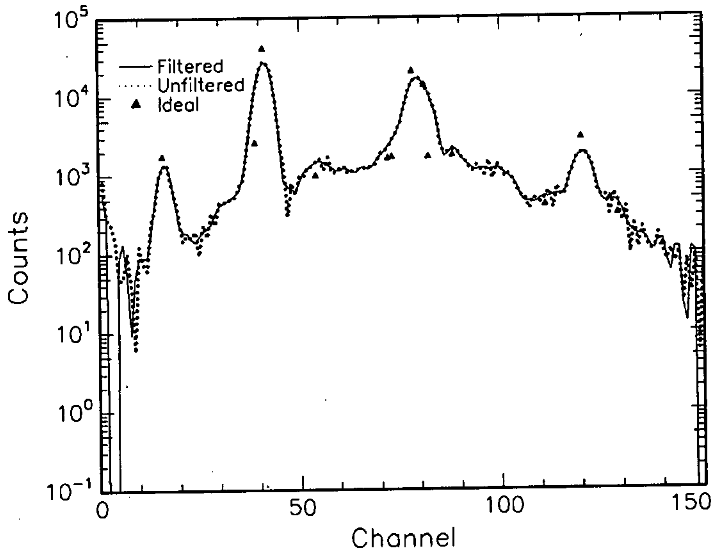

Figure 5.14. Fourier Transform reconstruction of PIXE spectrum with background. There is no noticeable improvement in the resolution of the spectrum when compared to *Figure 5.8*. The ▲ indicates the position and amplitude of the original peak.

The χ_{gof}^2 of 2.5 for the PIXE spectrum with a background (*Figure 5.15*) indicates that there is some under-fitting occurring but this occurs to a slightly lesser degree than in the spectrum without a background. As with the reconstruction of the spectrum without a background, the FFT reconstruction in *Figure 5.14* seems to agree quite well with the data in regions of high counts but not so well in regions of low counts. This is probably because the effect of the noise is more pronounced with a few counts than with a high number of counts since the variance of the noise in a particular channel is equal to the number of counts in that channel as in (3-52). This supports the argument that linear methods have no consistent method for eliminating noise in a spectrum.

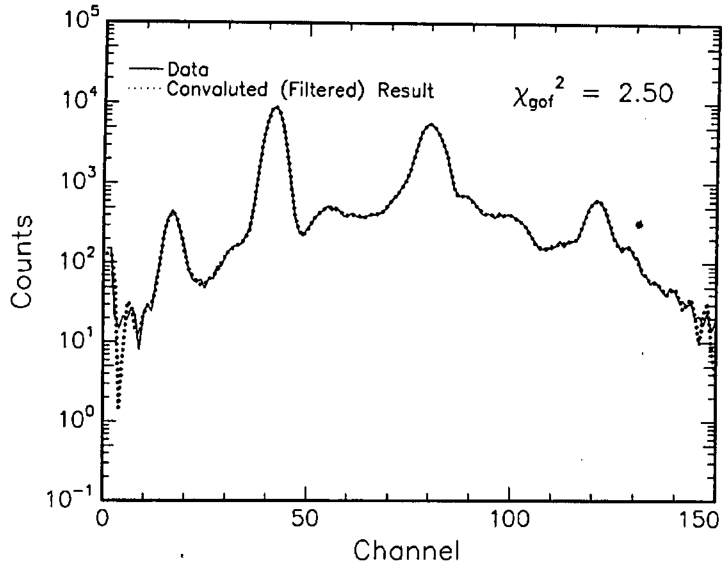

Figure 5.15. Comparison of the convoluted Fourier reconstruction and original data of the PIXE spectrum with a background. The value of χ_{gof}^2 implies that there is some under-fitting of the data but on average, the reconstruction is better than that of the PIXE spectrum without a background.

5.3.3. Summary

Fourier Transforms can be used to improve the resolution of RBS spectra, but there is no noticeable improvement in the resolution of PIXE spectra. The direct reconstruction, i.e. without applying a filter to eliminate the noise, has a noise component that is even greater than that of the original data. A filter must be applied to reduce the noise but this poses many problems. Firstly, the data can also contain high frequency components which are lost when a filter is applied, as there is no way of selectively filtering out the noise. Secondly, there is no formal or consistent method for choosing the parameters of the filter. The presence of negative components in the reconstruction is another undesirable feature of this method. Based on the values of χ_{gof}^2, the method seems to under-fit the data with the reconstruction agreeing with the data in certain regions of the spectrum while differing in others. The only advantage of using Fourier Transforms is that it is a single step process and therefore contains none of the problems associated with an iterative method.

5.4. Deconvolution Using Jansson's Method

Jansson's method, as described in §2.3.2 was coded into FORTRAN-77, using an algorithm described by Coote and Kwan [1995] and a more general relaxation function [Marchetti & Mignerey 1993] than that described in (2-20). The relaxation function had the form

$$\rho(\hat{f}^{(k)}) = \rho_0 \left[1 - \frac{2}{C_h - C_l} \left| \hat{f}^{(k)} - \frac{C_h + C_l}{2} \right| \right] \tag{5-2}$$

where the relaxation constant ρ_0 controls the amount of correction applied in each iteration and C_h and C_l are the upper and lower limits, respectively, of the reconstructed spectrum.

The reconstructed spectra depended strongly on the value of ρ_0 and the number of iterations, with noise in the reconstruction increasing as the number of iterations was increased. The values used in the final reconstruction were determined by trial and error, and computation times ranged from a few seconds to a few minutes depending on the number of iterations.

5.4.1. Deconvolution of RBS Spectra

Following the suggestion of Marchetti and Mignerey [1993], the experimental GaAs spectrum in *Figure 5.4* was smoothed with a Gaussian before applying Jansson's method. The "effective" width of the response function was then increased according to (2-3c). There is, however, no formal reason for smoothing the data. It could be similar to the concept of a pre-blur i.e. to introduce correlations to the reconstruction but there is no consistent method of determining the width of the smoothing function. The results can be seen in *Figure 5.16.*

Like Fourier Transforms, Jansson's method improves the resolution of the GaAs spectrum, with a clear step around channel 80. However, the method depends strongly on the value of the relaxation constant and the number of iterations. A value of ρ_0 = 0.1 was found to produce good results but this was determined by trying out various values. C_l and C_h were set to 0 and 350 respectively. Increasing the number of iterations improved the resolution of the reconstructed spectrum (*Figure 5.16*) but the level of noise also increased. This increase in noise is probably due to the fact that errors, due to the presence of noise, in the first iteration are propagated through to the last and there is no formal method of determining the number of iterations used in the deconvolution.

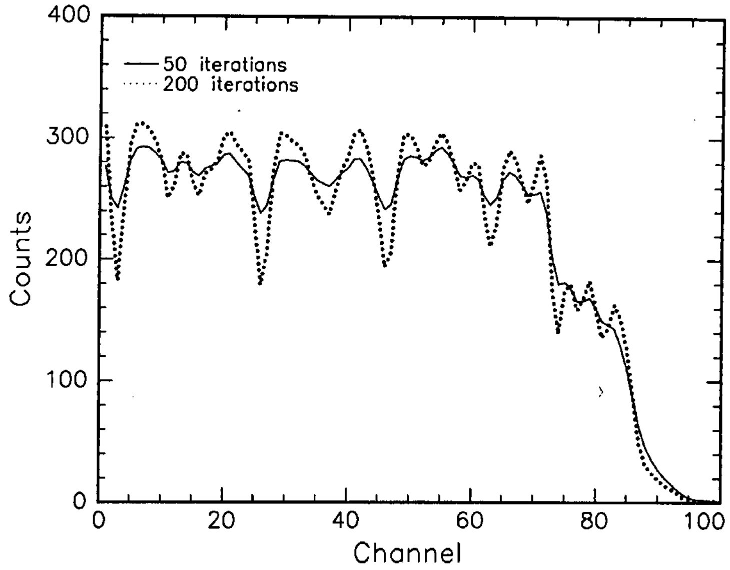

Figure 5.16. Deconvolution of GaAs using Jansson's method. The resolution increases as the number of iterations increase but this also increases the noise in the final reconstruction. This is a problem intrinsic to most iterative procedures since errors in the first iteration are propagated through to the last iteration. There is evidence of a step around channel 80.

The comparison of the convoluted results and the original GaAs data can be seen in *Figure 5.17* (50 iterations) and *Figure 5.18* (200 iterations).

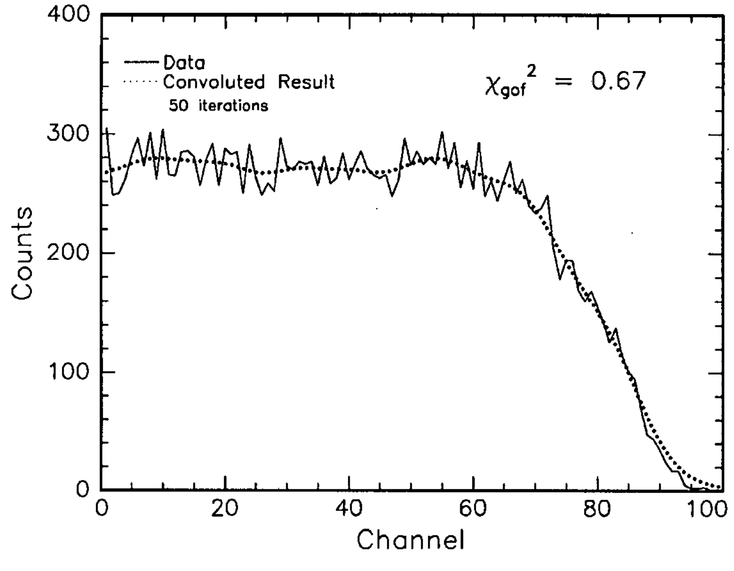

Figure 5.17. Comparison of the convoluted result from Jansson's method with 50 iterations and the original data of the GaAs spectrum. The value of χ^2_{gof} implies that there is some over-fitting of the data as the reconstruction tries to follow the data too closely.

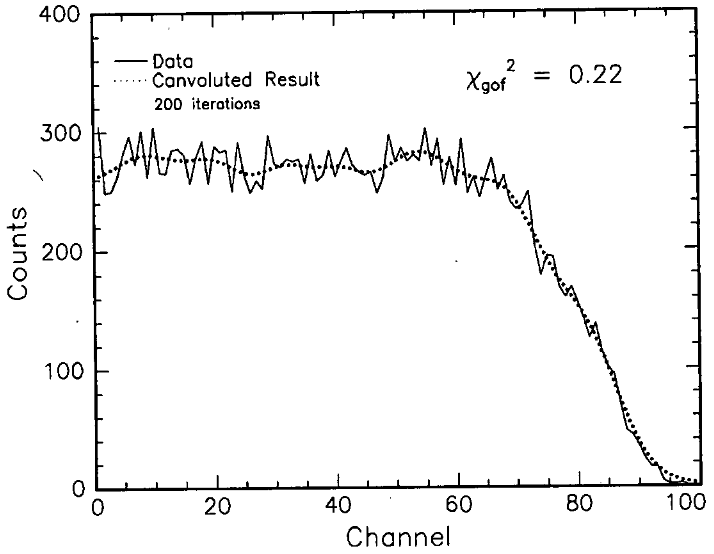

Figure 5.18. Comparison of the convoluted result from Jansson's method with 200 iterations and the original data of the GaAs spectrum. The value of χ^2_{gof} implies that the data is being over-fitted and the degree of over-fitting is greater than for 50 iterations. This confirms the greater amount of noise present in the reconstruction with 200 iterations.

The values of χ^2_{gof} for both reconstructions are less than 1 which is evidence of over-fitting. This can clearly be seen in *Figure 5.16* where there is a significant amount of noise in the reconstructions. The reconstruction of the GaAs spectrum is much better than that obtained using Fourier Transforms since there are no negative components nor are there any spurious peaks at the ends of the spectrum.

5.4.2. Deconvolution of PIXE Spectra

The PIXE spectra were not smoothed before applying Jansson's method since this made no noticeable difference to the final result. The Jansson method reconstruction of the PIXE spectrum without a background can be seen in *Figure 5.19* with the comparison of the convoluted result and the original data in *Figure 5.20.* The reconstruction of the PIXE spectrum with a background is shown in *Figure 5.21* and the comparison between the original data and convoluted result is shown in *Figure 5.22.* The reconstructions in *Figure 5.19* and *Figure 5.21* show an improvement in the resolution in the sense that many of the broad peaks from the original spectrum have been separated but there is also a high noise component which makes it difficult to distinguish between the real peaks and peaks due to noise.

An improvement in the resolution was obtained for relaxation constants of 0.8 (*Figure 5.19*) and 0.2 (*Figure 5.21*), after 10 000 iterations, in each case. Both parameters were determined by trial and error. C_l and C_h were set to 0 and 10^5, respectively. These values were obtained from the ideal PIXE spectra in *Figure 5.5* and *Figure 5.6.*

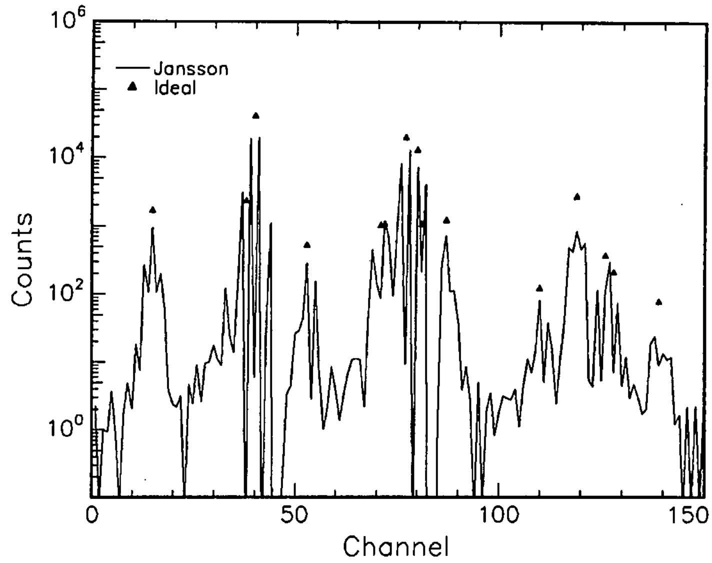

Figure 5.19. Reconstruction of PIXE spectrum without a background using Jansson's method. Although there is an improvement in the resolution, the result is noisy and one has difficulty distinguishing between the true peaks and artefacts of the reconstruction. A relaxation constant of $\rho_0 = 0.8$ was used and the result was obtained after 10 000 iterations.

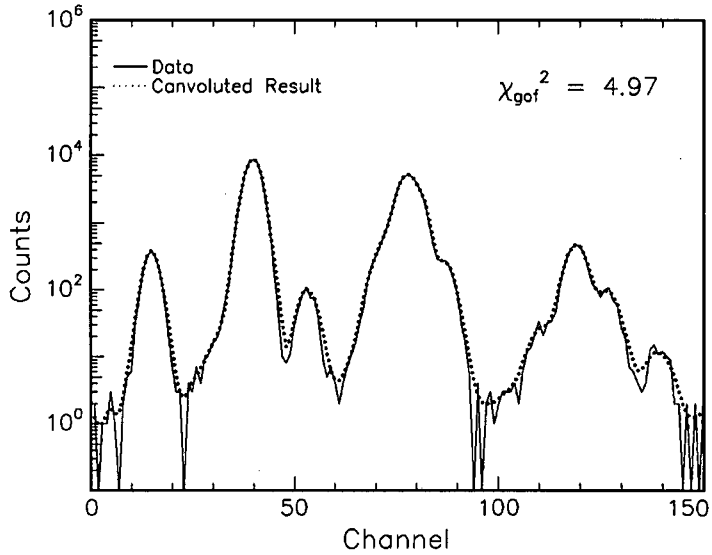

Figure 5.20. Comparison of convoluted reconstruction using Jansson's method and the original data of the PIXE spectrum without a background. The value of χ^2_{gof} implies that there is some under-fitting of the data.

Comparison of the convoluted result and the original data yields a χ^2_{gof} of 4.97 (*Figure 5.20*). This shows that there is some under-fitting of the data as the convoluted reconstruction is not close to the data in regions of the spectrum with low counts as seen in *Figure 5.20.* Like the FFT reconstructions of the PIXE spectra, Jansson's method fits the data well for high counts but performs poorly in regions of the spectrum with low counts.

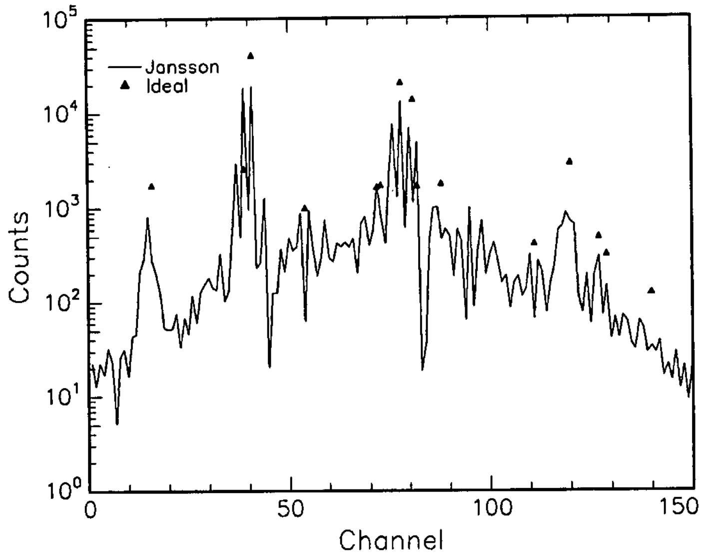

Figure 5.21. Reconstruction of PIXE spectrum with background using Jansson's method. There is an improvement in the resolution in the sense that the broad peaks in *Figure 5.8* have been separated but the result is very noisy. A relaxation constant of $\rho_0 = 0.2$ was used and the result was obtained after 10 000 iterations.

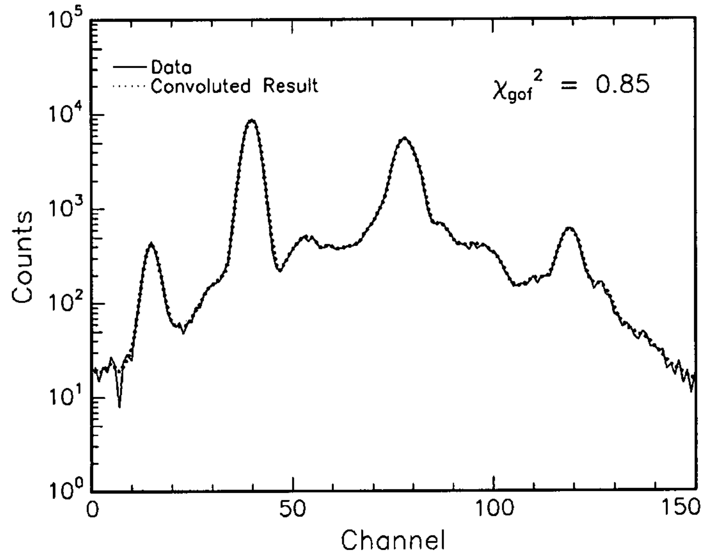

Figure 5.22. Comparison of convoluted Jansson reconstruction and original data of the PIXE spectrum with a background. The value of χ_{gof}^2 implies that there is some over-fitting of the data.

The χ_{gof}^2 of 0.85 shows that the data is being slightly over-fitted and the reconstruction follows the noise in the data. On average, the reconstructions obtained using Jansson's method are better than those obtained using Fourier Transforms. The degree of under-fitting for the PIXE spectrum without a background is less for Jansson's method than for Fourier Transforms while the degree of over-fitting in the Jansson reconstruction of the PIXE spectrum with a background is less than the degree of under-fitting in the Fourier Transform reconstruction.

5.4.3. Summary

Jansson's method improves the resolution of both RBS and PIXE spectra but there are many problems associated with the method. Firstly, it is an iterative procedure which means that errors generated in the first iteration are propagated through to the last. Secondly, there is no formal method of determining the number of iterations used in the deconvolution. Thirdly, Jansson's method requires knowledge of the physical bounds of the problem which means knowledge of the maximum and minimum values of the spectrum. The lower limit is simple to obtain because the solutions must be positive but the upper limit can pose problems. If the upper limit is set too low, features higher than the limit will have flat tops. The values of the upper limits used in the preceding reconstructions were obtained from the ideal spectra. Other parameters required are the relaxation constant and the number of iterations. There exists no formal procedure for determining these parameters. Jansson's method does, however, eliminate the negative components that were present in the Fourier Transform reconstructions.

Of the three spectra fitted, two (GaAs spectrum and PIXE spectrum with a background) were over-fitted while the PIXE spectrum without a background was under-fitted but the reconstructions were better than those obtained using Fourier Transforms.

5.5. Deconvolution Using Maximum Entropy with a Constant Pre-Blurring Function

For this portion of the investigation, the width of the pre-blurring function was set to a constant value. Computation times were of the order of an hour, running on a 120MHz PC.

5.5.1. Deconvolution of RBS Spectra

The reconstruction of the GaAs spectrum in *Figure 5.23* was obtained with a constant pre-blur width of 4 channels. The smoothing effect of the entropy is immediately obvious since the reconstruction exhibits none of the noise characteristic to Fourier Transforms and Jansson's method. The reconstruction shows a clear step around channel 80. There is, however, considerable ringing on the plateau, and although the improvement in resolution is sufficient to distinguish between the Ga and As signals, it is not enough to separate the isotopes of Ga. The ringing is due to the fact that there are too many degrees of freedom in the fit and the reconstruction is over-fitted on the plateau. The pre-blur width of 4 channels is also insufficient to smooth out the ringing. Reconstructions done for a fixed pre-blur width of 1 channel resulted in a very noisy spectrum.

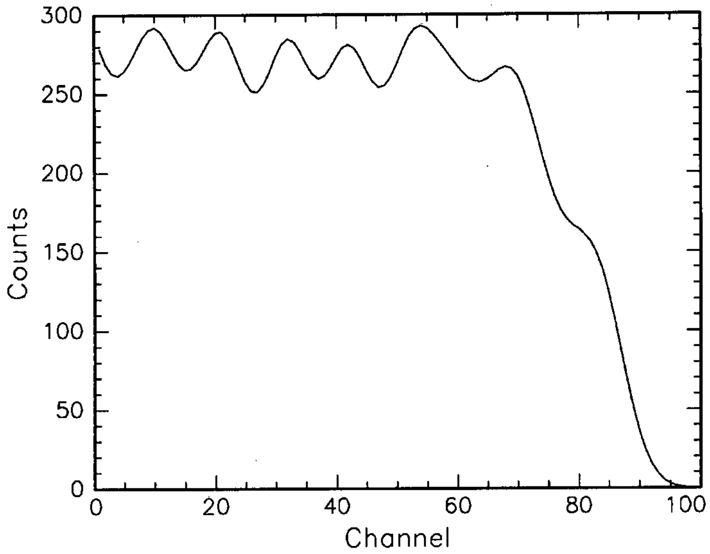

Figure 5.23. Reconstruction of GaAs using MaxEnt with a constant pre-blurring function. The reconstruction is smooth and shows a clear step around channel 80 but there is considerable ringing on the plateau.

The ringing on the plateau of the reconstruction leads to a χ^2_{gof} of 0.67 (*Figure 5.24*) which implies that the reconstruction is over-fitting the data. The value of χ^2_{gof} is equal to that of the Jansson reconstruction but the MaxEnt reconstruction is much smoother than the Jansson reconstruction.

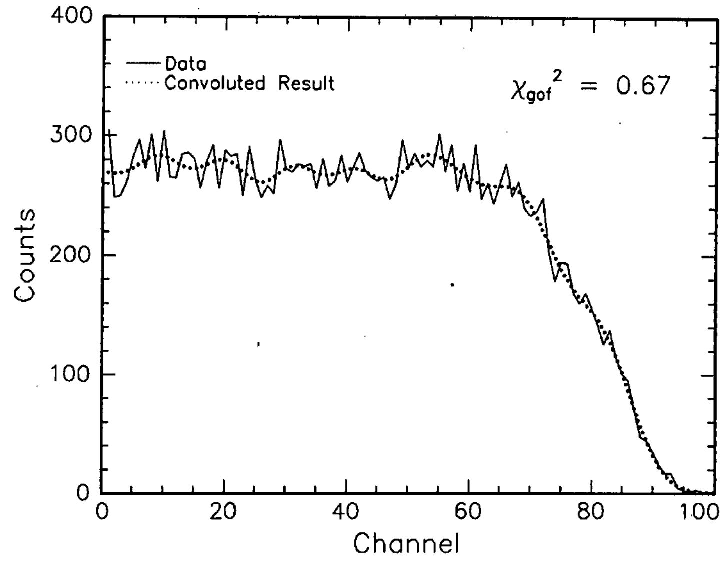

Figure 5.24. Comparison between convoluted result and original GaAs data. The value of χ^2_{gof} implies that the reconstruction over-fits the data. This is expected since there is a considerable amount of ringing on the plateau.

5.5.2. Deconvolution of PIXE Spectra

The reconstruction of the PIXE spectrum without a background, in *Figure 5.25*, was obtained with a constant pre-blur width of 0.2 channels. The reconstruction shows a marked improvement in resolution. All peaks have been extracted and although there are a few spurious noise sources, these peaks all contain around 10 counts per channel. The comparison of the convoluted result with the original data can be seen in *Figure 5.26.*

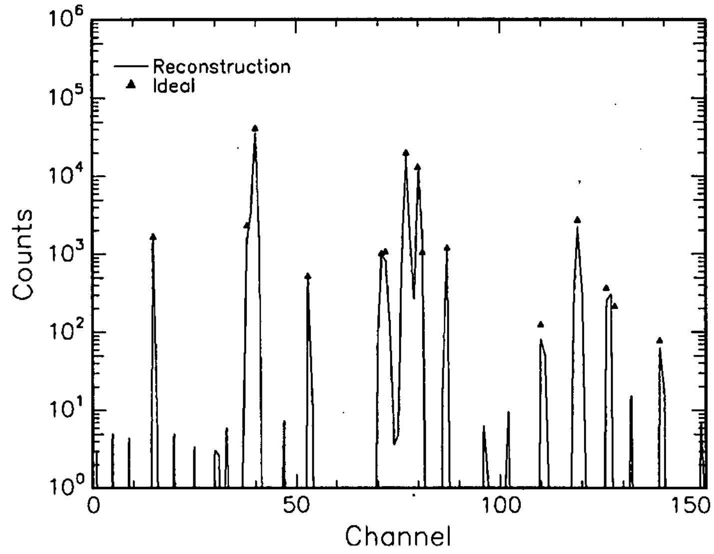

Figure 5.25. Reconstruction of PIXE spectrum without a background using MaxEnt with a constant pre-blur of width 0.2 channels.

The χ^2_{gof} of 1.74 in *Figure 5.26* indicates that, on average, there is a slight under-fitting of the PIXE spectrum without a background.

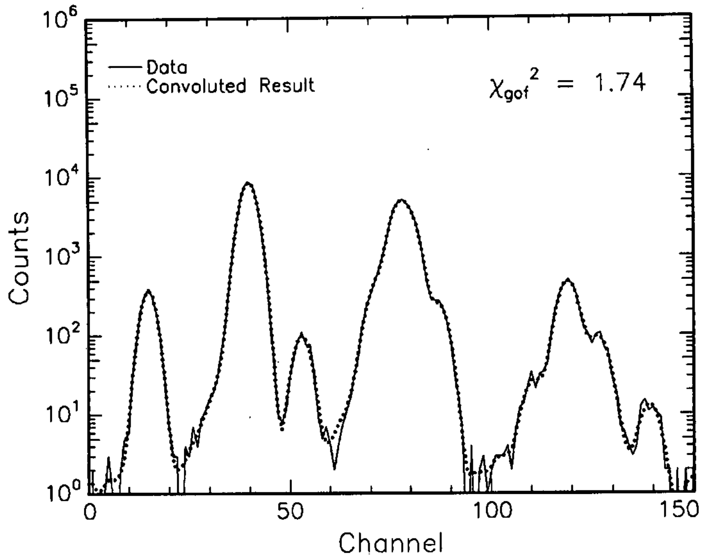

Figure 5.26. Comparison of convoluted result and original data of the PIXE spectrum without a background. The value of χ^2_{gof} implies that there is some under-fitting of the data.

The PIXE spectrum with a background was also reconstructed with a fixed pre-blur width of 0.2 channels (*Figure 5.27*).

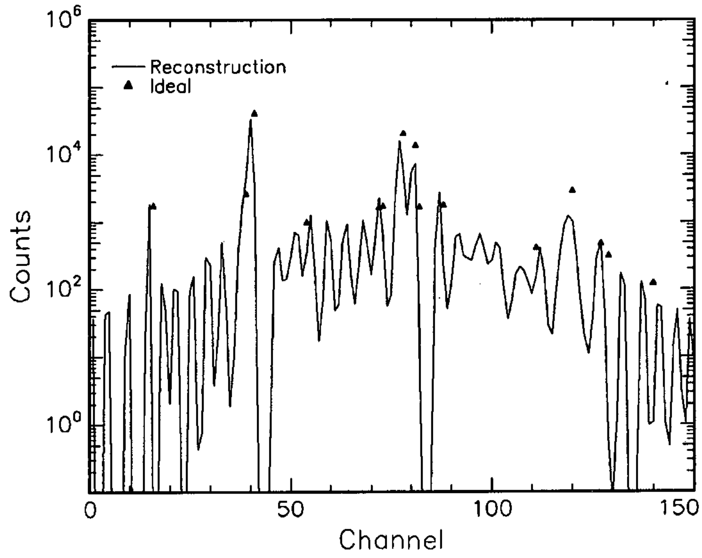

Figure 5.27. Reconstruction of a PIXE spectrum with a background, using MaxEnt with a constant pre-blur of width 0.2 channels.

All peaks have been extracted, but there is considerable ringing on the background portions of the spectrum, because the pre-blur width of 0.2 channels is insufficient to smooth out these spurious peaks. The spectrum was then reconstructed with a constant pre-blur width of 1 channel. The noise on the background portions were smoothed out but most of the peaks were not separated.

The χ^2_{gof} of the PIXE spectrum with a background indicates that the data is being over-fitted. This can be seen in *Figure 5.28* where the convoluted result closely follows the data.

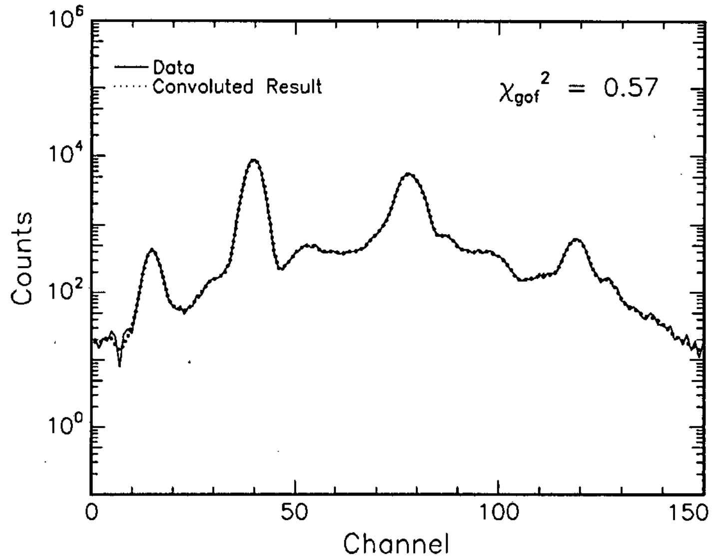

Figure 5.28. Comparison of convoluted result and original data of the PIXE spectrum without a background. The value of χ^2_{gof} implies that there is some over-fitting of the data.

5.5.3. Summary

Maximum entropy with a constant pre-blur exhibits different results for different situations. The reconstructed GaAs spectrum is smooth compared to the reconstructions obtained using Fourier Transforms or Jansson's method, and the step around channel 80 is clearly seen. There is, however, considerable ringing on the plateau which is non-physical – the reconstruction over-fits the data on the plateau. Using a constant pre-blur width means that the correlations are introduced over fixed regions of the spectrum irrespective of the information density. In the GaAs spectrum, the width of 4 channels was sufficient not to smooth out the step but not enough to eliminate the ringing on the plateau. If a smaller width were used, then this would have resulted in a greater amount of ringing on the plateau but the

Ga and As edges would still have been resolved. If a value larger than 4 were used then the ringing on the plateau would have been smoothed out but the step will also be smoothed out.

The reconstruction of the PIXE spectrum without a background is close to the ideal spectrum (*Figure 5.5*) with noise sources containing around 10 counts per channel. The χ^2_{gof} for this reconstruction is 1.74 which shows that the reconstruction is slightly under-fitting the data but all the peaks have been extracted. If a pre-blurring width larger than 0.2 were used, then the noise peaks would be smoothed out, but not all peaks will be resolved.

The reconstruction of the PIXE spectrum with a background shows an improvement in the resolution but the reconstruction has a high level of noise, and therefore a value of χ^2_{gof} which is less than 1, indicating over-fitting. These results can be understood on the basis that the pre-blur width of 0.2 channels in the reconstruction of the PIXE spectrum with a background (*Figure 5.27*) is too small to smooth out the noise peaks while it is sufficient for the reconstruction of the PIXE spectrum without a background (*Figure 5.25*). There are no correlations between peaks in the ideal PIXE spectrum without a background (*Figure 5.5*) and the pre-blur width of 0.2 channels is sufficient to extract the peaks. A reconstruction of the PIXE spectrum with a background, using a pre-blur width of 1 channel was less noisy but most of the peaks were not separated.

5.6. Multiresolution Pixons

A varying pre-blur width was used for this portion of the investigation. Computation times were of the order of a few hours for about 150 data points.

5.6.1. Deconvolution of RBS Spectra

The deconvoluted GaAs spectrum (*Figure 5.29*) shows a clear step between the Ga and As signals. The resolution has been improved from ~25 keV to ~10 keV, an improvement of 2.5 which is sufficient to separate the As from the Ga but insufficient to separate the isotopes of

Ga. For a 2 MeV incident beam, the energies of the backscattered particles off ^{71}Ga and ^{69}Ga are 1.596 MeV and 1.586 MeV respectively ($\Delta E_{^{71}Ga-^{69}Ga}$ = 10 keV).

The vertical position of the step is higher than that for the As yield but lower than that for the ^{71}Ga yield when compared to the ideal spectrum. The improvement in 'resolution' in the counts is sufficient for the algorithm to extract information about the ^{71}Ga yield but because the improvement in detector resolution is not sufficient to resolve the isotopes, the algorithm is 'unsure' as to where the ^{71}Ga yield should be and therefore "averages" the height of the ^{71}Ga and As yields.

As seen in *Figure 5.29* and *Figure 5.30*, the width of the pre-blur varies in such a manner so that it has a minimum in regions of the spectrum that contain the most detail, namely the step in the GaAs spectrum. The ringing on the plateau of the spectrum has been significantly reduced by the use of pixons, but there is a slight oscillation present. Although this oscillation is consistent with the experimental data (*Figure 5.4*) within the bounds imposed by the noise, it could also be due to the fact that the pre-blur widths are forced to vary by, at most, one pixel width from neighbouring pixels. This assumption leads to a slowly varying function of pre-blur widths but it leads to problems, since the pre-blur widths are *forced* to decrease slowly until areas of high information density are reached. For the GaAs spectrum, the ideal pre-blur width would have a constant value throughout the entire spectrum except at the step, where it would drop to its minimum value.

Examining the pre-blur widths more closely reveals that the width varies between about 4 (at the step) and 16 (on the plateau). This explains why the reconstruction using a constant pre-blur width of 4 was able to extract the step but was unable to smooth the slight ringing on the plateau.

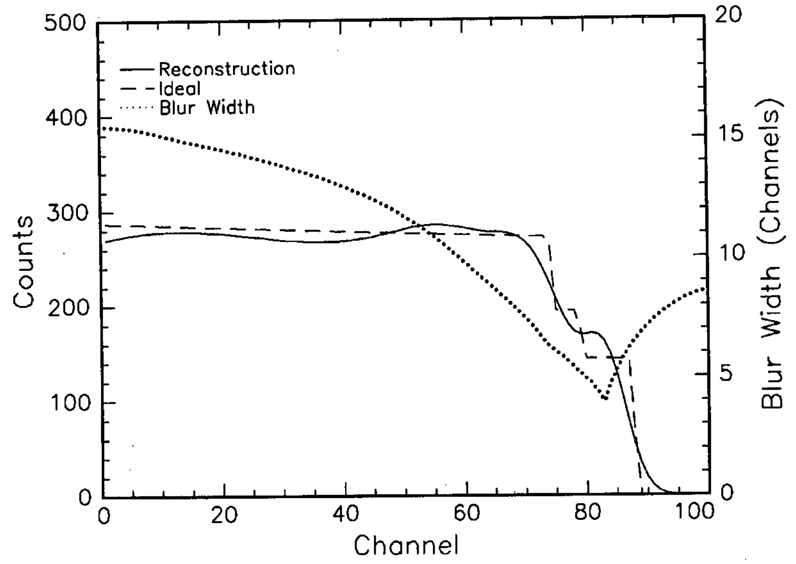

Figure 5.29. Multiresolution pixon reconstruction of GaAs. The spectrum is smooth and the ringing has been reduced on the plateau. The pre-blur width (shown as a dotted line) decreases until it reaches its minimum value at the region of the spectrum where there is greatest detail namely the step between the Ga and As signals.

The hidden image of the reconstruction (*Figure 5.30*) shows oscillations on the plateau which occurs because there are no correlations between neighbouring channels.

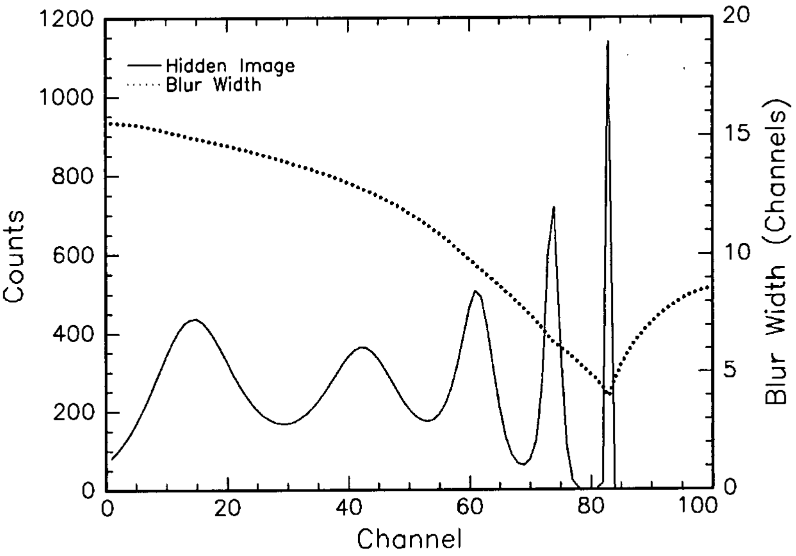

Figure 5.30. The corresponding hidden image of *Figure 5.29.* The hidden image represents the 'base' image with no correlations between neighbouring channels.

Figure 5.31 show the comparison between the original data and the convoluted result.

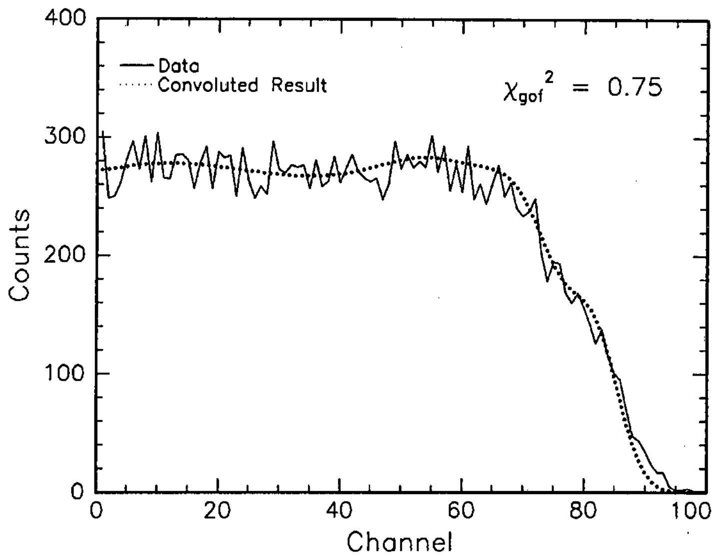

Figure 5.31. **Comparison of convoluted result and original GaAs data. The value of χ^2_{gof} implies that there is some under-fitting of the data, most likely on the plateau of the spectrum.**

The χ^2_{gof} of 0.75 indicates that, on average, the reconstruction slightly over-fits the data. The over-fitting occurs on the plateau which should ideally be flat, but instead shows traces of oscillations. The reconstruction is however still better than those obtained from Fourier Transforms, Jansson's method or MaxEnt with a constant pre-blur.

5.6.2. Deconvolution of PIXE Spectra

The reconstruction of the PIXE spectrum without a background in *Figure 5.32* shows a clear improvement in resolution. The improvement is, however, not sufficient to separate all the peaks (the peaks between channels 70 and 85 have not been separated). The reconstruction is smoother than those obtained with Fourier Transforms, Jansson's method and MaxEnt with a constant pre-blur.

The pre-blur widths behave as expected decreasing to their lowest values at positions of the true peaks and increasing in regions of noise. The requirement that the pre-blur widths vary slowly causes problems with the reconstruction of PIXE spectra. This constraint results in closely-spaced peaks not being completely resolved since the width cannot vary quickly enough between the zero level and the peaks. This is seen around channel 75 where the pre-blur width does not reach its minimum value and the peak is not completely resolved. If the behaviour of the pre-blur widths are compared to the peaks, it is easy to distinguish between the spurious noise peaks and the true peaks since the widths are broadest in regions of the noise peaks.

The pre-blur widths have a minimum value of around 2, which explains why the reconstruction with a constant pre-blur width had many spurious peaks; the constant width of 0.2 was insufficient to smooth out the noise.

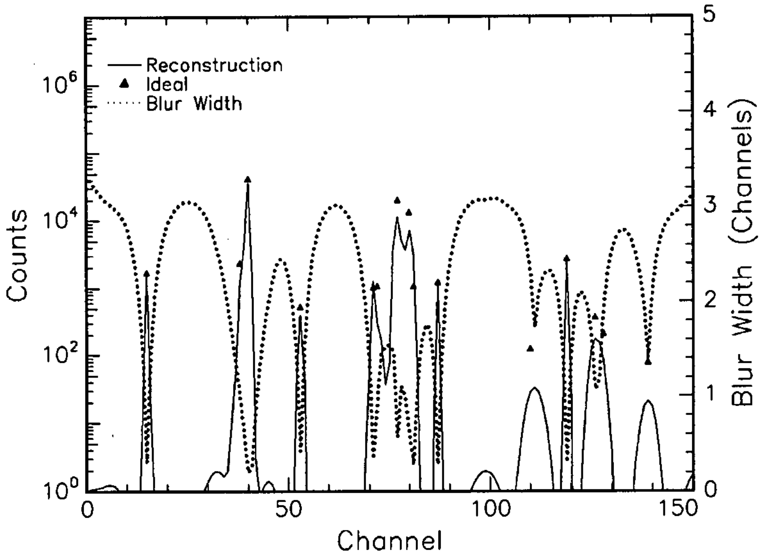

Figure 5.32. Reconstruction of PIXE spectrum without a background. All the peaks have been deconvoluted. The improvement in the resolution is however insufficient to separate all peaks. It can again be seen that the pre-blur widths have minimum values at the points in the spectrum that contain the most information, according to the multiresolution theory.

Figure 5.33 shows the hidden image corresponding to the reconstruction in *Figure 5.32*. The improvement in resolution in the hidden image of the PIXE spectrum without a background (*Figure 5.33*) is however, even better than that of the reconstruction with a constant pre-blur width of 0.2 channels (*Figure 5.25*). This is because the hidden image represents a reconstruction without correlations and this best describes the ideal spectrum in *Figure 5.5*. This makes sense since the ideal spectrum contains single energy peaks with no correlations between them, as in the hidden image. In this special case where each channel is independent of its neighbours, correlations can be ignored and the hidden image would be a more accurate reconstruction. This is however a *special* case since all real spectra have some form of correlation between neighbouring channels, usually in the form of a background.

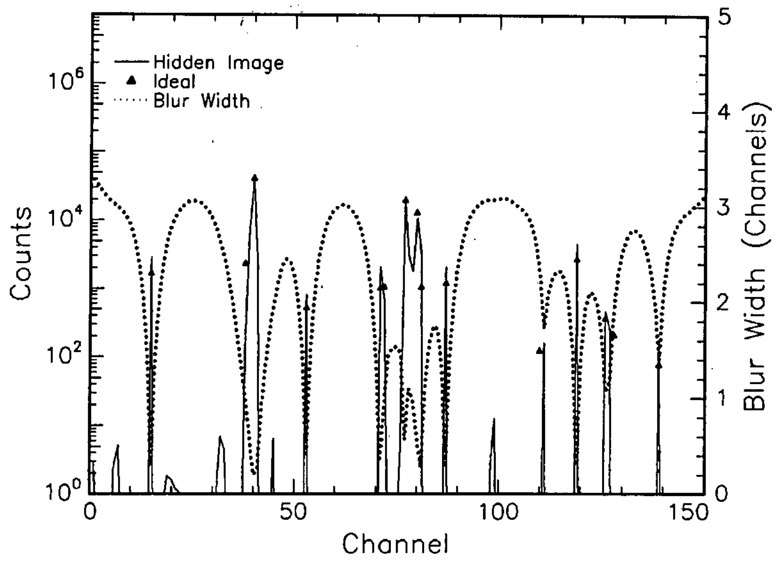

Figure 5.33. The corresponding hidden image of *Figure 5.32* The hidden image contains no correlations between neighbouring channels.

The value of 1.89 for the χ^2_{gof} indicates that there is some under-fitting of the data (*Figure 5.34*). The degree of the under-fitting is however greater than the MaxEnt reconstruction with a constant pre-blurring width which implies that the reconstruction of the PIXE spectrum obtained with a constant pre-blur width of 0.2 channels is better than that using multiresolution pixons.

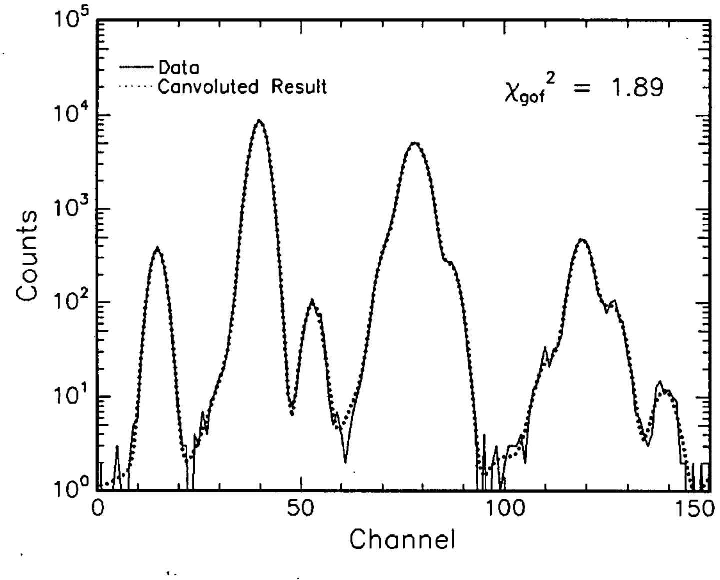

Figure 5.34. Comparison of convoluted reconstruction and original data of the PIXE spectrum without a background. The value of χ^2_{gof} implies that there is some under-fitting. The degree of under-fitting is however greater than that of the MaxEnt reconstruction with a constant pre-blur width.

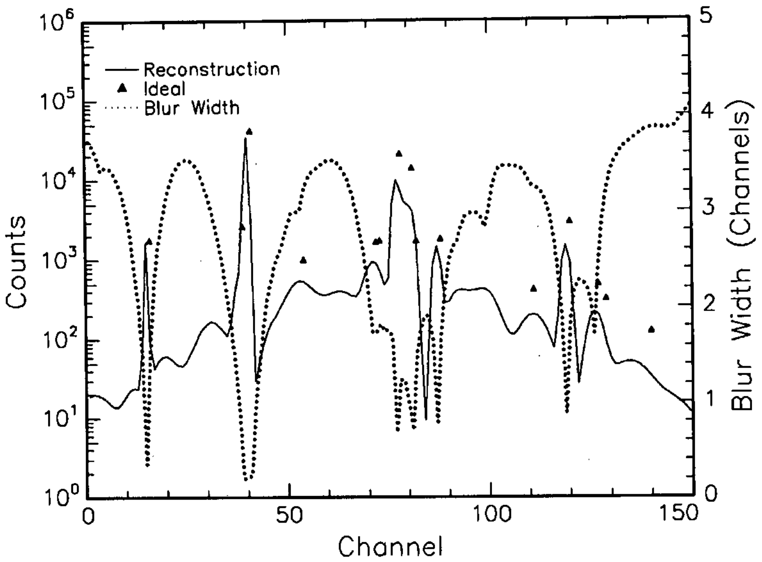

Figure 5.35. Multiresolution pixon reconstruction of PIXE spectrum with background. Most of the peaks have been deconvoluted but the improvement is not sufficient to completely separate the peaks. The widths of the pre-blurring function behave as expected decreasing to a minimum at regions containing the most informaion.

The reconstruction of the PIXE spectrum with a background in *Figure 5.35* shows that most of the peaks have been extracted but the improvement in resolution is insufficient to separate closely-spaced peaks. As explained previously this is due to the requirement that the pre-blur widths vary slowly. The ideal PIXE spectrum with a background (*Figure 5.6*) contains single energy peaks with values ranging from ~100 to ~10^4. These peaks are superimposed on a background and the pre-blur widths have to vary slowly between the background values and the peak values. Certain peaks are therefore broader than they ideally could be since the pre-blur width is not as small as it should be. Examination of the pre-blur widths shows that the width has values around 0.2 for the peaks around channel 15 and channel 40, around 0.7 for the peaks between channels 60 and 90 and around 0.8 for the peaks between channels 100 and 130. Since the peaks between channels 60 and 90 and between channels 100 and 130 are close to each other, the widths did not reach the minimum value of 0.2 in these regions of the spectrum.

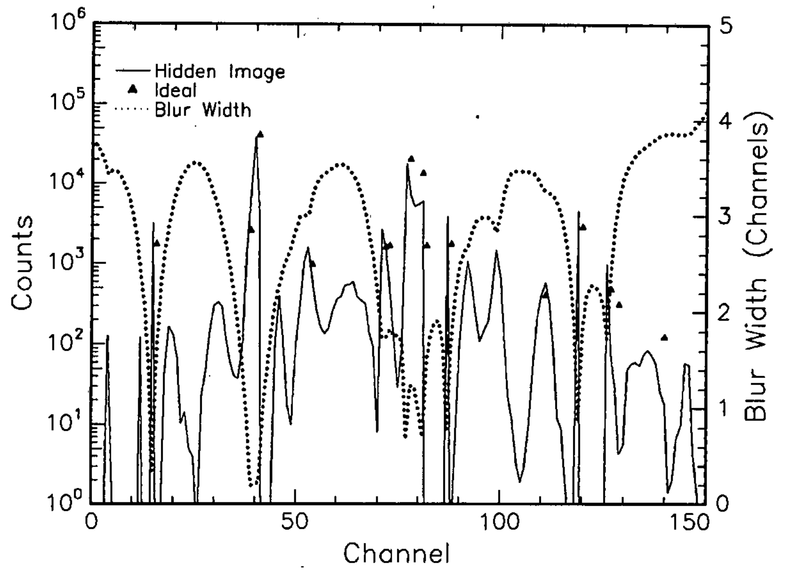

Figure 5.36. The hidden image corresponding to *Figure 5.35.*

The lack of correlations in the hidden image in *Figure 5.36* results in many spurious peaks which are smoothed out by the pre-blur in the final spectrum. The positions of the spurious peaks correspond to large pre-blur widths. In the case of the reconstruction of the PIXE spectrum with a background in *Figure 5.35* correlations between neighbouring pixels exist because of the background, and so the hidden image will not be an accurate reconstruction because of the lack of correlations.

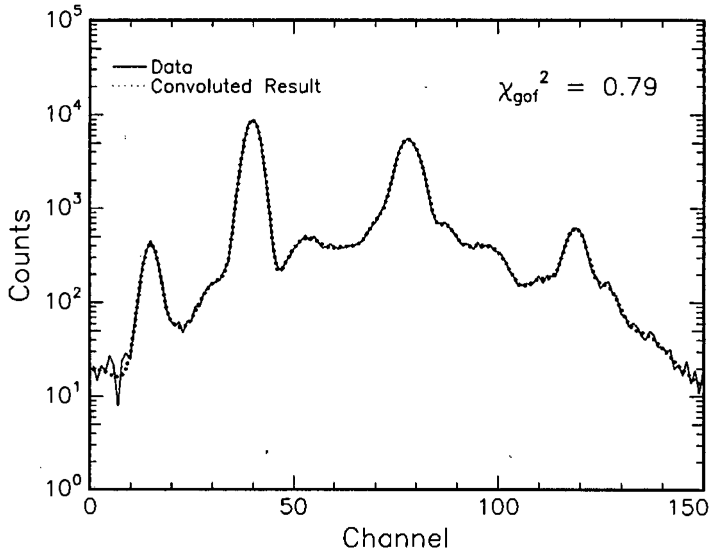

Figure 5.37. Comparison of convoluted reconstruction and original data of the PIXE spectrum with a background. The value of χ^2_{gof} implies that the data is being over-fitted, and features which are due to noise alone are included in the final reconstruction.

The multiresolution pixon reconstruction of the PIXE spectrum with a background over-fits the data (*Figure 5.37*). The reconstruction is, however, better than that obtained using a constant pre-blur width (*Figure 5.27*) and the ringing present in the constant pre-blur reconstruction *Figure 5.27* has been significantly reduced. This is because the minimum pre-blur width in the reconstruction in *Figure 5.35* is around 2 which is much larger than the constant width of 0.2 used in the reconstruction of *Figure 5.27*. The constant width of 0.2 was therefore too small to smooth out the ringing in *Figure 5.27*.

5.6.3. Summary

The reconstruction obtained using multiresolution pixons show around a factor 2.5 improvement in the resolution of the RBS spectrum. The spectrum is also smooth and there is very little trace of ringing on the plateau. The Ga and As edges are resolved but there is no separation of the isotopes of Ga. The pre-blur widths behave as expected, reaching a minimum value in the region of the spectrum with the greatest amount of detail, around channel 80.

The reconstruction of the PIXE spectrum without a background in *Figure 5.32* shows some evidence of peaks which are not completely resolved. This is due to the requirement that the pre-blur width vary slowly. In the PIXE spectrum without a background the values have to vary between the zero level and the peaks and the pre-blur width is unable to vary quickly enough for closely spaced peaks. The hidden image of the PIXE spectrum without a background in *Figure 5.33* shows much improvement in the resolution and all the peaks have been resolved. For the case of the PIXE spectrum without a background, the hidden image is a more appropriate reconstruction since there are no correlations between peaks in the original spectrum so there is no need for correlations in the reconstruction. There are some spurious peaks in *Figure 5.32* and *Figure 5.33* but these can be distinguished from the true peaks since the pre-blur width is narrow in regions of the true peaks and broad in regions of noise.

There is a slight improvement in the resolution of the PIXE spectrum with a background, compared to the original spectrum in *Figure 5.8*. Most of the peaks have been separated but the improvement is not sufficient to completely separate closely spaced peaks. This is due to the fact that the pre-blur widths vary are not allowed to vary quickly between the width required for the background and that required for the peaks.

5.7. Summary of Deconvolution Techniques

All techniques used show an improvement in the resolution of the RBS spectrum and all show a step around channel 80. None of the techniques increased the resolution to the extent that the isotopes of Ga could be separated. Only the MaxEnt techniques showed an improvement in the resolution of the PIXE spectra. In comparing the methods, there is an *ad hoc* nature to both the Fourier Transform method and Jansson's method in that there is no consistent method of determining the filter parameters in the Fourier Transform method, and no formal way of determining the relaxation constant and number of iterations required in Jansson's method. The multiresolution pixon technique, on the other hand, is completely general and all information about the reconstruction is included in the formalism. The use of MaxEnt yields smooth reconstructions unlike Fourier Transforms and Jansson's method which have high noise components.

MaxEnt with a constant pre-blur (i.e. constant correlation lengths) produces reconstructions that show a clear improvement in resolution of the GaAs spectrum but there are traces of ringing on the plateau of the reconstruction. There are also traces of ringing on the background of the PIXE reconstruction (*Figure 5.27*). The ringing is due to the fact that the value of the pre-blur width used in these reconstructions was too small to smooth out the ringing. The ringing is significantly reduced in the multiresolution pixon reconstructions where the size of the pre-blur width depends on the information density of the spectrum.

Table 5-3. Comparison of χ^2_{gof} values for each technique of deconvolution.

<table>
<tr><th></th><th>FFT</th><th colspan="2">Jansson</th><th>MaxEnt with constant pre-blur</th><th>Multiresolution Pixons</th></tr>
<tr><td rowspan="2">GaAs</td><td rowspan="2">225</td><td>50 iterations</td><td>0.67</td><td rowspan="2">0.67</td><td rowspan="2">0.75</td></tr>
<tr><td>200 iterations</td><td>0.22</td></tr>
<tr><td>PIXE without background</td><td>6.9</td><td colspan="2">4.97</td><td>1.74</td><td>1.89</td></tr>
<tr><td>PIXE with background</td><td>2.5</td><td colspan="2">0.85</td><td>0.57</td><td>0.79</td></tr>
</table>

The different values of χ^2_{gof} indicate that each method performs in a different manner for each spectrum. Fourier Transforms seem to constantly under-fit the data but this is expected since the filter to eliminate noise is applied in a very *ad hoc* manner and could possibly be eliminating some of the Fourier components of real data. These components will therefore be missing in the final reconstruction. Both Jansson's method and MaxEnt with a constant pre-blur over-fit the GaAs spectrum and the PIXE spectrum with a background but under-fit the PIXE spectrum without a background. The technique of multiresolution pixons under-fits the GaAs spectrum and the PIXE spectrum without a background and over-fits the PIXE spectrum with a background.

Deconvolution using MaxEnt (both with and without pixons) consistently produces good results and has the further advantage that the results are quantitatively correct since the conservation of the total counts in a spectrum is included in the pixon formulation of the deconvolution problem. All MaxEnt reconstructions, especially of the PIXE spectra, correspond closely to the ideal values of the peak height. In cases where it does not, the peak is usually broad and the area under the broad peak corresponds statistically to the ideal peak height. Neither Fourier Transforms nor Jansson's method has the option of preserving the counts in a spectrum.

The technique of maximising the entropy is very powerful and performs much better than conventional methods in deconvoluting spectra. The MaxEnt formalism uses all information *supplied* to it to perform a reconstruction and this means that it is highly sensitive to the assumptions made about the reconstruction. This is demonstrated, for example, by the fact that including correlations in a reconstruction reduces the ringing by reducing the number of degrees of freedom available for fitting the data. This implies that correlations do exist in certain cases. But including correlations in a reconstruction where none exist leads to a poorer reconstruction as seen in the reconstruction of the PIXE spectrum without a background where no correlations exist (*Figure 5.32*).

There are, however, problems associated with the reconstruction technique using the concept of a pre-blur to introduce correlations into a reconstruction since there is no formal reason for introducing a pre-blur. Ideally, one should use an expression for the entropy that includes correlations.

5.8. Depth Profiles

5.8.1. RBS Depth Profiles

Based on the good results of the MaxEnt formalism in deconvoluting spectra, it was applied to a direct calculation of depth profiles from an RBS spectrum. The standard method of determining depth profiles is to simulate a spectrum from an assumed profile, convolute and compare to the data using a least-squares criterion. The second method of determining depth profiles is to do a backward calculation using (4-5) – (4-7). The latter method uses only the spectrum and does not require any simulation. The latter method is quick and simple but the resulting depth information still contains the smoothing effect of the detector so small changes will be smoothed out in the data and therefore missed in the guessed profile. This has the additional disadvantage that non-zero concentrations may appear at negative depths.

In order to apply the Bayesian formalism to inverse problems, it is necessary to be able to simulate pseudo-data assuming the physics is known, in order to compare the reconstruction with the available data (§3.5.1). In the case of the deconvolution problem, this means convoluting the physics and then comparing the convoluted result to the data. For the case of extracting depth profiles from spectra, it is necessary to be able to simulate a spectrum from a given depth profile before the reconstructed depth profile can be compared to the data. The simulation of RBS spectra from a given depth profile was previously programmed into FORTRAN-77 by V.M. Prozesky (VDGG, NAC) and this module was added to the standard MaxEnt program and the modified program was used to extract depth profiles from RBS spectra.

The data consisted of a simulated spectrum of 2×10^{17} atoms/cm^2 of Ir on a Si substrate, which was convoluted with the detector response function of the SSB detector as described in §4.1.1 and Gaussian noise was added. This spectrum was then used to calculate depth profiles by the second method described above (i.e. extracting a depth profile directly from the data) (*Figure 5.38*) and in the Bayesian framework (*Figure 5.39*).

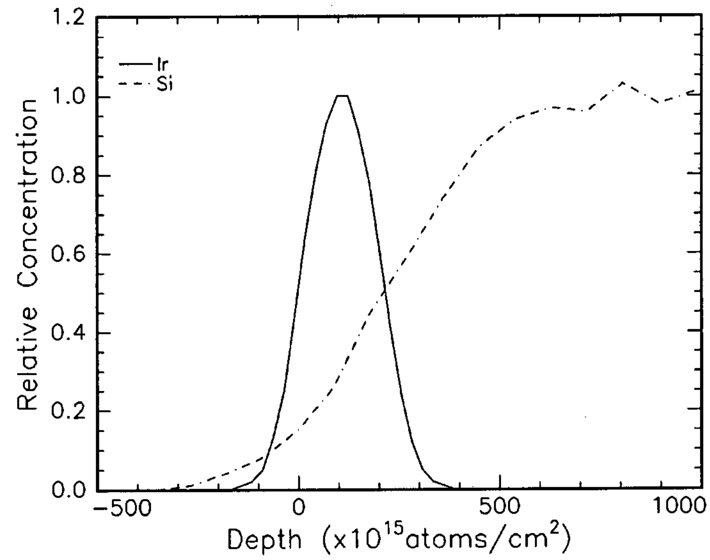

Figure 5.38. Depth profiles of Ir and Si extracted from an RBS spectrum using a direct calculation. There is no deconvolution of the detector response function and this has resulted in non-zero concentrations at negative depths.

The depth profiles in *Figure 5.38* show the non-zero concentrations at negative depths that is characteristic of a method that does not consider the effect of the detector response function.

The profiles also do not have sharp boundaries and there is evidence of noise on the plateau of the Si profile.

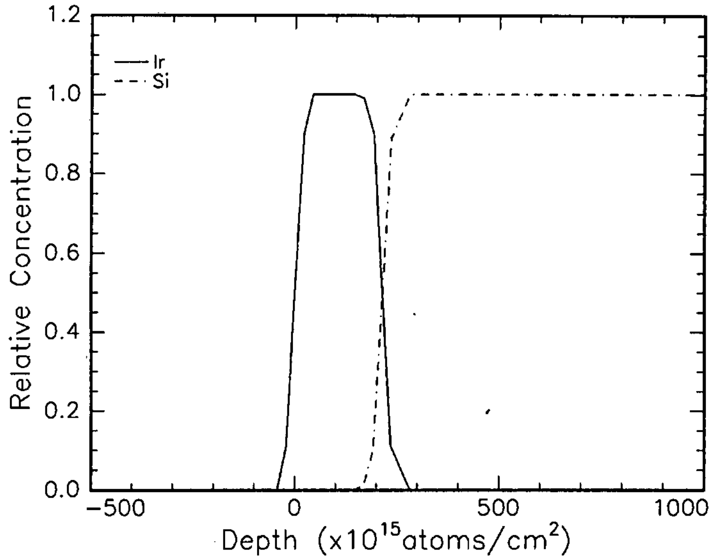

Figure 5.39. Depth profiles of Ir and Si extracted from a deconvoluted spectrum. The profile is smoother because noise has been taken into account.

The MaxEnt reconstruction in *Figure 5.39* shows an improvement of about a factor of 5 in resolution and clearly shows the flat tops of the Ir and Si distributions. The profiles are also smoother because the statistics of the noise have been taken into account. The MaxEnt reconstruction of the depth profile still shows non-zero concentrations at negative depths but this is to a much lesser degree than in *Figure 5.38.*

5.8.2. PIXE Depth Profiles

As discussed in §4.4.2.4, PIXE is not a good analytical technique for recovering depth information but some information may be extracted from PIXE spectra by changing the beam energy, changing the sample angle or by using two detectors. For this investigation, the two-detector set-up, shown in *Figure 4.10* was used.

As with extracting depth profiles from RBS spectra, it was necessary to simulate x-ray yields from a given depth profile. The simulation of x-ray yields from a given depth profile was previously programmed into FORTRAN-77 by V.M. Prozesky (VDGG, NAC) and this module was added to the standard MaxEnt program.

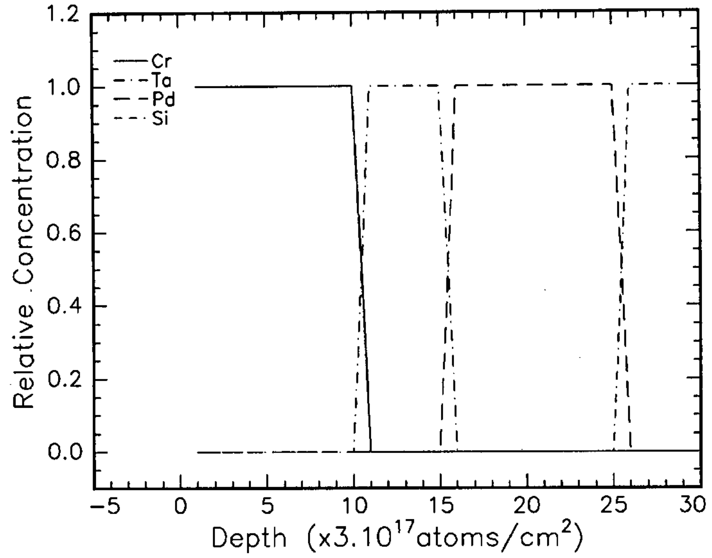

Figure 5.40. Ideal depth profile. The target consisted of 3×10^{18} atoms/cm^2 Cr on 1.8×10^{18} atoms/cm^2 Ta on 3×10^{18} atoms/cm^2 Pd on a Si substrate

The data was simulated for each of two detectors (situated at 45° and 60° with respect to the target normal) from a target consisting of 3×10^{18} atoms/cm^2 Cr on 1.8×10^{18} atoms/cm^2 Ta on 3×10^{18} atoms/cm^2 Pd on a Si substrate. Noise was added to the simulated yields to simulate counting statistics. *Figure 5.40* shows the depth profiles that were used to simulate the PIXE yields for each element, for each of the detectors in the two-detector set-up. The Bayesian formalism was applied to these simulated spectra in an attempt to recover the original depth profiles. The results can be seen in *Figure 5.41.*

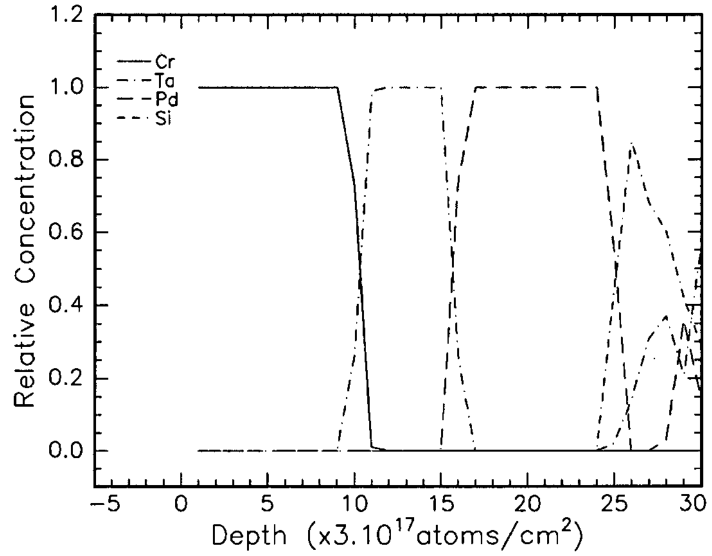

Figure 5.41. Reconstructed depth profile. The reconstruction compares very well to the theoretical profile up to the depth of the bulk Si. At this depth, the attenuation of the Si x-rays is quite strong, so that there was insufficient information in the spectra for the procedure to extract any more detail.

The information for 90 variables (30 depth intervals for each of 3 independent concentrations) was extracted from 8 x-ray yields – a very ill-defined problem. MaxEnt recovers the information quite well until it reaches the bulk Si. Si x-rays have a low energy (1.74 keV) and at this depth, the attenuation of the Si x-rays is quite strong, so that there was insufficient information in the spectra for the procedure to extract any more detail.

5.9. Deconvolution of the Beam Profile from 1-D NMP Scans

As discussed in §4.1.2, the NMP is ideal for obtaining lateral information about the distribution of elements in a target. The microbeam is scanned over the target and an image of the elemental distribution in the target is obtained. The use of the NMP for imaging has two disadvantages. Firstly, features smaller than the beam are smoothed out by the beam, which usually means a necessity for smaller beams which means new hardware and increased costs. Secondly the use of a small beam means that the beam currents must be small. This results in a low analytical precision because there will be few counts in the spectra which. With the MaxEnt technique of deconvolution, it should be possible to deconvolute the beam profile from the images and extract more information from images without increased hardware costs. At this stage, the MaxEnt algorithm is not fast enough to process two-dimensional images so one-dimensional scans were deconvoluted. For a 64 pixel by 64 pixel image, computation time would be of the order of a day.

1-D scans (*Figure 5.42*) of an integrated circuit were performed using *ion-beam-induced charge* (IBIC) microscopy [Breese *et al* 1991] with a 3 MeV proton beam with a FWHM of about 200 nm × 200 nm. The beam profile was assumed to be Gaussian. The deconvoluted spectra can be seen in *Figure 5.43.* In this case of deconvolution, the real beam profile was not known and had to be estimated using normal beam size measurements [Prozesky *et al* 1995]. The reconstruction is smooth and there is an enhancement of certain features with a small improvement in resolution. This is due to the fact that the description of the beam profile was not accurate enough (c.f. §2.1). In order to improve the resolution further, a more accurate description of the beam profile is required.

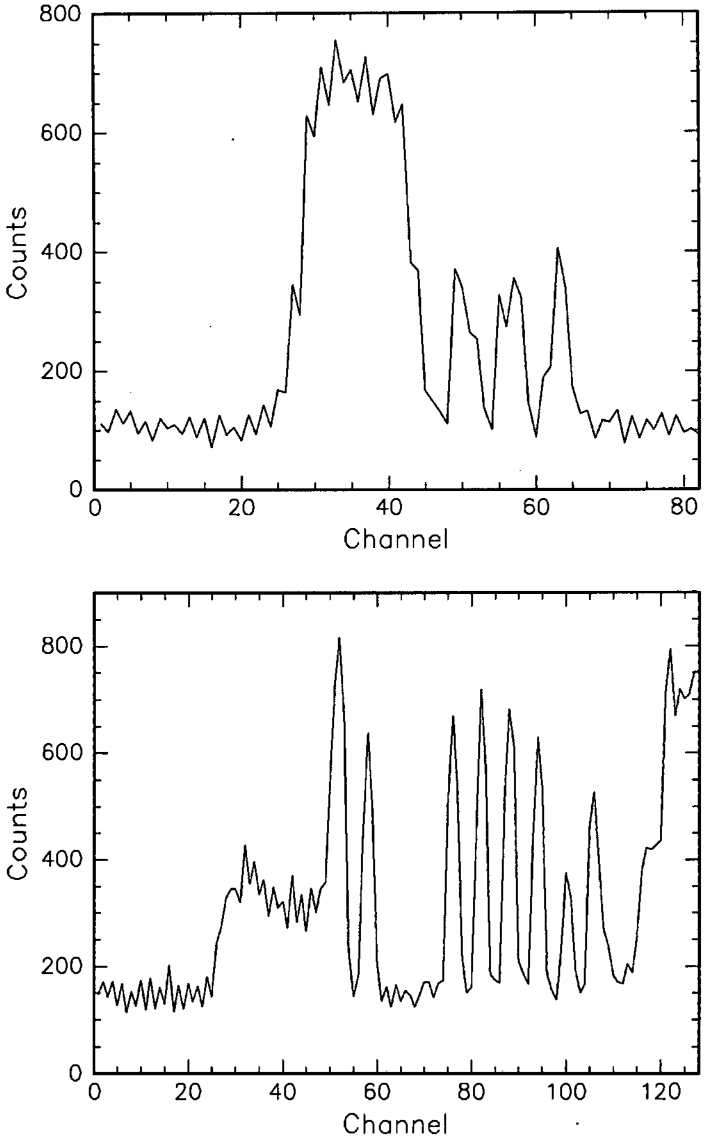

Figure 5.42. One-dimensional scans of a Si chip obtained using the IBIC technique. The beam size was about 200 nm × 200 nm.

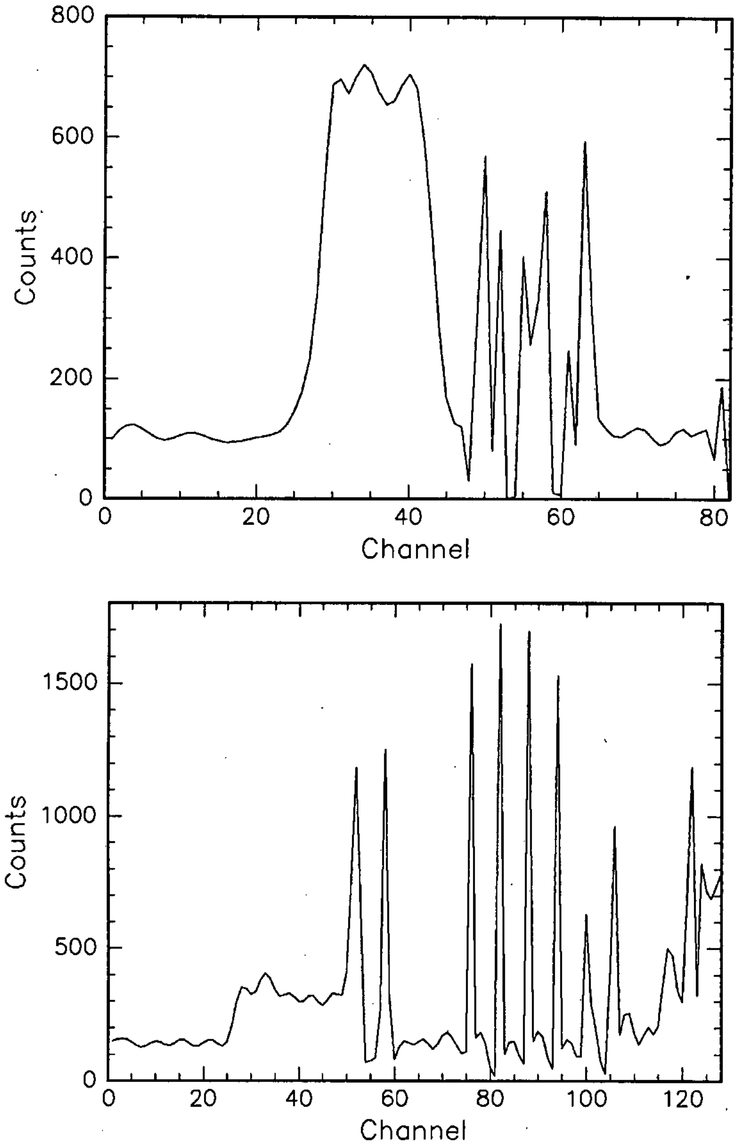

Figure 5.43. One-dimensional scans with the beam profile deconvoluted. A slight enhancement in the resolution of the beam is seen and the scans are much smoother than the originals.

Chapter 6

Conclusions and Future Work

6.1. Conclusions

The aim of this investigation was to demonstrate the efficiency of Bayesian Statistics using the concepts of maximum entropy and multiresolution pixons, in solving ill-defined inverse problems from IBA applications, the most important being the deconvolution of the detector response function from spectra.

Three methods of deconvolution (Fourier Transforms, Jansson's Method and MaxEnt with a constant pre-blurring function) were compared to deconvolution using multiresolution pixons. The various techniques were tested on RBS and PIXE spectra. The RBS spectrum of GaAs was flat with some detail in a small portion of the spectrum. In contrast, the PIXE spectra contained single energy peaks; one superimposed on a background and one without a background. Whereas the RBS spectrum had counts of the order of a few hundred counts per channel, the counts of the PIXE spectra varied between ~10 to ~10^4 counts per channel.

In general, all deconvolution techniques resulted in an improvement in the resolution of the RBS spectrum in that the Ga and As edges were resolved but each technique produced different results on the plateau of the spectrum. None of the techniques were able to improve the resolution to the extent that both isotopes of Ga were separated. Only the MaxEnt techniques resulted in an improvement in the resolution of the PIXE spectra.

Fourier Transforms can be used to improve the resolution of RBS spectra, but there is no noticeable improvement in the resolution of PIXE spectra. Deconvolution using Fourier Transforms requires a low-pass filter to eliminate the high frequency noise components, but there is no consistent method of choosing the parameters of the filter. A further disadvantage of using a filter is that high frequency components of the data are also filtered out. The presence of negative components in the reconstruction is another undesirable feature of this method. One advantage of using Fourier Transforms is that it is a single step process and therefore contains none of the problems associated with an iterative method. It also requires very little computation time. The use of Fourier Transforms (without a filter) in deconvolution demonstrates the weakness of linear deconvolution methods in treating noise as well as the lack of constraints on the reconstruction.

Jansson's method is a non-linear deconvolution technique and allows for a positivity constraint to be imposed on the reconstruction. The resolution of both RBS and PIXE spectra was improved but there are many problems associated with the method. Firstly, it is an iterative procedure, which means that errors generated in the first iteration are propagated through to the last. Secondly, there is no formal method of determining the various parameters required. Jansson's method does, however, eliminate the negative components that were present in the Fourier Transform reconstructions. Computation times depend on the number of iterations and are of the order of a few minutes for 10 000 iterations.

Maximum entropy with a constant pre-blur exhibits different results for different types of spectra. The reconstructed GaAs spectrum is smooth compared to the reconstructions obtained using Fourier Transforms or Jansson's method. The step around channel 80 is clearly seen but there is a considerable amount of ringing on the plateau which is non-physical. The reconstruction of the PIXE spectrum without a background is close to the ideal spectrum with noise sources containing around 10 counts per channel. The reconstruction of the PIXE spectrum with a background shows an improvement in the resolution but the reconstruction has a high level of noise. These differing results can be understood on the basis that the pre-blurring function introduces correlations into the reconstruction. For a fixed pre-blur width, correlations are introduced over fixed portions of the spectrum. For the GaAs spectrum the width of 4 is sufficient to resolve the Ga and As edges but is insufficient to smooth out the ringing on the plateau. For the PIXE spectra, the width of 0.2 channels is too small to smooth out the noise peaks in the spectrum with a background while it is sufficient for the spectrum without a background.

The reconstruction obtained using multiresolution pixons shows a marked improvement in the resolution of the RBS spectrum, when compared to the original spectrum. The reconstruction is smooth and there are slight traces of ringing on the plateau. The Ga and As edges are resolved but there is no separation of the isotopes of Ga. The pre-blur widths behave as expected, reaching their minimum value in the region of the spectrum with the greatest amount of detail, namely in the region of the step, around channel 80.

The multiresolution pixon reconstruction of the PIXE spectrum without a background in *Figure 5.32* shows some evidence of peaks that have not been completely resolved. This is due to the requirement that the pre-blur widths vary slowly and the widths are unable to vary quickly enough between the zero value and the peak value to separate all peaks completely. There are some spurious peaks in the reconstruction but these can be distinguished from the true peaks since the pre-blur width is narrow in regions of the true peaks and broad in regions of noise. The hidden image of the PIXE spectrum without a background (*Figure 5.33*) shows a significant improvement in the resolution and all the peaks have been separated. For the case of the PIXE spectrum without a background, the hidden image is a more appropriate reconstruction since there are no correlations between peaks in the original spectrum so there is no need for correlations in the reconstruction.

There is a slight improvement in the resolution of the PIXE spectrum with a background. Most of the peaks have been separated, but the improvement is not sufficient to completely separate closely spaced peaks. This is, again, due to the fact that the pre-blur widths are not allowed to vary quickly between the width required for the background and that required for the peaks.

Another application of the Bayesian formalism, of direct importance to the NMP is the deconvolution of the beam profile from images. The deconvoluted 1-D scan is much smoother but there is not much improvement in the resolution because the description of the beam profile is not accurate enough.

In general, a greater improvement in resolution can be obtained by improving the description of the detector response functions, since the amount of information that can be extracted from experimental data is dependent on the prior information supplied. For the deconvolution of the beam profile from NMP images, a more accurate description of the beam profile is required.

The multiresolution pixon technique is also ideal for extracting information from spectra. The MaxEnt extraction of depth profiles from an RBS spectrum of a thin layer of Ir on Si shows a distinct interface, unlike conventional methods. There is about a factor 5 improvement in resolution because the detector response function has been taken into consideration in the calculation. The flat tops of the Ir and Si distributions are clearly seen and the profiles are also smooth because noise has been taken into account.

It is not possible to extract depth information from conventional PIXE spectra but by using a two-detector set-up, depth information can be extracted. The results show that MaxEnt recovers depth information quite well from PIXE spectra.

More accurate depth profiles can be extracted from spectra by improving the simulation of the pseudo-data. This means that all possible processes that lead to the generation of a spectrum must be included in the simulation.

The technique of using Bayesian statistics with an entropic prior is ideal for solving ill-defined inverse problems in IBA. It is especially suited to the deconvolution of detector response functions from spectra where it extracts all possible information from the available data. The formalism allows for many constraints (like positivity, correlations, smoothness, etc.) to be imposed on the reconstructed spectrum and all the constraints are explicitly included in the formalism.

The theory incorporating the idea of a pre-blur to introduce correlations into the prior probability distribution performs well in the deconvolution of spectra from IBA techniques like PIXE and RBS. The resolution of these spectra are increased and more information can be extracted from existing spectra without additional hardware requirements. MaxEnt improves upon conventional deconvolution techniques, which do not usually consider prior information, and it has been shown that including prior information like correlations, results in a better reconstruction.

6.2. Future Work

Although the theory performs quite well in deconvolution, many questions have arisen during the course of this work which could be further investigated.

The main concern at this point is the *ad hoc* introduction of the pre-blur. The concept works well implying that correlations do exist but the question arises whether this could be incorporated into the form of the entropy. This leads one to consider *Tsallis' entropy* [Tsallis 1988] which implicitly includes correlations.

The formalism, as it stands at present, introduces correlations into the reconstruction via the pre-blurring function but the form of the likelihood still assumes that neighbouring data points are logically independent from each other. The effect, of using such a form for the likelihood with correlations introduced by the pre-blur, on the reconstruction is unknown.

The theory could be applied to experimental PIXE spectra although this would mean that the description of the response function must be improved. The theory of splines [von der Linden *et al* 1997a] could be used to obtain a more accurate fit to the detector response function.

Von der Linden *et al* [1997b] have developed a procedure based on Bayesian Statistics to subtract backgrounds from spectra. This leads one to the possibility of subtracting the background from PIXE spectra before processing. This may eliminate the need for a pre-blur when analysing PIXE spectra since these spectra consist of single energy peaks superimposed on a background with correlations existing *only* because of the background.

The speed of the MaxEnt algorithm could be increased so that NMP images could be processed. The theory of simulated annealing [Metropolis *et al* 1953, Kirkpatrick *et al* 1983] could be incorporated into the algorithm to improve processing.

Appendix A

Measure Theory

Measure theory is a branch of analysis with its most important application being integration. In order to define a *measure* it is necessary to define an *algebra*.

An algebra is a class $\mathcal{A}$ of subsets of a space $\mathcal{S}$ that satisfy the following conditions:

1. $\varnothing \in \mathcal{A}$
2. If $A \in \mathcal{A}$, then $A^c \in \mathcal{A}$
3. If $A_1, A_2, \ldots \in \mathcal{A}$ then $\bigcup_i A_i \in | \mathcal{A}$
4. If $A_1, A_2, \ldots \in \mathcal{A}$ then $\bigcap_i A_i \in \mathcal{A}$

Further, $\mathcal{A}$ is a σ-algebra if $\mathcal{A}$ contains the limit of every monotone sequence of its sets. The pair $(\mathcal{S}, \mathcal{A})$ is then a measurable space and the sets in $\mathcal{A}$ are measurable [Doob 1994, Parker 1984].

A *measure* μ is a function which assigns a positive real number to each set in $\mathcal{A}$ and satisfies the following conditions:

1. $\mu(\varnothing)=0$
2. If A and B $\in \mathcal{A}$ and $A \cap B = \varnothing$ then $\mu(A \cup B) = \mu(A) + \mu(B)$
3. If $A_1, A_2, \ldots \in \mathcal{A}$ then $\mu\left(\bigcup_i A_i\right) \leq \sum_i \mu(A_i)$

The integral of a function *f* on a measurable space $(\mathcal{S}, \mathcal{A})$ can then be written as:

$$\int f(x)\,\mathrm{d}\mu(x)$$

For example, consider the integral of a *simple* function *f* defined by:

$$f = \sum_{i=1}^{n} \alpha_i \chi_{A_i}(x)$$

where

$$\chi_A = \begin{cases} 1 & \mathit{if}\ x \in A \\ 0 & \mathit{if}\ x \notin A \end{cases}$$

is called the characteristic function of the set A and the αs are real numbers. The integral of *f* is then [Halmos 1950]

$$\int f(x)\,\mathrm{d}\mu(x) = \sum_{i=1}^{n} \alpha_i \mu(A_i)$$

since

$$\int \chi_A\,\mathrm{d}\mu = \int_A \mathrm{d}\mu = \mu(A)$$

Appendix B

Counting Statistics

The mechanism responsible for the generation of radiation when a nuclear interaction occurs is subject to inherent statistical fluctuations. These fluctuations represent an unavoidable source of uncertainty in all nuclear counting experiments.

In most cases the data can be described by a distribution function dependent on the experimental mean $\bar{d}_e$ and sample variance s^2 of the data set d.

$$\bar{d}_e = \frac{1}{N_{data}} \sum_{i=1}^{N_{data}} d_i \tag{B-1}$$

$$s^2 = \frac{1}{N_{data} - 1} \sum_{i=1}^{N_{data}} \left(d_i - \bar{d}_e\right)^2 \tag{B-2}$$

The sample variance gives a measure of the degree of fluctuation. The expressions given above are approximations, with the true values $\bar{d}$ and σ^2 obtained by letting $N_{\text{data}} \to \infty$, but since only a finite set of measurements is available, $\bar{d}_e$, derived from the data is used. The effect of using $\bar{d}_e$ instead of $\bar{d}$ is that the variance is smaller.

The distribution functions that can be derived are based on a measurement of the number of successes in a given number of trials t with each trial considered as a binary process – a trial is either a success or not. Assuming that the probability of success p is a constant, three distribution functions are applicable: the *binomial distribution*, the *Poisson distribution* and the *Gaussian distribution*.

B.1. The Binomial Distribution

The probability of counting s successes in t trials is described by the binomial distribution:

$$P_B(s\,|\,tp) = \binom{t}{s} p^s (1-p)^{t-s} \tag{B-3}$$

The binomial distribution is the most general distribution applicable to constant p processes but it is computationally intensive in nuclear experiments because of the large number of nuclei in nuclear counting experiments.

Properties of the Binomial Distribution

$$\sum_s P_B(s\,|\,tp) = 1 \tag{B-4}$$

$$\bar{s} = \sum_s sP_B(s\,|\,tp) = tp \tag{B-5}$$

$$\sigma^2 = \sum_s (s-\bar{s})^2 P_B(s\,|\,tp) = tp(1-p) \tag{B-6}$$

B.2. The Poisson Distribution

The Poisson distribution is an approximation to the binomial distribution for a low probability of success i.e. $p \ll 1$. The Poisson distribution is applicable to ion beam experiments where the number of atoms is large compared to the number of interactions yielding a positive result:

$$P_P(s\,|\,tp) = \frac{\bar{s}^s \exp(-\bar{s})}{s!} \tag{B-7}$$

Properties of the Poisson Distribution

$$\sum_s P_P(s\,|\,tp) = 1 \tag{B-8}$$

$$\bar{s} = \sum_{m} s P_P(s \mid tp) = tp \tag{B-9}$$

$$\sigma^2 = \sum_{s} (s - \bar{s})^2 P_P(s \mid tp) = tp = \bar{s} \tag{B-10}$$

$$\sigma = \sqrt{\bar{s}} \tag{B-11}$$

B.3. The Gaussian Distribution

Just as the Poisson distribution approximates the binomial distribution for $p \ll 1$, so the Gaussian distribution approximates the Poisson distribution for a large mean. Knoll [1979] suggests that $\bar{s}$ should be greater than say 20 to 30 for the Gaussian approximation to be valid:

$$P_G(s \mid tp) = \frac{1}{\sqrt{2\pi\bar{s}}} \exp\left(-\tfrac{1}{2} \left(\frac{s - \bar{s}}{\bar{s}} \right)^2 \right) \tag{B-12}$$

Properties of the Gaussian Distribution

$$\sum_{s} P_G(s \mid tp) = 1 \tag{B-13}$$

$$\bar{s} = \sum_{s} s P_G(s \mid tp) = tp \tag{B-14}$$

$$\sigma^2 = \sum_{s} (s - \bar{s})^2 P_G(s \mid tp) = tp = \bar{s} \tag{B-15}$$

$$\sigma = \sqrt{\bar{s}} \tag{B-16}$$

References

Beiser, A. (1987). *Concepts of Modern Physics.* McGraw-Hill. (4th edition). (New York).

Brandt, W. and Lipicki, G. (1981). *Phys. Rev. A.* **23**. 1717.

Breese, M.B.H., Grime, G.W. and Watt, F. (1991). Oxford Nuclear Physics Report OUNP-91-33.

Campbell, J.L., Millman, B.M., Maxwell, J.A., Perujo, A. and Teesdale, W.J. (1985). Analytic Fitting of Monoenergetic Peaks from Si(Li) X-ray Spectrometers. *Nucl. Instr. Meth. B* **9.** 71-79.

Carter, R.D., Damcott, D.L., Atzmon, M., Was, G.S., Bruemmer, S.M., and Kenik, E.A. (1994). Quantitative analysis of radiation-induced grain-boundary segregation measurements. *J. Nucl. Mat.* **211**. 70-84.

Chu, W.H., Mayer, J.W. and Nicolet, M.A. (1978). *Backscattering Spectrometry.* Academic Press. (New York)

Churms, C., Pilcher, J.V., Springhorn, K. and Tapper, U.A.S. (1993). A VAX and PC-based data acquisition system for MCA, scanning and list-mode analysis. *Nucl. Instr. Meth. B* **77.** 56-61.

Coote, G.E. and Kwan, B.P. (1995). Deconvolution of one-dimensional microprobe scans using the non-linear Jansson method. *Nucl. Instr. and Meth. B* **104.** 228-232.

Cox, R.T. (1946). Probability, Frequency and Reasonable expectation. *Amer. J. Phys.* **14.** 1-13.

Cox, R.T. (1961). *The Algebra of Probable Inference.* John Hopkins Press.

Doob, J.L. (1994). *Measure Theory.* Springer. (New York).

Doolittle, L.R. (1986). A semiautomatic algorithm for Rutherford Backscattering Analysis. *Nucl. Instr. Meth. B* **15.** 227-231.

Evans, B.J. and Smith, L.M. (1992). Determining Railgun Plasma Current Distribution Using Jansson's Method to Deconvolve B-Dot Probe Signals. *IEEE Transactions on Plasma Science* **20**. 432-438.

Fischer, R. , Von der Linden, W. and Dose, V. (1997). Enhancement of the energy resolution in ion-beam experiments with the maximum-entropy method. *Phys. Rev. E* **55**.

Frieden, B.R. (1972). Restoring with maximum likelihood and Maximum Entropy. *J. Opt. Soc. Am.* **62.** 511-518.

Garrett, A.J.M., Prozesky, V.M., Padayachee, J., and Mostert, F.J. (1997). Bayesian Cherry Picking in *Proceedings of the Maximum Entropy Conference 1996.* M. Sears, V. Nedeljkovic, N.E. Pendock and S. Sibisi (eds.). 31-40.

Gold, R. (1964). An Iterative Unfolding Method for Response Matrices. AEC Research and Development Report, ANL_6984, Argonne National Laboratory, Argonne, Illinois.

Goldman, S. (1953). *Information Theory.* Dover. (New York).

Gull, S.F. (1989). Developments in Maximum Entropy Data Analysis in *Maximum Entropy and Bayesian Methods.* J. Skilling (ed.). Kluwer. 53-71.

Gull, S.F. and Skilling, J. (1984). Maximum Entropy method in image processing in *IEE Proceedings*. **131**. 646-659.

Gull, S.F. and Skilling, J. (1985). The entropy of an image in *Maximum Entropy and Bayesian Methods in Inverse Problems.* C. Ray Smith and W.T. Grandy (eds.). Kluwer. 287-301.

Halmos, P.R. (1950). *Measure Theory.* Nostrand. (New York).

IDL (1993). *Interactive Data Language. Users' Manual.* Research Systems Inc. (Boulder).

Jansson, P.A. (ed.) (1984). *Deconvolution: with Applications in Spectroscopy.* Academic Press. (Orlando).

Jaynes, E.T. (1957). Information Theory and Statistical Mechanics. *Phys. Rev.* **106**. 620-630, *Phys. Rev.* **108**. 171-190.

Jaynes, E.T. (1978). *The Maximum Entropy Formalism.* R.D. Levine and M. Tribus (eds.). Kluwer. (Dordrecht).

Jaynes, E.T. (1995). *Probability Theory: The Logic of Science.* (in press)

Johansson, S.A.E. and Campbell, J.L. (1988). *PIXE: A Novel Technique for Elemental Analysis.* Wiley.

Johansson, S.A.E., Campbell, J.L. and Malmqvist, K.G. (1995). *Particle Induced X-ray Emission Spectrometry.* Wiley.

Johansson, T.B., Akselsson, R. and Johansson, S.A.E. (1970). X-ray Analysis: Elemental Trace Analysis at the 10^{-12}g level. *Nucl. Instr. Meth.* **84** 141-143.

Kirkpatrick, S., Gelatt, C.D. and Vecchi, M.P. (1983). Optimisation by Simulated Annealing. *Science* **220** 671-680.

Knoll, G.F. (1979). *Radiation detection and measurement.* Wiley. (New York).

Livesey, A.K. and Smith, G.C. (1994). The determination of depth profiles from angle-dependent XPS using maximum entropy data analysis. *J. Electron Spectroscopy and Related Phenomena.* **67**. 439-461.

Maj, B., Mutsaers, P.H.A., Rokita, E. and de Voight, M.J.A. (1996). Determination of microbeam profile using deconvolution procedures. *Nucl. Instr. Meth. B* **113**. 391-395.

Marchetti, A.A. and Mignerey A.C. (1993). Deconvolution of Mass Spectra. *Nucl. Instr. Meth. A* **324**. 288-296.

Maxwell, J.A., Campbell, J.L. and Teesdale, W.J. (1989). The Guelph Software Package. *Nucl. Instr. Meth. B* **43**. 218-.

Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N. and Teller, A.H. (1953). Equation of State Calculations by Fast Computing Machines. *J. Chem. Phys.* **21.** 1087-1092.

NAG (1993). *The NAG FORTRAN Introductory Guide Mark 16.*

Naylor, D.A., Lermer, N. and Furniss, I. (1991). Deconvolution of Fabry-Perot Spectra. *Infrared Phys.* **31**. 401-408.

Parker, S.P. (ed.) (1984). *Mc-Graw Hill Concise Encyclopaedia of Science and Technology.* Mc-Graw Hill.

Press, W.H., Teukolsky, S.A., Vetterling, W.T. and Flannery, B.P. (eds.) (1992). *Numerical Recipes in FORTRAN* (2nd ed.) Cambridge University Press. Cambridge.

Prozesky, V.M., Przybylowicz, W.J., van Achterbergh, E., Churms, C.L., Pineda, C.A., Springhorn, K.A., Pilcher, J.V., Ryan, C.G., Kritzinger, J., Schmitt, H., Swart, T. (1995). The NAC nuclear microprobe facility. *Nucl. Instr. Meth. B* **104.** 36-42.

Puetter, R.C. (1995). Pixon-based multiresolution image reconstruction and the quantification of picture information content. *Int. J. Image Sys and Tech.* **6.** p314.

Robinson, D.R.T. (1991). Maximum Entropy with Poisson Statistics in *Maximum Entropy and Bayesian Methods.* W.T. Grandy and L.H. Schick (eds.). Kluwer. 337-341.

Ryan, C.G. and Cousens, D.R. (1992). *GeoPIXE: PIXE Microanalysis Software System, Users Guide.*

Schaffer, J.P., Shaughnessy, E.J. and Jones, P.L. (1984). The Deconvolution of Doppler-broadened positron annihilation measurements using Fast Fourier Transforms and Power Spectral Analysis. *Nucl. Instr. Meth. B* **5.** 75-79.

Shannon, C.E. (1948). *The Mathematical Theory of Communication.* Bell Systems Technical Journal. **27**. 379.

Shannon, C.E. and Weaver, W. (1949). *The Mathematical Theory of Communication.* University of Illinois Press. (Urbana).

Skilling, J. (1988). The Axioms of Maximum Entropy in *Maximum Entropy and Bayesian Methods in Science and Engineering (vol. 1).* G. Erickson and C.R. Smith (eds.). Kluwer. 173-187.

Skilling, J. (1989). Classic Maximum Entropy in *Maximum Entropy and Bayesian Methods.* J. Skilling (ed.). Kluwer. 45-52.

Skilling, J. (1992). Fundamentals of MaxEnt in data analysis in *Maximum Entropy in Action.* B. Buck and V.A. Macaulay (eds.). Clarendon Press. 19-40.

Smith, G.C., Park, D and Cochonneau, O. (1995). Maximum entropy reconstruction of compositional depth profiles from electron microprobe microanalysis data. *J. Microscopy.* **178**. 48-55.

Sprenger, D and Anderson, O. (1991). Deconvolution of XPS Spectra. *Fresenius J. Anal. Chem.* **341**. 116-120.

Strauss, C.E.M. , Wolpert, D.H. and Wolf, D.R. (1993). Alpha, evidence and the entropic prior in *Maximum Entropy and Bayesian Methods.* A. Mohammad-Djafari and G Demoments (eds.). Kluwer. 113-120.

Tirira, J, Serruys, Y and Trocellier, P. (1996). *Forward Recoil Spectrometry: Applications to Hydrogen Determination in Solids.* Plenum. (New York).

Tribus, M. (1969). *Rational descriptions, decisions and designs.* Pergamon Press. (USA).

Tribus, M. and McIrvine, E.C. (1971). Energy and Information. *Scientific American* Sept. Iss. 179-184.

Tsallis, C. (1988). Possible Generalization of Boltzmann-Gibbs Statistics. *J. Stat. Phys.* **52.** 479-487.

Van Cittert, P.H. (1931). *Z. Phys.* **69** 298-308.

Von der Linden, W. (1995). Maximum Entropy data analysis. *Appl. Phys. A* **60**. 155-165.

Von der Linden, W. Dose, V and Fischer, R. (1997a). Spline based adaptive resolution image reconstruction in *Proceedings of the Maximum Entropy Conference 1996.* M. Sears, V. Nedeljkovic, N.E. Pendock and S. Sibisi (eds.). 154-163.

Von der Linden, W. Dose, V and Fischer, R. (1997b). How to separate the signal from the background in *Proceedings of the Maximum Entropy Conference 1996.* M. Sears, V. Nedeljkovic, N.E. Pendock and S. Sibisi (eds.). 146-153.

Watt, F. and Grime, G.W. (eds.) (1987). *Principles and applications of high-energy ion microbeams.* Adam Hilger. (Bristol).

Wolpert, D.H. and Strauss, C.E.M. (1995). What Bayes has to say about the evidence procedure in *Maximum Entropy and Bayesian Methods.* G. Heidbreder (ed.). Kluwer. 61-78.

XSYS (1985). A General Purpose Data Acquisition System, Indiana University Report.