The hybrid designer: Testing, evaluation, and identification of additional features

FRANK HOCHMUTH

April 1998
ENERGY & DEVELOPMENT RESEARCH CENTRE
University of Cape Town

Contents

2.2.1 10 individuals

Individuals	10	Duration (days)	1	Max number components allowed	D: 2 C:2
Mutation rate	0.4	Convergence crit.	0.01		B: 20 P: 20 I: 2 W: 2
AC demand day	0	DC demand day	150	Components in database	D: 2 C:1
AC demand night	0	DC demand night	150		B: 3 P: 2 I: 1 W: 2

Results

Time (secs): 49	Generations: 40	Sizing: OK	Components: not OK

Result discussion

HD found a range of solutions during various runs. These comprise PV systems, PV/wind systems, and also systems with chargers for no reason. It was impossible to get the same system twice. This proves that the population size of 10 is too small. Setting the convergence rate down to 0.0004 (instead of 0.01) had no major effect.

The lowest life-cycle costs encountered were around R33 000.

2.2.2 50 individuals

Individuals	50	Duration (days)	1	Max number components allowed	D: 2 C:2
Mutation rate	0.4	Convergence crit.	0.01		B: 20 P: 20 I: 2 W: 2
AC demand day	0	DC demand day	150	Components in database	D: 2 C:1
AC demand night	0	DC demand night	150		B: 3 P: 2 I: 1 W: 2

Results

Time (secs): 1'45"	Generations: 37	Sizing: ok	Components: ok

Result discussion

Several runs produced various systems, which indicates that the population size was still somewhat small for the problem applied to it. The systems, however, were all much more sensible than they were with the population size of 10. The systems costs were all between R29 300 and R33 000. Most systems were chosen to have three modules and between 11 and 18 batteries. The cheapest system however was a 600 W wind turbine with one battery only.

The process was much faster than with a population size of 250.

2.2.3 250 individuals

Individuals	250	Duration (days)	1	Max number components allowed	D: 2 C:2
Mutation rate	0.4	Convergence crit.	0.01		B: 20 P: 20 I: 2 W: 2
AC demand day	0	DC demand day	150	Components in database	D: 2 C:1
AC demand night	0	DC demand night	150		B: 3 P: 2 I: 1 W: 2

Results

Time (s): 11'42"	Generations: 43	Sizing: ok	Components: ok

Result discussion

Several runs produced the same system. The system uses three PV modules and quite a high number of batteries (ten). The batteries are cycled only very shallowly and are replaced every four to five years. The reason can be found in the type of battery used: the number of full cycles the battery can take is only 111.

The small number of modules chosen is due to the fact that the irradiation data for the day chosen was very good. The same run carried out for a longer period (with spells of worse irradiation conditions) produced systems with more modules. This testing, however, was not intended during the testing sequence 'variation of the population size'.

2.2.4 Summary

While a population size of ten is definitely too small for the problem it was applied to, higher population sizes produced good results. There are many debates internationally about what population sizes are required for certain problems. At population sizes of 250 the software produced good and reliable results.

The readership is herewith referred to the publications of Prof D E Goldberg form the Illinois Genetic Algorithms Laboratory (http://www-illigal.ge.uiuc.edu/). The institute has produced a number of papers and reports on this matter, most of which can be downloaded from their server.

2.3 Complex optimisation runs

2.3.1 Long-term sizing

The run was made over a simulation period of three months, using a population size of 300. Furthermore, an AC demand was introduced. Otherwise the parameters remain as the above.

This run took several hours (up to 24) using the computing equipment described in this report. Ideally, it would be useful to run the optimisation over one or several years, which should be possible in the near future considering the technical improvements taking place in personal computers.

Individuals	300	Duration (days)	90	Max number components allowed	D: 2	C:2
Mutation rate	0.4	Convergence crit.	0.01		B: 20 I: 2	P: 20 W: 2
AC demand day	300	DC demand day	150	Components in database	D: 2	C:1
AC demand night	300	DC demand night	150		B: 3 I: 1	P: 2 W: 2

The system came up with a PV-wind hybrid system with reasonable sizing. No useless components were chosen.

It should be noted that while complex systems involving diesels as well can be optimised, the simulation and optimisation of systems including diesels takes much more time than systems consisting of renewable energy sources and batteries only. This can largely be ascribed to the additional number of parameters coming from the diesel operating strategies to be optimised

2.3.2 Diesel operating strategies

To test the functionality of the diesel operating strategies, a simulation with a manually given system configuration was completed first. The upper boundary value for the diesel to be switched off was 90% of the battery SOC, while the lower boundary value was 60% SOC.

While much more comprehensive tests were performed, for simple demonstration purposes the simulation here was done over one day which takes less that one second of simulation time. The system consist of four 105 Ah batteries, a charger, an inverter, and one diesel of 10 kW.

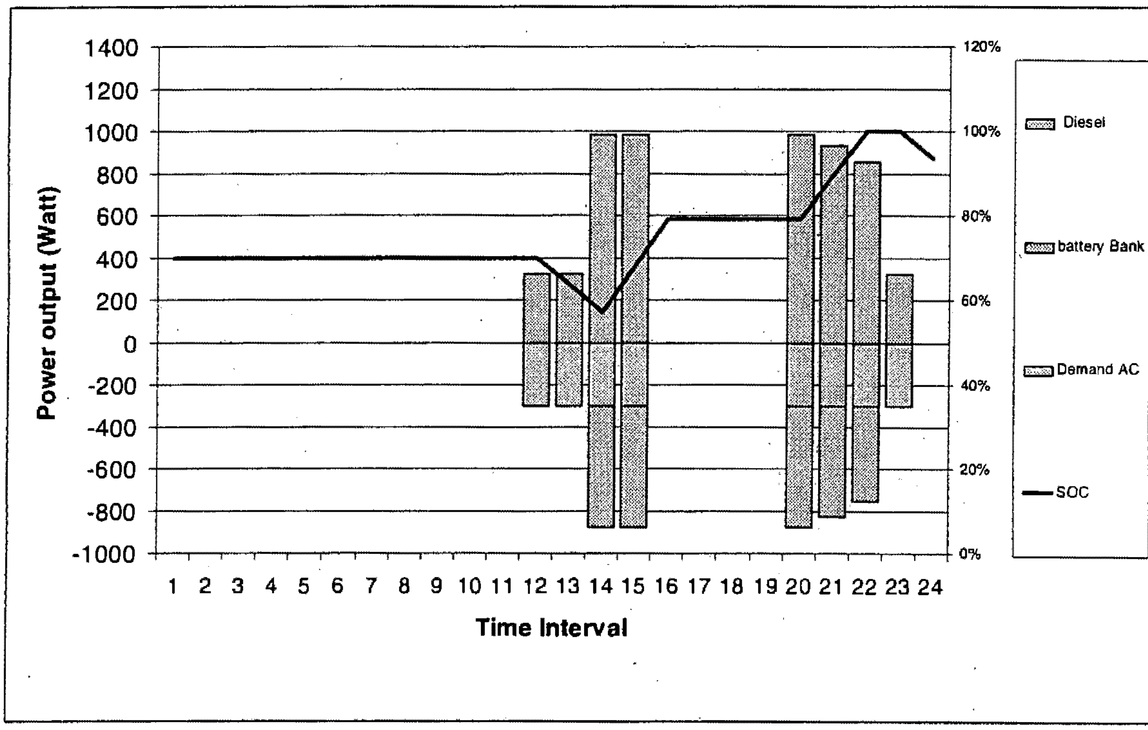

Figure 1: Demonstration of the operation strategies through simulation of a given system.

Figure 1 shows that when the first demand occurs (noon) it is covered initially by the batteries exclusively. This causes the SOC to drop until at 14:00 the SOC is below the 60% threshold which causes the diesel to be switched on. Thereupon the diesel not only covers the demand but also charges the batteries at the maximum allowed current (in this case 12A). It should be noted that the diesel for some reason did not switch off between the two demand batches but ran idle without creating any electricity output. (The problem was caused by a minor bug in the software that was subsequently removed.)

The next step was to let the software find the best operation strategy by itself as part of the design process. A run was performed over one entire week for that purpose. The system selected by HD consists of 11 batteries, one diesel generator and the necessary amount of chargers and inverters. The upper boundary value for the diesel to be switched off was 72% of the battery SOC, and the lower boundary value was chosen to be 58% SOC.

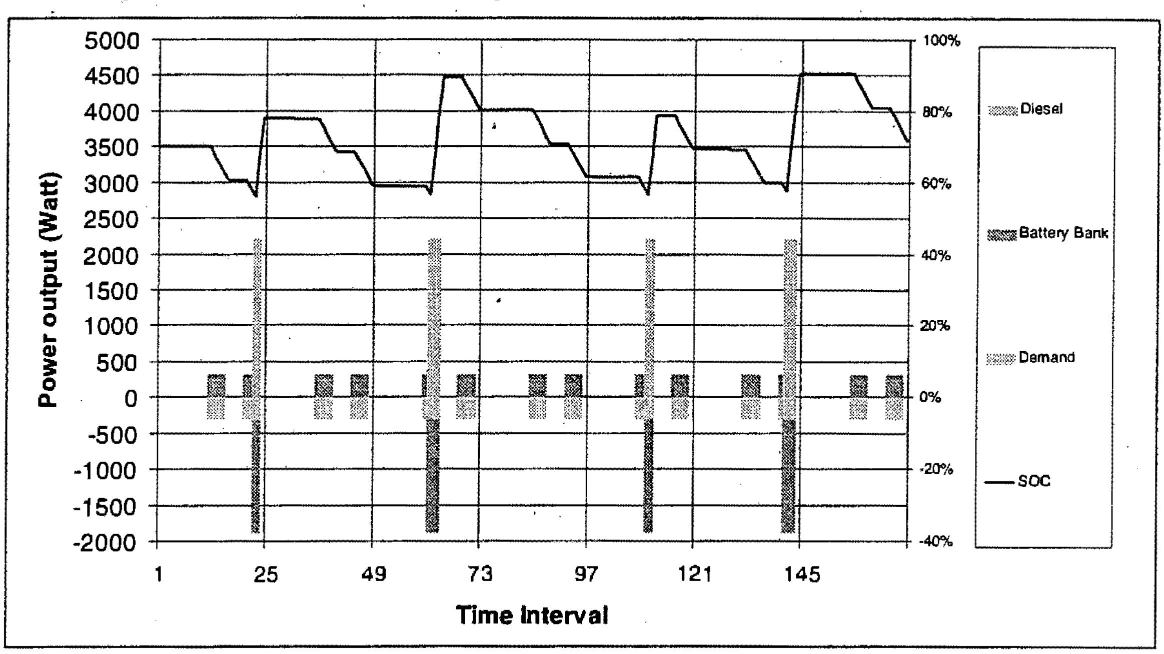

Figure 2: Optimisation of the operation strategies over a 7 day period

Figure 2 shows the effect of those boundaries on the system operation. At the beginning the demand is covered by the batteries, until the SOC drops down to 56% which causes the diesel to switch on. The diesel then stays on until the 72% boundary is reached (because of the one-hour time step of the simulation, the SOC at the end of the hour is actually 78%).

During the whole week of the simulation/optimisation cycle, the diesel is switched on four times only. These all occur at times when demand is not zero – i.e. the diesel is being used to cover the demand and recharge the batteries at the same time.

2.4 Repeatability of the results

If the chosen population size is high enough and the convergence criterion is chosen as zero, the systems can be reproduced exactly. However, the last generations of the optimisation will only develop marginal differences to their predecessors and therefore the optimisation can be stopped before that in order to save time.

It would be helpful for the user to have a better overview and knowledge about the validity of the results in relation to the input parameters. This applies particularly to the population size. At present, the only way for the user to assess the validity of the results is to compare the so-called 'winning solutions' within one run, or to have several consecutive runs which should all bring the same results.

The time required for the simulation can also vary substantially in this regard, depending on the randomly selected initial population and on the random mutation and recombination during the run.

2.5 Speed

The processor of the computer used for the simulations was a Pentium 150 Mhz CPU. This technology has already been overtaken by much faster processors that are commercially available at reasonable prices.

Using the Pentium 150, it was possible to perform sizing/simulation runs using a part of the database with 8 different components (two wind turbines, two diesels, one inverter, one battery charger, one battery, one PV module). A total number of 88 components (two wind turbines, two diesels, two inverters, two battery chargers, 40 batteries, 40 PV modules) were allowed.

The simulations gave reasonable results after about 40 generations, with population sizes of between 300 and 500. The runs took about 36 hours with the hardware equipment mentioned above.

Pentium 300 MHz chips are already commercially available. Intel has presented a 700 MHz Pentium II chip and has announced the release of a 450 MHz chip later this year. The Hybrid Designer software will be very suitable for these computers.

The system runs substantially faster towards the end of the optimization procedure, which means that redundant systems are not being calculated again and again.

The speed is also greatly influenced by the type of components that the algorithm considers for the system configuration. While batteries and renewable energy sources take less time, diesels, chargers and inverters appear to need much more time.

2.6 Stability

The latest versions of the HD ran with very good stability through the test period. Old solution files sometimes had to be deleted prior to opening the scenario file, but this is not a serious problem.

For unknown reasons, it was not possible to create a smaller version of the HD at first. The versions were about 5 Mb in size and contained the debug information which is not necessary for the end user at all. However, a program file of this size does not present a problem.

A smaller version of about 1.8 Mb is now available and runs without problems.

2.7 Test with large numbers of components

Tests with large numbers of components should be avoided where possible. One can easily imagine that if some 50-100 different components are being selected, the number of permutations can easily go into billions of different system configurations. The only way to reduce these to a set of optimal solutions is to have huge population sizes. This is, however, very time-consuming and not recommended here.

It is recommended that the user select from the modular components one PV panel type and one battery type only. For the non-modular components such as wind and diesel, several components with various sizes in sensible ranges should be selected. By selecting fewer than 10 to 15 different components, HD will be able to optimise the system using population sizes of 300-1000, depending on the type of components and other factors.

3. Problems and faults

A range of problems and faults was identified during the testing. All of these have now been isolated and fixed. Amongst these are:

- a large gap between in-depth costing and quick costing – this problem has been fixed;
- differing results when updating the component database – this problem has also been fixed.

Some minor problems are not fixed as yet. These are predominately the fault of the C++ development kit, or can be considered as 'cosmetic':

- The 'Save as' option in the file pull-down menu is not always clickable.
- The progress indicators for population evaluation and simulation continue when the 'in-depth analysis of output' is being performed.
- Some checkboxes only show grey if checked. This does not affect the performace of the software.

4. Additional features

Recommended features (additions and alterations) that it would be worthwhile to implement in the future include:

- Demand entry window: make ADD button the default button.
- A penalty for battery SOC being higher or lower than initial (between certain ranges).
- Output: Number generations, time used.
- Battery lifetime to be calculated according ti SOC levels.
- Title of the scenario should be the filename, not a separate title which causes confusion .
- Listing of components in the database should be grouped according to types of components.
- Display (or saving to a special file for later processing) the evaluation of the LCC and the system components as they develop over the iterations.

5. Progress

This section aims to verify the progress that has been made according to the specifications given (terms of reference for van Kuik).

Context-sensitive help

This part of the project was more difficult and required more time as expected. Van Kuik experienced major problems with the documentation for creating help files using the C++ development kit. Correspondence with experts via the Internet assisted in the creation of help files which are available now.

Speeding up of the code

The code was accelerated through the elimination of various unnecessary program instructions and the simplification of some objects. In particular the handling of input and efficiency data by the inverters, battery chargers and batteries was improved.

Documentation and User's Manual

The draft of the manual was handed in by E van Kuik after some delay on 15th April.

Adding additional features

The following features as outlined in the TOR were added:

- Hellman factor/height of wind turbine: has been added and tested. The impact of the rotor height works satisfactorily.
- PV temperature dependence: has been added and tested. The problem was to obtain a formula that would give the output of the PV module in dependence on the irradiation, the environment temperature. Such a formula was found through a listserver participant responding to a request send out (Thanks to Joe McCabe).

Furthermore, the following additional features that have not been outlined in the TOR were added:

- File opening and storing progress windows.
- The operating strategies for the diesel generator take substantial time to be optimised. It was therefore decided to limit the OPS to only adopt values in certain steps (e.g. 10%), and to limit the range the values can take to reasonable limits.

6. Conclusion

The software works well with some minor problems still to be sorted out and some additions still to be made. Hybrid designer can be used to design and simulate hybrid energy systems using genetic algorithms.

It is recommended that the system be used by trained experts only. While it has been made very user-friendly using a Windows 95 interface, the degree of sophistication and the nature of the genetic algorithm make some training for the user necessary in order to be able to fully exploit the program's abilities. Alternatively, the user's manual could include an in-depth section about the impact of the various options on the validity and the time required by the run.

The system is not supposed to be used with more that one battery type, or more than one panel type.

7. Recommendations

Convergence

The convergence criterion does not seem to be very well defined since, after the convergence has been reached, the winning solutions do sometimes still differ substantially from each other. It would be useful to give the users some means of obtaining an overview of components the individuals of the last generation (i.e. when the criterion was met) consisted of.

Furthermore, it is essential for the user to have some indication about the relationship between and the impact of various parameters (e.g. the convergence criterion, the population size, the maximal allowed components, the number of selected components, etc.) on the accuracy, validity, and the repeatability of the optimisation runs.

File storage

For longer period runs (e.g. three months), the software needs a lot of time for saving the solution data. This can take up to 20 minutes and more. It is recommended that a more efficient file format and handling will be used in future work on the software.

One way to shorten the time for storing the solution file is to reduce the number of winning solutions from say, five, to one. This, however, would mean that the user of the software could not verify the degree of convergence by comparing the life-cycle costs of the winning solutions. In case a more efficient file saving procedure cannot be found it is recommended that the values of the life-cycle costs be stored without the entire data related to the solution.

Weather data

It would be advantageous if the simulation period could be selected at any time within the weather data. At the moment, this can only be achieved by creating new weather data files specifically for this purpose.

User instructions

The user should be instructed specifically about how to use the program and to explore its abilities to the highest levels possible. This applies particularly to the genetic options (e.g. what population size is sensible in relation to the given problem), and also for selecting components. It would be advisable to include a brief tutorial in the user's manual.

The user should also be advised about not using the program for too difficult tasks, such as running it with a large database over extended time periods. It is advisable to only use one type of battery and one type of module in the database at a time.

Stop now button

It is recommended that the 'Stop Now' button of the progress indicator window should not complete the current population evaluation cycle, but rather discard it. Long runs can take a long time for the cycle to be completed, which the user might not want. On the other hand, if the cycle time is rather short, the user can simply wait until the current cycle is completed.

Interim results

It would be useful if the interim results for all the best individuals could be viewed while the simulation is still running.

Time required

The user should be given an indication about the time that will be required for the optimisation run in advance. This could be accomplished by a rough estimate provided by the programme. The user manual should, furthermore, indicate the exact relationships between time required for the run, the population size chosen, the number of components selected from the database, the maximum number of components that is allowed to be used, the simulation period, and other parameters that have an impact on the running time.

Operating strategies

The operating strategies consist of five parameters. While the first two (diesel switch in dependence on SOC) are being used and greatly influence operation strategies, the other three take much time to be 'optimised' but are not really used since they hardly ever come into effect. It might, therefore, be worth investigating the possibility that these three parameters are not optimised but set by the user manually.

EDRC REPORT SERIES

The hybrid designer: Testing, evaluation and identification of additional features.

Frank Hochmuth

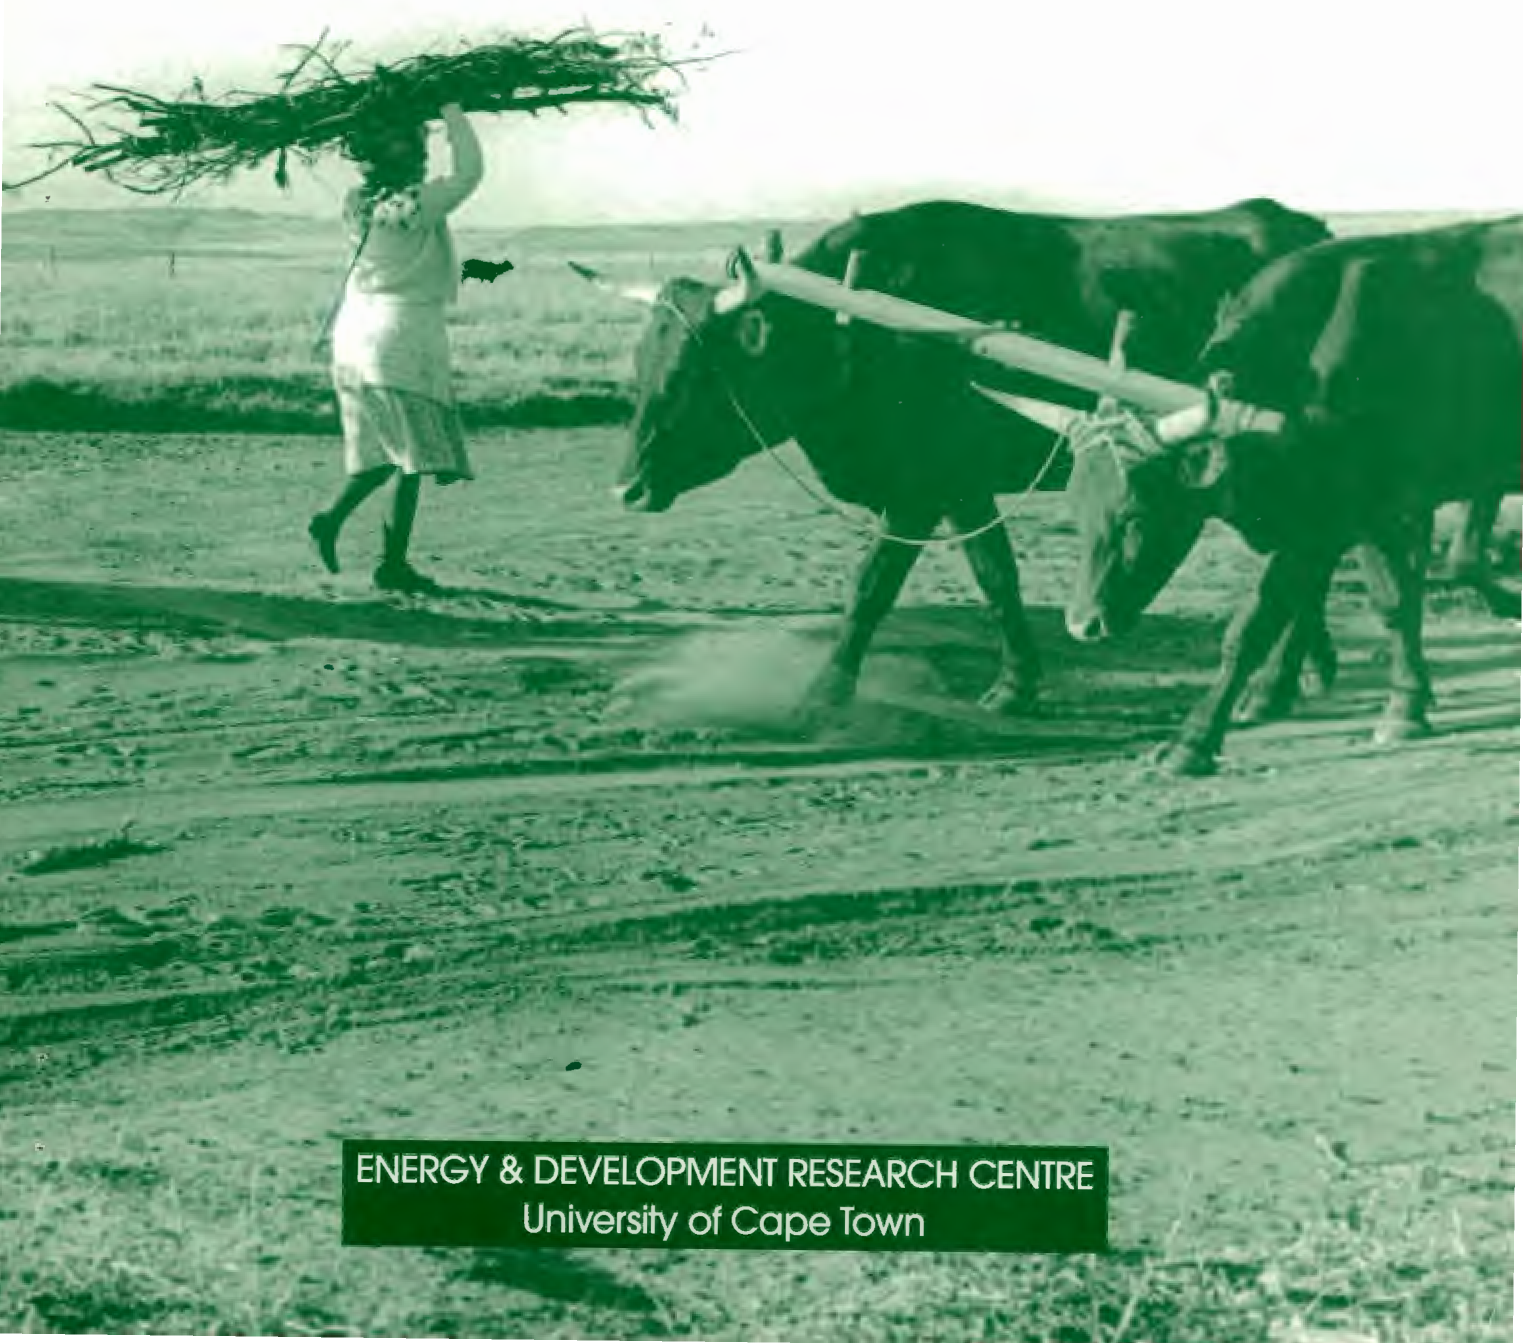

ENERGY & DEVELOPMENT RESEARCH CENTRE
University of Cape Town