AN ASSESSMENT OF THE MIKE SHE HYDROLOGICAL MODEL FOR APPLICATION IN SOUTH AFRICAN CATCHMENTS

Final Report to the Water Research Commission

by the

Institute for Water Research
Rhodes University
Grahamstown
South Africa

D A Hughes
July 2000

IWR Report No. 4/2000
WRC Report No. 747/1/00
ISBN 1 86845 624 2

Table of Contents **Page**

EXECUTIVE SUMMARY

MIKE SHE is the commercial version of the distributed, physically-based, surface-ground water interaction model developed over the last 15 to 20 years by various European hydrological and hydraulic research agencies. It is marketed by the Danish Hydraulic Institute (DHI) and operates under either UNIX or Windows operating systems. This project was carried out over a two year period from January 1996 to December 1997 and the main reasons for the delay in producing the report were related to the difficulties in generating any sensible simulations with the MIKE SHE model. It became clear that the project team under-estimated the resources and time that would be required to become familiar with this model and to be able to make use of it in a practically (from a water resource management point of view) beneficial way.

The objectives of the project were to develop experience in the use of the model, to evaluate the model with respect to the general availability of information to establish parameter values and the applicability of the model formulation to South African conditions and to assess the applicability of the model to the solution of a range of water resource related problems in South Africa. The project teams ultimate task was to make recommendations regarding the future use of the model in South Africa.

The evaluations were based upon several example applications, ranging from a small catchment in the Eastern Highlands of Zimbabwe where afforestation had been demonstrated by the use of other models to have reduced runoff, to a 766 km^2 catchment in the headwaters of the Sabie River, Mpumalanga. The latter represented a reasonable test of a potentially typical application in South Africa, where a good understanding of the hydrology had already been gained through the application of other models, but there is no detailed, experimental information available on soils, geology and vegetation, etc. (i.e. the only data available are what might be expected to be generally available for most South African catchments).

One positive conclusion reached is that the model, and associated software, contains certain features that make it relatively easy to establish a model set-up for a particular catchment. These include graphical display and editing features that are fairly flexible and allow spatial distributions of either time series input data or model parameter values to be established relatively efficiently. This is quite an important feature in a fully distributed model. However, being able to set-up a model in a relatively efficient way does not necessarily mean that it will generate useful results.

One of the main problems experienced was with the run control parameters, which are not part of the hydrological set-up of the model, but rather part of the mathematical processing - something that is very difficult for a new user to develop an understanding of, yet is critical if the model is to run successfully and not take an excessive amount of time. A further problem is that there are a number of components of the model that are essentially empirically based and that under certain conditions the physically-based options in the model are not valid and it is almost essential to resort to the alternative empirical methods. Advice from DHI suggests that this type of approach is required when the topography is steep, while many of the questions related to land use change impacts on water resources in South Africa are associated with afforestation in steep topography areas. It is difficult to see how this model offers South African hydrologists any advantage over methods that have already been developed locally.
A major issue that is related to the points raised in the previous paragraph is that the manual

does not give any adequate guidelines to either setting the run control parameters or when to use the empirical formulations rather than the physically-based ones (i.e. the range of conditions over which the model assumptions are valid). The project team did not find that the suppliers of the software were very quick to respond to questions and therefore some problems took a great deal of time to resolve - a factor that was not considered when designing the project to run over a two year period. It is possible that attending a training course with DHI would have prevented some of these problems from arising, but that would make the establishment of the model in any organisation a great deal more expensive than the already high cost of purchasing the software.

From a hydrological point of view, the model seems to illustrate the problems of scale in hydrological modelling that were clearly demonstrated by Dr Keith Beven in the mid 1980s. The problems are related to the application of physically-based hydraulic equations at the catchment scale that describe flow through porous media and can be readily proved to work at small scales. It is almost impossible to adequately define the characteristics of the media at the larger scale and therefore very difficult to apply the equations as they were designed to be applied. A comparison between applying the model to the small Zimbabwe catchment (1 km^2) and the much larger Sabie catchment (766 km^2) further illustrated this problem. While, being far from straightforward to set-up, the model was found to be relatively successful in the small catchment. However, no sensible simulations could be achieved for the Sabie without resorting to a completely empirical method of representing the catchment hydrology.

The development of surface hydrology models in South Africa has followed the more conceptual modelling route, where the distribution system has been based on catchments rather than grids and detail has been added through the representation of internal catchment process variation through statistical distributions, rather than explicit modelling of those variations. These models have worked very well and have been demonstrated to be sensitive to the types of issues that are important in South African water resource management (land use changes, climate variations, patterns of water use, etc.). It has also been demonstrated that these models can be developed to keep pace with the changing demands of water resource management and that other individuals and organisations than the those involved in their development can be trained in their use. While no models have been developed in South Africa that attempt to explicitly link surface and ground water in the same way as the SHE model, there is also a great deal of experience in the country of the development and use of ground water models.

With respect to an organisation that is seeking a model to use for long term water resource planning and management of a relatively large drainage basin (a catchment management agency, for example), the choice of any model will require a large investment of time, money and expertise. It is possible that the relatively high cost and training requirements of the SHE model could be justified under such situations and it is very difficult to make a firm recommendation on the basis of this projects experience. This is largely because the project team under-estimated the resources that would be needed to carry out such a complete evaluation - a further illustration of the capacity that would be required to effectively install the model for future use. There is little doubt that the model has the potential to be a valuable tool for water resource management in that it incorporates most of the components that would be required. However, local developments in South Africa over the last 20 years or so have also generated similarly useful tools and there is a stronger, and existing, experience base in these tools. The question then becomes one of opting to develop new experience, and relying upon an overseas agency to continue with developments that are relevant to South African water resource management problems, or to direct and support further developments in local products, learning from South African experience as well as overseas developments. If the latter is the preferred option, it should be noted that it will be necessary to maintain South African capacity in model development (as well as application). This can only be achieved through support of the research organisations that have a successful record in this area and through cooperative research efforts between the developers and the end users (the Department of Water Affairs and Forestry and the future Catchment Management Agencies).

ACKNOWLEDGEMENTS

The project was funded by the Water Research Commission of South Africa under the dircetion of Mr Hugo Maaren, to whom the project team are grateful for the support and guidance given during the project and for his chairmanship of the Steering Committee. The other members of the Steering Committee are also thanked for their contributions and constructive comments :

Mr T Coleman, formerly of the University of the Witwatersrand
Prof. CA James, of the University of the Witwatersrand
Mr J Smithers, of the University of Natal, Pietermaritzburg
Mr C Swiegers, of the Department of Water Affairs and Forestry

Dr Vladimir Smakhtin and Mr Karim Sami were both part of the project team during the time that they were part of the Institute for Water Research, while Ms Deidre Watkins of the Institute contributed through assistance with GIS and the conversion of ARCINFO data to SHE spatial coverages. Mr Andrew Groom, an honours student in the Department of Geography at Rhodes University contributed to the project by testing the aquifer/river exchange module of the model on the Crocodile River.

Finally, Henrick Sørensen, Thomas Clausen and Anders Refsgaard from the Danish Hydraulic Institute provided some support and suggestions when the project team had problems.

1. INTRODUCTION

This project was proposed in 1995 and designed to be carried out over a two year period from January 1996 to December 1997. The main reasons for the delay in generating the report were related to the difficulties that the project team had with getting any sensible simulations using the model and the relatively slow response of the model developers to questions. There was therefore little to report at the end of the project, apart from the information that was contained within the Progress Report submitted during May 1997. Since then there have been some further developments, which although relatively minor in terms of the final conclusions, nevertheless allow the project to be concluded and this report finalised.

1.1_ Project Motivation

MIKE SHE is the commercial version of the distributed, physically-based, surface-groundwater interaction model developed over the last 15 to 20 years by various European hydrological and hydraulic research agencies. It is marketed by the Danish Hydraulic Institute (DHI) and operates under either UNIX or WINDOWS operating systems. It is designed to represent an example of the category of models that are complex, strongly-physics based and fully distributed (using a grid approach) and is therefore quite different to models that have been developed within South Africa. The model incorporates detailed surface water components with unsaturated zone functions and full three dimensional saturated sub-surface components. Additional functionality is added through GIS based pre- and post-processing routines for data input and results analysis, as well as incorporating (as optional extras) water quality and sediment production routines.

It has been tested and validated in a number of different climate and physiographic zones worldwide and there has been a body of opinion in South Africa that this type of model could be the answer to many of the countries water resource information requirements in preference to some of the simpler approaches that have been used in the past.

Because the model is complex, it has large data requirements and can take a substantial amount of time to set up, even for relatively small catchments. It is also expensive software to purchase if it is to be used commercially (i.e. by consultants). It was therefore considered essential that the usefulness of the model, in the context of the South African situation, be evaluated before the model is considered for widespread use within the region. The usefulness should be evaluated not only in terms of its applicability to South African hydrological and water use conditions, but also in terms of the general availability of the required data and the resources needed to collect and use such data.

It was always accepted that the large resources required to successfully set up a model of this type will mean that existing modelling approaches, developed and used within South Africa, will continue. However, with an increasing recognition of the need for integrated catchment planning and management and the sensitivity of our water resources to adverse impacts, there are likely to be future situations where the use of more detailed modelling approaches is justified. Some potential future uses of the model were identified at the time of proposing this project:

- ❑ For use where less physically-based, calibration type models have such wide potential error margins that the results may be inconclusive (e.g. some lad use change assessments situations, assessing the impacts of rainfall stimulation, etc.).
- ❑ For use where complex surface/groundwater interactions are important.
- ❑ For water quality (surface and groundwater) modelling.
- ❑ To provide a standard against which simpler and quicker to establish models could be evaluated.

Inevitably, these potential uses rely on the model being truly successful and applicable within the South African context.

1.2 Project Aims and Objectives

The objective of the project were originally stated as follows :

- ❑ To develop experience in the use of the MIKE SHE model
- ❑ To evaluate the MIKE SHE model with respect to :

  General availability of information to establish parameter values.
  Applicability of the model formulation to South African conditions.

- ❑ To assess the applicability of the MIKE SHE model to a range of water resource related problems in South Africa and to make recommendations regarding its future use.

1.3 Intended Work Programme

A major part of the project was designed to concentrate on determining the amount and type of information that is required to set up the model, to what extent there is sufficient information available in South Africa to satisfy the requirements and determining the influence on the simulation results of not providing sufficiently accurate parameter data. Associated with this part of the programme, it was expected that the usefulness and general accessibility of existing databases for setting up the model would be evaluated.

Some of the catchments that were to be included for assessment were intended to be selected from those that the IWR already had some experience of applying other models to. Others would be selected on the basis of their suitability for testing specific components of the model and because of the availability of the required data.

It was initially uncertain whether the water quality and sediment modules should be evaluated under the proposed project. As will be seen from later comments in this report, problems with achieving a satisfactory level of success with the water movement modules precluded any attempts to evaluate the quality and sediment routines. The sediment production module had not reached a very advanced stage of development at the time of the project and therefore was not considered for assessment.

2. MIKE SHE - SUMMARY OF TECHNICAL DETAILS

2.1 General

It is not the intention of this report to provide all the details of the operation of the model and of the model algorithms. However, it is important to establish the technical context in which the discussions about model testing in South Africa can be viewed. Most of the descriptions provided in this report relate to the version (5.2) that the IWR purchased for testing. Versions 5.3 and 1999 are now available. While the IWR is able to run the version 5.3 with the hardware key supplied for Windows 95 operation, it was not possible to run the 1999 version. No comments can therefore be made about the latest version.

2.1.1 Summary of Overall Structure

There are three main manuals for the model; User Guide, Pre- and Post-processing Module, User Guide, Water Movement Module and Technical Reference Manual. The structure of the manuals is indicative of the structure of the modelling package, in that setting up the model and viewing the results are quite separate processes from running the model.

MIKE SHE is a fully distributed model operating with a horizontal grid that can be established by the user, but is fixed in size over the area of a specific project. The pre-processing module includes GIS type facilities to establish grid coverages of various model parameters and variables.

2.1.2 Hardware Platforms

The original version of the model that was received by the IWR was designed to operate on a UNIX workstation and the Windows version (95 and NT) only became available towards the end of the project. It is very unfortunate that the Windows version was not available earlier as this would have avoided the relatively high cost of the UNIX hardware. The Windows version of the model has not been specifically written for Windows but executes under a X-server program eXceed and therefore looks identical to the UNIX version. This could be something of a problem for some users who are familiar with the Windows environment, which is quite different in some respects to the UNIX environment.

2.2 Interception and Evapotranspiration

The interception function is based on the Rutter model using canopy drainage and storage parameters that are linked to leaf area index (LAI) values. The parameters can vary over space and time using ground cover indices. There are two evapotranspiration models referred to in the manual, the Kristensen-Jensen model and the Penman-Montieth. However, the K-J model is the only one that is available in the version of the software provided to the project team. The K-J model is a relatively standard approach to estimating evapotranspiration relying upon a series of empirical constants, root depth and distribution, LAI, potential evaporation and soil moisture content. Root depth can be made to vary over time, in the same way that many of the SHE model parameters can.

2.3 Surface Runoff

Surface runoff takes place via ponding when the net rainfall rate exceeds the infiltration capacity of the soil. There is a detention storage capacity parameter as well as roughness coefficients for both overland and channel flow. The flow routing is controlled by the surface topography (through the elevation of each grid square) and is modelled by approximations of the St. Venant equations of continuity and momentum. Flow into channels is determined where a river segment occurs between two overland flow grids. The relationships between river channel information and the surface topography of the grid is discussed in section 2.8.2 below.

2.4 Unsaturated Zone

The unsaturated zone modelling is through solutions to the one-dimensional Richards' equation using parameters for the soil moisture tension/content relationship, unsaturated hydraulic conductivity as a function of soil moisture and an empirical bypass ratio of net rainfall (to simulate macro-pore flow).

Soil types and their hydraulic properties are initially entered into a soils database, from which the user may select when defining soil (or unsaturated) profiles as soil types over a depth range. These properties include the saturated hydraulic conductivity, the exponent in the relationship between hydraulic conductivity and capillary pressure, volumetric soil moisture contents at saturation and effective saturation, residual soil moisture content, capillary pressure at field capacity and wilting point, threshold value of capillary pressure (a computational control parameter) and the moisture contents at various points on the moisture retention curve. The ACRU manual could be useful for establishing a set of soil properties for commonly occurring South African soils, while at present only a few relatively simple soil definitions have been entered. It is also possible to define how many vertical calculation grids are used in each soil profile. It is also possible to define the spatial extent over which defined soil profiles apply and whether or not the unsaturated zone model is calculated in all the grids, at specific spatial points, or using an automatic scheme that attempts to optimise the number of calculation points with respect to modelling efficiency. Further details can be found in the user manual.

2.5 Saturated Zone

The saturated zone is defined through a number of vertical layers and lenses in much the same way as surface topography (through elevations of the vertical boundaries in a spatial coverage). The saturated zone is modelled in fully three-dimensional flow through the non-linear Boussinesq equation which is solved using an iterative finite difference technique. One of the problems experienced is that the definition of the unsaturated zone is through various 'soil' profiles, while the saturated zone is defined through geological layers and their associated hydrogeological parameters. When the phreatic surface is expected to fluctuate over a wide range of depths, it would seem to be important to ensure that the two zones have compatible hydraulic parameters, but they are defined in very different ways in the model. This makes setting the model up quite complex, if the two definitions are to be reasonably consistent with each other. If they are not, it is apparent that iteration problems can arise during periods when the water table level is changing quite rapidly and the model tries to decide how far to move the water table upwards (into the UZ) or downwards (out of the UZ) for a given increase or decrease in water volume. It is thought that, in some of the examples tested, the number of iteration necessary to perform these calculations is partly responsible for excessive run times.

The vertical discretization in the model can be controlled in several different ways, either using the geological layers, uniform vertical discretization, or through explicitly defined lower levels. There are also methods of defining various boundary conditions.

Drainage flow can also be simulated using an empirical formula, a level for the drain and a routing time constant. The drainage function allows more rapid drainage from the saturated zone than though the hydraulic conductivity and is suggested as a way of expanding the channel network beyond that which is specified directly.

2.6 Aquifer-River Exchange

The aquifer-river exchange component allows for two options where the river is in full contact with the aquifer, or where a low permeability river bed separates the river from the aquifer. The correct operation of this component clearly relies upon an adequate definition of the elevation of the river bed with respect to the surrounding topography and phreatic levels.

2.7 Channel Hydraulics and Routing

Channel cross-sections can be added at a number of points on the channel network and roughness coefficients defined for each of these. These cross-sections are then used with approximations of the St. Venant equations to carry out the flow routing. The newer versions of the SHE model appear to include the use of MIKE 11 (the channel routing package) for dealing with channel flows.

2.8 Input of Data

In general terms it was found to be quite straightforward to understand the formats of the files required by the MIKE SHE model and write small conversion programs to translate existing information into the required formats. This applies to time series data as well as most of the spatial data.

2.8.1 Catchment Configuration and Topography

Although a digitising program is supplied with the model, it was found to be easier to digitise contours (or surface elevation, or geological layers) using ARCINFO and then to export the data in a suitable format (which then had to be further converted using an additional simple program). Facilities provided within the pre-processing module of MIKE SHE allow point data to be converted to matrix (or grid) data and for the grid data to be displayed and edited. These facilities were found to be quite valuable and relatively easy to use. There are a number of facilities that allow new surfaces (grid values) to be established from existing ones using a variety of operators.

The topographic surface is commonly defined by the input of raw elevation coordinates for a series of points within the area of interest, a pair of origin coordinates, a grid size and a search radius that the interpolation algorithm uses to determine the elevation to represent each grid point. The specification of the grid size is usually a compromise between including sufficient detail and the amount of time taken to run the model (see comments about runtime for typical applications).

The grid size may also be related to the original source of information, of which there are several possibilities in the South African context. Perhaps the most obvious, but also the most labour intensive and time consuming, is digitised coordinates of contour lines taken from 1:50 000 or 1:25 000 (or other scale) topographic maps. There are facilities provided by MIKE SHE to carry this out on a DOS based PC and the IWR has also added its own facility. The final result, in terms of the 'accurate' representation of the topography in the SHE grid, will clearly depend upon the relationship between the grid size, the scale of the source data and the contour interval(s) used. This procedure has been used for the Erin and Bedford catchment setup using 1:5 000 and a grid size of 50m and 1:50 000 and a grid size of 100m, respectively, for the two catchments.

One alternative that was considered for relatively large areas was the use of the 1'x1' elevation coverage (available from the CCWR) for the whole of South Africa. This was initially used for the Sabie setup, after a program had been written to convert the longitude-latitude data to x, y coordinates in m or km from a defined origin. Final interpolated grid sizes of 1 and 1.5 km were tried (an increase in resolution from the original data). The result appeared quite sensible when viewed on its own, but when overlayed with the channel network, digitised from the 1:50 000 topographic maps, revealed some serious problems. One of the more obvious problems was that the channel network did not always follow the lowest grid squares and in places was up to two squares away. A further problem was that the catchment boundary grids were not always higher than other close-by grids falling within the catchment. This would be likely to cause problems with boundary effects and generate surface outflows from the catchment at headwater boundary points. Yet another problem occurred with respect to the relative elevation of the channel network and the topography (discussed later).

Two other sources of gridded data are available. The first is the 200 x 200m DTM available from the Trig. Survey (at a cost) and the 800 x 800m DTM available through the Internet from the USGS. Mr S Lynch of the Department of Agricultural Engineering, University of Natal, kindly passed on to the IWR the USGS data for the Sabie area and this was used to create a SHE topographic coverage with a grid size of 1km. The data were also supplied to us in a latitude-longitude format and required conversion to LO coordinates, but a similar program as used for the 1'x1' grid data was used. The improvement was very obvious and when the channel network was overlaid, there were far fewer points where the topography and channel were not in phase. It is clear that the 200m grid data would also have worked successfully.

2.8.2 Channel Position, Cross-Sections and Slope

The channel pattern is established through a set a digitised points and associated elevations (which was generated using ARCINFO and then converted in this project). The MIKE SHE software then allows the user to establish points at which cross-sections are entered, node points that define the connectivity of tributaries, boundary condition points and flow stations which can be for monitoring simulated flow, or for entering observed flows.

There are some good points about the facilities for establishing the river network, but there are also some weak points. The biggest problem experienced by the project team was that the channel network and the topography are established independently and there are no automatic checks to ensure that they are compatible with each other. It is therefore quite possible to have a river elevation point that lies above the surrounding topography. While the model reports this as a problem when the set up data files are created before running the model, there is no straightforward way of correcting it. For example, it would be very useful to be able to overlay the channel network with the topography grid and identify channel points that are above the adjacent ground surface or far below it.

Any of the river-aquifer exchange procedures will be strongly affected by the head differences between the water in the river and that within the aquifer. Given that the topography and the channel network can be defined from different sources and are almost inevitably defined with different resolutions, there are distinct possibilities of problems occurring in this part of the model.

If a relatively coarse grid is used in an area that has quite steep topography, particularly in the vicinity of rivers and their valleys, the differences in elevation between the channel bank points and the surface topography (derived from averaging elevations over the grid square) can be quite substantial. In the case of the Sabie when the 1'x1' grid source data were used, several points were found where the channel bank elevation was above the surface. In all cases, in the Sabie setup, there are long sections of channel which are as much as 100m and more below the defined surface topography. The extreme cases occur where valleys are narrow, relative to the size of the grid, and valley sides are steep. It should be pointed out that there are no easy to use facilities in the MIKE SHE pre-processing module to check these type of situations. The potential for a problem to arise was mainly revealed when the model was run and rather strange ground water responses resulted. It was initially not clear why the ground water was being drawn down to 100m below the surface, until it was realised that the channels were at such levels in many parts of the area.

A comparison between the Sabie and Erin examples can be used to illustrate the problem of scale and the way in which the model is setup. The Sabie was established with a 1km grid, while the in the Erin example a 50m grid was used. Typical relief variations within a grid square in the vicinity of river channels in the Sabie catchment are between 80 and 200m, while in the Erin catchment they are less than 15m. If it assumed that the ground water level is close to the surface in the near channel grids in both cases, then the gradients of flow to the Erin channel are greater than the Sabie. However, the elevation differences are far greater, suggesting that the problems occur with the compatibility of the channel elevations with respect to the ground surface and the definition of the geological layering. In the Erin setup the best results were obtained with a relatively highly transmissive surface layer of about 5m (weathered granites) underlain by a layer with much lower hydraulic conductivities. This allowed a rapid response of baseflow to occur during the wet season when ground water levels were elevated and a slower, more sustained response to occur during the dry season. Applying a similar approach to Sabie did not work and this is thought to be related to the fact that the ground water level is always drawn down into the lower, less transmissive, layers in the grid squares adjacent to the channels. This seems to limit the possibility of achieving a relatively rapid wet season baseflow response.

2.8.3 Time Series Data

Rainfall data, for as many stations or areas as required, are input as intensities (mm h^{-1}) with an associated date and time in a relatively simple text format. A program has been written to re-format HYMAS distributed rainfall data files for use with SHE and can be used for fixed interval data (usually daily), or variable interval, where available. The areas over which the individual rainfall records apply are usually input by digitising the boundaries of the relevant polygons. Alternatively, a grid editing program supplied with SHE can be used to create the coverage of rainfall distribution codes. There is no procedure for catchment, or sub-area, averaging of station rainfall data within SHE and therefore this must be done externally (the HYMAS routines have been used in all examples established so far). While it is accepted that accurate reproduction of short term runoff responses are unlikely without the disaggregation of daily rainfall data, no attempt has yet been made to develop such a methodology. The only example where less than daily interval data has been used is the Bedford (Eastern Cape) catchment.

Potential evaporation data files have the same format as the rainfall data files and as many records (individual stations or areas) as required can be used together with a coverage of the grid squares over which they apply. All the input data used to date have been taken from mean monthly S-pan values.

2.8.4 Parameter Data

Input and editing of parameter data within the model is relatively efficient and can be handled through the spatial distribution facilities, or through simple tables on the screen.

2.9 Run Control Parameters

These are used to control the iterations that are used within the model to solve the differential equations. However, there are no clear guidelines provided in the manual and it is not at all clear how these should be set to achieve a compromise between accurate modelling (in terms of limiting water balance errors) and run time length. The setting of these values has turned out to be one of the most difficult problems to resolve, particularly given the lack of information provided. It is worth noting that a user would have to have a relatively detailed knowledge of the internal workings of the model to be able to understand how best to set these values.

2.10 Output and Analysis of Results

There are a number of different options available for assessing the results, including both the retrieval of input data and results to the screen or to output files as well as graphical display of the input data or results. While in general terms, a large number of options are available that might be considered to cater for most users requirements, many of them end up being slow and clumsy to use. There are facilities to save graphical layouts and retrieve them for later use, which helps a lot given the number of options that have to be set to establish a graphical presentation. Part of the problems lies with the amount of time taken to display a graph.

2.11 Manual and Support Material

There are parts of the manual that are quite good and more than adequate to gain a sound understand of how to operate the system. However, there are also many gaps and weak points that make it very difficult for the user to find the information that is required to carry out a simulation exercise successfully. Most of the problems relate to tracing errors which are not necessarily fatal, but certainly mean that the results are worthless. The most difficult (and frustrating) aspect of applying the MIKE SHE model was the fact that very few guidelines exist for correcting a model set-up that does not run properly. Given the amount of time taken for even relatively short time series to be simulated, a trial and error approach to changing run-control parameters or model definition parameters is simply not feasible.

3. ASSESSMENT OF SIMULATION CAPABILITIES

The first progress report of the project referred to some of the problems that the project team were experiencing with establishing a satisfactory model setup for the Erin catchments in the Eastern Highlands of Zimbabwe (the first example tried). While a more satisfactory result was eventually obtained (largely through trial and error evaluations) for these catchments, the team did not get very much closer to a real understanding of the way in which the model operates. Part of the problem lies in the lack of detail provided in the manual, part with the amount of time the model takes to complete a run for even a few years and part with the relatively slow response to questions put to the developers at DHI.

Despite these problems it has been possible to approximately reproduce the general characteristics of the response for the main Erin catchment, largely through trial and error changes to the geological model setup and the associated hydrogeological parameters. It has not proved to be simple to adjust the parameters in the smaller gauged sub-catchment to simulate the somewhat different baseflow response that occurs there. The fact that the simulated flows do not adequately reflect the short term peak responses is hardly surprising given that the input rainfall data used have a daily resolution.

The report will attempt to highlight some of the characteristics of the model setup and results achieved for the examples used. These are the Erin catchments (Eastern Highlands, Zimbabwe), a channel and floodplain reach of the Crocodile River (Northern Province), Sabie River (X3H006, Mpumalanga) and one of the Bedford catchments (Eastern Cape).

3.1 Erin catchments - Zimbabwe

The Erin catchments (total area of about 1 km^2) were setup using a digitised contour map at a scale of 1:2500 and the grid size was set to 50m. Two geological layers were established, an upper one of a uniform 5m depth to represent weathered granite and a lower one of 110m to represent less weathered granite. The depth of the lower one was set large enough to ensure that this layer never dries out. It was found necessary to set the hydraulic conductivities (horizontal and vertical) of the upper layer to very high values to achieve a reasonably sensible response. The horizontal value was set at 36 and the vertical at 180 mm h^{-1}, while the specific yield was set at 0.03. The equivalent values for the lower layer were 0.72, 0.72 and 0.001. The vegetation parameters were not properly calibrated but probably reflect grassland (the conditions at the start of the simulation period) rather than forestry (planted over about 0.7 km^2 in 1980/81). A single soil profile was established with 0.3 m of sandy clay loams overlying 1.5 m of sandy clay. Although some soil descriptions are available, they are not really detailed enough to allow the spatial variations to be adequately defined.

Figure 1 illustrates the results for the main catchment outlet as well as the pattern of aquifer flow to the river and overland flow for a single valley bottom grid square. The general pattern of baseflow response has been reasonably well simulated, except in wet years where a prolonged response occurs throughout the year and the seasonal recession is over-simulated. The model is also not able to simulate the many short term events and this is mainly due to the lack of finer time resolution rainfall data than 1 day.

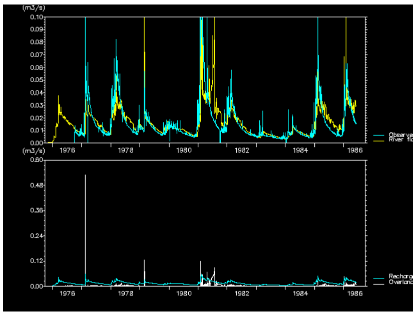

Figure 1 Example of Erin catchment results

There is no indication in the results that the afforestation programme has much effect, although most of the effect can only be seen later in the record. The main problem with the simulation is the time taken for the model to run for a relatively small catchment and only 400 grid squares. Most of the time appears to be in the saturated zone calculations and the number of iterations is very high, with the maximum number exceeded regularly. This also seems to lead to quite high water balance errors. The problem may be related to the large differences in hydraulic conductivity between the two geological layers, but the research team are not certain about this. The next step in this example is to try and decrease the number of iterations and attempt to achieve a similar result with more realistic upper layer hydraulic conductivities. After that it should be possible to concentrate on other aspects of the model, such as the vegetation parameters.

Some feedback was obtained from DHI on this model setup and they are summarised below.

- They suggested that the iteration stop criteria for the saturated zone was too low and that a higher value (0.001 m) is more appropriate to a transient model.
- It seems important to set the maximum amount of rainfall in a single time step to a value such as 5 or 10mm to prevent too many time steps being used.
- There were apparently some problems with the original definitions of the unsaturated zone soil parameters.
- They recommended the use of drains to simulate some fast runoff.

- DHI changed the downstream boundary condition in the channel component.

- They suggest that the unsaturated zone classification that was used was too coarse.

- They commented on the small values that had been used for the lower layers and yet these were considered to be relatively un-weathered granites and typical values were obtained from experienced ground water hydrologists.

- One of the critical issues that was raised was that they indicated that they would not like to use MIKE SHE for a catchment with such steep slopes (15-20%) and that this catchment would be on the limit for the validity of the St. Venant equations. They suggest that when they use the model in steep areas they do not try and apply the model in a physically-based way, but use the drains component.

3.2 Crocodile River - Northern Province

The data for setting up this example was taken from a WRC Report (113/3/87) which provided some information about the character of the alluvial sediments adjacent to a 50 km reach of the Crocodile River in the Northern Province and some cross sections of the valley bottom and channel. An upstream gauging station provided a time varying boundary condition for the channel flow calculations and the model was run with the express purpose of learning more about the operation of the river-aquifer exchange component. The example formed part of a project by a Geography Department Honours student (Mr Andrew Groom). Initially the model was setup with all components operational (including the UZ, evaporation, rainfall input, overland flow, etc.), but problems with getting the model to run to completion under such circumstances led to a simplification. The model setup was then confined to the UZ, SZ and the channel flow, with zero rainfall input and no evaporation. Figure 2 illustrates the results for four locations along the channel (a full colour version will be available at the meeting)..

The alluvial material consists of layers of clay material in a generally sandy gravel surround and the purpose was largely to see how the spatial position of these affected flows to and from the channel. In general terms, the results appeared to be satisfactory (figure 2), although the lack of a downstream gauge prevented the total channel losses simulated by the model (which were quite small) from being assessed. The base of the channel is frequently close to the bedrock, which has been assumed to be largely impermeable. There is also no allowance for abstraction from the alluvial material. It is therefore not surprising that very little water is actually lost from the channel reach as a whole, it is mainly redistributed in time and space. The spatial redistribution is largely controlled by the position of the clay lenses, which restrict exchanges in both directions. It was found that there was not really enough information at a satisfactory resolution provided in the report to readily set the model up (neither the geological/alluvial layers, nor the channel cross-sections) and that a great deal of subjective interpolation between the available data points was necessary. This suggests that even relatively detailed survey information is not sufficient to fully describe the physical characteristics of such an area for the purposes of running SHE with even a moderate degree of confidence.

Despite these problems, a variety of tests run over a 3 year period with different hydraulic conductivities for the clay and sandy alluvium materials illustrated the operation of the aquifer-river exchange component.

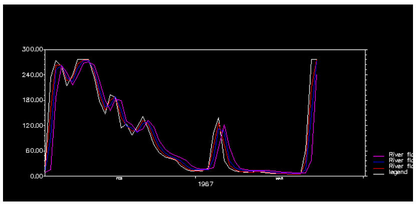

Figure 2 Streamflow at four locations within the Crocodile River channel

3.3 Bedford catchment - Eastern Cape

One of the Bedford catchments (about 40 km^2) has been set up using the contour data that were digitised during an earlier WRC project and a topography coverage created with a 100 m grid. The rainfall data were extracted from the data files that form part of HYMAS and are based on variable interval data (5 minutes to 1 day) generated for 9 sub-areas based on some 5 raingauges. The soils information was based upon the rough surveys carried out by the IWR and includes an approximate indication of the spatial variation in depth. The model was run for a five year period but the results were not examined in detail as there is very little observed flow data (few events and some data missing). However, the model certainly simulates very few events and the scale of streamflows appears to be approximately correct. The next stage in the testing of the model in this catchment was going to be to decide how best to set up the geological layers for an area of alluvium in the bottom end of the catchment which is known to absorb a great deal of runoff generated upstream. However, problems experienced in other test cases suggested that this would not provide the team with any conclusive results.

3.4 Sabie catchment - Mpumalanga

The Sabie catchment (X3H006 - 766 km^2 with three other gauged sub-catchments) was set up as a test of a typical medium sized catchment within South Africa where only commonly available data would be available, but where the project team have quite extensive experience of applying other models. The topography data for this example have already been discussed in section 2.8, while the geological layers were defined on the basis of 1:50 000 geological maps. Unfortunately, such maps provide no information about the subsurface extent and characteristics of the rocks and it is therefore difficult to decide how to establish layering. The present set up assumes that the surface lithology extends vertically downwards (certainly not the case in reality) and that the three different rock types (two quartzites and an area of dolomite) each have three different vertical layers with decreasing transmissivity and storativity with depth. At this stage, only a single soil profile has been defined as it was considered more important to establish the saturated zone component of the model and a reasonable reproduction of the long term baseflow pattern, before attempting to add more detail.

Despite many attempts with different combinations of parameter values, it was not possible to generate results which are even close to the observed flows. Part of the problem is that the model took between 6 and 12 hours to run over a 5 year period, making it extremely time consuming to calibrate. For example, most of the runs were carried out with a single computational layer for the SZ and a moderately sensible pattern of baseflow response began to emerge after some manipulation of the hydrogeological parameters. Subsequently, the number of computational layers was increased to four (uniform vertical distribution). The model took over 12 hours to run and the streamflow response was much higher than it should have been and more or less constant. However, on a more positive note, the project team learned a great deal about the problems and pitfalls of setting up the model in this type of situation.

4. CONCLUSIONS AND RECOMMENDATIONS

The MIKE SHE model is a commercial product that is relatively expensive to purchase for groups within South Africa and requires quite a high level of investment of time and effort to generate any kind of acceptable results. It is therefore reasonable to assume that potential users should be able to expect quite a lot in return. Unfortunately, the project teams experiences suggest that the return is not as high as might be expected. This conclusion has been reached for various reasons, not all of which can be blamed on the model itself. However, there are certainly deficiences in the model and its packaging that would perhaps not be expected given the way in which it is promoted. Some of these relate to the model formulation, while others relate to the support material (manuals, etc.) that is available to assist the user. Some more specific conclusions are highlighted below :

- One critical issue is the influence of the run control parameters on the results. These are part of the mathematical formulation of the model, which standard users should not be expected to have to fully understand. An experienced hydrologist may be able to define the parameters of the various hydrological equations that make up the model, but cannot really be expected to understand the influence of the run control parameters without very clear guidelines. These guidelines are, unfortunately, not available in the manuals. It is possible that by attending a training course on the model, such guidelines could be passed on to the users, but this serves to increase the costs of establishing the model.
- There appear to many situations where empirical parameters are used, either as a standard in the component equations, or as an alternative to the more physically-based approaches in data poor situations, or where the model assumptions are not valid (see example of the steep Erin catchments). Once again, guidelines on this issue are not available in the manual and the user has to rely upon input directly from DHI. The initial experience of the project team was that this input was not forthcoming and in fact it was only well after the project had finished that some guidelines were offered by DHI.

- Given the number of empirical type parameters that may be required, it is clear that the model becomes a 'calibration' type model, rather than a truly physically-based model. This is related to the first point above about the run control parameters and becomes a real problem when the run time for even relatively short time series is excessively long - it is simply not practical to make many repeated runs of the model to achieve an acceptable fit with a sensible set of parameters.

- The model formulation is clearly based on hydraulic rather than hydrological principles and everything within the model is forced to fit in with the hydraulic approach. Clearly, this is based on trying to establish a physically-based model that attempts to simulate the physics of water movement over surfaces and through porous media. However, its application then suffers from the scale problems often referred to in the modelling literature and related to the inability of the user to obtain sufficiently detailed data about the surface or sub-surface characteristics to enable the model to simulate conditions accurately. Less rigorous models of catchment hydrology (conceptual, semi-distributed models) have recognised this constraint and many have incorporated components (often invisible to the user) to handle issues of scale and spatial variations that are impossible to account for explicitly. Some of these issues appear to have been accepted by the MIKE SHE model developers and have added empirical components. However, this just seems to make things even more confusing because now we have a physically-based model with additional components to use when the physically-based components don't work. The difficult issue for the user is to decide when this occurs. A comparison of the levels of difficulty in getting reasonable results between the very small Erin catchments and the Sabie are believed to illustrate the scale problem very well. The grid size used for the Erin was 50m, while for the Sabie it was 1000m. At the Sabie scale there appeared to be too many problems related to defining the hydraulic setup of the catchment, but going to a smaller grid size would have meant that the model took far too long to run.

- There are no really obvious indications for the methods that should be used to incorporate water resource usage into the model. The manual does refer to various sources and sinks (return flows and abstractions) that can be added, but it is not very clear how these operate.

- There is little doubt that the model can be made to perform adequately, or even very well, However, for a new user, a very high investment of time and effort is required. While, attending some kind of training course with the developers might help (and would be recommended if a group were to seriously consider the model for use), it is the conclusion of the project team that existing models used in South Africa should be able to offer the same results for less effort. At present there do not seem to be any applications in South Africa where the MIKE SHE model is likely to offer additional benefits. However, this conclusion should be qualified by taking into account the future requirements of catchment management agencies. If they are likely to require an integrated model of surface water, ground water and channel flow routing for long term future use in catchment management, then the investment required might be considered worthwhile. It is not totally clear how reservoirs, water use and abstractions are included within the model, but it is understood that this is possible.

- It is worthwhile noting that the software environment of the Windows version is a conversion from the UNIX version, which will not be considered very 'friendly' by users who are only familiar with the standard Windows environment. This is a relatively minor point, but one that does influence initial impressions.

- Graphics displays and file access seems to be very slow compared to other systems that the project team are familiar with. While it is accepted that this may be less of an issue if the model is to be set up for long term application at a single site, it was a major factor in the ability of the project team to evaluate the model in the limited amount of time that was originally budgeted for.

- One of the mistakes that the project team made was that the evaluation was designed to be operated on a part time basis and with limited time resources. It became apparent after the first year of the project that it would have been much better to have at least one dedicated individual to concentrate on learning the use of the model.

- A final conclusion can be reached in the form of a more general warning about the selection of this type of model. It is relatively easy for any model developer to demonstrate the value of their model when they are using an example that has been established by themselves with their detailed understanding of its operation and internal structure. It is a very different matter to expect a new user to be able to achieve the same level of success, regardless of that individuals expertise in hydrology or other models. The main point is that at this level of modelling (rather than the use of simpler models) the model specific experience of the user becomes more important than the model itself in determining the success of a modelling exercise. This is a conclusion that is often neglected and means that a great deal more emphasis has to be placed on the learning capacity of the organisation proposing to adopt the model.