Reservoir System Operational Optimisation

NJ Manson

Report to the Water Research Commission
by the
Water Systems Research Group
University of the Witwatersrand

WRC Report No 757/1/98

Reservoir System Operational Optimisation

by

NJ Manson

Report to the Water Research Commission

by the

Water Systems Research Group
University of the Witwatersrand
Private Bag 3
WITS
2050

Head of Department
and Project Leader: Prof D Stephenson

Report No 757/1/98
ISBN 1 86845 452 5
ISBN SET 1 86845 454 1

NB. THE REPORT SUPPORTED BY SOFTWARE CALLED "RSPO"

The software emanating from this study called "RSPO"-Reservoir System Pumping Optimisation is available only from the WRC internet site, http:/www.wrc.org.za. Should you not have access to internet facilities or encounter problems in downloading the software, kindly inform us by email or fax, and a copy of the software will be forwarded to you:

Water Research Commission
PO Box 824
PRETORIA, South Africa
0001
Tel.: (012) 330-0340
Fax: (012) 331-2565
email: wrcit@wrc.org.za

EXECUTIVE SUMMARY

Introduction

Service and distribution reservoirs are used for a number of purposes, such as balancing storage, emergency storage and fire storage. In terms of the daily running of the reservoir, the most important component is the balancing storage. This provides a buffer between the consumers' demands and the supply equipment. The simplest form of reservoir operation would be to supply water at the average demand, and provide a reservoir large enough to meet any fluctuations in the demand. This is unfortunately not feasible for the following reasons.

- The average demand is not known sufficiently accurately, if at all.
- Meeting a full years demand fluctuation would usually require a very large reservoir.
- The demand changes with changes, such as urbanisation, in the supply area.

The above reasons, and others, mean that a reservoir system cannot just be built and left to run by itself. Rather, it must be operated on a daily basis.

Problem Statement

Experienced human operators run most reservoir systems in South Africa, and throughout the world, without any formal operation policy. This allows very flexible operation, but may not result in the minimum operation cost. Labour costs are also continuing to rise relative to computer and control equipment costs, so many systems are tending towards automated operation. To automate a system, it is necessary to develop a formal, mathematical operation policy, and this needs to be optimised so that operating costs will be a minimum.

Aim

The aim of this project is to develop a model that can be used to calculate an operation policy for a reservoir system that will minimise the cost of operating the reservoir system. These costs are largely the cost of energy for pumping.

Literature Review

Reservoir system optimisation can be looked at from two different viewpoints. Capital cost or design optimisation looks at finding the combination of pump, supply and reservoir capacities that result in the lowest lifetime cost of the system. This aspect was investigated in depth in the report entitled "Capital Cost Optimisation of Pumping and Reservoir System Design", which forms part of this project.

Operation Optimisation looks at the day-to-day operation of the reservoir system, specifically the control of the pumping equipment, and attempts to minimise the running costs while meeting the consumer demands. This is the focus of this report.

System Components

To optimise the operation of a reservoir system, it is necessary to model three main components. The first component that must be modelled is the hydraulic network. Hydraulic network models fall into three categories.

1. Mass balance models assume that the supply rate from the pumps is constant and not affected by head against which they must pump. They also assume that system pressure requirements are met if the reservoir volumes remain within certain specified ranges. This is the simplest model and therefore the easiest to implement and quickest to calculate, but it is the least accurate.
2. Regression models fit non-linear curves to the system hydraulics. These models are relatively simple and fast to calculate, but provide greater accuracy than mass-balance models. This is the model used in this project.
3. Hydraulic simulation models use a calibrated network model of the system. These are most accurate, but require more computational effort, and significantly more system data than the above models.

The second component of a reservoir system that must be modelled is the consumer demands.

- Lumped models assume all the demands occur at a single point. This is the model used in this project.
- Proportional demand model assumes that each location in the demand draws a fixed proportion of the total demand.
- Distributed demand models attempt to simulate the spatial and historical variation in demands.

The third component of a reservoir system that must be modelled is the control policy. These vary greatly depending on the physical system complexity, the data available, and the objective of the optimisation.

Optimisation Models

The optimisation model consists of an objective function that defines exactly what is the aim of the optimisation process. Most models attempt to minimise the cost of pumping, while other try to minimise the cost of both pumping and switching between settings. The model used in this project allows the costs of pumping, switching and storage to be minimised.

Engineering optimisations are always subject to a number of constraints that can be classified as physical system limitations, governing physical laws and externally imposed constraints. The physical constraints on this model are the available pumping capacities and the reservoir sizes. The model is also constrained by the conservation of mass and energy laws. Minimum and maximum allowable volumes are often modelled as external constraints, but in this model they are not. Rather, financial penalties are imposed for exceeding these limits. Very large penalties would, however, function as constraints.

An optimisation model consists of a number of decision variables. The process of optimisation seeks values for these variables that minimises the objective function.

In this model the decision variables are the reservoir volumes at which to change from one pump setting to another.

Reservoir System Analysis

The reservoir system is modelled as two main components.

Physical System Model

The physical reservoir system is modelled as a pump station supplying a reservoir from which the consumers' demands are met. The demands are lumped into a single global demand, and all the available reservoirs in the system are lumped into a single storage.

Reservoir Sizes and Costs

The reservoir is then modelled by defining maximum and minimum volumes. These limit the amount of water that can be put into the distribution system. Financial penalties are also defined for exceeding these limits. These are physical limits imposed by the existing system, and cannot be changed without major capital expense. High and low warning levels, and associated penalties, are also defined. These should be calculated from the costs and risks of failure, but are usually policy decisions made by the reservoir operators. A storage cost function, which is a linear function of the volume of water in the reservoir, is also defined.

Pumping Capacities and Costs

The pumping capacity of the system is modelled by defining a number of independent pump settings. For example, a system that contains 2 pumps, each of which can be either 'On' or 'Off' will have four possible settings, defined as:

0: Both off
1: Pump 1 on, pump 2 off
2: Pump 1 off, pump 2 on
3: Both on

Variable speed pumps would be modelled as a number of discrete settings.

For each of these settings the output of the pumps is modelled as a function of the volume in the reservoir. The system head against which the pumps must operate is given by $H_{sys} = H_{elev} + \frac{V_R}{A_R} + k_S Q^2$, where k_s is a function of the friction losses in the pipeline, and needs to be fitted by regression. Each pump setting will also have a pump curve given by $H_{pump} = f_p(Q)$. At the operating point $H_{sys} = H_{pump}$, so by combining these two equations an expression of the form $Q = f_Q(V_R)$ can be obtained. This is usually a cubic polynomial.

Similarly, a relationship between the flow delivered and the energy consumed by the pumps can be found. The cost of pumping is calculated by defining a number of time zones throughout the day, each with its own tariff. The amount of energy consumed is multiplied by the applicable tariff to find the cost of pumping.

An important cost in running a reservoir system is the cost of controlling the pumping rate. This is modelled by defining a switching cost for each possible pair of settings.

Operating Policy

The operation policy model used in this project is based on the reservoir volume at the start of a time step, the estimated demand for the current time step, and the pump setting used during the previous time step. The *Expected Volume* is defined as the current volume minus the expected demand. A change level is defined for each pump setting, which specifies the maximum value of the *Expected Volume* for which that pump setting should be used. A typical operation policy is shown in the graph below, which indicates that pump setting 2 should be used if the *Expected Volume* is between 25% and 50% of the reservoir capacity.

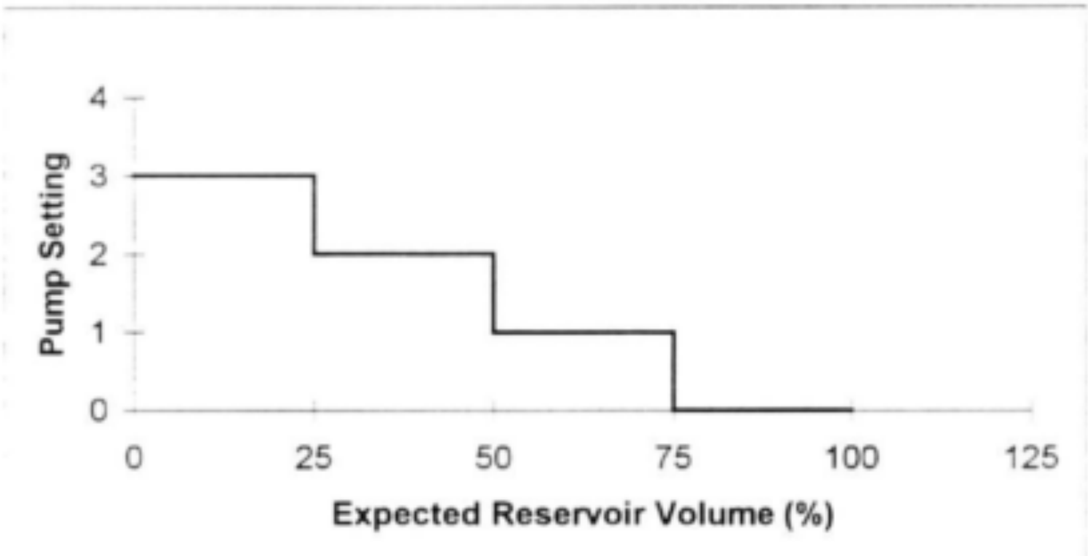

Figure 1: Typical Operation Policy

To minimise the cost of switching from one setting to another, a *Don't Change* variable is also defined. This means that if the new *Expected Volume* is only just outside the range of the current setting, then the setting should not be changed.

Demand Data Model

A statistical model of the consumers' demands is used to estimate the expected demand for the next time step, and to generate synthetic demand data. The total demand for each time step is made up of a number of components. To find the parameters of the model a set of historically measured data is required. It should be at least a year long to model the seasonal variations, but 3 to 5 years is even better.

Secular Trends

These are the trends that continually grow or decline through the whole data set. The model uses a least-squares regression to fit both a linear and a logarithmic trend. The best fit is then used.

Periodic Trends

The periodic trends represent the components that have a regular variation. The model uses a harmonic analysis to fit periodic trends to the seasonal, weekly and daily variations.

Autoregressive Component

The residual component that remains after the secular and periodic trends have been removed is random in nature, but may not be serially independent. A linear autoregressive model is used to represent the serial correlation between a particular demand value and those of the previous few time steps. The order of the model is based on the significance level chosen by the user.

Independent Random Component

The model estimates the mean and standard deviation of the independent random component by finding the values the result in the best fit between the historical data and the calculated data. In this case the fit is measured using a χ^2 test.

Reservoir System Operation Optimisation

The model developed during this project was coded into a software package called *Reservoir System Pumping Optimisation* or *RSPO* for short. The program consists of three main functions, and a number of utility functions. The three main functions are discussed below.

Analysis

The analysis function takes a set of historically measured demand data and fits a demand data model as discussed in the previous section.

Simulation

Simulations are useful to generate synthetic demand data for periods different to the historically available data, and to test the effects of various system changes.

The reservoir system being simulated is described in a *System Description File*, which contains the following information:

- The minimum and maximum reservoir volumes.
- The high and low warning levels for the reservoir.
- The penalties for exceeding the above limits.
- The storage cost function.
- The number of available pump settings.
- The flow rate function for each pump setting.
- The power consumption function for each pump setting.
- The time zones and electrical tariffs associated with each zone.
- The cost of switching between any two pump settings.

The simulation uses this information as well as either historical or synthetic demand data to calculate the pump setting, pump rate, power consumption and reservoir volume for each time step. A report is produced summarising any limit violations, the reservoir and pumping costs and the power consumed.

Optimisation

The optimisation function finds the values for the *Change Volume* and *Don't Change* parameters of the operation policy that result in the lowest total cost. It does this using a technique known as Downhill Simplex Optimisation. This runs repeated simulations, adjusting the operation policy between each run, until a minimum cost is found.

Dynamic Programming can also be used to find the best possible way of operating a reservoir for a particular set of data, but it cannot calculate an optimum policy. This method is provided as a comparison to check the results of the Downhill Simplex optimisation.

Case Study Results

The Rand Water distribution system was chosen as a case study because of its size and proximity. Two simple sub-systems were chosen for initial monitoring, these being the Libanon/Driefontein and the Bloemendal/Wildebeesfontein sub-systems. A third, more complex system, the Mapleton sub-system, was also analysed.

Daily reservoir level and pumping rate information was available for each of these systems. From this data the daily demands for a period of six months was calculated.

Simulation Results

The six months of historical demand data was used to fit a stochastic model. This model was then used to generate six, thirty and sixty months of synthetic demand data. Analysis results on the synthetic data show a good correlation with the analysis results of the historical data.

Optimisation Results

The main purpose of the program is to calculate optimum operation policies, but these calculations can be adversely affected by a number of factors. The effect of each of these factors on the case study optimum was investigated.

Local and Global Minima

Because the Downhill Simplex method is a search technique it stops when it finds a minimum, although that minimum may not be the global minimum. This is a problem with all searching optimisation techniques. It was found that although there are local minima in the solution space, these are small, and the optimums found from different starting points were within 1% of each other.

Effect of Data Set Length

In many optimisation techniques the length of the data available affects the accuracy of the result. To test this, a set of data with 1800 points was generated. From this, four sets were extracted, with 225, 450, 900 and 1800 points respectively. Optimums were calculated from each of the four sets, and the results were found to be within 0.05% of one another.

Effect of Cost Functions

Different optimum policies were also calculated based upon different cost functions, to check that the individual costs had the expected effect. These gave good results, showing that the model reacted as expected to different costs.

A Comparison with Historical Results

The two optimisation techniques were used to calculate optimums using the all the system cost functions and compared with the historical costs. The optimum pumping rates found are shown in the figure below.

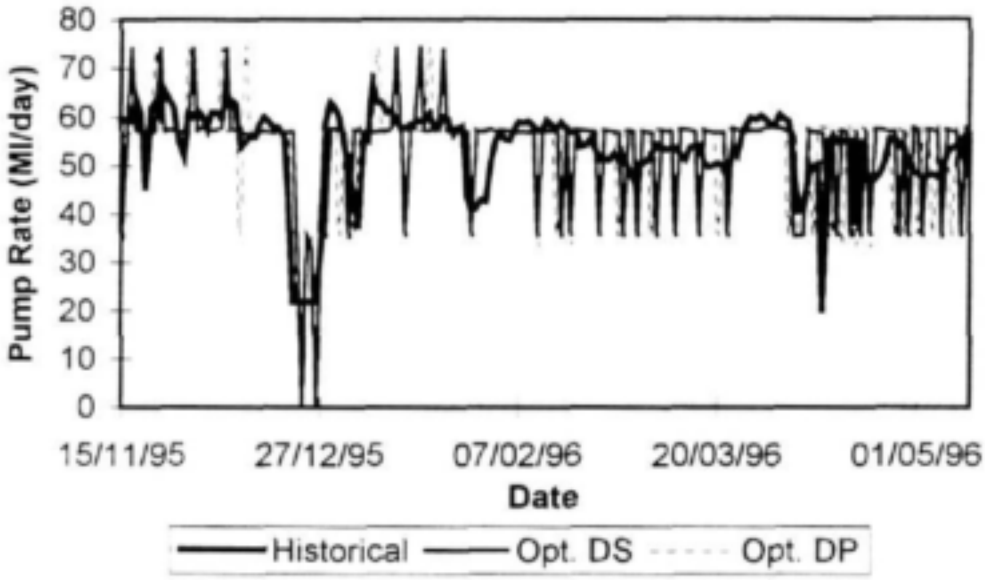

Figure 2: Pumping Rate

The main difficulty in comparing optimum costs to the historical costs is the lack of historical cost data. The only data available was the pumping rate for each day and the reservoir content at the end of each day. The pump setting for each day was estimated from the system head, and this was used to calculate the power consumed and the cost. The comparison between these estimated historical costs and the optimum costs are shown in the table below.

Table 1: Power Consumption's and Costs

Operation Policy	Power Consumption kWh	Cost R	Saving %
Estimated Historical	129113.18	165626.39	
Downhill Simplex Optimisation	126638.31	162451.62	1.92%
Dynamic Programming Optimisation	126080.40	161735.94	2.35%

Although the cost savings were not found to be highly significant, the results do indicate that a simple automated procedure can operate the system as well as an experienced human operator.

Conclusions and Recommendations

The aim of the project was to develop a dynamic computer program to simulate and optimise the operation of a general reservoir system. This was achieved with the completion of *Reservoir System Pumping Optimisation*.

The Analysis option of *RSPO* fits a demand model to the secular, periodic, autoregressive and independent random components of the historical demand. This model is then used to predict the demand one time step into the future, and also to generate synthetic demand data.

The simulation option of *RSPO* allows the user to compare the effects of different operation policies and costs on the total running cost of the system. It also allows the user to generate sets of synthetic demand data for periods different to the historical data.

The simulation option of *RSPO* calculates an optimum operation policy using the Downhill Simplex technique, and calculates the optimum pump settings for a set of data using the Dynamic Programming technique. The effect of local minima on the finding of an optimum solution was shown to be less than 1% of the total costs. It was also found that the optimum solution found was hardly affected by the length of the data set at all.

A comparison with the historical operating costs showed that, although only small savings were achieved, the optimum policy performed at least as well as an experienced human operator.

ACKNOWLEDGEMENTS

The research report emanated from a project funded by the Water Research Commission entitled:

"Development of a Model to be used for the Economic Optimisation of the Pumping and Design Policies of Reservoir Systems".

The steering committee responsible for this project consisted of the following persons:

Mr JN Bhagwan	W R C (Chairman)
Mr HC Chapman	W R C
Dr TC Erasmus	W R C
Prof. EA Moss	University of the Witwatersrand
Mr JJ van der Walt	Magalies Water
Mr I Buchan	Umgeni Water
Mr G Hough	Bloem Water
Mr J Frame	City Council, Cape Town
Mr C Coetzer	Dept. Water Affairs and Forestry
Mr G Constantinides	Rand Water
Mr R Thomson	Rand Water

The financing of the project by the Water Research Commission and the contribution of the members of the Steering Committee is gratefully acknowledged.

This project was only possible with the co-operation of many individuals at Rand Water. The authors therefore wish to thank the following:

Mr KT Owen, Chief Electronic Engineer

Mr M Lines, Regional Superintendent, Operations, East

Mr R Barry, Regional Superintendent, Operations, North

Mr J Kruger, Senior Administration Assistant, Central Depot

Mr D Leach, Metering, Central Depot

Mr M Aspeling, Superintendent, Bloemendal Pump Station

CONTENTS

LIST OF FIGURES

LIST OF TABLES

1. INTRODUCTION

Water supply authorities throughout the world use balancing reservoirs as a buffer between the users' demands, and the pumping required [Nel and Haarhoff,1996]. If there is no reservoir then the pumping rate must continually, and instantaneously, meet the demand. In most municipal areas the demand is never constant, which requires that the pumping rate be monitored and controlled at all times. Traditional water pumps were not very easy to control, and usually only have two settings, 'ON' or 'OFF'. Modern pumps, and electronically controlled motors, are more flexible, but are still expensive to control. With a balancing reservoir in place, the pump can be set to pump at the average demand. Any excess demand is drawn from the reservoir, and if the demand is less than the average, the excess pumping fills up the reservoir.

For this system to work correctly, it is necessary to know the average demand very accurately, and to have a large enough reservoir. If the pump is not set exactly equal to the average demand, the reservoir will eventually either run dry, or overflow. If the reservoir is too small, it will also run dry and overflow on occasions. A reservoir that is too big will not cause any operational problems, but will be unnecessarily expensive to construct.

There are a number of problems with this ideal situation. One is in finding the "average" demand. When a new reservoir is being designed, it is not possible to measure the water demand being drawn from that reservoir. Even if there is an existing reservoir, the average demand that can be measured there will not be quite the same as for the new reservoir. This average demand also varies continually as the number of users grows, and the types of user changes. The average demand also changes from season-to-season, and differs from weekdays to weekends. All these factors mean that there is no way to accurately know what the average demand will be by the time the reservoir is built, and certainly not 5 or 10 years into the future.

Because the demand varies from season to season, it would be necessary to build a very large reservoir if one wanted to keep the pumping rate constant for an entire year. For a summer rainfall area, this would imply storing the rain from summer until it could be used in winter. This is often one of the functions of supply reservoirs, but it would be very expensive to do for a municipal balancing reservoir. Thus a compromise must be made. It will be necessary to change the pumping rate occasionally, but one would like to keep the pumping rate as constant as possible.

Much work has been done on how one selects the optimum size of a new reservoir. The actual size of a reservoir is affected by more factors than just the balancing volume discussed so far. A municipal reservoir also has to provide emergency storage, in case of a supply failure upstream of the reservoir; and it usually has to provide water for fire-fighting. Most reservoirs also have a volume of 'dead water' or bottom storage. This is water at the bottom of the reservoir that is never pumped into the supply system. This is done to ensure that any contamination in the water that has settled to the bottom is not put into the supply. *Nel* [1993] has looked at a number of ways of calculating the sizes of the various components of a reservoir and gives references for further reading.

It is also necessary to design a reservoir for a certain 'design horizon' [Stephenson, 1989]. If one estimates the demand too far into the future, the reservoir will be too big for much of its life, and this will dramatically increase the construction costs. If the demand estimate is too small, or made for too short a time, the reservoir will only have a short useful life. Another reservoir will have to be built in the near future. South African municipal reservoirs are generally designed to store anything from 24 to 76 hours worth of storage of the annual average daily water demand [Nel, 1993]. For example, Rand Water generally design their reservoirs for 36 hour's storage [Stallard, 1980].

Once the reservoir has been designed and built, it has to be continually operated. This operation can be either manual or automatic. The function of reservoir operation is essentially to ensure that the reservoir never overflows, and that there is always sufficient water in the reservoir to meet users' demands and emergency storage requirements. Practically, operation of a municipal balancing reservoir is the process of deciding when, and by how much, to change the pumping rate supplying the reservoir. The set of rules or guidelines that are used to make this decision is known as the pumping or operation policy, or operating rule.

1.1. Problem Statement

Rand Water are currently using pumping guidelines developed in 1980 [Stallard, 1980]. Since that report was written, a number of factors have changed; the network has been upgraded, electrical tariff structures have changed, user requirements have changed, and there have been a number of advancements in the science of optimisation and the available computer technology.

That report also focused on Rand Water's global demand pattern, and tried to optimise the pumping at the primary pump stations of Vereeniging, Zuikerbosch and Zuurbekom. Secondary and tertiary pump stations and reservoir systems were ignored.

Pumping costs represented 16.3% of Rand Water's total expenditure in 1995 [Rand Water, 1995], therefore minimisation of this component is of vital importance. In fact the cost of pumping is one of the largest single components of energy use by local governments. Sadowski, Nitivattananon and Quimpo [1995] estimate that the city of Pittsburgh spent 25% of their total electricity cost on pumping. Ormsbee and Reddy [1995] state that the optimisation of daily pumping schedules is one of the greatest potential areas for cost saving. Reduction in operating costs will result in lower water tariffs. These will affect everybody, as everyone needs to use water. The cost of water also affects the cost of almost everything produced, so a decrease in the cost of water will result in a general decrease in the cost of living for all people.

1.2. Aim

The aim of this project is to develop a model that can be used to calculate a pumping policy for a reservoir sub-system that results in the minimum operating cost for the system. This will be done by developing a mathematical model of the reservoir sub-system, and then incorporating the model into a computer program. It is intended that the program will be flexible and allow the user to modify all the input parameters. This will mean that the program will be useful to many water supply authorities, and not just Rand Water.

The optimum policy generated by the program will tell the reservoir operator which of the available pump settings to use depending on the expected demand and the volume of water in the system reservoirs. It will be an optimum policy in that it minimises the total running costs of pumping water to the system reservoirs to meet the consumers' demands.

1.3. Technology Transfer

The research results of this project have been compiled into a suite of computer programs, namely ***Reservoir System Pumping Optimisation*** and ***Reservoir System Design Optimisation***. The Water Systems Research Group in the Department of Civil Engineering at the University of the Witwatersrand will present a training workshop on these two programs. The Water Systems Research Group will also keep records of all the data used in this project, and this can be made available on request.

A paper based on this research was also presented to the Biennial Congress of the Southern African Division of the IAHR in Sun City in 1996 (Manson, 1996).

1.4. Project Report Overview

This chapter introduced the project, presented a statement of the problem, and detailed the aim of the project. Chapter 2 will look at the background theory of optimisation and what it means to optimise a system. It will also look at the methods that have been presented in the literature for solving similar problems. Chapter 3 covers a number of mathematical techniques that have been used to optimise systems. Chapter 4 discusses the analysis of reservoir systems and discusses how the various components can be modelled. Chapter 5 presents the development of a computer program that analyses, simulates and optimises reservoir systems based on the model developed in Chapter 4. Chapter 6 discusses the results of applying the computer program to a specific reservoir sub-system in the Rand Water distribution area. The results of the optimised operating rules are also compared with historical operating information. Chapter 7 presents the conclusions and some recommendations for further study.

2. LITERATURE REVIEW

2.1. Reservoir System Optimisation

Optimising reservoir systems means different things to different people. The design engineer will think about selecting the optimum capacities for the reservoir storage, the supply pipelines and the pumping equipment. Here the objective would be to minimise the capital cost while ensuring that the constraints are met with a reasonable level of security. The operations engineer, on the other hand, would see reservoir optimisation as the minimising of the cost of operating an existing reservoir system, while still meeting the requirements of the distribution network being supplied.

The word reservoir also has different meanings to different people. The dictionary defines a reservoir as "a place where a great stock of anything is accumulated" [McLeod, 1986]. This can mean dams that are used to provide water for agricultural, industrial, residential, recreational use as well as for flood control. These are usually referred to as supply reservoirs. Reservoirs can also refer to the steel or concrete tanks that are used by local water supply authorities to provide buffer storage between pump stations and the users' demands. These reservoirs are usually referred to as service reservoirs, balancing reservoirs or service tanks.

2.1.1. Design Optimisation

Much work has been done on the optimisation of the design of balancing reservoirs. The traditional method of selecting the size of the reservoir has been to apply, almost blindly, the standards set by water supply authorities [Loubser, 1986]. More recently, however, these standards have been subject to renewed scrutiny. Smook [1985] looked at the daily and hourly demand patterns and their effect on the size of the balancing storage required. Grosman and Basson [1985] also used stochastic modelling and risk analysis to establish size-risk-cost planning aids for the design of balancing storage.

Municipal service reservoirs are also not only used for balancing storage. Nel and Haarhoff [1996] used Monte Carlo simulation to calculate the probability of failure of a service reservoir of a given size. They included demand storage, fire storage, emergency storage, a control provision, bottom storage and operational freeboard in their size calculations.

2.1.2. Operation Optimisation

Increasing urbanisation and demand for water has led to increasing complexity of supply networks [Ormsbee and Lansey, 1994]. This has caused the water utility industry to begin investigating on-line control and optimised pump scheduling [Lansey and Awumah, 1994; Ormsbee and Reddy, 1995]. Because the cost of pumping is one of the largest single components of energy use by local governments, optimising the pumping policies used is one of the areas of greatest potential saving [Jowitt and Germanopoulos, 1992; Sadowski, Nitivattananon and Quimpo, 1995; Ormsbee and Reddy, 1995].

Optimising the operation of reservoir systems, specifically the pumping policy, is the focus of this project, and the work currently available in this area will be discussed in detail in the following sections.

2.2. System Components

The aim of operational optimisation of reservoir systems is to find an operational policy that will result in the lowest operation cost, while still meeting the requirements of the supply system. The requirements would include meeting the consumption demands, maintaining sufficient pressures, providing water for emergencies and others. To optimise a reservoir system, three main components need to be modelled. These are the distribution network, the consumers' demands and the actual control policy. Ormsbee and Lansey [1994] give a good review of the different models that have been proposed for each of these three components.

2.2.1. Hydraulic Network Models

A model of the distribution network is necessary to evaluate the effects of the demands and the pumping policies imposed on the system. These models vary from a simple mass-balance model to full hydraulic simulation models

Mass-balance models

These models generally assume that the pumps in the system have a constant delivery that is unaffected by the head against which it is pumping. This means that the demand is equal to the change in reservoir volume plus the pumping rate. System pressure requirements are assumed to be met if the reservoir volumes are within the required ranges. The main advantage of this model is its simplicity and speed of calculation.

Regression models

These models use regression analysis to fit non-linear curves to the system hydraulics. The data used to perform the regression analyses can be obtained from calibrated simulation models or from measured historical data. These models are also relatively simple and allow fast calculation, but they represent the system hydraulics with greater accuracy than the mass-balance models.

Hydraulic network simulations

The most accurate way of modelling the distribution network is with a full network simulation that has been calibrated to the actual system. This allows the greatest flexibility, and provides the best accuracy, but at the expense of considerably increased data storage and computational requirements.

The exact model used is very dependent on the complexity of the system being optimised, and the formulation of the objective function.

2.2.2. Demand Models

To calculate an optimum operation policy it is necessary to know the demand that must be met. A system can either be optimised for a set of historical demands, or the historical demands can be used to forecast future demands. The demand forecast can be a function of season of the year, the day of the week, the time of day, the forecasted weather conditions and many other factors.

The forecasted demand can be incorporated into the hydraulic network model in one of three ways, lumped, proportional, or distributed.

Lumped demands

In this case the complete demand on the entire network is assumed to occur at a single point. This form of demand incorporation is most suitable for mass-balance hydraulic models.

Proportional demands

In this case the demand at each point in the network is assumed to be a fixed proportion of the total demand. This is suitable for incorporation into a regression type hydraulic model.

Distributed demands

These demands are incorporated into a full hydraulic simulation at each of the nodes from which water is removed from the network. They are allowed vary both in space and in time.

Techniques for forecasting the overall demand on the network are well established, but methods for distributing the demands across the network and throughout the operational period are limited due to the lack of data.

2.2.3. Control Policies

The control policies used in the network determine the model used to optimise the system. These vary greatly depending on the physical system complexity, the data available and the objective of the optimisation. Details of these models are presented in the next section.

2.3. Optimisation Models

2.3.1. Objective Functions

The most important part of the optimisation model is the objective function. This defines exactly what is meant by an optimum for the current system. Because the cost of the energy required to drive the pumps is such a major component of the operating cost of water supply utilities, most models have as their objective the minimisation of this cost [Sadowski, Nitivattananon and Quimpo, 1995]. In many places the cost of energy is made up of two components, a charge per unit of energy consumed (kWh) and a charge for the maximum demand (kVA) for the billing period [Jowitt and Germanopoulos 1992]. Some models, therefore, try to minimise both of these components of the cost. In some areas the electricity tariff is lower during off-peak times of the day, so some models attempt to schedule pumping during these times.

Lansey and Awumah [1994] proposed a model whose objective was to minimise the number of pump switchings as well as the energy cost. Very little work has been done on the effect that optimum pumping policies have on the cost of maintenance, but minimising this cost could also be a model objective [Ormsbee and Lansey, 1994].

2.3.2. Constraints

The optimisation model will always be subject to a number of constraints. These can be classified as physical system limitations, governing physical laws and externally imposed requirements.

Physical system limitations

These include the capacities of the pumps and pipelines in the system, the amount of water available from a source, the feasible pump and valve settings and others [Ormsbee and Lansey, 1994]. Many models include the maximum and minimum reservoir capacities as constraints, but these can also be included as penalty costs.

Governing physical laws

The main laws applicable to hydraulic networks are the conservation of mass at each node, and the conservation of energy around each loop. The equations giving pressure losses along pipes or across valves are also governing laws.

External Requirements

The main external requirement is to meet the demand from the users, although this could also be formulated as a penalty for any under supply. External requirements can also include restrictions on reservoir volumes, restrictions on network pressures, limitations on withdrawals from sources and others.

2.3.3. Decision Variables

The optimisation model attempts to select the values for certain variables that best meet the objective over the optimisation horizon. These variables are known as the decision variables. The choice of which variables are to be the decision variables has a significant effect on the formulation and solution of the optimisation problem.

Some models use a direct approach to pump scheduling. They divide the optimisation period into intervals, and they calculate which pumps should run during each interval, and for what proportion of the interval. Other models use a two-stage approach. A surrogate variable is chosen as the decision variable, for example, tank level or overall pump station output. This method requires that the costs be expressed in terms of the surrogate variable, and that the resultant optimum values of this variable be translated back into pump schedules[Ormsbee and Lansey, 1994].

Roman and Chandramouli [1996] proposed a model that directly calculated the optimum values of the decision variable, and then used these values to train a neural network to derive a general operating policy.

The fact that there are so many different possible objective functions and choices of decision variables results in the proliferation of different model formulations seen in the literature.

2.4. Other possibilities for reducing operating costs

This report has focused on optimising the pumping policy to achieve a reduction in the cost of pumping, but this is not the only route to improved operating efficiency. Walski [1993] presents a number of other possible cost savings techniques.

Minimise wasted water

Water that is pumped from a source into a network where it leaks out and performs no useful purpose is not only a waste of water, but it is a waste of energy as well. It is therefore doubly necessary to minimise leaks and water wastage.

Reduce Heads

Keeping the reservoir levels as low as is acceptable and by keeping the suction wells as full as possible can reduce the pumping heads. Care must be taken that this does not cause the pump to operate at a point further from its optimum point.

Avoid Peak-time pumping

If the tariff for electricity varies during the day, it will be naturally be better to pump as much as possible during the lowest tariff period. This is only applicable if there is sufficient storage in the system.

Check pump efficiency

Even though a pump is operating satisfactorily, in that it is providing sufficient flow and pressure, does not mean that it is operating efficiently. As pumps wear their efficiency drops, and this needs to be periodically checked.

3. OPTIMISATION TECHNIQUES

The previous chapter discussed the formulation of the optimisation problem, looking at the choice of objective function, and selection of decision variables. Once the optimisation has been formulated, it still needs to be solved, which is a non-trivial problem in itself.

To date, much of the work that has been done on developing solution techniques for this kind of problem has been done in the context of optimising the release rules for large supply reservoirs. Municipal balancing reservoirs are very similar to supply reservoirs except that they are the "other way around". Balancing reservoirs have a controllable input, the pumping rate, and a random output, the user demand. Supply reservoirs, on the other hand, have a random input, the rainfall run-off or stream-flow, and a controllable output, the amount of water released. The following sections will evaluate the work that has been published on the solution of operating policy optimisation for supply reservoirs, and relate it to the control of balancing reservoirs.

3.1. Linear Programming

Linear Programming (LP) is a widely used and very popular technique for optimising an objective function, subject to a set of constraints. In the case of a reservoir, the constraints will include the maximum and minimum storage volumes, the maximum and minimum release rates, continuity, equipment limitations and legal or contractual obligations [Yeh, 1985]. The objective function will depend on the main function of the reservoir, for example, water supply, hydropower, recreation, flood control, etc., and will need to be either maximised or minimised.

A typical LP model would have an n-dimensional vector of decision variables. Let this be $\boldsymbol{X}$. There would also be $\boldsymbol{C}$, an n-dimensional vector of objective function coefficients. The constraints on the system would be given by $\boldsymbol{A}$, an $m \times n$ matrix of constraint coefficients; and $\boldsymbol{B}$, an m-dimensional vector of constraint limits. The aim would then be to minimise $\boldsymbol{Z} = \boldsymbol{C}^{\mathrm{T}}\boldsymbol{X}$ for all values of $\boldsymbol{X}$, subject to $\boldsymbol{AX} \geq \boldsymbol{B}$ and $\boldsymbol{X} \geq \boldsymbol{0}$ [Yeh, 1985]. The 'T' indicates the transpose operation. Some models may want to maximise $\boldsymbol{Z}$, in which case the constraints would be given by $\boldsymbol{AX} \leq \boldsymbol{B}$.

LP is a fairly simple, but a powerful technique for solving linear systems. A number of methods have been developed to enable the application of LP to more complex, or non-linear systems. *Meier and Beightler* [1976 (Yeh, 1985)] proposed a branch decompression technique to simplify a system with multiple parallel reservoirs. *Parikh* [1966 (Yeh, 1985)] proposed a system of spatial decomposition and combination with dynamic programming, which was extended by *Roefs and Bodin* [1970 (Yeh, 1985)] to include decomposition over time. *Hall and Shephard* [1967 (Yeh, 1985)] also developed a method that combined LP and dynamic programming. They divided a multi-reservoir system into a master problem, solved by LP, and sub-problems, one for each reservoir, which were solved by dynamic programming. *Windsor* [1973 (Yeh, 1985)] developed a technique that used recursive LP to optimise the flood control of a multi-reservoir system.

There are two main types of LP models, namely deterministic and non-deterministic. Deterministic models calculate a single, optimum result, but do not take into account any uncertainty in the input variables or constraints. They generally use average or mean values, which result in optimistic designs or policies, overestimating benefits and under estimating costs and losses [Reznicek and Cheng, 1991]. The models mentioned above are all deterministic LP models.

Unfortunately, most hydrologic systems contain a significant amount of uncertainty in their parameters. Deterministic models can be useful to evaluate the system for specific values of certain input variables. The uncertainties can also be taken into account through sensitivity analyses, but LP does not explicitly consider uncertainty, and thus may lead to unsatisfactory results [Yeh, 1985]. A number of models have been developed that can be used when parameters, such as stream-flows, are non-deterministic.

3.1.1. Stochastic Linear Programming

In a simple, single reservoir system, where the reservoir has a single main function, the main element of uncertainty is the inflow to the reservoir. These inflows are assumed to be random and serially correlated. *Loucks* [1968] used a first order Markov chain to model these inflows. The probability of the inflow in a particular period changing to another given inflow in the next time period was calculated from the historical information. From these inflow probabilities, the probability of the reservoir volume changing from the current value to a new value in the next period can be calculated from the continuity equation. This information can then be used to choose the optimum release rule for the reservoir, depending on the function of the reservoir. The main shortcoming of this model is the problem of dimensionality, as the number of constraints can exceed several thousands in real situations [Loucks, 1968].

3.1.2. Stochastic LP with Recourse

This is a slightly more complex model than the one described above, but it permits the incorporation of random variables in the constraint set, and allows the decision to be made in multiple stages, usually two. *Prekopa* [1980, (Reznicek and Cheng, 1991)] describes the model as follows:

$$\underset{X}{\text{minimize}} \quad Z = C^{T}X + E\left(\underset{Y}{\min}\,\{G^{T}Y\}\right) \qquad \text{(Eqn 3.1)}$$

subject to $\boldsymbol{AX} = \boldsymbol{B}$,

$\boldsymbol{TX} + \boldsymbol{WY} = \boldsymbol{P}$, and

$\boldsymbol{X} \geq \boldsymbol{0}$, $\boldsymbol{Y} \geq \boldsymbol{0}$,

where $E()$ is the expected value operator,

A is the matrix of constraint coefficients,
B is the vector of constraint limits,
X and Y are deterministic vectors of decision variables,
C and G are vectors of cost coefficients,
T and W are coefficient matrices that may contain random elements, and
P is a vector of random right hand sides.

To start with, a vector X^* is chosen such that $AX^* = B$, where X^* is chosen from a set K of X vectors which all have at least one corresponding Y vector that satisfies

$TX + WY = P$ for whatever T, W and P are realised. The cost of the decision X^* is $C^T X^*$.

The second stage of the solution takes place after the random events T, W and P have been realised. A recourse action Y^* is chosen such that $G^T Y^* \leq G^T Y$ for all $Y > 0$, subject to $WY^* = P - TX^*$ for $Y^* > 0$. *Loucks et al.* [1981, (Reznicek and Cheng, 1991)] show that, if each random variable has a discrete distribution, the first and second stage optimal solutions can be obtained simultaneously, since the problem resolves to a deterministic LP problem. The main limitation to this method is the difficulty of estimating the losses resulting from recourse actions [Yeh, 1985].

3.1.3. Chance-Constrained LP

Chance-constrained LP was first applied to reservoir problems by *ReVelle et al.* [1969, (Reznicek and Cheng, 1991)], and is a methodology that allows the uncertainty to be incorporated into the system constraints. A typical chance-constrained LP model would be the following [Reznicek and Cheng, 1991]:

$$\underset{X}{\text{minimise}} \quad Z = C^T X \qquad \text{(Eqn 3.2)}$$

$$\text{subject to } AX = B,$$
$$P[TX \geq P] \geq \alpha,$$
$$X \geq 0,$$

where $P[]$ denotes the probability; α is a constant vector with elements $0 \leq \alpha_i \leq 1$; and the other symbols have the same meaning as above.

If the probability distribution of the vector P is known, or can be approximated from historical data, then this model can be converted into a deterministic model. If $F_P(TX)$ is the cumulative probability distribution function then the constraint $P[TX \geq P] \geq \alpha$ can be written as $F_P(TX) \geq \alpha$. If F_P is known, and can be inverted, then the constraint can be written in the form $TX \geq F_P^{-1}(\alpha)$, which is a deterministic constraint.

Chance-constrained formulations neither penalise constraint violations, nor provide any form of recourse action [Yeh, 1985]. For this reason it is a difficult task to assign probability values to α.

3.1.4. Linear Decision Rules

Linear decision rules (LDR) are the rules used to relate the release from a reservoir to the storage, inflow and decision parameters of the reservoir. *ReVelle et al.* [1969 (Yeh, 1985)] proposed the original LDR for reservoir design and operation, which has the form $R_t = S_{t-1} - b_t$, where R_t is the release during period t, S_{t-1} is the storage at the end of period $t-1$, and b_t is the decision parameter to be determined. This rule is useful as it allows the release to be determined at the beginning of the period, and it eliminates the mathematical difficulties of formulating chance constraints. It has two basic limitations, it yields conservative, that is, unnecessarily large reservoir capacities, and it is not guaranteed to result in an optimal solution. A number of other LDR's have been formulated in an attempt to reduce these limitations, and *Yeh* [1985] gives a good review of these.

3.2. Dynamic Programming

Dynamic programming (DP) is a popular method for optimising complex multistage problems that was largely formulated by *Bellman* [1957 (Yeh, 1985)]. This method is popular in solving water supply problems because it can incorporate both the non-linear and stochastic properties of these problems. It also effectively decomposes large complex problems, with a large number of variables, into a series of simpler sub-problems.

DP is used to optimise a system over a number of stages. In most problems that are solved using DP, the stages are periods of time, but this is not necessarily the case. At each stage there will be a number of variables that describe the state of the system and these vary with the stage of the problem. These state variables are combined into a state vector $\boldsymbol{X}_i$. At each stage a decision must be made as to the value of the decision variable $\boldsymbol{D}_i$. This decision is based on the state vector and the stage of the problem. A stage-to-stage transformation equation relates the value of the state variable at a particular stage to the state at the previous stage, the decision taken, and the stage. At each stage a return can be calculated from the state vector, the decision variable and the stage [Jacobs, 1967]. The aim of DP is to maximise (or minimise) the total return over all the stages of the problem. Often the range of the state and decision variables are restricted by certain constraints. These may be physical, contractual or institutional.

A simple example of DP for reservoir operation would be as follows. The state variables would be the volume of storage, S_i, in the reservoir at the end of the stage, and the inflow I_i during the stage. The decision variable would be the amount of water to release, R_i, from the reservoir. The stage-to-stage transformation would be $S_{i-1} = S_i + I_i - R_i$, where evaporation has been ignored for clarity. There would be a restriction on the value of S_i given by $S_{min} \leq S_i \leq S_{max}$. A certain economic return would be obtained from the water released from the reservoir, either from the sale of the water, or from the hydropower generated. This might be the value of power generated in a hydroelectric power plant, or it might be the value of crops grown by irrigating with the water released. An economic return may also be assigned to the volume of water stored in the reservoir, depending on the use of the reservoir. Thus, a function $\boldsymbol{B}(R_i, S_i)$ would be the return function, or benefit obtained from a release R_i and a storage S_i for this problem.

The aim of this problem would be to maximise the total return over all the stages. This is done recursively by defining the Optimal Return Function

$$f_n(S_n) = \max_{R_n} \left\{ \boldsymbol{B}(S_n, R_n) + f_{n-1}(S_{n-1}) \right\} \quad \text{[Williams, 1970]} \qquad \text{(Eqn 3.3)}$$

To solve this recursive equation, a boundary condition is required. This would normally be the current value of the reservoir storage if a forward formulation is being used, or the desired storage at the end of the problem is a backward formulation is being used.

This formulation of the DP solution assumes that the state variables and the time stages are discrete. It also assumes that the return function and the stage-to-stage transformation equation are completely known, that is, that they are deterministic. If these properties are not true, it is possible to use other formulations of the DP solution as discussed below. Many different practitioners have applied DP to the solution of reservoir problems. Some examples are *Little* [1955], *Hall and Buras* [1961], *Young* [1967], *Meier and Beightler* [1967], *Hall et al.* [1968], *Schweig and Cole* [1968], *Fitch et al.* [1970], *Liu and Tedrow* [1973], *Opricovic and Djordjevic* [1976], and *Collins* [1977] [Yeh, 1985].

3.2.1. Incremental DP (IDP) and Discrete Differential DP (DDDP)

One of the biggest problems with DP is the "curse of dimensionality" [Yeh, 1985]. For example, if the first reservoir in a two reservoir system can take on 20 different states and the second reservoir can take on 40 different states, the DP algorithm would have to be evaluated for 800 points in each stage. If it is feasible to change from any state at the beginning of a stage to any other state at the end of the stage, 800×800 transitions would have to be evaluated. The number of calculations and the amount of storage space required for the solution quickly gets out of hand for a multi-reservoir system with many possible states at each reservoir.

The IDP and DDDP techniques were developed to relieve this problem. One starts by assuming a set of feasible state vectors, known as the state trajectory. This gives an initial policy. The set of states that are just above and below the state trajectory are then examined. If any of them give a better value for the return function, it replaces the state trajectory. This process continues until convergence occurs. For the two reservoir problem mentioned above, the use of IDP reduces the number of state transitions to 81 [Yeh, 1985]. Both *Hall et al.* [1969 (Yeh, 1985)] and *Turgeon* [1982 (Yeh, 1985)] discuss methods for determining the best increments in the state variables to ensure rapid convergence.

3.2.2. Incremental DP with Successive Approximations (IDPSA)

Another way of reducing the dimensionality problem is to use Bellman's original concept of successive approximation. This method decomposes a complex problem that has many state variables into a series of sub-problems, each with only a single state variable. The solutions of these sub-problems converge to the solution of the original problem. This method is discussed in *Bellman* [1957] and a number of other references.

3.2.3. Stochastic DP

Stochastic DP is a particularly important form of dynamic programming for reservoir operations, as the inflows to, or demands of a reservoir are almost never known deterministically. A typical formulation for a stochastic DP problem would be the following:

$$f_t(S_t, I_{t+1}) = \max_{R_t} \left\{ \sum_{I_t=0}^{I_{t,max}} P\left[I_t | I_{t+1}\right]\left[B(R_t) + f_{t-1}(S_{t-1}, I_t)\right] \right\} \qquad \text{(Eqn 3.4)}$$

subject to $S_{t-1} = S_t + I_t - R_t - e_t$, and

$$f_1(S_1, I_2) = \max_{R_1} \left\{ \sum_{I_1=0}^{I_{1,max}} P\left[I_1 | I_2\right]\left[B(R_1)\right] \right\}$$

where all the symbols have the same meaning as previously defined and $P[I_t|I_{t+1}]$ is the probability that an inflow of I_t will occur during period t, given that the inflow I_{t+1}

occurred during period $t+1$. e_t is the evaporation during period t and the time periods, t, are numbered from the end of the planning horizon [Reznicek and Cheng, 1991].

Stochastic DP also suffers from the problem of dimensionality, but it is not possible to use an incremental DP type of decomposition because the probability of each state must be considered in the optimisation. *Arunkumar and Yeh* [1973 (Reznicek and Cheng, 1991)] used a heuristic decomposition method in their stochastic DP optimisation of the firm power output of two reservoirs. *Stedinger et al.* [1984] also developed a stochastic DP algorithm for the optimisation of the operation policies of a reservoir at Aswan on the Nile. Instead of using the previous period's inflow as a state variable, as shown above, they used the best available forecast of the current period's inflow. This resulted in a substantial improvement in the simulated reservoir operations.

3.2.4. Reliability-Constrained DP

Long-term reservoir planning has to make a compromise between maximum return and the risk associated with achieving that return. Generally, the returns are included in the return function used to optimise the system, and the risks are incorporated in the constraints. *Askew* [1974 (Yeh, 1985)] developed a probabilistic DP model that included a penalty function to account for the reliability of the system. The model does not consider the serial dependence of the inflows, and discounts the returns to present value using an assumed interest rate. This model and some variations of it are discussed in *Yeh* [1985].

3.3. Reliability Programming

This is a programming technique based on chance-constrained programming. Chance-constrained programming, whether linear or dynamic, requires the specification of the probability levels of the inputs to the model. These levels are often not easy to determine. Chance-constrained models also often do not explicitly penalise any constraint violations, they only predict the probabilities of violations.

Sengupta [1972 (Reznicek and Cheng, 1991)] introduced a stochastic reliability programming approach, where the chance-constrained reliability's are not fixed beforehand, but are included in the model as extra decision variables. These extra decision variables are included in the objective function through the use of risk-loss functions. The major drawback of this method is the non-linearity of the programmes.

3.4. Non-linear Programming

Non-linear programming has not been as widely used in water resource problems as LP and DP. This is mainly because the optimisation process is slower and requires more resources to solve than DP or LP. The mathematics involved are also more complicated, but it does have the advantage that it allows more general mathematical formulations [Yeh, 1985]. These restrictions will become less significant as computers become more powerful and cheaper. *Yeh* [1985] discusses some of the recent work that has been done on non-linear programming in the field of water resources.

There are a number of non-linear optimisation algorithms available. The choice of algorithm will depend on a number of factors such as if the derivatives of the function are known or not, the number of independent variables, etc. A simple form of non-linear programming, such as the Downhill Simplex method can be used with a simulation routine to optimise a reservoir system. The simulation routine would take a number of parameters that define the operation policy, and calculate a running cost for the system simulated. The algorithm can then be used to find the set of parameters that give the lowest simulated running cost.

There are two main problems with this sort of optimisation. The first is that it would require repeated simulation of the system. As computers become more powerful, cheaper and faster, this becomes less and less of a problem. The other problem is that non-linear programming techniques can usually only find a local optimum and not the global optimum. The final solution may depend on the starting guess used by the algorithm [Press et al., 1986].

3.5. A Min-Max Operation Policy

Orlovski et al. [1984] have proposed a deterministic 'Min-Max' approach to reservoir operation. The problem with most of the stochastic optimisation procedures is that they have received very little support from reservoir managers. Often, managers will focus their attention on avoiding dramatic system failures, even at the expense of the systems average performance.

The solutions presented by *Orlovski et al.* [1984] have the following advantages over stochastic optimisation procedures:

- they do not require complex algorithms or on-line optimisation,
- they can be visualised in terms of classical storage allocation zones,
- they make reasonable use of real-time forecasts of inflows, and
- they suggest a range of possible releases instead of a single value.

This last property is valuable as it allows some flexibility in the decision making process. An experienced manager may use this flexibility to accommodate secondary objectives or unexpected system demands. It would also be possible to use a stochastic, or other, optimisation method to optimise the long-term operation within this flexible region. This would ensure that the managers risk averse attitude is satisfied, while still optimising the operation as much as possible.

The solution is based on a set of yearlong daily inflow sequences. These are either recorded or synthetically generated, and are hopefully the set of 'worst-case' inflow sequences. These are usually approved by the manager as being particularly troublesome. The solution provided by this method is a region on the Release versus Storage graph that will minimise the system failures.

3.6. Simulation Models

Simulation modelling is slightly different from the above methods of finding optimum reservoir operation policies. Simulation modelling is not designed to explicitly find the optimum policy, but rather as a way of testing different operation strategies and the effect of various inputs on the system.

A simulation model is a mathematical representation of the real system. It is designed so that it will give the same outputs as the real system for the same inputs. The simulation model can then be used to test the operation strategies that have been developed using any of the above techniques. The simulation can also be used to analyse the real system's performance under any conditions. This is particularly useful for testing the system under extreme events such as floods or droughts. It is possible to apply extreme inputs to a simulation model that would have a very low probability of occurring in the real system, and would probably have serious consequences on the real system if they did occur.

Simulation models have been used in the past to develop operation policies by a trial-and-error method. An operation policy is proposed and run on the model. The results are then analysed and an improved policy is then developed. This process will continue until a policy that is satisfactory is developed. This method will not necessarily lead to an optimum policy, and often depends on the initial policy that is proposed. This is the method that was used to develop the current policy used by Rand Water to operate their Zuikerbosch and Vereeniging pump stations [Stallard, 1980].

The main disadvantage of simulation models is that they can almost never represent the real system with 100% accuracy. It is usually necessary to make certain assumptions about the response of the system in certain areas. Simulation models also have to use historical stream-flow data as the input to the model. Stochastic modelling can be used to extend the historical data, and to generate feasible extreme events, but this still assumes that what has happened in the past will be repeated in the future. This problem is not restricted to simulation modelling, but is true of all the optimisation techniques discussed. It is, unfortunately, not possible to see into the future.

A number of simulation models have been developed in the past, and many of these are discussed in *Yeh* [1985]. Some of these were developed to model specific projects, and others have been developed as general models. For example, the *Hydrologic Engineering Centre* has developed two different generalised models. The *HEC-3* model [Hydrologic Engineering Centre, 1971] is designed to analyse reservoir systems to improve conservation, and the *HEC-5* model [Hydrologic Engineering Centre, 1979] simulates flood control and conservation systems.

Monte Carlo Simulation is a technique that is fairly popular in water systems research. In this technique a system is modelled in the same way as discussed above, but instead of using historical stream-flow sequences to drive the model, a set of random numbers are generated and used for the stream-flow values [Hammersley and Handscomb, 1964]. Of course the statistical properties of the random numbers must agree with the statistical properties of the historical stream-flow data, to ensure that the numbers are representative of stream-flows that can realistically occur in a particular system. *Matalas and Wallis* [1976] discuss a number of stream-flow generating techniques that can be used to drive a simulation model.

3.7. Real Time Operation

The real time operation of existing reservoirs is usually performed based on forecast information for the stream-flow and other inputs. The accuracy of these forecasts deteriorates with increasing time spans. The actual models used to calculate the optimum operation procedures, both in the long- and short-term, are usually deterministic and use the best available forecasts, rather than stochastic optimisation techniques.

The operation models are often decomposed into a long-term and a short-term model. The usual planning horizon for the long-term model will be a year, and will include seasonal and monthly variations. This would be updated each year before coming to the end of the current planning year. The short-term model will be based on forecasts for the next 7 to 30 days with daily increments. If the reservoir supplies water to a hydropower scheme, the increments may be as short as hours.

The objective function that will be used to optimise the operation will be dependent on the uses of the particular reservoir, but will often have to include assumptions and simplifications to allow for practical solution. The objectives will also often be framed in the context of long-term contracts, institutional agreements and even legislation. These objectives will be very difficult, or even impossible to change in the short-term. Usually, the only short-term benefits available to the reservoir operator will be hydropower generation [Yeh, 1985].

3.8. Moran's Theory of Storage

Moran [1959] developed the 'Theory of Storage' that is widely used today in the analysis of reservoir operations. This theory calculates the probability distribution of the storage volume, given the distribution of the inflows, the capacity of the reservoir, and the release rule used.

If the inflow to a reservoir is random and serially independent, and if the probability distribution of the initial storage volume is known, then a stable, stationary distribution that is independent of the initial conditions can be found for the volume of storage in the reservoir. From this, the distribution of the release and any overflows can also be calculated. These distributions are dependent on the release rule used for the reservoir, and the capacity of the reservoir. These can then be studied to find optimum values.

3.9. Neural Network Control

Neural network control is a system of control that is completely different from the control and optimisation methods discussed thus far. The control method is not based on any form of model of the system, but rather on an 'intelligent' model of the biological nervous system as it is currently understood.

Neural networks are made up of a large number of very simple processors, as opposed to traditional computing machines, which use a few complex processors. The processor in a neural network is known as a neuron, and consists of a number of inputs, with a weighting factor for each input, a transfer function and, usually, a single output. The inputs are positive for an excitatory input and negative for inhibitory inputs. The transfer function can be any function, but is often a sigmoidal function that returns '*1*' for large positive inputs and '*0*' for large negative inputs. The output is calculated by applying the transfer function to the weighted sum of the inputs.

The property of neural networks that make them significantly different from mathematical models is that the weights do not have to be explicitly determined. The network 'learns' by being exposed to sets of known input-output pairs. During this learning phase, the weighting factors in each neuron are adjusted until the output for a given input set matches the known output. Once the network is 'trained', it can recognise and classify input sets that were not in the training set, and produce the correct output. Neural networks are good at recognising patterns, noise reduction, and can also produce correct outputs from incomplete input data.

There are a number of different neural network models, which differ in architecture, method of learning, means of handling data, and the use of time. *Shaw* [1993] compared a back-propagation neural network, a mathematical model developed by *Dold* [1982] and an experienced human operator in the control of a balancing reservoir for a wastewater treatment plant. He reported that the network produced more constant outputs than either the human operator or the mathematical model, with no under- or overflow of the reservoir.

3.10. Summary

A number of methods for finding the optimum operating rule for a reservoir, or system of reservoirs, have been discussed above. These methods are also applicable to smaller balancing reservoirs, as these are very similar to supply reservoirs. The main difference is that the inflows to a supply reservoir are random, whereas for a balancing reservoir, the outflows, or demands, are random.

Balancing reservoirs are also often simpler systems than supply reservoirs. Supply reservoirs often have to satisfy multiple, competing, objectives such as flood control, water supply, hydropower, recreation, navigation and others. Balancing reservoirs, on the other hand, usually have a main objective of water supply and a secondary objective of emergency storage. Supply reservoirs also usually have more numerous and more complex constraints than balancing reservoirs.

The computer program developed during the course of this project uses two methods to optimise the operating policy of a balancing reservoir. The main method is a simulation model with a controlled search technique. The parameters of the policy are estimated, and a simulation is run. The parameters are then adjusted and the simulation is run again. This process is repeated until the results obtained from the simulations are sufficiently close to an optimum. The computer program can also find the optimum using Dynamic Programming, and this can be used to check that the simulation process finds a result that is near the global optimum.

4. RESERVOIR SYSTEM ANALYSIS

The actual physical reservoir system consists of the installed equipment, and the running costs associated with that equipment. These include the pumping equipment, the pumping costs and the policy used to control the pumps. The other main component is the set of demands placed on the system by the consumers. This will consist of a statistical description of the variation of demand with time. Each of these components are discussed in detail below.

4.1. Reservoir System Model

To optimise a system, it is necessary to fully understand the system, and it is often necessary to try out different operation policies on the system. This is generally not practically possible with the real system, so some form of model needs to be built. In this case the model will be a mathematical model that will form the basis of a computer program. The program can then be used to simulate various solution scenarios, and can also perform direct optimisation of the system.

The mathematical model consists of a number of variables that define aspects of the model, and a number of equations that can be used to calculate simulation steps.

4.1.1. Reservoir Volumes, Demands and Pumping Rates

For a model to be successful, it must incorporate all the parts of the real system that have an impact on the optimum solution, and it must leave out any unnecessary details.

The basic system that this project is looking at consists of a pump station, a reservoir, and a set of user demands. These demands are random and cannot be calculated deterministically. This demand is the uncontrolled input to the model. The pumping rate delivered by the pump station is the controlled input to the model. The total cost of running the system, including pumping, switching and penalty costs, is the output from the model. An optimum solution will be one that minimises the total running cost for a given period of time, while satisfying all the constraints of the system.

Generally, real systems are not as simple as this. For example, one of the sub-systems monitored for the case study consisted of a single pump station containing five pumps, three reservoirs and approximately one hundred consumers. This system can still be modelled as a single pump station, single reservoir system, and a single set of consumer demands. The total capacities of the three reservoirs are lumped together to make up the capacity of the model reservoir, and no concern is given to exactly where the demand is drawn from. All the users' demands are lumped together and assumed to act at a single point in the system. Water is put into the system by the pump station, and drawn from the system be the users. The difference between the input and the output is taken up by the reservoir.

To model this system, it is essential to get all the necessary details for each of the components of the system. These include the pumping capacities and operating costs of the pumps, the sizes and allowable limits on volumes of the reservoirs, the relationship between pressure and flow rate for the system, and statistical information about the historical demand. Once this information has been obtained, a mathematical model can be constructed. An important part of the construction process is to verify the model. This will be done by applying known historical inputs to the system, and ensuring the outputs agree with what really happened, to a reasonable degree of accuracy.

Many reservoir operators do not monitor the demands on a daily basis, but rather monitor only the reservoir levels and the pumping rates. Demands are usually only measured monthly for billing purposes. Therefore, to optimise on a daily basis, the daily demands have to be calculated from the daily pumping rates and reservoir contents. This is done using a simple mass-balance equation as follows:

$$D_t = P_t + V_{t-1} - V_t \quad \text{(Eqn 4.1)}$$

where D_t is the demand for a particular time period; P_t is the volume pumped for the same period; V_{t-1} is the volume in the reservoir at the beginning of the period; and V_t is the volume in the reservoir at the end of the period. This equation assumes that there are no leaks in the system, or that leaks are included in the demand, and that the reservoir never overflows. These assumptions must be kept in mind when obtaining data for the model, and interpreting the results from the simulations.

The mathematical model will also simulate the operation of the reservoir system using the mass-balance equation above. The simulation is broken down into time steps, called stages, depending on the data available. For example, the system monitored for the case study reports their pumping rates on a daily basis, so the time step would be 24 hours. The volume at the end of each time step, V_t, is calculated from:

$$V_t = V_{t-1} + P_t - D_t. \quad \text{(Eqn 4.2)}$$

To start a simulation, it is necessary to know a starting volume V_0. For each stage the demand will be calculated from either historical measurements or from synthetically generated data. The pumping rate can also be calculated from historical information or by using a mathematical operation policy. Once these are known, the new volume can be calculated. For each time step, the cost of operation for that period is also calculated. These costs will be summed over the entire period of the simulation, to give the total operating cost.

Sometimes the reservoir system may fail. A failure is defined as either a failure to meet the users' demands, or overflow of the reservoir. If the demand is greater than the pumping rate and the available volume in the reservoir then the water supplied is equal to the available volume in the reservoir plus the pumping rate, and the reservoir volume at the start of the next stage is the absolute minimum. If the pumping rate minus the demand is greater than the available space in the reservoir then an overflow failure will occur. The volume in the reservoir at the start of the next stage will be equal to the maximum volume in the reservoir. Failures of this kind will usually have financial implications, and are incorporated into the model as penalties.

4.1.2. Reservoir Sizes and Costs

Each reservoir system will consist of one or more reservoirs that will have been designed to store a certain volume of water. These reservoirs are modelled by defining that maximum and minimum volumes available. The maximum volume is defined as the absolute maximum storage capacity available in the system. Any attempt to store more water than this maximum will result in the system overflowing, with the resultant loss of water and possible secondary damage. The minimum volume is defined as the volume of water that cannot be used out of the reservoir system. It will not be possible to meet a demand that is greater than the difference between the volume available and the minimum volume. If a demand is not met, there will be both direct and indirect costs to the system operators. Direct costs will occur if the system operators have contractual obligations to supply water to the consumers, or if the unmet demand has to be made up from other more expensive sources. The indirect costs will be due to problems such as loss of production, depending on the nature of the consumers.

The model makes provision for these costs by defining a Maximum Penalty and a Minimum Penalty. The maximum penalty is the cost, in Rands per megalitre, that will be incurred if the reservoir system overflows. The minimum penalty is the cost of not meeting a demand, also in Rands per megalitre. These costs should be as accurate as possible and include both the direct and indirect costs.

The Maximum and Minimum volumes are structural limits that are imposed by the physical reservoir system. They cannot be changed, except by spending major capital on physical upgrading of the system. The model also defines High and Low Warning limits. These are policy decisions made by the operators of the reservoir system, and are not physical limits. The operators can also impose financial penalties on the system if these limits are not met. In the model, these are called the High and Low limit penalties respectively. For example, the Low limit can be used to ensure that the reservoir system always contains a certain volume of water for emergencies, such as fire fighting or breakdown requirements. The Low limit penalty would then be calculated on the estimated cost risk of not having water available in an emergency situation.

The model also defines a storage cost function, based on the volume in the reservoir at each stage. This is a straight line function made up of the Storage Cost Slope (in Rands per megalitre) and the Storage Cost Offset (in Rands). This cost can be based on the cost of the original construction of the reservoir, or on the cost of unused capacity in the reservoir.

4.1.3. Pumping Capacities and Costs

The reservoir system will also contain a number of pumps that will need to be controlled. The model defines the pump settings as the possible ways in which a set of pumps may be operated. A simple system with a single pump would have two settings, either 'ON' or 'OFF'. A system with two different pumps would have the following 4 possible settings:

- Both off
- Pump 1 on, Pump 2 off
- Pump 1 off, Pump 2 on
- Both on

If the system contains pumps that have variable settings it will be necessary to approximate the available pumping rates by a number of discrete settings. The more settings that are chosen, the more accurate the result of the simulation, but the more time it will take to calculate an optimum operation policy.

The range of settings available will be entirely dependent on the physical system being modelled. For each setting the actual output of the pumps will be dependent on the head against which the system is pumping. This is a function of the volume in the reservoir and the demand throughout the system, and the difference in elevation between the pump station and the reservoir. Because the exact distribution of the demand affects the flow rate within the network, and therefore the pressure losses due to friction, and the demand is modelled as a single lumped demand, the pumping rates are approximated as follows.

The system head, against which a pump must work, for a single pipeline, is given by:

$$H_{sys} = H_{elev} + \frac{V_R}{A_R} + \frac{8\lambda LQ^2}{\pi^2 gD^5} + k_L\frac{8Q^2}{\pi^2 gD^4} \quad \text{[White, 1986]} \qquad \text{(Eqn 4.3)}$$

where: H_{sys} is the total system head,
H_{elev} is the difference in elevation between the reservoirs and the pump station,
V_R is the volume of water in the reservoir,
A_R is the surface area of the water in the reservoir,
λ is the Darcy coefficient of friction,
L is the total length of the pipeline,
Q is the flow rate in the pipeline,
D is the diameter of the pipeline,
g is the acceleration due to gravity, and
k_L is the coefficient of head loss due to the local losses.

For a network that currently installed, all the above terms are constant, except for the volume in the reservoir and the flow rate. This means that Eqn 4.3 can be simplified as:

$$H_{sys} = H_{elev} + \frac{V_R}{A_R} + k_S Q^2 \qquad \text{(Eqn 4.4)}$$

where k_S is a constant that relates head loss to Q^2. This constant would need to be estimated for the system from historical pumping data.

For each available pump setting there is also a pump curve that relates the increase in head across the pump to the flow rate through the pump. This can be written as:

$$H_{pump} = f_P(Q) \qquad \text{(Eqn 4.5)}$$

The function $f_P(Q)$ is usually approximated as a cubic polynomial. When a pump is operating in a particular system the H_{sys} given by Eqn 4.4 must be equal to the H_{pump} given by Eqn 4.5. By eliminating the heads between these two equations and rearranging, it is possible to get an expression of the form:

$$Q = f_Q(V_R) \qquad \text{(Eqn 4.6)}$$

The function $f_Q(V_R)$ is also usually a cubic polynomial

Thus, for each available pump setting, it is possible to find an equation that relates the pumping rate to the volume in the reservoir.

For each of these settings, the pumps will also require a certain power input to deliver the calculated flow. Therefore, for each setting, the model also requires an equation that gives the power requirement as a function of the flow being delivered. These equations are available directly from the pump curves.

To calculate the cost of pumping, it is also necessary to know the cost of energy available to the pump station. Often power is available at different tariffs, depending on when the power is used. The model allows the 24 hour day to be broken down into a number of power tariff zones, each of which has its own tariff. The cost of pumping for each time step is calculated by finding the energy consumed during the step, and multiplying it by the applicable tariff.

Another major cost in running reservoir systems is the cost of controlling the pumping rate. The system can be either manually or automatically controlled. If it is manually controlled, each time a change in pumping rate is required, someone would have to go out to the pump and physically change the setting. If the system is automatically controlled, then electronic control and switching gear would need to be installed. Both of these options cost money. To take account of this cost, the model defines a switch cost for each possible pair of pump settings. This will be the cost of switching from setting 1 to setting 2. If it is not possible to switch from a particular setting to another setting, this cost should be set infinitely high.

4.1.4. Operating Policy

The operating policy of a reservoir system is the set of rules that indicate what the pump setting should be, based on a number of factors such as the content of the reservoirs, the current pump setting and the expected demand. An optimum policy is one that results in the lowest overall operating cost for the system. The model uses the following operating policy. The *Expected Volume* at the end of a time step is calculated from

$$EV_t = V_{t-1} - ED_t, \qquad \text{(Eqn 4.7)}$$

where EV_t is the expected volume at the end of time step t, V_{t-1} is the volume in the reservoirs at the end of time step $t-1$, and ED_t is the *Expected Demand* during time step t.

A set of *Change Volume* variables are defined, one for each available pump setting, except for the setting of all pumps off. Each of these variables indicate the range of volumes for which a pump setting should be used. For example, the typical operation policy shown in Figure 4.1 below would use pump setting 0, that is, all pumps off, if the expected reservoir volume is greater than 75% of its maximum capacity. The maximum setting would be used whenever the expected volume fell below 25% of maximum capacity. The process of optimising the operation policy involves finding the values of these *Change Volume* parameters that minimises the total operating cost of the system.

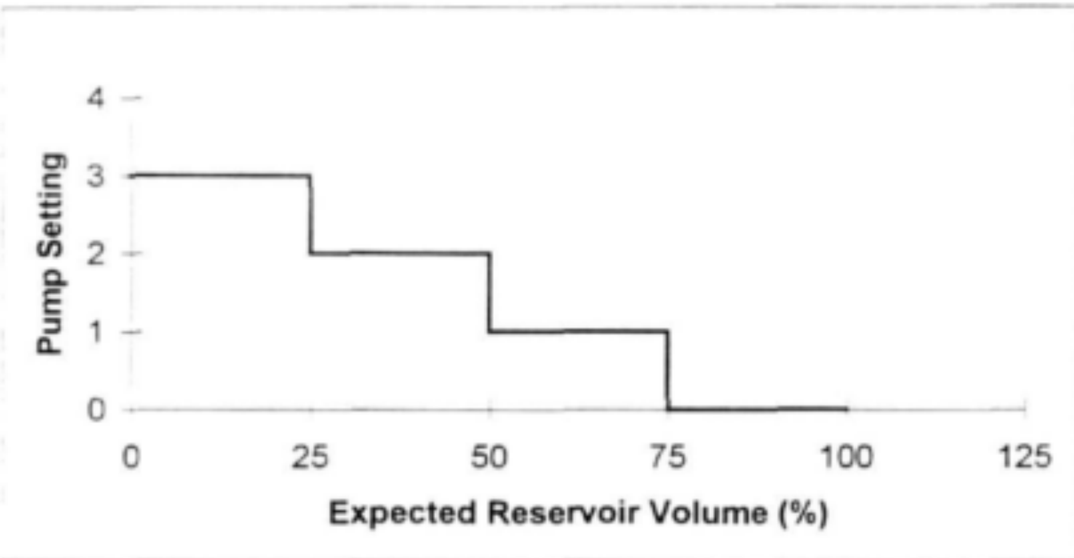

Figure 4.1: Typical Operation Policy

As mentioned earlier, one of the costs of operating a reservoir system is the cost of switching from one setting to another. To try to minimise the impact of this cost the model also defines a *Don't Change* variable. This variable governs whether or not to change the setting if the new setting is sufficiently close to the old setting. For example, if the current setting is setting 0 and the new expected content is 73%, then the new setting should be setting 1. But if the *Don't Change* value is 10% then the *Don't Change* range will be 10% of the difference between 50% and 75%, that is 2,5%. The new expected contend is within 2,5% of the *Change Volume* so no change in setting will occur. This variable helps to reduce the cost of running a reservoir system when the cost of switching from one setting to another is relatively high.

4.2. Demand Data Model

An important part of the model is the demand data that is used in the simulation and optimisation processes. It is not good enough to just use data that has been measured from the existing system as this is only one possible realisation of a random function. The user demands placed on a reservoir system are random in nature, and cannot be predicted explicitly. It is therefore necessary to be able to generate synthetic data sets that represent different possible realisations of the same random processes that generated the measured data. This will allow the system to be optimised and tested under as wide a range of conditions as possible.

To generate synthetic demand data it is necessary to get some idea of the statistical processes that generated the measured data, and to fit a model to this data. This model can then be used both to generate synthetic data sets, and to predict the next time step's *Estimated Demand*. The process of calculating the properties of the historical demand data, and constructing a model that will generate statistically similar data sets is known as *Stochastic Modelling*. Once a stochastic model has been constructed, it can be used to generate as many data sets as are required. All the data sets will be statistically similar to, but different from, the original historical data set.

To construct an accurate stochastic model, it is necessary to break the demand into various components. There are five basic components that can be divided between *Systematic* and *Non-systematic* components. The systematic components are the *Secular*, *Periodic* and *Cyclical Trends*. The non-systematic components are chance, or chance-dependent effects. These are *Episodic Events* and the *Stochastic Component* [McCuen 1993].

4.2.1. Secular Trends

The secular trend is that part of the data that continually increases or decreases through the whole data set. This can be linear, logarithmic, or a number of other forms. It should be possible to find a logical explanation for a secular trend in demand data. An example would be an exponential growth in the population being supplied will lead to an exponential growth in the demand.

4.2.2. Periodic Trends

Periodic trends are those effects that have a definite period and recur at fixed, regular intervals. For example, residential demand varies between a minimum in the early hours of the morning to a maximum in the early evening. This has a fixed period of 24 hours, and occurs every day. There should always be a logical explanation for the periodic trend.

4.2.3. Cyclic Trends

Cyclic trends are those factors that tend to recur, but at irregular intervals. For example, it is said that drought years occur approximately every 7 years, but this is definitely irregular.

4.2.4. Episodic Events

These are events that are once-off in nature. It should also be possible to trace the cause of episodic events, especially in municipal demand data. An example would be an extremely high demand for a few hours cause by a major pipe break.

4.2.5. Stochastic Component

The stochastic component consists of random effects that cannot be explained by physical logic. These require probabilistic concepts for description, and are caused by the fact that the demand is made up by many individual users, each acting independently.

Not all of these components will be present in a particular data set, and generally, demand data contains only secular and periodic trends, and a stochastic component. Any episodic events that have occurred during the measurement period should be removed. The process of stochastic modelling consists of finding values for the parameters that define each of these components.

The details of how this model is calculated are discussed in the following chapter.

5. RESERVOIR SYSTEM OPERATION OPTIMISATION

Once the reservoir system and demand data models had been constructed they were coded as a computer program. This allows the models to be fully tested and to be used by reservoir system operators.

5.1. Computer Program Description

A program was written using Borland's® Delphi™ programming language. This language was chosen as it allows the rapid development of programs for the Microsoft® Windows™ operating system, and it is based on the Pascal language, which is the language with which the author is most familiar.

The program, called *Reservoir System Operation Optimisation*, or *RSPO* for short, has four main options:

- Analysis - which analyses data measured from the existing system and calculates the Stochastic model,
- Simulation - which runs simulations on the modelled system to calculate the total operating cost, and
- Optimisation - which finds the optimum operating policy.
- Pump Setting – which calculates the expected demand and optimum pump setting for the current time step.

There are also a number of utility functions that help the user to manage their reservoir system models. These include file management functions and functions to view and edit the input and output files of the models. Full details of how to install and use the program are given in the user manual in Appendix C. A set of extension routines that include import and export functions were written specifically for the data used in the case study. These are fully described in Appendix D: Rand Water Extensions User Manual. Similar extensions could be written for any specific data file format.

5.2. Analysis

The analysis option of the *RSPO* program takes a data file containing the date and time, reservoir volume, demand and pumping rate for each time step and produces an *Analysis Report File* and a *Stochastic Model File*. Details of all the files used by *RSPO*, and the user options on the Analysis Dialog screen are given in the user manual in Appendix C.

5.2.1. Secular Trend

Two possible secular trends are fitted to the demand data using regression analysis, namely a linear trend and a logarithmic trend. The linear trend has the form $D_{s,t} = mt + c$, where $D_{s,t}$ is the secular trend component of the demand at time t, m is the slope and c is the intercept. The values m and c are calculated from the formulae shown below. The correlation coefficient, R, is also calculated to indicate of how well the trend fits the data.

$$m = \frac{N\sum(Dt) - \sum D \sum t}{N\sum t^2 - \left(\sum t\right)^2}$$

$$c = \frac{\sum D}{N} - m\frac{\sum t}{N} \qquad \text{(Eqn 5.1)}$$

$$R_{lin}^{\ 2} = m\frac{N\sum(Dt) - \sum D \sum t}{N\sum D^2 - \left(\sum D\right)^2}$$

Similarly, a logarithmic trend of the form $D_{s,t} = ab^t$ is also fitted. The coefficients a and b, and the correlation coefficient are found using the following formulae:

$$b' = \frac{N\sum(\ln(D)t) - \sum \ln(D)\sum t}{N\sum t^2 - \left(\sum t\right)^2}$$

$$a' = \frac{\sum \ln(D)}{N} - b'\frac{\sum t}{N}$$

$$R_{log}^{\ 2} = b'\frac{N\sum(\ln(D)t) - \sum \ln(D)\sum t}{N\sum \ln(D)^2 - \left(\sum \ln(D)\right)^2} \qquad \text{(Eqn 5.2)}$$

$$a = e^{a'}$$

$$b = e^{b'}$$

The trend that has the greatest absolute value of R^2 is then used as the trend component of the stochastic model. This secular trend is also subtracted from the demand data before the periodic trend is fitted [Larson, 1982].

5.2.2. Periodic Trend

Municipal demand data typically has three basic periodic trends. These are the annual trend of the seasons, the weekly trend where weekend consumption is different to weekday consumption, and the daily trend, where the consumption differs from night to day. Each of these periodic trends are analysed and reported on by *RSPO*.

The basic period of the trend is known as the fundamental wavelength. In the case of the weekly trend, the fundamental wavelength will be 7 days. The actual shape of the periodic trend can be constructed by summing a number of sinusoidal curves with frequencies of $1/T$, $2/T$, $3/T$,..., where T is the fundamental wavelength. The calculation of these sinusoidal curves is known as harmonic analysis. The sinusoidal curve with frequency $1/T$ is called the first harmonic, the curve with a frequency of $2/T$ is called the second harmonic, and so on. The harmonic representation of the periodic component is

$$D_{p,t} = \mu_D + \sum_{i=1}^{L} \lambda_i \sin\left(\frac{2\pi t}{T} i + \phi_i\right), \qquad \text{(Eqn 5.3)}$$

where $D_{p,t}$ is the periodic component of the demand at time t, μ_D is the mean of the demand data, and λ_i and ϕ_i are the amplitude and phase of the frequency ith harmonic. Because at least three points are required to draw a curve, the shortest possible wavelength a harmonic could have is $2\Delta t$, where Δt is the time step of the demand data. L is the maximum number of harmonics that can be fitted to a set of data, and will be equal to $N/2$ if N, the number of data points in the set, is even, or $L = (N-1)/2$ if N is odd.

The fundamental period can be broken down into a number of intervals. *RSPO* breaks the year into 52 weeks, the week into 7 days, and the day into 24 hours. The mean, m_τ, of all the data points that fall into each interval is then found, where τ is the interval number. For example, when calculating the weekly periodic trend, m_1 would be the average of all the data points that occur during the first day of the week.

It is now necessary to fit a set of harmonic sine curves to these average points. To do this, Eqn 5.3 can be rearranged as follows:

$$D_{p,\tau} = \mu_D + \sum_{i=1}^{p/2} \alpha_i \, sin\left(\frac{2\pi\tau}{p} i\right) + \sum_{i=1}^{p/2} \beta_i \, cos\left(\frac{2\pi\tau}{p} i\right), \qquad \text{(Eqn 5.4)}$$

where $D_{p,\tau}$ is the harmonically fitted, periodic component of demand for the interval τ, p is the number of intervals in the fundamental wavelength, and A_l and B_l are coefficients to be determined.

Kottegoda [1980] shows that the coefficients can be determined from:

$$\begin{aligned}
\alpha_i &= \frac{2}{p}\sum_{\tau=1}^{p} m_\tau \, sin\left(\frac{2\pi\tau}{p} i\right), \qquad i = 1,2,\ldots,\frac{p}{2}-1 \\
\alpha_{p/2} &= 0 \\
\beta_i &= \frac{2}{p}\sum_{\tau=1}^{p} m_\tau \, cos\left(\frac{2\pi\tau}{p} i\right), \qquad i = 1,2,\ldots,\frac{p}{2}-1, \\
\beta_{p/2} &= \frac{1}{p}\sum_{\tau=1}^{p} m_\tau (-1)^\tau
\end{aligned} \qquad \text{(Eqn 5.5)}$$

Figure 5.1 below shows the weekly periodic trend for the Libanon/Driefontein subsystem. The open points (**O**) show the calculated means m_τ for each day, the dashed lines show the first and second harmonics, and the solid line indicates the fitted periodic trend, made up of three harmonics.

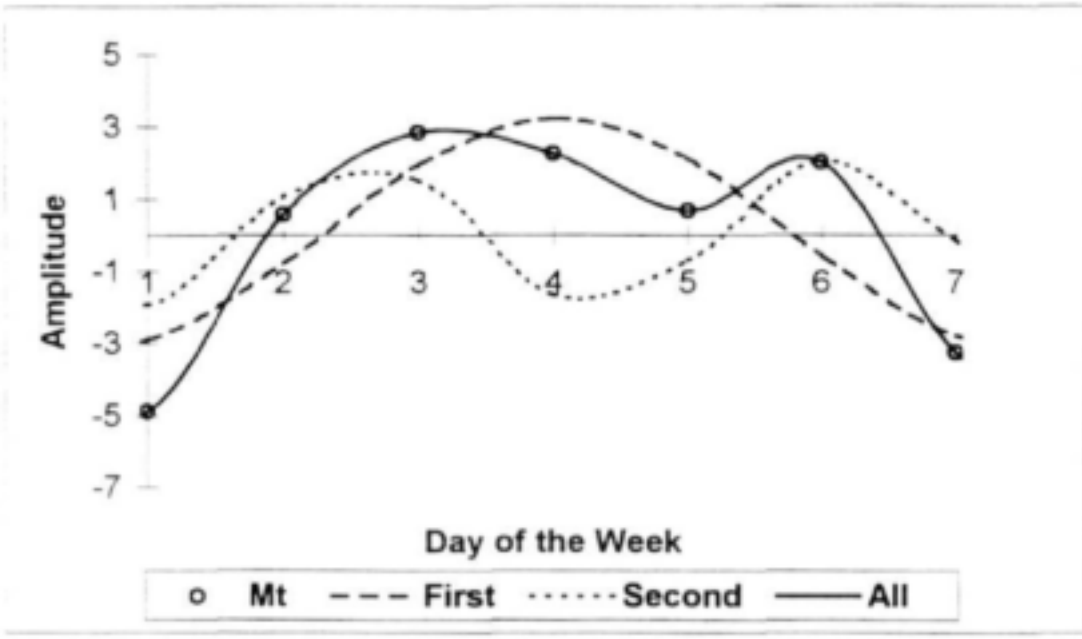

Figure 5.1: Harmonic Analysis

The significance of each of the harmonics can be tested as follows. If the significance is higher than that specified by the program user, then the harmonic is subtracted from the demand data. If the significance is not high enough, the variance caused by the harmonic is passed onto the random component of the data [Kottegoda, 1980]. To test the significance of each harmonic, a null hypothesis, which states that the variance explained by harmonic i is zero, is assumed. The variance explained by harmonic i is $(N/2)(\alpha_i^2 + \beta_i^2)$. Each harmonic is tested, starting with harmonic $p/2$ and proceeding down to harmonic 1, using the F-test for analysis of variance [Larson, 1982]. If the null hypothesis is not rejected then the sum of squares of the harmonic is added to the residual sum of squares. For the example shown in Figure 5.1, harmonics 3, 2 and 1 had significance's of 1.7%, 96.3% and 99.9% respectively. If the significance level chosen by the user was anything less that 96.3%, harmonics 1 and 2 would be subtracted from the data set, and harmonic 3 would be ignored.

5.2.3. Residual Stochastic Component Parameters

Once the secular and periodic trends have been removed from the demand data, all that is left is the residual stochastic component. To start with, *RSPO* calculates the first four moments of this data, and converts them to the mean, Standard Deviation, Skewness and Kurtosis using the following formulae:

$$\mu_D = \frac{\sum D_{r,t}}{N}$$
$$\sigma_D^2 = \frac{\sum (D_{r,t} - \mu_D)^2}{N}$$
$$Skew_D = \frac{\sum (D_{r,t} - \mu_D)^3}{N\sigma_D^3} \qquad \text{(Eqn 5.6)}$$
$$Kurt_D = \frac{\sum (D_{r,t} - \mu_D)^4}{N\sigma_D^4}$$

where $D_{r,t}$ is the residual demand at time t, N is the total number of data points and the summations are over all the data points. These are calculated using two passes through the data, the first to calculate μ_D, and the second to calculate σ_D, $Skew_D$ and $Kurt_D$.

5.2.4. Serial Autocorrelations

On the Analysis Dialog screen, the user can also specify a value for the Maximum Lag for Correlations. If this value is set to any value greater than zero then *RSPO* will calculate serial autocorrelations with lags from zero to the number specified. The autocorrelations are calculated from the formulae specified in *Kottegoda* [1980], as follows:

$$r_l = \frac{N}{N-l} \frac{\sum_{t=l}^{N-l} \left\{ (D_{r,t} - \mu'_D)(D_{r,t+l} - \mu''_D) \right\}}{\sqrt{\left\{ \sum_{t=l}^{N-l} (D_{r,t} - \mu'_D)^2 \right\} \left\{ \sum_{t=l}^{N-l} (D_{r,t+l} - \mu''_D)^2 \right\}}}, where$$
$$\mu'_D = \frac{\sum_{t=l}^{N-l} D_{r,t}}{N-l}, and \qquad \text{(Eqn 5.7)}$$
$$\mu''_D = \frac{\sum_{t=l}^{N-l} D_{r,t+l}}{N-l}$$

Here r_l is the serial autocorrelation at lag l. This value is also calculated using two passes through the data, the first to find μ'_D and μ''_D, and the second to find r_l.

5.2.5. Autoregressive Component

The residual stochastic component that remains after the secular and periodic trends have been removed is random in nature, but it may not be serially independent. To model the correlation between a particular demand value and the previous few demands, a Linear Autoregressive model is fitted to the remaining residuals [Nel and Haaroff, 1996].

In a linear autoregressive model, the current value of the demand is the weighted sum of the previous p demand values and an independent random value. The number p is known as the order of the model. This can be written mathematically as

$$D_{r,t} = \sum_{i=1}^{p} \phi_{p,i} D_{r,t-i} + D_{n,t} \qquad \text{(Eqn 5.8)}$$

where the $\phi_{p,i}$ values are the weights of a pth order model, which need to be determined, and $D_{n,t}$ is the independent random component of the demand at time t.

To fit a linear autoregressive model, it is necessary to choose the order of the model, and then to find the best weights for that model. *Kottegoda* [1980] derives the following recursive formulae for estimating the weights:

$$\phi_{p,p} = \frac{r_p - \sum_{j=1}^{p-1} \phi_{p-1,j} r_{p-j}}{1 - \sum_{j=1}^{p-1} \phi_{p-1,j} r_j}$$

$$\phi_{p,j} = \phi_{p-1,j} - \phi_{p,p}\phi_{p-1,p-j} \qquad \text{(Eqn 5.9)}$$

$$j = 1,2,3,\ldots,p-1$$

r_p is the serial autocorrelation with a lag of p, and $\phi_{p,p}$ is known as the Partial Autocorrelation Coefficient of order p. The set of partial autocorrelation coefficients for p = 1, 2, 3, ..., is known as the Partial Autocorrelation Function or PAF.

For an autoregressive process of order p, the variance of the partial autocorrelations is given by:

$$\mathbf{var}(\phi_{k,k}) \approx 1/N, \text{ for } k > p \qquad \text{(Eqn 5.10)}$$

This information can be used to construct a test by calculating confidence limits that are $\frac{x}{\sqrt{N}}$ above and below the horizontal axis of the PAF, where x is a value such that the standard normal distribution evaluated at x is equal to the significance level chosen by the program user. *RSPO* evaluates the partial autocorrelation coefficients using Eqn 5.9, and uses the last coefficient that is greater than this limit as the order of the autoregressive model.

The weights found using Eqn 5.9 are then used as a first guess to find a set of weights that minimise

$$\sum_{t=1}^{N} \left\{ abs\left(D_{r,t} - \sum_{i=1}^{p} \phi_{p,i} D_{r,t-i} \right) \right\}, \qquad \text{(Eqn 5.11)}$$

where $D_{x,t-i} = 0$ if $t\text{-}i < 1$.

RSPO uses a Downhill Simplex algorithm to find the set of weights, $\phi_{p,i}$, that minimise Eqn 5.11. The *Analysis Report* shows the values of the initial guesses used and the value of Eqn 5.11 for this first guess. It also shows the number of iterations taken to find a best fit, the equation value and the values of the weights that give this best fit.

5.2.6. Independent Random Component

After finding the best fit autoregressive weights, the linear autoregressive model is subtracted from the residual demand, leaving only the independent random component $D_{n,t}$. Firstly, *RSPO* estimates the mean, standard deviation, skewness and kurtosis of the random component using Eqn 5.6, and these values are shown in the *Analysis Report*. Next, a χ^2 test is used to find the normal distribution that best fits the independent random component of the demand. In this case 'best fit' means the parameters that minimise the χ^2 value where

$$\chi^2 = \sum_{i=1}^{I} \frac{O_i - E_i}{E_i}, \quad \text{(Eqn 5.12)}$$

O_i is the number of data points observed in interval i, and E_i is the expected height of interval i. *RSPO* calculates the boundaries of the intervals such that each interval contains the same number of observed data points, and selects I, the number of intervals such that there are at least 5 data points in each interval.

If the user chooses a verbose *Analysis Report*, then the boundaries and O_i and E_i for each interval are shown in the report. The χ^2 value of the first guess and the best fit are also shown, along with the number of iterations taken to find the best fit set of parameters.

This, then, completes the analysis of the data and the construction of the Stochastic model. The values of all the parameters that make up the stochastic model are saved in a *Stochastic Model File* that is used during simulations to generate synthetic demand data, and in both simulations and optimisations to forecast demand values.

5.3. Simulation

RSPO allows the user to run simulations of the reservoir system being studied. Simulations are useful for two reasons; the first is to generate synthetic data files, and the second is to test various operating rules. When the Simulation option is first chosen, the user is presented with a choice of three options:

- Existing Demand data - the simulation uses demand data from an existing *System Run File*, and calculates the pumping rates from the given operating rule. This option is useful for comparing the effects of different operating rules on the same set of demand data.
- Synthetic Demand Data - the simulation generates synthetic demand data from the stochastic model, and calculates the pumping rates from the operating rule. This option would be used to generate further data sets on which to run analyses and optimisations.

- Existing Demand and Pumping data - this option takes both the demand and pumping data from an existing *System Run File*. This allows the user to calculate the running costs of the system using the historical data and the cost information stored in a *System Description File*. This is useful to compare the running costs of the system using its existing operation policy to the running costs using a newly generated operation policy.

Running Simulations will result in the creation of two new files, a *System Run File*, containing the reservoir volume, demand and pumping rate information for each time step in the simulation, and a *Simulation Report File*, containing information about the simulated running costs and any reservoir limit violations. The input files required by the simulation will depend on the simulation option chosen by the user.

- Existing Demand data - requires an existing *System Run File* for the demand data, a *Stochastic Model File* to forecast demands, and an *Operating Rule File* to calculate the pumping rate at each stage.
- Synthetic Demand Data - requires a *Stochastic Model File* to generate demand data and to forecast demands, and an *Operating Rule File* to calculate the pumping rates.
- Existing Demand and Pumping data - requires an existing *System Run File* for the demand data and the pumping rates.

All options will require a *System Description File*. This file is fundamental to the simulation as it describes the capacities of the reservoir and pumping equipment, and gives details of all the cost functions involved. A detailed description of this file is given in the next section.

5.3.1. The *System Description File*

A typical *System Description File* is shown below. The file consists of a number of sections, each with a section heading enclosed in square brackets "[...]". Each section then contains a number of variable descriptions. Each description has a variable name, followed by an equal sign, "=", followed by the variable's value. For example, in the [Reservoir] section, there is a line that says "Max Volume=124.354". The variable name is "Max Volume", and describes the maximum available volume in this reservoir. The value after the equal sign is 124.354 that tells the program that this reservoir can contain a maximum of 124.354 Ml of water.

In the *System Description File*, the semi-colon character ";" is used as a comment character. Any information after the semi-colon is ignored and does not need to be in the file. It is just there for information.

Each section of the *System Description File* will be described separately below:

```
[Header]
System=C:\RSPO\EXAMPLES\EXAMPLES.RSPO
Comment=Example Reservoir System to Demonstrate the use of RSPO
Time=01:00
Date=01/06/97
```

The [Header] section describes the system to which the *System Description File* applies. The System variable must contain the full path of the *System Initialisation File*. The Comment variable can contain any information you like, up to 255 characters. The Time and Date variables are useful to keep track of when the file was last edited.

```
[Reservoir]
Start Volume=104.10     ; Simulation initial volume
Max Volume=124.354      ; Maximum possible volume in the reservoir
High Volume=120.900     ; High limit volume
Low Volume=62.177             ; Low Limit volume
Min Volume=0            ; Absolute minimum volume, usually zero

; Each of the following penalties must be given in Rands per Ml
Max Penalty=100               ; Penalty for exceeding the maximum
                              ; volume. Should include the cost of
                              ; lost water.
High Penalty=10               ; Penalty for exceeding the high limit
                              ; volume
Low Penalty=10          ; Penalty for not meeting the low limit
                              ; volume
Min Penalty=100               ; Penalty for running the reservoir dry

; Storage costs
Storage Cost Slope=10 ; Basic cost of storage, in Rands per Ml
Storage Cost Offset=10 ; Reservoir fixed costs in Rands
```

The [Reservoir] section describes the constraints and costs that are applicable to the reservoir in the system. The Min Volume and Max Volume variables describe the absolute minimum and maximum volumes allowable in the reservoir. All the volumes must be in megalitres (Ml). The High Volume and Low Volume variables are warning limits on the volume in the reservoir. The low volume warning is often used to ensure that the volume in the reservoir is always sufficient to meet emergency demands. The Start Volume variable describes the volume that is assumed to be in the reservoir at the start of a simulation. This must be within the Max Volume and Min Volume limits.

The next four variables allow the user to impose financial penalties if the limits are not met. The penalties must be in Rands per Ml over or under the limit. If the Max Volume limit is exceeded, both the Max Penalty and the High Penalty will be applied. Similarly if the Min Volume limit is not met, both the Low Penalty and the Min Penalty will be applied. If it is not acceptable for a reservoir to overflow, for example if it is situated on dolomitic bedrock, then a very large value must be used for the Max Penalty.

The storage costs allow the user to apply a straight line cost function to the cost of storing water in the reservoir. The cost at each time step is calculated as

Cost = Storage Cost Slope * Reservoir Volume + Storage Cost Offset

A positive slope will cause the optimum operation policy to keep the reservoir as empty as possible, while a negative slope will cause the optimum operation policy to keep the reservoir as full as possible.

```
[Pumping]
; Number of possible pump settings, excluding all pumps off
Settings=2

; Equations giving flow rates in Ml/d as a function of
; reservoir content in Ml for each pump setting
; Content is the variable x
Flow1=-0.0250*x+31.8165
Flow2=-0.0200*x+58.1333

; Equations giving the power consumption in kW as a function of
; Flow rate in Ml/d for each pump setting
; Flow is the variable y
Power1=170.532+10.467*y-0.112*y^2
Power2=11.043+29.725*y-0.199*y^2
```

The [Pumping] section describes the pumping that is available in the system. Settings gives the number of different settings that are available. This will include all the combinations of all the pumps that can be run simultaneously. If there is a variable speed pump available in the system, it must be approximated by a number of individual settings. The larger the number of settings, the more accurate the optimum results will be, but the program will take longer to find a solution. The setting that has all pumps turned off is always assumed, and therefore is included in the *System Description File*.

For each of the settings there must be a Flow? variable. This gives an equation that describes the flow rate as a function of x, the reservoir content. These flow rates must be in increasing order, that is, Flow1 must deliver a smaller flow rate than Flow2 over the entire range of reservoir content. The reservoir content, x, will be in Ml and the flow rate must be given in Ml/d.

The Power? variables are similar to the Flow? variables. For each of the settings, there must be a Power? equation the can be used to calculate the power consumption for a given flow rate. The flow rates, y, will be in Ml/d, and the power must be in kW.

```
[Energy]
; Number of Power tariff zones
; The zones MUST total 24 hrs, and must have NO overlaps and NO gaps!
Zones=2

; Zone start and end times
Zone1 Start=06:00
Zone1 End=18:00
Zone2 Start=18:00
Zone2 End=06:00

; Zone power cost in cents per kWh
Zone1 Cost=5.94
Zone2 Cost=4.75
```

The energy section describes the cost of energy that will be used for pumping that is available to the system. Provision is made for multiple tariff zones. For each of the zones given by the Zones variable, there must be a start time and an end time, as well as a cost, in cents per kWh. The zones must be continuous, and total exactly 24 hours. If they do not, the results of the simulation will be unpredictable and inaccurate.

```
[Switching]
; Cost of switching from one pump setting to another, in Rands.
; There should be up to (Settings+1)^2 switches. Default value for
; switches not listed is R0,00
From0to0=0.00
From0to1=41.04
From0to2=102.96
From1to0=-41.04
From1to1=0.00
From1to2=61.92
From2to0=-102.96
From2to1=-61.92
From2to2=0.00
```

The [Switching] section describes the cost, in Rands, of switching from one pump setting to another. It is not necessary to list every single switch cost, as shown above, but any costs not listed are assumed to be zero.

5.3.2. Demand Forecasting

As mentioned previously, the calculation of a pump setting for a particular stage of the simulation requires an estimate of the demand for that stage. This forecasting is done using the stochastic model calculated during the analysis process, and the following formula:

$$D_{f,t} = D_{s,t} + D_{p,t} + \sum_{i=1}^{p} \phi_{i,p} D'_{r,t-i} + \mu_{D_r} \qquad \text{(Eqn 5.13)}$$

where $D_{f,t}$ is the forecasted demand for stage t, $D_{s,t}$ is the secular trend component of demand for time t, $D_{p,t}$ is the periodic trend component, $\phi_{i,p}$ are the autoregressive weights, p is the order of the linear autoregressive model, and μ_{D_r} is the mean of the independent random component. The $D'_{r,t-i}$ values are the i previously observed demand values after removal of the secular and periodic trends. If $t-i < 1$, then these values are set to zero [Kottegoda, 1980].

This formula is used in both the simulation and optimisation processes to estimate the demand for the current stage based on the stochastic model and the demands that were actually realised during the previous i stages.

5.3.3. Synthetic Demand Data Generation

A very similar process to the above is used to generate synthetic demand data. Instead of using the previously observed demand values for $D'_{r,t-i}$, however, the previously calculated values are used. Also, instead of using the mean of the independent random component, random numbers are generated using the following algorithm.

```
1     function NormGen(ParmArray: TMatrixClass): double;
2     var
3        fac,r,v1,v2: double;
4     begin
5        if (gliset = 0) then
6        begin
7           repeat
8                 v1:= 2.0*random-1.0;
9                 v2:= 2.0*random-1.0;
10                r:= sqr(v1)+sqr(v2);
11          until (r < 1.0);
12          fac:= sqrt(-2.0*ln(r)/r);
13          glgset:= v1*fac;
14          result:= v2*fac;
15          gliset:= 1;
16       end
17       else
18       begin
19          result:= glgset;
20          gliset:= 0
21       end;
22       Result:= Result * ParmArray.GetElement(0,2) +
         ParmArray.GetElement(0,1);
23    end;
```

The variables glgset and gliset must be declared externally to this function as they need to be preserved from one function call to the next. The random function call in lines 8 and 9 use the built in random number generator to generate a uniformly distributed random number between 0 and 1.

The algorithm generates two uniform variates and transforms them into two normal variates. One is returned to the calling function, and the other is preserved and returned the next time function NormGen is called. The calls to ParmArray.GetElement(0,2) and ParmArray.GetElement(0,1) in line 23 return the standard deviation and the mean of the required normal distribution. A χ^2 test on 1000 data points generated with this algorithm indicated that they fitted a standard normal probability distribution with a significance of 97.6%. The absolute maximum serial autocorrelation on lags up to 250 was found to be 0.110.

5.3.4. Running the Simulation

The actual simulation is calculated as follows. The user options are first checked, and the *System Description File* is read. If one of the existing demand data options was chosen, then the existing data is read, otherwise the synthetic demand data is generated as described above. If the existing pumping data option was chosen, then the pumping data out of the existing data file is used, otherwise the *Operation Policy File* is read. The variables are then initialised and the actual simulation begins.

For each stage, the expected demand is calculated from Eqn 5.13. This is used in the operation policy to calculate the pump setting. From the pump setting and the volume in the reservoir at the start of the stage, the actual pump rate can be calculated from the Flow equations in the *System Description File*. The new volume in the reservoir is then calculated, and the limits checked. If any of the limits are violated, then the penalties are calculated, and added to the running total for the applicable penalty. The storage cost is also found and added to its running total. If the maximum volume was exceeded then the new volume is reset to the maximum volume. This is also done if the new volume is below the minimum volume.

The Power equations, together with the pump setting and the pump rate, are used to calculate the power requirements for the stage. From this, the energy requirements and costs are found, and added to the running total. Finally, the current setting is compared to the previous setting and the switch cost calculated.

This process is repeated for each stage in the simulation. When the simulation is complete, the data generated is saved to a *System Run File*, and a *Simulation Report File* is written. This shows the names of all the input and output files, the period of the simulation, the number of times limits were violated, and a cost breakdown. The cost breakdown includes each of the penalty costs, the storage cost, the power cost and the switching cost.

5.4. Optimisation

The main reason that this program was created was to assist reservoir system operators in improving the efficiency of their systems. One of the ways of doing this is to find operation policies that are better than the ones currently being used, in other words, to optimise the operation policy. In this context, an optimum operation policy is one that minimises the total running cost for a reservoir system, for a given set of demand data and a given system description.

RSPO provides two different ways to do this, *Downhill Simplex Optimisation* and *Dynamic Programming Optimisation*. Downhill Simplex Optimisation finds the set of *Change Volume* and *Don't Change* parameters that minimise the running cost, whereas Dynamic Programming Optimisation finds the best possible pump setting for each stage of the given data set. Dynamic Programming Optimisation provides the best possible way to operate a system for a given set of demand data, but does not give any help in running the system in the future. Downhill Simplex Optimisation, however, provides a policy that can continually be used until some component of the system changes.

5.4.1. Downhill Simplex Optimisation

The Downhill Simplex Algorithm is a technique for finding the minimum of a function that has more than one independent parameter. In this context the parameters are each of the *Change Volumes* and the *Don't Change* value. To describe how this algorithm works, a simple example system will be used. This system has only two pump settings, *High* and *Low*, so it will have only two *Change Volume* parameters. This example system will also not use the *Don't Change* parameter. Because there are only two independent parameters, this can be described as two dimensional optimisation.

If x is used to represent the maximum reservoir volume at which the *High* pump setting is used, and y represents the maximum volume at which the *Low* pump setting is used, then the operation policy will be as shown in Figure 5.2 below.

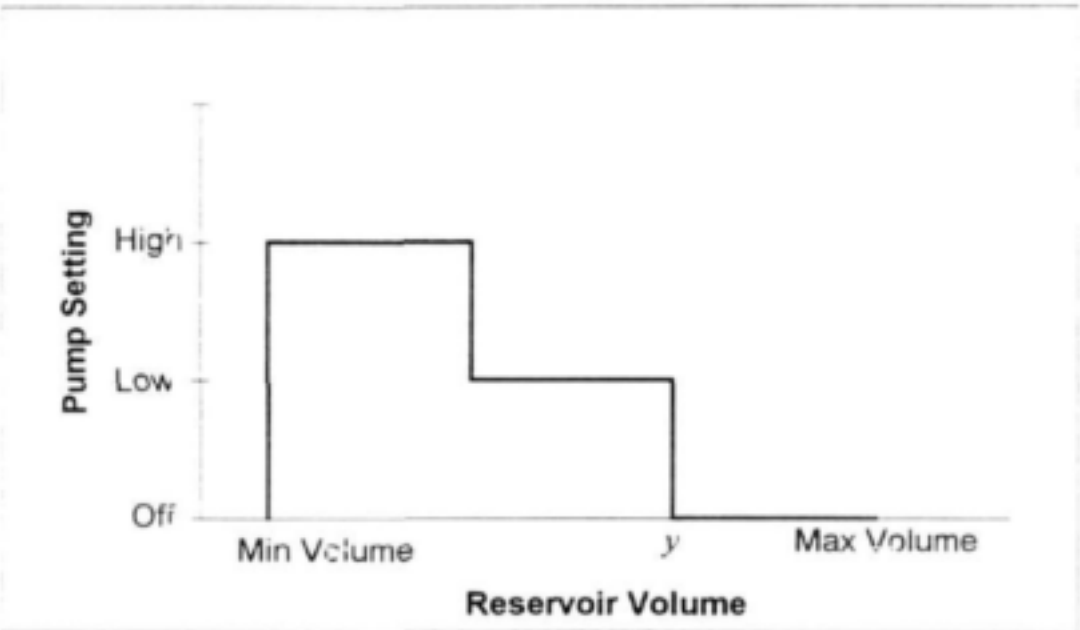

Figure 5.2: Example Operation Policy

To optimise this operation policy, it is necessary to find the values of x and y that result in the lowest operating cost for the system. Each pair of values $(x;y)$ represent a point on a plane, and the system operating cost for those two values represents a height above the plane. There are also constraints on the values that x and y can take on. They cannot be more than the maximum reservoir volume, and x must be less than y. The heights of all the points within the constraints define a cost surface. The object is to find the lowest point of this surface.

A simplex is defined as a figure with one more vertex than there are dimensions, so for this example system there are two dimensions so the simplex has three vertices, that is a triangle. The Downhill Simplex Algorithm finds the lowest point on the cost surface by "walking" the simplex downhill until it can go no further, using one of the four steps shown below:

1. Reflection - the highest point is reflected about the point that lies at the centre of the other two points. If this point is lower than the lowest point, then it replaces the highest point as one of the vertices of the simplex. The closed dots (●) are the old simplex and the open dot (O) is the new vertex.

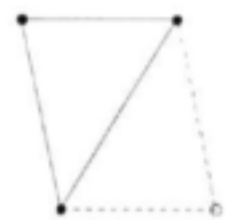

Figure 5.3: Reflection

2. Reflection and Extension - the highest point is reflected about the centre of the other two points. If this new point is lower than the current lowest point, then a further extension is tried. If the further point is lower than the current lowest point then this new point replaces the current highest point in the vertex. If it is not lower than the reflected point, then the original reflected point replace the highest in the simplex.

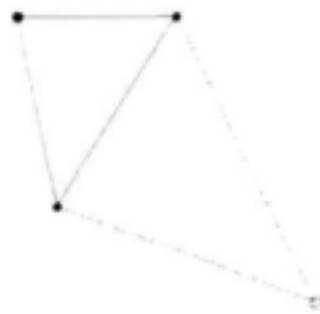

Figure 5.4: Extension

3. Single Axis Contraction - if a reflected point is not lower than the lowest point, then a point halfway between the highest point and the centre of the other points is tried. If this is lower than the lowest point then it replaces the highest point in the simplex.

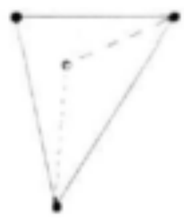

Figure 5.5: Single Axis Contraction

4. All Axes Contraction - if no point can be found that is lower than the lowest point, using any of the steps above, then the simplex is contracted about the lowest point.

Figure 5.6: All Axes Contraction

The Downhill Simplex Algorithm is an iterative technique that repeatedly applies one of the steps above until a simplex that is sufficiently small encloses the lowest point. One of the options on the Optimisation Dialog screen is the Tolerance value. When the cost evaluated at each of the vertices in the simplex are within the Tolerance value of each other the Downhill Simplex Algorithm stops iterating. The point returned as the solution to the optimisation is the point that lies at the centre of the vertices of the simplex [Press et al., 1986].

The other option on the Optimisation Dialog screen is the Maximum Iterations value. This value is used to ensure that the Downhill Simplex Algorithm does not continue iterating indefinitely. This can sometimes happen if a sequence of steps returns the simplex to one of its previous positions. If this happens, the simplex will "walk" around in a circle forever. If the algorithm stops because it reached the maximum number of iterations, it is recommended that the optimisation be restarted using a different starting position.

The algorithm used by *RSPO* includes a slight modification, which is necessary to take account of the constraints on the values of x and y. Each time a new point is calculated, it is checked to ensure it is within the boundaries. If it is not, the nearest point on the boundary is used instead.

The main problem with the Downhill Simplex Algorithm, as with most iterative optimisation algorithms is that it can only find a local minimum, and it is not possible to know if the solution is the global minimum. It is therefore recommended that the optimisation be started from a number of different starting points. If all the optimisations converge to the same point, then it is fairly likely that this is the global minimum. If they do not all converge to the same point, then the lowest solution is most likely to be the global optimum. The recommended starting points for the example system described above would be at each of the vertices of the constraints and in the centre of the constraints, as listed in the following table:

Table 5.1: Recommended Starting Points

Starting Point	x	y
1	Minimum Volume	Minimum Volume
2	Minimum Volume	Maximum Volume
3	Maximum Volume	Maximum Volume
4	$\frac{1}{3}$(2Min Vol + Max Vol)	$\frac{1}{3}$(Min Vol + 2Max Vol)

The *Optimisation Report* produced when a Downhill Simplex Optimisation is run shows the number of iterations taken to find the solution and the total running cost at the optimum point. It then lists the values for each of the dimensions in the solution with a plus and a minus value. These are the distances from the centre point of the simplex to the two points furthest from the centre. This information provides an indication of how close the returned point is to the actual minimum point. For example, if the final simplex in the example system was made up of the points *(1;2)*, *(1;5)* and *(4;5)* then the resulting values for x and y would be:

Table 5.2: Resultant Optimum Point

Dimension	Value	Plus	Minus
x	*2*	*2*	*-1*
y	*4*	*1*	*-3*

The above gives a simplified explanation of how the Downhill Simplex Algorithm works in two dimensions. The algorithm works in the same manner for as many dimensions as are necessary, although it is not possible to visualise the process for more than three dimensions.

5.4.2. Dynamic Programming

The Optimisation Dialog screen of *RSPO* also offers the user the option of optimising the system using Dynamic Programming (DP). DP is an entirely different method of optimising a system to the Downhill Simplex method. DP does not need the user to define an operating policy for the system, the parameters of which need to be optimised. Rather, DP looks at the total period of data that is being optimised, and finds the optimum pump setting for each stage. This will usually give a lower total operating cost than the Downhill Simplex method, but it can only optimise the current data set. DP gives no help on how to run the system in the future.

How DP works is described in Chapter 3.2 Dynamic Programming. *RSPO* defines the decision variable as the new pump setting required for the current stage. The state variable is the volume of water in the reservoir at the end of the stage. The DP algorithm requires that the state variable can only take on a number of discrete values. For this reason, the available volume in the reservoir must be broken down into a number of states. The number of states used is chosen by the user, and entered on the Optimisation Dialog screen. The higher the number of states, the better the resolution of the solution, but the longer it will take to find a solution. The number of states is also limited by the programming language used so that the maximum number of states is $\frac{8192}{N}$, where N is the number of stages in the data set.

The stage-to-stage transformation equation is

$$V_t = V_{t-1} + \boldsymbol{P_r}(P_{s,t}, V_{t-1}) - D_t, \qquad \text{(Eqn 5.14)}$$

where V_t is the volume of water in the reservoir at the end of stage t, V_{t-1} is the volume in the reservoir at the end of the previous stage, D_t is the demand for the stage, $P_{s,t}$ is the pump setting, and $\boldsymbol{Pr(\,)}$ is the function that calculates the pumping rate for the stage from the pump setting and the reservoir volume at the start of the stage. The volume calculated using Eqn 5.14 will fall into one of the reservoir states. The volume that is the centre of this state will then be used as the new volume.

The Optimal Return Function then becomes:

$f_t(V_t) = \min_{P_{s,t}} \left\{ C(V_t, P_{s,t}) + f_{t-1}(V_{t-1}) \right\}$, where $C(\,)$ is the cost incurred at state V_t using a pump setting of $P_{s,t}$. The boundary condition used to solve this is the starting reservoir volume given in the *System Description File*.

The optimum found using either of the above methods is based on the historical data used in the optimisation. If the optimisation is repeated at regular intervals then the optimum policy will remain up-to-date. This way, any fluctuations in weather patterns can be accounted for. The optimisation process takes account of the periodic nature of the weather and the demands, but it cannot make allowance for sudden radical changes in demand. An example of this would be a period of dry weather, followed by a few days of rain. As the dry weather continues, people will start to water their gardens, pushing up the demand. As soon as the dry period ends, people will stop watering their gardens, leading to a sunned drop in demand. The optimum policy can unfortunately not predict, or make allowance for this situation. An experienced human operator will be aware that as soon as the dry spell ends the demand will drop. He should therefore override the setting suggested by the optimum policy, and use a lower pump setting. As this example demonstrates, this software cannot replace an experienced human operator for running a reservoir system, but is designed to complement the operator.

6. CASE STUDY RESULTS

To fully test the model and the simulation and optimisation results, an existing system was analysed and modelled. A description of the system, and the results obtained are presented in this chapter.

6.1. System Description

6.1.1. Rand Water's Supply System

The Rand Water system was chosen as an initial monitoring site because of its convenience, and the fact that it is the largest bulk water supplier to urban areas on the continent. In their 1995 financial year, Rand Water supplied a record annual supply of 1 033 953,7Ml, or a daily average of 2 832.75Ml/day [Rand Water, 1995].

Rand Water's overall system consists of 2 major purification works, 10 pump stations, 49 reservoirs and 2 584km of pipelines with a diameter larger than 155mm. The reservoirs have a total capacity of 4 034Ml, which is equivalent to approximately 34 hours of storage. The limits of supply cover an area of 16 807km^2, and contain a population of approximately 9 million people. These limits are shown in Figure 6.1. Water is abstracted from the Vaal river at the Vereeniging and Zuikerbosch works. A small amount of water is also abstracted from boreholes at the Zuurbekom pump station.

Rand Water is the intermediate supplier in a three tier supply system. Rand Water are allocated their water from the Department of Water Affairs and Forestry (DWAF). The DWAF is a government ministry that is responsible for the management of the entire country's water resources. Rand Water, in turn, manages the water resources for its region of supply, which is approximately the Gauteng Province. Rand Water then supplies water to the local authorities and bulk consumers in this area. These form the third tier in the system, and supply water to individual consumers.

Because of the size and complexity of the entire Rand Water system, and the expenditure required to monitor the system adequately, it was felt that it would not be feasible to optimise the entire system as a single entity. A number of subsystems have been identified for initial simulation and optimisation. Once the functioning of these systems has been fully understood, further work can be done on analysing the inter-relationships between subsystems and optimisation of the entire system. The subsystems identified for monitoring, simulation and optimisation are discussed below.

Figure 6.1: Rand Water's Limits of Supply

6.1.2. Sub-systems chosen for Modelling

It was decided that initially a simple system would be analysed and modelled. Once this was complete, and a better understanding of the requirements and complexities of the optimisation process was gained, a larger and more complex subsystem could be investigated.

The initial system should contain a single pump station and a relatively small number of reservoirs. It should also be possible to isolate the subsystem so that a complete water balance could be calculated for the period of monitoring. This would allow for a detailed description of what is happening to the water in the system at all times, and an accurate understanding of the user demand patterns. Once the user demands have been calculated, they can be analysed and used to optimise the pumping policy.

Discussions with Rand Water indicated that two subsystems fulfil these requirements and would be suitable for initial monitoring. The first of these systems is the Libanon/Driefontein system. It contains the Libanon pump station, which supplies the Driefontein reservoir. This subsystem is located on the far West Rand, as shown in Figure 6.2.

Figure 6.2: Libanon/Driefontein Subsystem

The second subsystem is the Bloemendal/Wildebeesfontein subsystem. This contains the Bloemendal pump station, and three reservoirs known as the Wildebeesfontein Reservoirs. This subsystem is located in the far East Rand. This subsystem is shown in Figure 6.3.

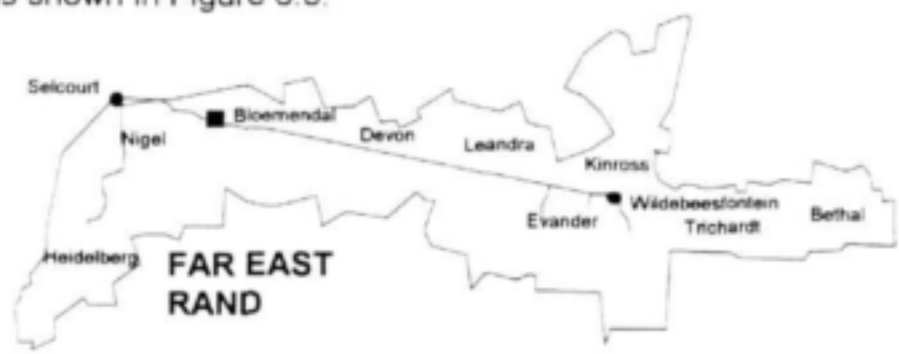

Figure 6.3: Bloemendal/Wildebeesfontein subsystem

A further subsystem was identified that could be used for a second stage of analysis and optimisation. This is the Mapleton subsystem. It includes the Mapleton pump station and 6 reservoirs. This subsystem also supplies water to the Bloemendal pump station, so the Bloemendal/Wildebeesfontein subsystem forms part of the larger Mapleton subsystem. Once the analyses of these two subsystems are complete, it will prove valuable to compare them, and see how they relate to one another.

It was decided to use the Libanon/Driefontein subsystem as the initial system for monitoring and model verification, mainly because this subsystem contains more monitoring equipment already installed. It will also be easier to isolate this system to ensure that accurate results are obtained.

6.2. Available Data

This section gives a description of the data that is monitored and recorded on a daily basis by Rand Water's Operations Department.

Rand Water reports the output from each pumping station on a daily basis at 06:00. The levels and contents in the reservoirs are also reported at this time, as well as at midnight and 16:00 each day. Currently, decisions to alter pumping rates are made in the morning, based on the reports of pumping rates and reservoir levels. The reservoir levels are also checked again in the afternoon, and adjustments can be made to the pumping rates at this time, but this is usually not necessary.

6.2.1. Morning Daily Distribution Log

This report is made available at 06:00 every morning and is used by the Operations department to decide whether to change the pumping rates from the previous day. An example of this report is shown in Appendix A: Available System Data. It includes the following information:

<u>Reservoir levels and content</u>

Each reservoir in the system is listed, and the level and content in the reservoir are shown for 06:00 on the current date, and 06:00 and 16:00 on the previous date. Any increase or decrease in the content since the previous day is also shown. The total available space, that is, the difference between the current content and the maximum available content, is also shown.

A similar list also shows the same reservoir information for 00:00 on the current and the previous days.

<u>Weather conditions</u>

The rainfall at 14 stations is recorded on a daily basis, and this is shown for the previous day on the morning report. The barometric pressure, temperature and humidity at Johannesburg International Airport at 06:00 on the current and previous mornings are also shown.

<u>Pumping rates</u>

The actual volumes pumped to and away from each pump station for the previous 24 hours up to 00:00 are recorded. The suction and delivery heads, and the average load, maximum demand, load factor and other information are also shown for some of the pump stations. The pumping section of the report also includes any known pipeline or reservoir losses. An 'assumed consumption' is calculated from the change in total reservoir content and the total volume pumped.

6.2.2. Monthly User Consumption Report

This report is broken down by pipeline. All the users on a particular pipeline are listed, with their monthly consumption over the past 6 months. The total for each pipeline is also calculated. An example of this report is shown in Appendix A: Available System Data.

6.2.3. Pumping Policy

Discussions were held with Rand Water concerning their current pumping policy. It was felt that, in order for any optimised policy to be accepted, it should fit in as well as possible with the existing system.

The daily operation policy used by Rand Water is currently based entirely on the operator's experience. There is no formal written policy used by the operators to help them decide when, and by how much, to change the pumping rates. At present, the operator receives a report at 06:00 and 16:00 each day. He will then decide, based on the total volume of water in all the system reservoirs, the contents of specific reservoirs, the current pumping rates and the current weather conditions; whether or not to change the pumping rates. His decisions are limited by a monthly quota that is given to him by the planning department in Head Office. These monthly quotas are based on historical averages for the month in question, and predictions of demand growth.

A set of extension routines was written for the computer program that allows the importing of this available data directly from the Rand Water file formats. The extensions also include a routine to export the weather data to a text file. For details on these extensions, see the Appendix D: Rand Water Extensions User Manual. Similar extensions could be written for any Water Supply Authority after examining the file formats of the available data.

6.3. Analysis

As discussed earlier, three subsystems were chosen for initial analysis. Six months worth of data was extracted from Rand Water's records and analysed using *RSPO*. The results are presented below:

Table 6.1: Analysis Results

System	Libanon/Driefontein Subsystem	Bloemendal/ Wildebeesfontein Subsystem	Mapleton Subsystem
Input Data File:	RSPO\LIB\LIB.RUN	RSPO\BW\BW.RUN	RSPO\MPL\MPL.RUN
Analysis Report File:	RSPO\LIB\LIB.ANR	RSPO\BW\BW.ANR	RSPO\MPL\MPL.ANR
Stochastic Model File:	RSPO\LIB\LIB.MDL	RSPO\BW\BW.MDL	RSPO\MPL\MPL.MDL
Data Range			
Start Date & Time	15/11/95 06:00	15/11/95 06:00	15/11/95 06:00
End Date & Time	12/05/96 06:00	12/05/96 06:00	12/05/96 06:00
Data Points	180	180	180
Minimum Demand	19.16 Ml/day	33.43 Ml/day	46.03 Ml/day
Maximum Demand	69.68 Ml/day	59.37 Ml/day	448.11 Ml/day
Minimum Pumping Rate	19.65 Ml/day	20.95 Ml/day	36.40 Ml/day
Maximum Pumping Rate	66.76 Ml/day	65.21 Ml/day	428.81 Ml/day
Secular Trend Parameters			
Type	Linear	Linear	Linear
Slope	-0.0337752708416917	0.025977746226735	0.211962612014774
Intercept	56.2718311847758	45.4248250460405	329.078124002455
R^2	0.0478160553849891	0.069607585855137	0.0495562818090543
Periodic Trend Components			
Significance Level	95%	95%	95%
Annual Cycle:	No Cycle Calculated	No Cycle Calculated	No Cycle Calculated
Weekly Cycle	0.0185984126984373	0.0066611721611768	0.129409035409103
Period of 7 Days			
Alpha	-1.49565921055988	-1.28594347735551	None
Beta	-2.83107663492425	-2.28715988505167	
Significance	99.9636137683068%	99.9998610885 18%	
Period of 3.5 Days			
Alpha	-2.03839456339651	None	None
Beta	-0.161319693733 24		

System	Libanon/Driefontein Subsystem	Bloemendal/ Wildebeesfontein Subsystem	Mapleton Subsystem
Significance	96.37201863337765%		
Daily Cycle	No Cycle Calculated	No Cycle Calculated	No Cycle Calculated
Estimated Parameters for Residual Stochastic Component			
Mean	0.00138188586646471	-0.0024503463152552	-0.129409035409036
StdDev	7.35387825011152	4.5704022296834	48.233445674355
Skew	-2.22349057666834	-0.197846105903257	-2.94279314409572
Kurt	11.227225833007	3.18317459732598	19.4165231681271
Serial autocorrelations			
Correlation Coefficients at Lag:			
0	1	1	1
1	0.701062386158384	0.083933790144456	0.257317292014276
2	0.548878474945079	0.11951918694044	0.189597461140394
3	0.463176962038485	-0.0797039098621735	0.102154528947518
4	0.316486163596267	0.0570633873592156	0.289680625642453
5	0.183619960720542	0.042232127133103	-0.0283529933976727
6	0.0944963105141409	0.147349770426641	0.0696094252983905
7	0.0714724566914906	-0.0445390404088936	0.0708067020242583
Autoregressive Component Parameters			
Initial $f_{1,1}$	0.701062386158384	No parameters with a	0.257317292014276
Initial Value	710.886922634639	significance greater	5155.73258415214
Best fit $f_{1,1}$	0.585310014373619	than 73%	0.357703648367105
Iterations:	34		32
Best Value	700.867418176444		5115.17150530555
Estimated Parameters for Independent Random Component			
Mean	-0.000235858848842080	-0.000188343959082691	-0.0214326403613605
StdDev	5.33644556519468	4.55455321443536	46.8844994301
Skew	-0.585781382189186	-0.174181880210224	-3.16311903940227
Kurt	5.08183373853231	3.23162880546901	22.6604454873142
Initial c^2	10.8799	20.4050	42.0722
Iterations	127	98	114
Minimum c^2	3.2747	18.8356	4.5519
Best Fit Parameters for Independent Random Component			
Mean	0.102745122164684	-0.16618387650817	2.50310780071958
StdDev	4.45537476662109	4.24987764286005	30.0367123452254
c^2 significance	99.3274285495774%	9.25738245409982%	97.1300014455736%

All three systems have a linear secular trend over the six months of available data. The Libanon/Driefontein Subsystem has a negative slope that indicates that the demand decreased slightly over the analysis period.

Annual cycles were not calculated because at least two years worth of data would be required for this. Daily cycles were also not calculated as this is only possible if the data available has a time step of 12 hours or less, and Rand Water monitor their system on a daily basis. Both the Libanon/Driefontein and Bloemendal/Wildebeesfontein Subsystems evidenced periodic demand patterns with a high significance level (greater than 95%). This was expected because of the difference in demand between weekends and weekdays. The Mapleton Subsystem only showed periodic trends at much lower significance levels. This is probably because the Mapleton Subsystem is a much larger subsystem than the other two, so differences in demand are less pronounced.

The Libanon/Driefontein Subsystem shows serial correlations of greater than 30% at lags up to 4 days, while the Bloemendal/Wildebeesfontein Subsystem has relatively low serial correlations, not rising above 15%. These are reflected in the autoregressive parameters. The low serial correlation in the Bloemendal/Wildebeesfontein Subsystem results in no significant autoregression in the stochastic model, while the other two subsystems use an autoregressive model of order 1.

The parameters of the independent random component are estimated using the method of moments on the data after the autoregressive model has been removed. Best fit parameters are then found by minimising the χ^2 value. A Normal distribution fits both the Libanon/Driefontein and the Mapleton Subsystems very well, but the Bloemendal/Wildebeesfontein Subsystem shows a low χ^2 value, indicating a poor fit.

6.4. Simulation

There are two reasons for running simulation, namely to generate synthetic demand data, and to compare the running costs of different operation policies. Cost comparison simulations will be discussed later in the section on Optimisation.

The historical demand data available spans a period of 6 months starting from 15/11/1995. below shows the analysis results for the historical data as well as for 6, 30, and 60 months worth of synthetic demand data.

The demand data sets generated by these simulations are also shown in Figure 6.4, with the historical demand for comparison. Each data set is displaced by 50 Ml/day on the ordinate axis for clarity. The decrease in demand with time is a result of the fact that only six months of data were used to fit the stochastic model used to generate these results. The six months of historical demand data start in summer and end in winter, giving a decrease in average demand. If the full years worth of data were used, this slope would not be evident.

In the Secular trend parameters, the intercept value is based on the start date of the data set, so the earlier the data set starts, the greater the intercept value will be, for a negative slope. The longer the synthetic data set, the more significant the secular trend becomes, as evidenced by the increasing R^2 values. In the historical data set, the downward trend is part of a seasonal cycle. This becomes exaggerated when the trend is extrapolated to 60 months.

The periodic trend generated remains consistent as the length of the synthetic data set grows. The serial correlations also appear consistent at lag 1. The serial correlations of the synthetic data sets decrease more rapidly after lag 1, but this is because the model only uses a first order autoregressive model. If a lower significance level were used when analysing the historical data, a higher order model may be calculated.

The calculated best fit $\phi_{1,1}$ parameter is sufficiently close in all four data sets, as are the parameters of the normal distribution used to generate the independent random component of the demand data.

These results show that the stochastic model used to generate synthetic data does in fact generate data that is statistically consistent with the historical data set. This then allows the use of synthetic data to test different operating policy scenarios over time periods longer than the available historical data. Also, by changing certain parameters of the stochastic model, for example the linear trend slope, the effects of changes within the system can also be modelled.

Table 6.2: Synthetic Demand Data

Data Type and Length	6 Month Historical	6 Month Synthetic	30 Month Synthetic	60 Month Synthetic
Output Data File:	RSPO\LIB\LIB.RUN	RSPO\LIB\LIB06.RUN	RSPO\LIB\LIB30.RUN	RSPO\LIB\LIB60.RUN
Analysis Report File:	RSPO\LIB\LIB.ANR	RSPO\LIB\LIB06.ANR	RSPO\LIB\LIB30.ANR	RSPO\LIB\LIB60.ANR
Stochastic Model File:	RSPO\LIB\LIB.MDL	RSPO\LIB\LIB.MDL	RSPO\LIB\LIB.MDL	RSPO\LIB\LIB.MDL
Data Range				
Start Date & Time	15/11/95 06:00	15/11/95 06:00	20/11/94 06:00	27/08/93 06:00
End Date & Time	12/05/96 06:00	12/05/96 06:00	07/05/97 06:00	31/07/98 06:00
Data Points	180	180	900	1800
Minimum Demand	19.16 Ml/day	37.042 Ml/day	21.977 Ml/day	7.847 Ml/day
Maximum Demand	69.68 Ml/day	68.863 Ml/day	80.340 Ml/day	99.293 Ml/day
Secular Trend Parameters				
Type	Linear	Linear	Linear	Linear
Slope	-0.033778	-0.037523	-0.03424	-0.03452
Intercept	56.27183	56.24333	68.67061	84.19499
R^2	0.04782	0.08946	0.68097	0.89401
Periodic Trend Components				
Significance Level	95%	95%	95%	95%
Annual Cycle:	No Cycle Calculated	No Cycle Calculated	No Cycle Calculated	No Cycle Calculated
Weekly Cycle	0.01860	0.01860	0.00117	0.00088
Period of 7 Days				
Alpha	-1.49566	-1.76970	-1.90074	-1.521733
Beta	-2.83108	-3.19194	-2.48056	-2.92221
Significance	99.96361%	99.99999%	1	1
Period of 3.5 Days				
Alpha	-2.03839	-1.76259	-2.15493	-2.09319
Beta	-0.16132	-0.04434	0.00340	-0.10845
Significance	96.37202%	98.71412%	99.99999%	1
Daily Cycle	No Cycle Calculated	No Cycle Calculated	No Cycle Calculated	No Cycle Calculated
Estimated Parameters for Residual Stochastic Component				
Mean	0.00138	-0.00056	-0.00003	0.00008
StdDev	7.35388	5.51948	5.46564	5.52613
Skew	-2.22349	-0.07489	-0.09998	-0.03374
Kurt	11.22723	3.37912	2.84839	2.96620
Serial autocorrelations				
Correlation Coefficients at Lag:				
0	1	1	1	1
1	0.69717	0.56922	0.61742	0.57763
2	0.54278	0.34641	0.40871	0.31779
3	0.45546	0.25846	0.23693	0.18413
4	0.30945	0.21528	0.16712	0.12759
5	0.17852	0.11115	0.11096	0.10045
6	0.09135	0.01927	0.08020	0.05074
7	0.06869	-0.06156	0.00072	0.00846

Data Type and Length	6 Month Historical	6 Month Synthetic	30 Month Synthetic	60 Month Synthetic
Autoregressive Component Parameters				
Best fit $\phi_{1,1}$	0.58531	0.56578	0.60580	0.59188
Estimated Parameters for Independent Random Component				
Mean	-0.00024	0.02299	-0.00305	-0.00176
StdDev	5.33645	4.55518	4.30321	4.51307
Skew	-0.58578	-0.01616	0.00790	-0.03135
Kurt	5.08183	3.69900	3.12702	2.87637
Best Fit Parameters for Independent Random Component				
Mean	0.10275	0.03672	-0.02939	0.00396
StdDev	4.45537	4.39060	4.33559	4.52903
χ^2 significance	99.32743%	97.72985%	99.27867%	97.47465%

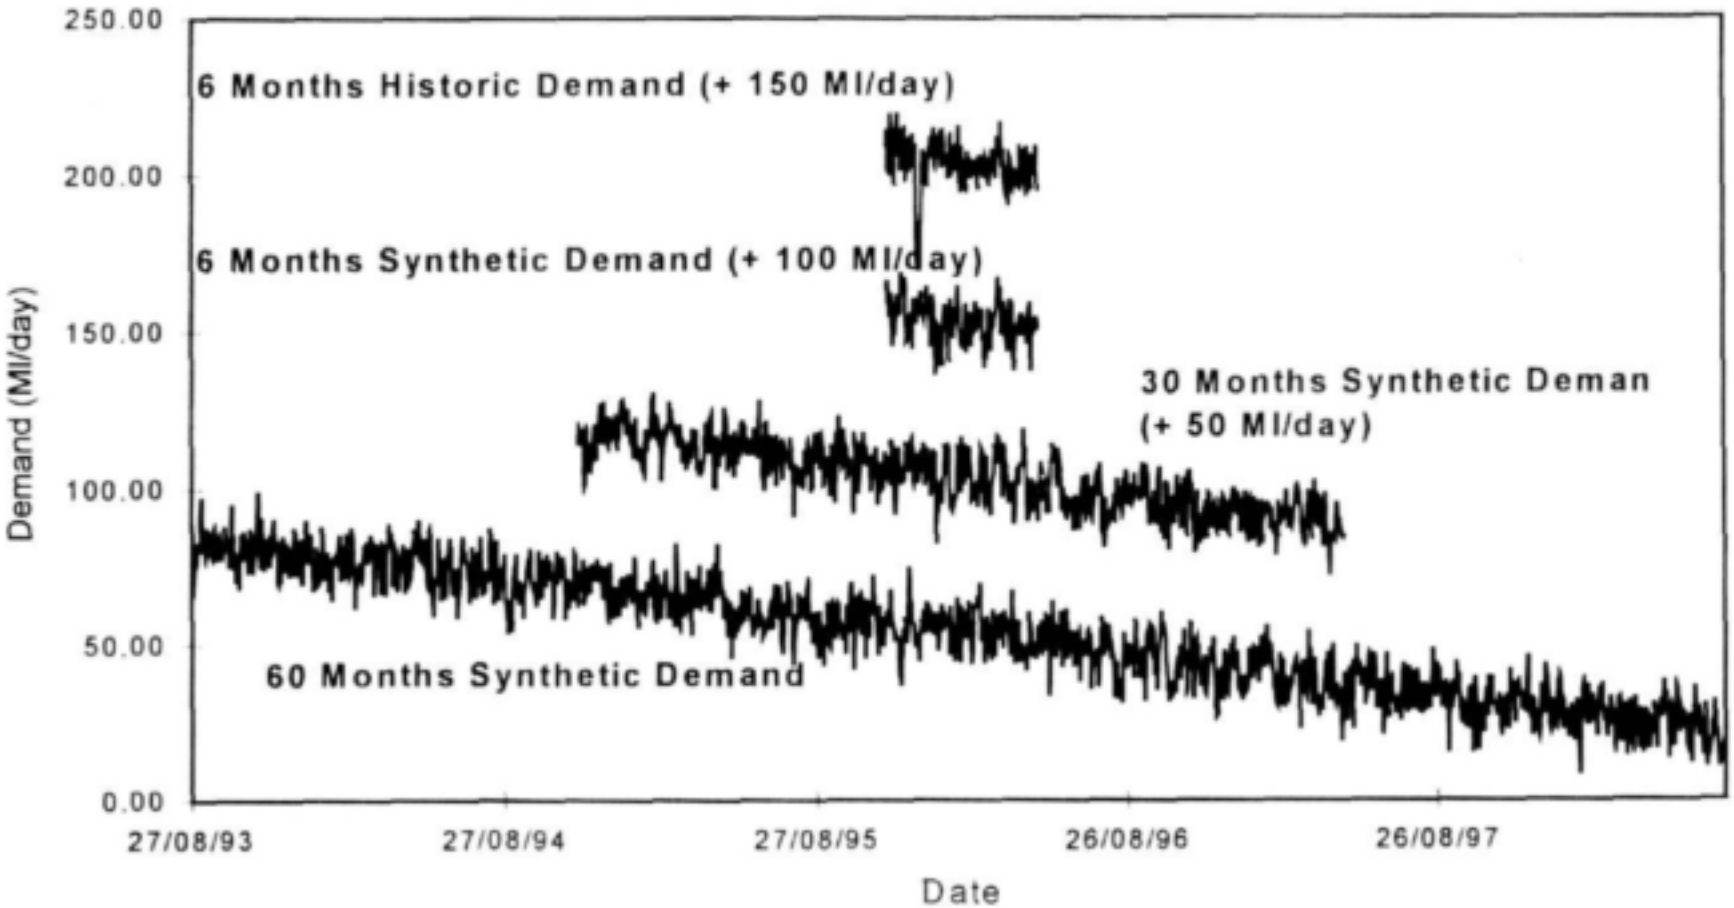

Figure 6.4: Synthetic Demand Data

6.5. Optimisation

The main objective of the *RSPO* program is to find an optimum operation policy. Two methods are provided for doing this, the Downhill Simplex method, and the Dynamic Programming method. The Downhill Simplex method is used to calculate an operating policy that is optimised for a particular set of data, but that can be used for the day-to-day operation of the reservoir system. The Downhill Simplex method is a non-linear search technique, and suffers from the disadvantage of all non-linear search techniques, namely that they can only find local optimum points which are not necessarily the global optimum point.

Dynamic Programming, on the other hand, will always find the global optimum for a particular set of data, but it cannot calculate a policy that can be used for the day-to-day running of a reservoir system. The Dynamic Programming method is provided as an option in the *RSPO* program so that the solution found with the Downhill Simplex method can be compared to the global minimum solution.

In the next sections the effect of local minima on the solution will be examined, as well as the effects of the length of the data set and the cost functions used in the optimisation.

6.5.1. Local and Global Minima

If the system being optimised has only a few local minima, then the chance of finding the global minimum using a non-linear search technique is good. The standard way of trying to find the global minimum in the presence of local minima is to start the search from a number of different points. Each search will find the local minimum closest to its starting point. If more searches are run than there are minima, then it is highly likely that the global optimum will be found. The problem with most systems is that the number of local minima are not known in advance, and with some systems there are very many local minima. This usually makes it very difficult to find the global minimum.

The Libanon/Driefontein subsystem was examined to check the effect of local minima on the determination of an optimum operation policy. Optimisations were run using four different starting points. The starting point labelled "Even Spread" in Table 6.3 below is the default used by *RSPO* if no initial starting point is given. The starting levels are evenly spread between the maximum and minimum volume values for the reservoir. This point is essentially the point at the centre of the constraints.

The points labelled "Minimum" and "Maximum" have all their starting levels equal to the minimum volume and maximum volume of the reservoir, respectively. For the Libanon/Driefontein subsystem these values are 20Ml and 124.354Ml. The point labelled "Centre" has all levels set to the average of the minimum and maximum volume values.

Table 6.3: Optimisation Results from Different Starting Points

	Even Spread	Centre	Minimum	Maximum
Starting Operating Rule				
Rule File	EVEN.ROR	CENT.ROR	MIN.ROR	MAX.ROR
Level 1	111.310	72.177	20.000	124.354
Level 2	98.266	72.177	20.000	124.354
Level 3	85.221	72.177	20.000	124.354
Level 4	72.177	72.177	20.000	124.354
Level 5	59.133	72.177	20.000	124.354
Level 6	46.089	72.177	20.000	124.354
Level 7	33.044	72.177	20.000	124.354
Don't Change	0.2500	0.250	0.000	0.500
Running Cost	R757 814.06	R894 317.83	R391 504.72	R1 353 463.18
Optimum Operating Rule				
Rule File	EVEN_O.ROR	CENT_O.ROR	MIN_O.ROR	MAX_O.ROR
Level 1	79.470	30.838	23.760	122.723
Level 2	67.557	26.095	23.760	121.092
Level 3	2.677	0.857	22.555	119.462
Level 4	-64.809	-65.892	-39.897	117.831
Level 5	-94.346	-103.763	-58.127	116.201
Level 6	-124.079	-145.735	-86.008	114.570
Level 7	-152.781	-192.305	-119.004	112.940
Don't Change	0.009	0.000	0.000	0.500
Running Cost	R322 081.54	R321 186.18	R317 506.20	R1 353 463.18

The table shows that each of the four optimisations found a different optimum operating policy. This indicates that there are local minima in the system, but the three best optimum points are close to one another, the total running costs are within less than 1,5% of one another. This indicates that there is a general global minimum, but there are small areas of local minima around it. While the simplex is sufficiently large, it moves "downhill" towards the global optimum, but when the simplex becomes too small, it becomes trapped in one of the local minima. Although the optimum solutions found are only local optima, they are very near to the global optimum.

The optimisation started at the point labelled "Maximum" did not progress to a minimum because all the points in the simplex returned the same value, as they were all very similar. This indicates that there exists a "Flat" region near the maximum values of the constraints, which makes it impossible to find the global minimum from this starting point.

A dynamic programming optimisation performed on the same input information as the above optimisations returned a total operating cost of R317 415.54 for the period optimised. This indicates that the downhill simplex method finds a very similar optimum point as the dynamic programming method, both of which will be very close to the global optimum of the system.

6.5.2. Effect of Data Set Length

In many optimisation situations, the length of the available data affects the accuracy of the optimum result obtained. To test this in the case of the Downhill Simplex optimisation method a data set of 60 months was synthetically generated from the 6 months of available historical data. Three sets of data were then extracted from this, each of differing lengths. The lengths and start and end dates of the four data sets are shown in Table 6.4 below. Optimum operation policies were then calculated using each of the four data sets, and these policies are also shown in the table.

Table 6.4: Optimisation Results with Different Length Data Sets

File	LIB_8.RUN	LIB_4.RUN	LIB_2.RUN	LIB_1.RUN
Data Points	225	450	900	1800
Start Date	24/10/1995	03/07/1995	20/11/1994	27/08/1993
End Date	04/06/1996	24/09/1996	07/05/1997	31/07/1998
Optimum Operating Policy				
Level 1	58.837	81.365	86.078	83.752
Level 2	49.675	63.779	68.388	66.444
Level 3	26.965	14.119	24.661	21.183
Level 4	-43.785	-43.131	-65.220	-56.593
Level 5	-64.626	-64.755	-96.145	-92.661
Level 6	-88.127	-88.687	-129.362	-129.529
Level 7	-110.257	-111.890	-161.734	-164.844
Don't Change	0	0	0	0
Running Cost	R3 312 086	R3 311 340	R3 311 890	R3 312 507

Once the optimum policies had been calculated, they were used in simulations with the previously generated 60 months of demand data, and the total running costs calculated are shown in Table 6.. There is less than 0.05% difference between the total running costs of all four operation policies, which indicates that the optimum policy calculated is not very sensitive to the length of data used to calculate it.

The results of this test also show that a policy optimised for a specific set of measured historical data will still be near optimum when used to operate the system in the future when the demands are not the same as the historical data, but are statistically similar. This indicates that it would not necessary to recalculate the policy every time step, but rather only when there are changes in system that would cause changes in the statistical properties of the demand.

6.5.3. Effect of Cost Functions

The most important factor that affects the optimum operating policy will be the combination of costs in the system. For a particular system, these costs will be described in the *System Description File (*.SDC)*. The optimisations calculated in the two previous sections all used the same file called *LIB/LIB.SDC* so that the results could be compared.

To examine the effects of each of the different costs described in this file, a number of optimisations were run, the results of which are shown below. In each case, all the costs except the one being discussed are set to zero. Each of the optimum policies was calculated for the same set of historical demand data, which is shown in Figure 6.5 below.

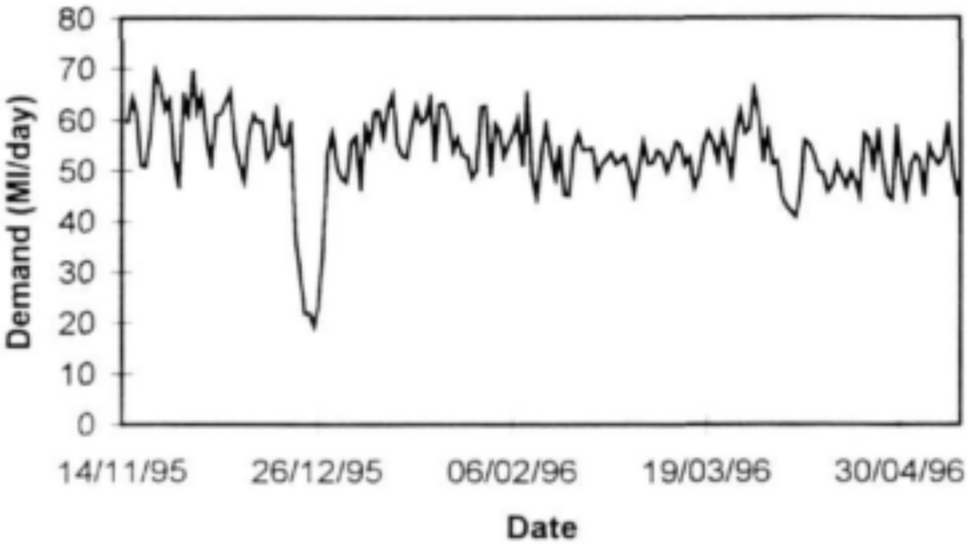

Figure 6.5: Historical Demand Data Graph

Pumping Costs

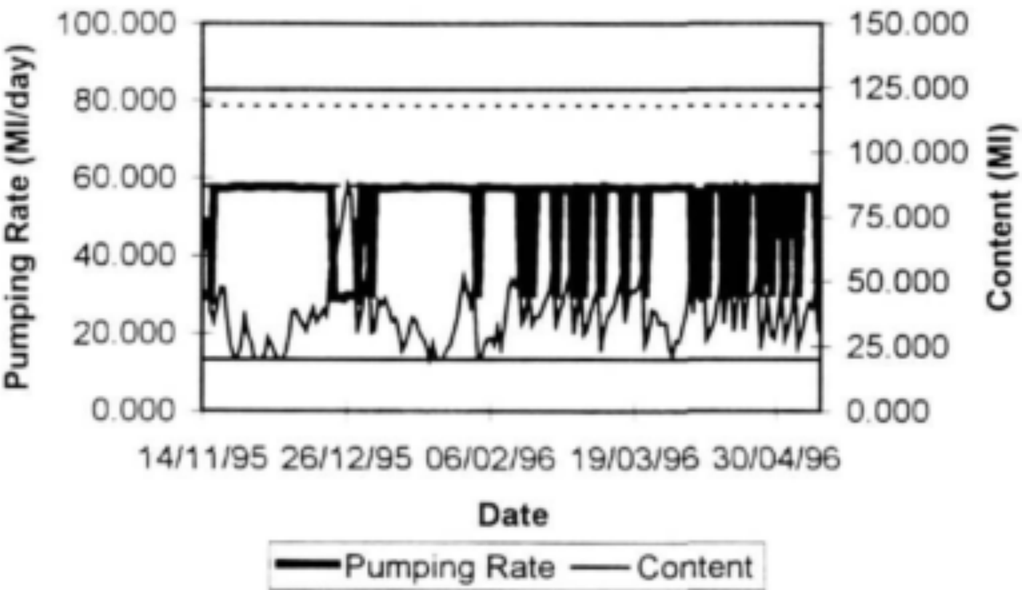

Figure 6.6: Minimum Pumping Cost Operation

In most systems the main cost of operating the system is the cost of the energy required to pump the water from the source to the reservoir that supplies the consumers. If this were the only cost, then the cheapest way of running the system would be to not pump at all, and this is the policy found by *RSPO*. This would naturally cause the reservoir to run dry almost immediately. To ensure the reservoir does not run dry the costs used in this optimisation included a penalty for running the reservoir dry. The operation of the reservoir with the optimum policy found when the only costs considered are the pumping cost and the penalty for running dry is shown in the graph above.

Figure 6.6 shows that the pumping rate, and therefore the level in the reservoir is kept as low as possible. This is what would be expected from a policy designed to minimise the pumping costs. The solid horizontal lines in Figure 6.6 indicate the Minimum and Maximum reservoir volumes, and the dashed lines indicate the Low and High Warning levels.

Switching Costs

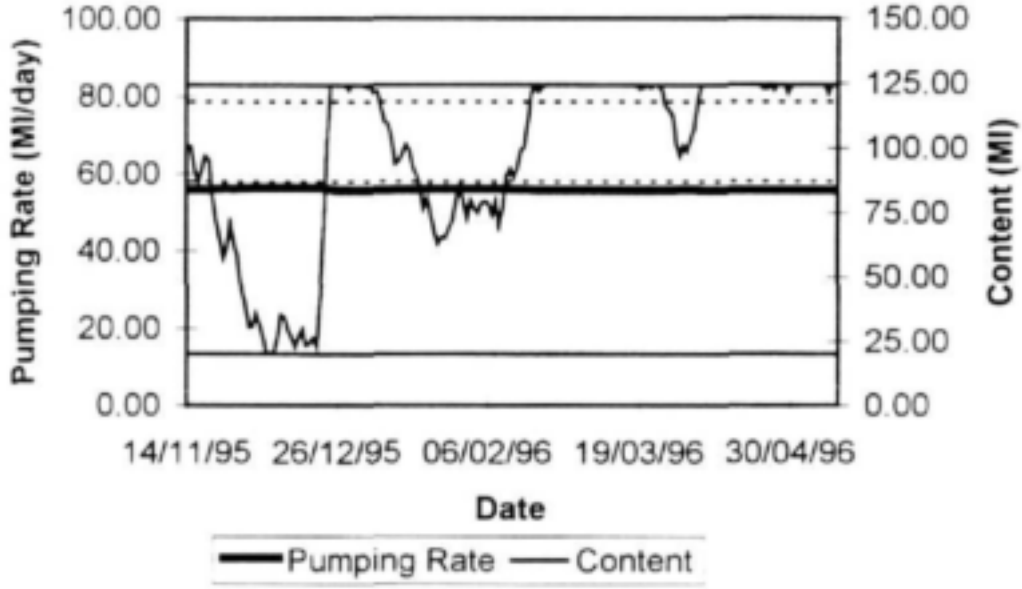

Figure 6.7: Minimum Switching Cost Operation

Another important cost in the operation of a reservoir system is the cost of controlling the system. This is modelled as a cost of switching the pumps from one setting to another. An operation policy that is optimised to minimise the switching costs will try to keep the pump rate as constant as possible. This is what is shown in Figure 6.7, where the pump rate is at the maximum for the entire optimisation period. There is no penalty included for overflowing the reservoir, so this happens quite frequently, but the switching cost is a minimum.

Penalty Costs

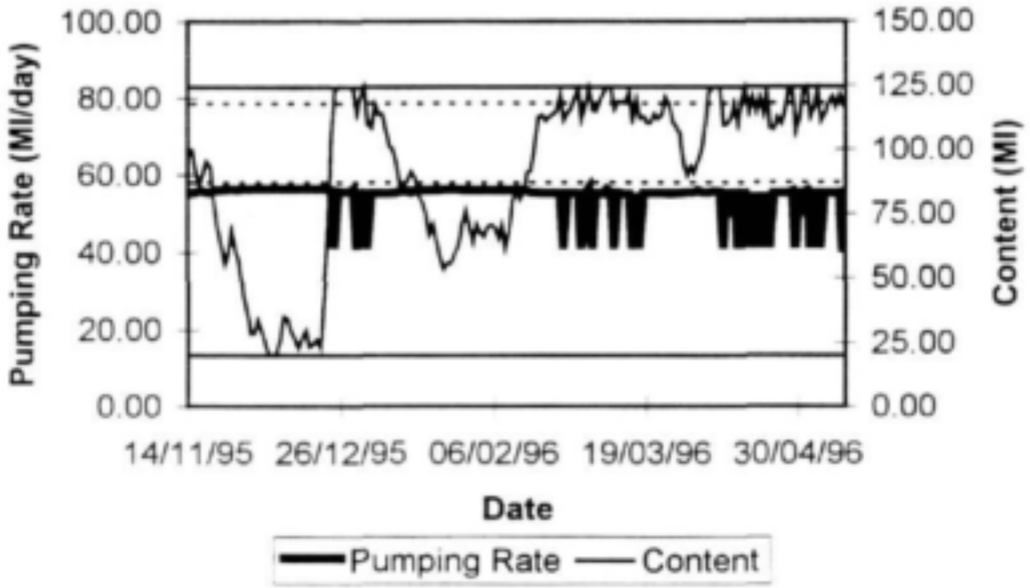

Figure 6.8: Minimum Penalty Cost Operation

In many reservoir systems it is desired that the volume not reach 100% full or completely empty. *RSPO* allows this to be modelled by defining High and Low warning levels, and associating penalties with these limits. The values for these penalties can either be policy decisions made by the reservoir operators, or can be calculated from the probability and cost of reservoir failure associated with particular volumes in the reservoir. The following operation policy was calculated using penalties for exceeding the High and Low Warning levels, and for overflowing or running the reservoir dry. All other costs were set to zero. The resultant policy is shown in Figure 6.8.

During the first half of December 1995, the Low Warning level was exceeded, and in fact the reservoir ran dry on two days. The reason for this is not because the policy was not adequate, but that the demand was greater than the pumps could supply, as shown in Figure 6.5. Figure 6.8 shows that the pumping rate was a maximum for most of December. For the rest of the period, the volume in the reservoir remained mostly within the High and Low Warning limits.

Reservoir Costs

RSPO also allows a cost to be associated with the volume of water stored in the reservoir at each stage. This can be used to model the cost of the construction of the tank, or it can be used as a policy variable to control how the operation policy uses the available storage in the reservoir. If the reservoir operators require that the reservoir be kept as full as possible, within the other constraints, they can assign a negative cost to the storage cost. This will result in a policy similar to that shown in Figure 6.9, where the volume in the reservoir is as high as possible. In this case there is no penalty for overflowing the reservoir, so this happens frequently. There is also no cost associated with the pumping, so the pumping rate is maximum for the entire period.

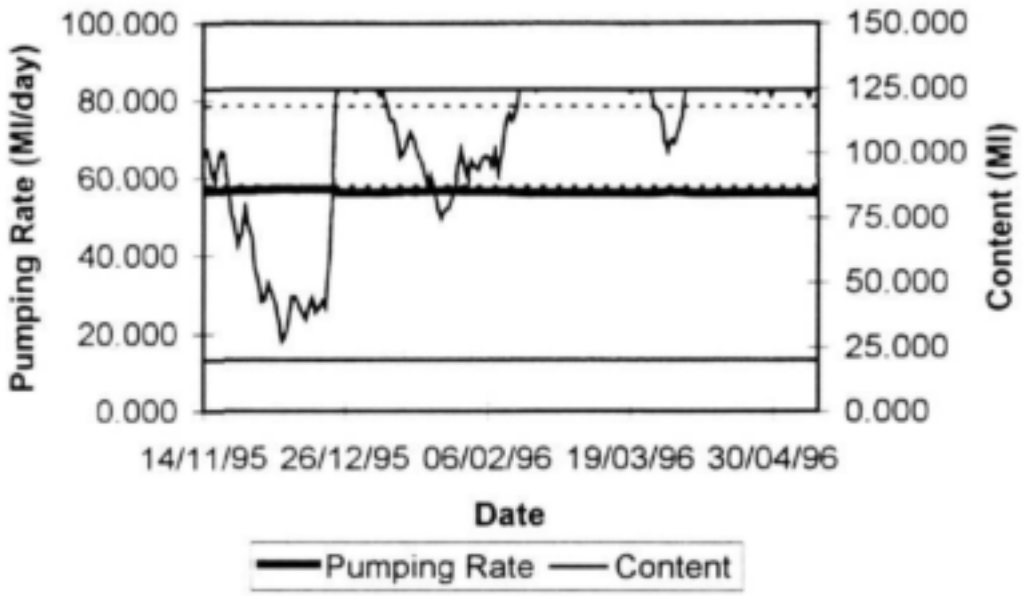

Figure 6.9: Negative Storage Cost Operation

A positive storage cost would have the opposite effect, that is, to keep the reservoir as empty as possible. If no other costs were taken into account, then the optimum policy would be to allow the reservoir to run dry, and leave it like that for the entire period. For Figure 6.10 a policy was calculated using a positive storage cost, as well as a penalty for allowing the reservoir volume to fall below the Low Warning level.

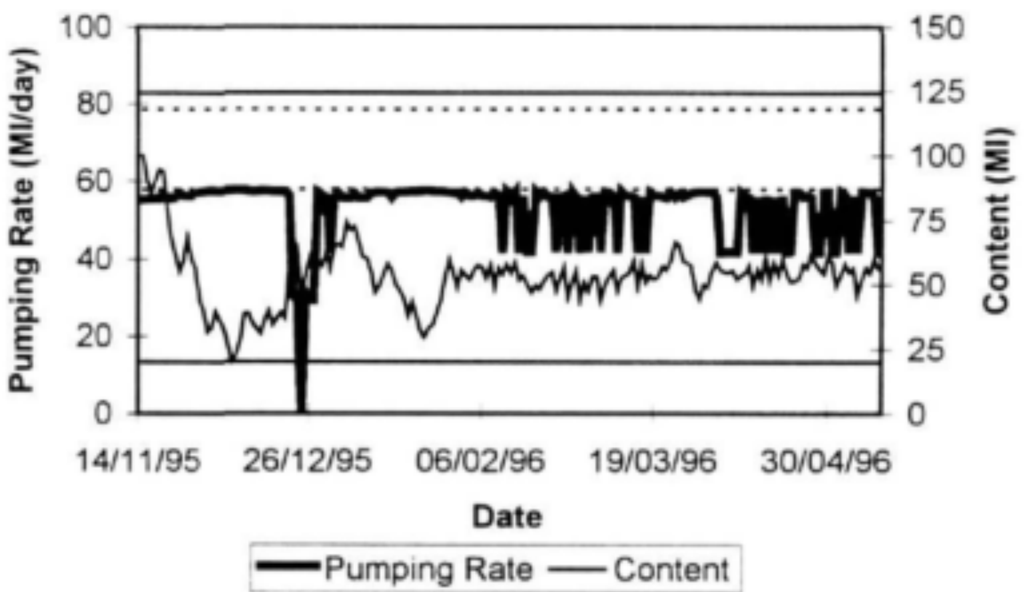

Figure 6.10: Positive Storage Cost Operation

Figure 6.10 shows that the reservoir volume was kept as close to the Low Warning level as possible. If the penalty had been set higher, there would have been fewer violations of the Low Warning level. If the storage cost had been higher, the reservoir level would have been lower, with more violations.

In real reservoir systems the operating policy must take into account all the costs discussed above simultaneously, and not just one at a time. The policy will also not only be affected by the absolute value of the cost functions, but also the relationship between the different costs. A typical operation is shown in Figure 6.12 and Figure 6.13, which was optimised using all the costs discussed above.

This operation graph shows how the policy attempts to keep the reservoir volume within the limits, while at the same time minimising the pumping, switching and storage costs.

In general the costs above should be calculated as accurately as possible from the actual reservoir system. The pumping costs are usually the easiest to obtain. The switching costs are also usually fairly easy to calculate. As discussed earlier, the storage cost can be either a real cost, or an operational policy cost. The penalty costs for overflowing the reservoir would include the cost of the lost water and subsequent damage around the actual reservoir. The penalty costs of running the reservoir dry would usually include either the cost of supplying the required water from some other, more expensive, source, or the cost of the losses incurred by the consumers. The legal responsibilities of the reservoir operators towards their consumers would have to be taken into account when calculating these penalties. The penalties imposed for violating the High and Low Warning limits, and the values of these limits, are more difficult to determine. They should be based on the risk of reservoir failure and the costs of such a failure for different limits, but they are more often chosen as policy decisions by the reservoir operators. Policy decisions of this type tend to be conservative as reservoir operators generally have a risk averse attitude [Orlovski et al. 1984], which could increase the actual running cost of the reservoir unnecessarily.

6.6. A Comparison with Historical Results

Two further optimisations were run for the historical data available for the Libanon/Driefontein sub-system, both using all the costs discussed in the above sections. The complete System Description is contained in the file 'LIB\HIST.SDC'. The operation policy found by the Downhill Simplex Optimisation method is shown in the following table and graph.

Table 6.5: Optimised Operation Policy

Change Level	Value
Level1	69.666
Level2	61.997
Level3	45.250
Level4	17.068
Level5	-2.860
Level6	-24.942
Level7	-40.023
Don't Change	0.000

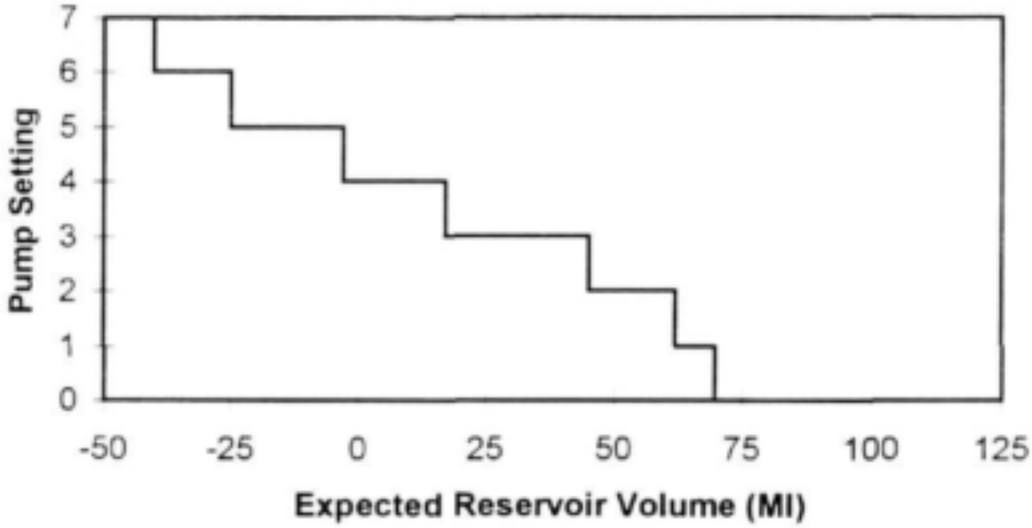

Figure 6.11: Optimised Operation Policy

The expected reservoir volume is the difference between the reservoir volume at the start of the day and the expected demand for the day. This policy indicates that if the expected volume is between 69.666 Ml and 61.997 Ml then the pumps should be set to setting 1. If the expected volume is over 69.666 Ml then the pumps should be set to setting 0, that is, all pumps turned off. If the expected volume is below -40.023 Ml then all pumps should be turned on.

The results of these optimisation runs are shown in the following two graphs.

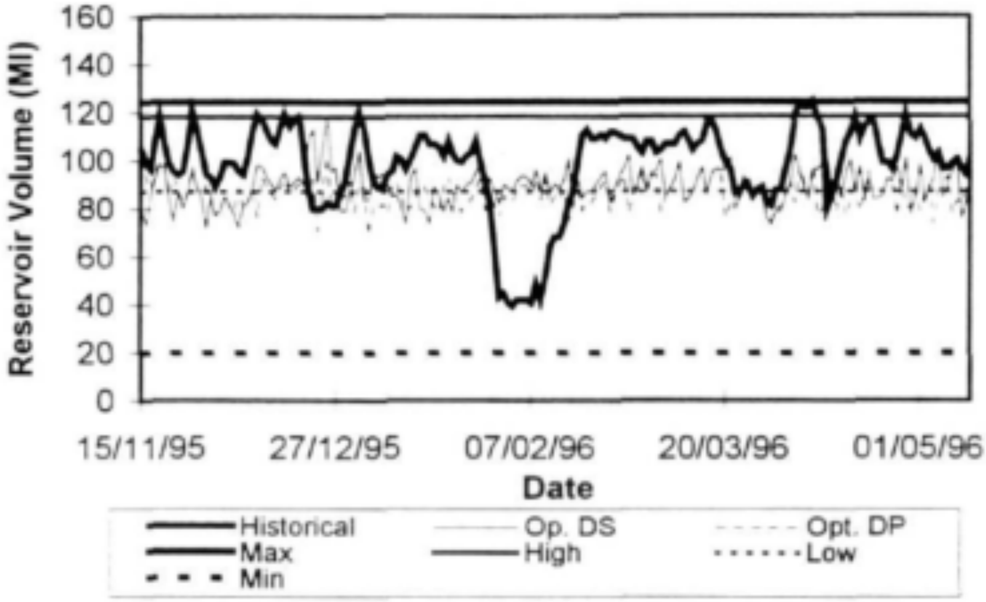

Figure 6.12: Reservoir Content

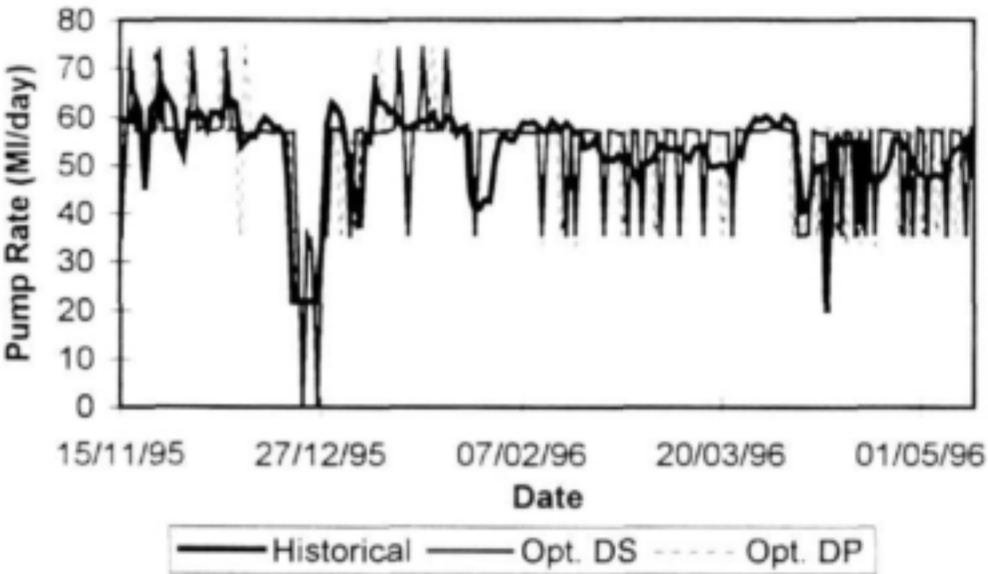

Figure 6.13: Pumping Rate

The main difficulty with attempting to compare the optimised operation with the historical operation is the lack of data available for the historical period. No information is kept regarding which pumps are running at any particular time, or when they were switched on or off. The only information available is the actual output from the pumps for each day, and the reservoir volumes at the start of the day.

To estimate the historical operation policy, the value of k_S in Eqn 4.6 was estimated by finding the value that minimised the difference between the system head and the pumping head of the closest pump setting. The value was found to be 0.00548 m/(Ml/day)^2. The pump setting that gave the pump head closest to the system head was then assumed to be the pump setting used for the entire day.

The power consumption's and costs obtained using this approximated historical operation policy, as well as those obtained from policies calculated using the Downhill Simplex Optimisation and Dynamic Programming Optimisation methods are shown below in Table 6.6.

Table 6.6: Power Consumption's and Costs

Operation Policy	Power Consumption kWh	Cost R	Saving %
Historical	129113.18	165626.39	
Downhill Simplex Optimisation	126638.31	162451.62	1.92%
Dynamic Programming Optimisation	126080.40	161735.94	2.35%

As can be seen from the table, the optimised policies obtained give results that are better than those obtained by an experienced human operator, although not by much. This indicates that even this fairly simple optimisation procedure can operate the system as well as the experienced operator.

The advantage of this procedure is that it gives a formalised policy that does not require the lengthy training of an operator to achieve. It also shows that an automatic control system, using a simple policy like this one, could successfully control the reservoir system.

This optimisation only attempts to minimise the cost of the power consumed. No information on the cost of switching from one pump setting to another was available, so this aspect was not optimised at all. It is expected that more significant savings could be achieved if this were optimised.

7. CONCLUSIONS AND RECOMMENDATIONS

The aim of this project was to develop a dynamic computer simulation model that could then be used to optimise the rules for pumping operation used by the operators of reservoir storage systems with respect to demand patterns, pumping costs, and storage volumes. This aim has been fulfilled with the completion of the computer program *Reservoir System Operation Optimisation*, which allows historical demand data to be imported and analysed, a stochastic model of that data to be fitted, and optimum operation policies to be calculated.

RSPO analyses the available historical data by fitting either a linear or exponential trend, which ever give the best fit. This trend is then removed from the data, and annual, weekly and daily periodic trends are fitted. To model the serial autocorrelations of the data, a linear stochastic model is fitted. A normal probability distribution is fitted to the remaining independent random component, and a χ^2 test shows that this generally gives a fairly good fit, but not in all situations.

The simulation option of *RSPO* allows the user to calculate the total running costs of a given reservoir system using either historical, or synthetically generated, demand data. The user can specify both the cost functions and the operating policy used in the simulation. This allows the user to test the effect of different cost and operation policies on the system, over any period the user chooses. The synthetic demand data generated using the simulation option has been shown to have a high statistical similarity to the historical data modelled.

RSPO also allows the user to calculate optimum operating policies for the system using either the Downhill Simplex method or the Dynamic Programming method of optimisation. The optimum policy calculated by the Downhill Simplex Optimisation method tells the operator which of the available pump settings to use for the current time period, based on the volume in the reservoirs at the start of the period, and the demand expected for the period.

The policies calculated by these methods have been shown to both be very close to optimum, and useful for the continued operation of the reservoir system. The effect of local optimum points in the cost space was investigated. It was found that although there are local optimum points that do tend to prevent the system from finding the global optimum solution, these tend to be small and not reduce the accuracy of the solution by more than 1%. It was also found that the accuracy of the optimum solution found is not very sensitive to the length of the data set used in the optimisation process.

A comparison with historical operating data reveals that the optimum policy can operate the reservoir system as well as an experienced human operator, and even achieve a small saving on the amount of power consumed.

7.1. Recommendations for Minimising Pumping Costs

Experienced operators operate most reservoir systems, both throughout the world, and especially in South Africa, with no formal, optimised policy. Because the cost of pumping forms such a large part of any water supply authorities running costs, it is vital that the operation of the system be optimised. To do this it is necessary to accurately analyse the system, construct the relevant models, and calculate the optimum policy.

Once an optimum policy has been derived, it is necessary to implement it. Ideally, this would involve an automatic monitoring and control system throughout the distribution network. To date, very few water utilities have this kind of system in place, although they are much more common in power supply utilities. As the cost of computer hardware and software continues to decline, and the costs of skilled labour, power and water continue to increase, automatic control systems will become both more economically feasible, and necessary.

7.2. Suggestions for further research

One of the main difficulties of this project was found to be obtaining accurate and detailed cost information of the actual running of reservoir systems. Also, during this project, only a few small subsystems in the Rand Water distribution area were analysed. It would be useful to extend this research to other water supply authorities to get more detailed and general cost information.

Another area of further research would be in the modelling of the distribution of consumer demands. In this project the consumer demands have been lumped together, and no account has been taken of the location and breakdown of the different demands that make up the total demand on the reservoir system. The difficulty with modelling the individual demands is the large number of demands on the system.

It would be ideal to model and optimise the entire system rather than single reservoir sub-system. This requires much more work in terms of data gathering and computational effort, but the results would be more comprehensive and give a better overall optimisation.

There are a number of areas related to reservoir operation that could have significant impacts on the total operating costs that have received very little attention in the literature. The first of these is the impact that 'optimum' operation policies have on the maintenance costs of the reservoir systems. The biggest problem with this field is that to get an indication of the effect on maintenance costs, it would be necessary to determine and implement an optimum policy at a specific system and allow it to run for a fairly long period of time. To date there are very few reservoir systems that have actually implemented optimum policy controls.

Another area is the relationship between network design and the use of optimum operation policies. Reservoir design is currently based on the demands that must be met, and very little attention is given to the way in which the final design will be operated.

Very little work has also been done on the robustness of optimum operation policies and how they affect the reliability of the supply to consumers. Currently, reliability is incorporated into reservoir optimisation by specifying a minimum constraint on the reservoir volumes, and assuming this will be sufficient.

REFERENCES

Arunkumar, S. and Yeh W.W.-G. [1973] Probabilistic models in the design and operation of a multi-purpose reservoir system, *Contributions 144*, California Water Resources Centre, University of California, Davis, Dec. 1973.

Askew, A.J. [1974] Optimum reservoir operating policies and the imposition of a reliability constraint, *Water Resources Research*, vol. 10, no. 1, 1974, pp 51-56.

Bellman, R.E. [1957] *Dynamic Programming*, Princeton University Press, Princeton, N.J., 1957.

Collins, M.A. [1977] Implementation of an optimisation model for operation of a metropolitan reservoir system, *Water Resources Bulletin*, vol. 13, no. 1, 1977, pp 57-70.

Dold, P.L. [1982] *Design and control of equalization reservoirs*, Research report W42, Cape Town, University of Cape Town.

Fitch, W.N., King, P.H., and Young, G.K. Jr. [1970] The optimisation of the operation of a multi-purpose water resource system, *Water Resources Bulletin*, vol. 6, no. 4, 1970, pp 498-518.

Grosman, D.D, ,and Basson, ,M.S. [1985] Stochastic modelling and risk analysis applied to urban reservoir design, *The Civil Engineer in South Africa*, vol 27, no. 8, Aug. 1985, pp 429-438

Hall, W.A. and Buras, N. [1961] The dynamic programming approach to water resources development, *Journal of Geophysics Research*, vol. 66, no. 2, 1961, pp 510-520

Hall, W.A., and Shephard, R.W. [1967] Optimum operations for planning of a complex water resources system, *Tech. Rep. 122 (UCLA-ENG 67-544)*, Water Resour. Cent., Sh. of Eng. and Appl. Sci., University of California, Los Angeles, Oct. 1967.

Hall, W.A., Butcher, W.S., and Esogbue, A. [1968] Optimisation of the operations of a multi-purpose reservoir by dynamic programming, *Water Resources Research*, vol. 4, no. 3, 1968, pp 471-477.

Hall, W.A., Harboe, W., Yeh, W.W.-G., and Askew, A.J. [1969] Optimum firm power output from a two reservoir system by incremental dynamic programming, *Contrib. 130*, Water Resour. Cent., University of California, Los Angeles, Oct. 1969.

Hammersley, J.M. and Handscomb, D.C. [1964] *Mote Carlo Methods*, London: Methuen & Co. Limited.

Hydrologic Engineering Centre, [1971] *HEC-3 reservoir system analysis*, Technical report, US Army Corps of Eng., Davis, California, 1971.

Hydrologic Engineering Centre, [1979] *HEC-5 reservoir system operation for flood control and conservation user manual*, Technical report, US Army Corps of Eng., Davis, California, June 1979.

Jacobs, O.L.R. [1967] *An Introduction to Dynamic Programming*, London: Chapman and Hall Limited.

Jowitt, P.W., and Germanopoulos, G. [1992] Optimal pump scheduling in water supply networks, *Journal of Water Resources Planning and Management*, vol. 118, no. 4, Jul/Aug 1992, pp 406-422.

Kottegoda, N.T. [1980] *Stochastic Water Resources Technology*, London: The MacMillan Press Ltd., 1980.

Lansey, K.E. and Awumah, K. [1994] Optimal pump operations considering pump switches, *Journal of Water Resources Planning and Management*, vol. 120, no. 1, Jan/Feb 1994, pp 17-35.

Larson, H.J. [1982] *Introduction to Probability Theory and Statistical Inference*, 3rd ed., New York: John Wiley & Sons, 1982, pp 463-480.

Little, J.D.C. [1955] The use of storage water in a hydroelectric system, *Operations Research*, vol. 3, no. 2, 1955, pp 187-197.

Liu, C.S. and Tedrow, A.C. [1973] Multi-lake river system operating rules, *J. Hydraul. Div., Am. Soc. Civ. Eng.*, vol. 99, no. HY9, 1973, pp 1369-1381.

Loubser, B.F. [1986] Discussion on Bepaling van die optimale toevoer- en opgaar-kapasiteite vir munisipale diensreservoirs, *The Civil Engineer in South Africa*, vol 28, no. 9, Sept. 1986, pp 352-359.

Loucks, D.P. [1968] Computer models for reservoir regulations, *Journal of Sanitary Engineering Division, Am. Soc. Civ. Eng.*, vol. 94, 1968, pp 657-669.

Loucks, D.P, Stedinger, J.R. and Haith, D.A. [1981] *Water Resource Systems Planning and Analysis*, Englewood Cliffs, NJ: Prentice-Hall.

Manson, N.J. [1996]. Optimizing Water Supply and Distribution, *Proceedings IAHR Southern African Division Biennial Congress*, Sun City, SA.

Matalas, N.C. and Wallis, J.R. [1976] Generation of Synthetic Flow Sequences, In: Bisuras, A.K. ed. *Systems Approach to Water Management*, New York: McGraw-Hill.

McCuen, R.H. [1993] *Microcomputer Applications in Statistical Hydrology*, Englewood Cilffs, New Jersey, 1993.

McLeod, W.T. ed, [1986] *The Collins Paperback English Dictionary*, London: Collins.

Meier, W.L. Jr., and Beightler, C.S. [1967] An Optimization Method for Branching Multi-Stage Water Resources Systems, *Water Resources Research*, vol. 3, no. 3, Mar. 1967, pp 645-652.

Moran, P.A.P. [1959] *The Theory of Storage*, London: Methuen & Co. Limited.

Nel, D.T. [1993] *Bepaling van die Optimale Stoorkapasiteit van Twee Johannesburgse Munisipale Diensreservoirs*, MSc Dissertation, Randse Afrikaanse Universiteit, Johannesburg.

Nel, D.T. and Haarhoff, J. [1996] Sizing municipal water storage tanks with Monte Carlo simulation, *Aqua*, vol. 45, no. 4, pp 203-212.

Opricovic, S. and Djordjevic, B. [1976] Optimal long-term control of a multi-purpose reservoir with indirect users, *Water Resources Research*, vol. 12, no. 6, 1976, pp 1286-1290.

Orlovski, S., Rinaldi, S. and Soncini-Sesssa, R. [1984] A Min-Max Approach to Reservoir Management, *Water Resources Research*, vol. 20, no. 11, Nov. 1984, pp 1506-1514.

Ormsbee, L.E. and Reddy, S.L. [1995] Non-linear heuristic for pump operations, *Journal of Water Resources Planning and Management*, vol. 121, no. 4, Jul/Aug 1995, pp 302-309.

Ormsbee, L.E. and Lansey, K.E. [1994] Optimal control of water supply systems, *Journal of Water Resources Planning and Management*, vol. 120, no. 2, Mar/Apr 1994, pp 237-252.

Parikh, S.C., [1966] Linear decomposition programming of optimal long-range operation of a multi-purpose reservoir system, *Rep. ORC 66-28*, Oper. Res. Cent., University of California, Berkeley.

Perkopa, A. [1980] Network planning using two-stage programming under uncertainty, in *Recent Results in Stochastic Programming*, edited by P.Kall and A.Perkopa, Springer-Verlang, New York, 1980.

Press, W.H., Flannery, B.P., Teukolsky, S.A. and Vetterling, W.T. [1986] *Numerical Recipes, The Art of Scientific Computing*, Cambridge: Cambridge University Press, 1986.

Raman, H. and Chandramouli, V. [1996] Deriving a gerneral operating policy for reservoirs using neural networks, *Journal of Water Resource Planning and Management*, vol. 122, no. 5, Sep/Oct 1996, pp 342-347.

Rand Water [1995] Annual Report, Johannesurg: Rand Water.

ReVelle, C., Joeres, E. and Kirby, W. [1969] The linear decision rule in reservoir management and design, 1, Development of the stochastic model, *Water Resources Research*, vol. 5, no. 4, 1969, pp 767-777.

Reznicek, K. and Cheng. T.C.E. [1991] Stochastic Modelling of Reservoir Operations, *European Journal of Operational Research*, vol. 50, Issue 3, 1991, pp. 235-248.

Roefs, T.G., and Bodin, L.D. [1970] Multi-reservoir operation studies, *Water Resources Research*, vol. 6, no. 2, Feb. 1970, pp 410-420.

Sadowski, E., Nitivattananon, V. and Quimpo, R.G. [1995] Compute generated optimal pumping schedule, *Journal of the American Water Works Association*, vol. 87, no. 7, Jul 1995, pp 53-63.

Schweig, Z. and Cole, J.A. [1968] Optimal control of linked reservoir, *Water Resources Research*, vol. 4, no. 3, 1968, pp 479-497.

Sengupta, J.K. [1972] *Stochastic Programming, Methods and Applications*, Amsterdam: North-Holland.

Shaw, I.S. [1993] Neural network-based controller for the balancing reservoir of a wastewater treatment plant, *J. S.A. Inst. Civ. Eng.*, vol. 35, no. 4, Fourth Quarter 1993.

Smook, A.W. [1985] Bepaling van die optimale toevoer- en opgaar-kapasiteite vir munisipale diensreservoirs, *The Civil Engineer in South Africa*, vol 27, no. 8 Aug. 1985, pp 429-438.

Stallard, K. [1980] *Economic Optimisation of the RWB System by Simulation Modelling*, Johannesburg: Rand Water Board, pp. 33.

Stedinger, J.R., Bola, F.S. and Loucks, D.P. [1984] Stochastic Dynamic Programming Models for Reservoir Operation Optimisation, *Water Resources Research*, vol. 20, no. 11, Nov. 1984, pp 1499-1505.

Stephenson, D. [1989] *Pipeline Design for Water Engineers*, 3rd ed., Amsterdam: Elsevier.

Turgeon, A. [1982] Incremental dynamic programming may yield non-optimal solutions, *Water Resources Research*, vol. 18, no. 6, 1982, pp 1599-1604.

Walski, T.M. [1993] Tips for saving energy in pumping operations, *Journal of the American Water Works Associations*, vol. 85, no. 7, July 1993, pp 49-53.

White, F.W. [1986] *Fluid Mechanics*, Iternational ed., Singapore: McGraw-Hill.

Williams, K [1970] *Dynamic Programming: Sequential decision making*, London: Longman Group Limited.

Windsor, J.S. [1973] Optimization model for the operation of flood control systems, *Water Resources Research*, vol. 9, no. 5, May 1973, pp 1219-1226.

Yeh, W.W.-G. [1985] Reservoir Management and Operations Models: A State-of-the-Art Review, *Water Resources Research*, vol. 21, no. 12, Dec. 1985, pp 1797-1818.

Young, G.K. [1967] Finding reservoir operation rules, *J. Hydraul. Div., Am. Soc. Civ. Eng.*, vol. 93, no. HY6, 1967, pp 297-321.

APPENDIX A: AVAILABLE SYSTEM DATA

Example Morning Daily Distribution Log

WD930RP

DAILY DISTRIBUTION LOG - CENTRAL DEPOT

24 HOURS BEGINNING 06:00 ON THE 01 01 1996

PAGE NO: 1
96-05-13 09:25:42

Level 06h00	Content 06h00	Level Noon	Level 16h00	Content 16h00	Incr.	Decr.	Top Levels	Reservoir Name	Top Content	Level 06h00	Content 06h00	Incr.	Decr.
9.90	101.39	.	8.90	89.91		11.48	10.825	LANGERAND	111.956	8.90	89.91		
6.10	31.27	.	5.80	29.73		1.54	9.850	DALESIDE	45.429	6.70	34.35	4.62	
2.50	11.10	.	2.70	11.99	0.89		4.525	ZWARTKOPJES BALANCING	20.103	2.50	11.10		0.89
12.40	20.90	.	11.50	19.37		1.53	13.725	PALMIET NO 1	23.149	13.40	22.61	3.24	
12.40	20.90	.	11.50	19.37		1.53	13.725	PALMIET NO 2	23.149	13.40	22.61	3.24	
12.50	18.43	.	12.80	18.88	0.45		13.570	EIKENHOF NO 1	20.025	13.10	19.32	0.44	
12.50	18.43	.	12.80	18.88	0.45		13.510	EIKENHOF NO 2	20.025	13.10	19.32	0.44	
8.20	12.55	.	7.70	11.78		0.77	13.980	MAPLETON	19.959	9.30	14.25	2.47	
0.80	4.18	.	1.20	6.30	2.12		4.575	FOREST HILL NO 1	24.153	3.40	17.96	11.66	
5.30	54.04	.	5.60	57.11	3.07		8.850	FOREST HILL NO 2	90.213	7.70	78.57	21.46	
6.10	68.96	.	6.40	72.40	3.44		10.075	FOREST HILL (3)	113.880	8.70	98.48	26.08	
3.60	144.55	.	3.60	144.55			13.480	KLIPRIVIERSBERG (1)	587.427	3.60	144.55		
5.80	157.35	.	6.60	179.12	21.77		10.075	MEYERS HILL	273.211	8.30	225.38	46.26	
11.20	96.50	.	11.70	102.46	5.96		13.425	MEREDALE	122.426	13.40	122.30	19.84	
11.50	165.58	.	11.70	169.18	3.60		13.400	MEREDALE (2)	200.888	12.00	175.81	6.63	
12.80	187.80	.	12.60	184.20		3.60	13.425	KLIPFONTEIN (1)	198.750	12.90	189.60	5.40	
12.90	248.97	.	12.50	241.25		7.72	13.400	KLIPFONTEIN (2)	258.700	12.90	248.97	7.72	
9.90	89.50	.	9.90	89.50			10.075	GERMISTON	90.935	9.90	89.50		
8.60	87.63	.	8.60	87.63			8.850	NORTHBRIDGE	90.076	8.60	87.63		
5.70	39.21	.	5.60	38.52		0.69	7.320	AIRFIELD (COMBINED)	50.438	5.60	38.52		
		.					11.200	VLAKFONTEIN (1)	210.617				
7.20	132.70	.	6.90	126.75		5.95	11.200	VLAKFONTEIN (2)	210.617	7.00	128.73	1.98	
1.50	4.41	.	1.50	4.41			4.275	BENONI (1)	12.611	1.50	4.41		
8.50	43.60	.	8.00	41.03		2.57	8.850	BENONI (2)	45.429	7.80	40.00		1.03
8.50	117.19	.	8.10	111.94		5.25	8.840	BENONI (3)	121.630	7.90	109.30		2.64
4.50	11.20	.	4.50	11.20			4.600	MODDER EAST	11.388	4.50	11.20		
3.10	15.87	.	2.60	13.30		2.57	8.850	BRAKPAN NO 1	45.429	3.60	18.43	5.13	
3.00	41.62	.	2.30	31.43		10.19	8.840	BRAKPAN NO 2	121.630	3.40	47.39	15.96	
3.70	18.37	.	4.00	19.86	1.49		4.575	SELCOURT	22.708	3.70	18.37		1.49
12.40	21.09	.	12.20	20.75		0.34	13.425	WILDEBEESTFONTEIN (1)	22.826	12.90	21.95	1.20	
12.10	41.42	.	12.00	41.07		0.35	13.430	WILDEBEESTFONTEIN (2)	45.959	12.80	43.82	2.75	
12.40	46.62	.	12.30	46.24		0.38	13.500	WILDEBEESTFONTEIN (3)	50.800	13.10	49.28	3.04	
		.					8.850	WATERVAL (1)	45.429				
10.90	61.17	.	10.80	60.48		0.69	10.975	WATERVAL (2)	61.672	10.90	61.17	0.69	
2.70	13.81	.	2.60	13.30		0.51	8.850	WITPOORTJIE (1)	45.429	2.70	13.81	0.51	
2.50	21.68	.	2.50	21.68			8.825	WITPOORTJIE (2)	75.162	2.50	21.68		
2.80	6.94	.	3.20	7.94	1.00		4.575	LIBANON (1)	11.374	3.40	8.44	0.50	
2.80	6.94	.	3.20	7.94	1.00		4.575	LIBANON (2)	11.347	3.40	8.44	0.50	
9.50	114.30	.	9.60	115.60	1.30		10.300	DRIEFONTEIN	124.354	10.00	120.60	5.00	
7.60	89.03	.	7.40	86.81		2.22	8.230	BARNARDSVLEI	95.965	7.70	90.14	3.33	
8.30	42.57	.	8.30	42.57			8.850	ROODEPOORT (1)	45.429	8.30	42.57		
4.10	5.46	.	4.20	5.59	0.13		4.575	KRUGERSDORP (1)	6.042	4.20	5.59		
3.90	19.37	.	3.90	19.37			4.575	KRUGERSDORP (2)	22.708	4.00	19.86	0.49	
5.80	41.39	.	5.90	42.11	0.72		6.425	KRUGERSDORP (3)	45.693	6.00	42.82	0.71	
8.20	95.68	.	8.70	101.18	5.50		9.900	KRUGERSDORP (4)	112.400	8.10	94.35		6.83
9.10	88.73	.	9.10	88.73			10.640	SASOLBURG	103.419	9.20	89.70	0.97	
		.					8.850	BLYVOORUITZICHT	22.790				
11.90	84.52	.	11.30	78.98		5.54	13.500	HARTEBEESTHOEK	99.317	12.40	89.15	10.17	
TOTAL INCR. / DECR.					52.89	65.42						212.43	12.88

TOTAL RESERVOIR CAPACITY	4212.276			
RESERVOIR CONTENT 06H00 (Y/DAY)	2764.920	DIFFERENCE OVER 1st 12 hours	12.530	DECREASE
RESERVOIR CONTENT 16H00 (Y/DAY)	2752.390	DIFFERENCE OVER 2nd 12 hours	199.550	INCREASE
RESERVOIR CONTENT 06H00 (T/DAY)	2951.940	DIFFERENCE OVER 24 hours	187.020	INCREASE
AVAILABLE SPACE - 12 HRS.	1459.886	AVAILABLE SPACE - 24 hrs.	1260.336	

Level 06h00	Content 06h00	Level 16h00	Content 16h00	Top Levels	Reservoir Name	Top Content	Level 06h00	Content 06h00
9.40	31.39	10.30	34.40	12.100	MEYERSDAL (MUN)	40.000	9.70	32.39

A2

DAILY DISTRIBUTION LOG - CENTRAL DEPOT

24 HOURS BEGINNING 00:00 ON THE 01 01 1996

Level 24h00	Content 24h00	Incr.	Decr.	Top Levels	Reservoir Name	Top Content	Level 24h00	Content 24h00	Incr.	Decr.
9.90	101.39			10.825	LANGERAND	111.956	8.60	86.46		14.93
4.50	23.06			8.850	DALESIDE	45.429	5.60	28.70	5.64	
2.50	11.10			4.525	ZWARTKOPJES BALANCING	20.103	2.70	11.99	0.89	
9.90	16.65			13.725	PALMIET NO 1	23.149	10.70	18.01	1.36	
9.90	16.65			13.725	PALMIET NO 2	23.149	10.70	18.01	1.36	
10.50	15.46			13.570	EIKENHOF NO 1	20.025	9.30	13.69		1.77
10.50	15.46			13.570	EIKENHOF NO 2	20.025	9.30	13.69		1.77
7.20	11.01			13.980	MAPLETON	19.959	8.40	12.86	1.85	
0.20	1.02			4.575	FOREST HILL NO 1	24.153	2.10	11.07	10.05	
3.30	33.60			8.850	FOREST HILL NO 2	90.213	6.40	65.29	31.69	
4.10	46.33			10.075	FOREST HILL (3)	113.880	7.30	82.61	36.28	
3.60	144.55			13.480	KLIPRIVIERSBERG (1)	587.427	3.60	144.55		
4.50	121.97			10.075	MEYERS HILL	273.211	7.40	200.89	78.92	
9.50	75.82			13.425	MEREDALE	122.426	12.70	114.21	38.39	
11.30	161.99			13.400	MEREDALE (2)	200.888	11.80	171.00	9.01	
12.50	182.39			13.425	KLIPFONTEIN (1)	198.750	12.60	184.20	1.81	
12.50	241.25			13.400	KLIPFONTEIN (2)	258.700	12.70	245.11	3.86	
9.90	89.50			10.075	GERMISTON	90.935	9.90	89.50		
8.60	87.63			8.850	NORTHRIDGE	90.076	8.60	87.63		
5.60	38.52			7.320	AIRFIELD (COMBINED)	50.438	5.60	38.52		
				11.200	VLAKFONTEIN (1)	210.617				
7.40	136.65			11.200	VLAKFONTEIN (2)	210.617	6.60	120.78		15.87
1.50	4.41			4.275	BENONI (1)	12.611	1.50	4.41		
8.10	41.54			8.850	BENONI (2)	45.429	7.70	39.49		2.05
8.20	113.26			8.840	BENONI (3)	121.630	7.90	109.30		3.96
4.50	11.20			4.600	MODDER EAST	11.388	4.50	11.20		
2.70	13.81			8.850	BRAKPAN NO 1	45.429	2.70	13.81		
2.50	34.36			8.840	BRAKPAN NO 2	121.630	2.50	34.36		
4.00	19.86			4.575	SELCOURT	22.708	4.00	19.86		
12.20	20.75			13.425	WILDEBEESTFONTEIN (1)	22.826	12.60	21.43	0.68	
11.70	40.04			13.430	WILDEBEESTFONTEIN (2)	45.959	12.40	42.45	2.41	
12.10	45.48			13.500	WILDEBEESTFONTEIN (3)	50.800	12.80	48.14	2.66	
				8.850	WATERVAL (1)	45.429				
10.80	60.48			10.975	WATERVAL (2)	61.672	10.90	61.17	0.69	
2.70	13.81			8.850	WITPOORTJIE (1)	45.429	2.70	13.81		
2.50	21.68			8.825	WITPOORTJIE (2)	75.162	2.50	21.68		
3.40	8.44			4.575	LIBANON (1)	11.374	3.60	8.94	0.50	
3.40	8.44			4.575	LIBANON (2)	11.347	3.60	8.94	0.50	
9.10	109.20			10.300	DRIEFONTEIN	124.354	9.70	116.80	7.60	
7.50	87.92			8.230	BARNARDSVLEI	95.965	7.60	89.03	1.11	
8.30	42.57			8.850	ROODEPOORT (1)	45.429	8.30	42.57		
4.20	5.59			4.575	KRUGERSDORP (1)	6.042	4.20	5.59		
4.00	19.86			4.575	KRUGERSDORP (2)	22.708	4.00	19.86		
6.00	42.82			6.425	KRUGERSDORP (3)	45.693	6.00	42.82		
8.70	101.18			9.900	KRUGERSDORP (4)	112.400	8.60	100.05		1.13
9.10	88.73			10.640	SASOLBURG	103.419	9.20	89.70	0.97	
				8.850	BLYVOORUITZICHT	22.790				
11.50	80.83			13.500	HARTEBEESTHOEK	99.317	11.70	82.68	1.85	

TOTAL INCR. / DECR. 240.08 41.18
TOTAL RESERVOIR CAPACITY 4212.276
RESERVOIR CONTENT 24H00 (Y/DAY) 2608.260
RESERVOIR CONTENT 24H00 (T/DAY) 2806.860 DIFFERENCE OVER 24 hours 198.600 INCREASE
AVAILABLE SPACE - 24 hrs. 1405.416

Level 24h00	Content 24h00	Top Levels	Reservoir Name	Top Content	Level 24h00	Content 24h00
9.40	31.39	12.100	MEYERSDAL (MBU)	40.000	9.50	31.73

DAILY DISTRIBUTION LOG - CENTRAL DEPOT

24 HOURS BEGINNING 06:00 ON THE 01 01 1996

WEATHER CONDITIONS

STATIONS	RAINFALL mm
BENONI	NIL
BLOEMENDAL	NIL
CENTRAL DEPOT	1,2
DALESIDE	NIL
ESSELEN PARK	NIL
FOREST HILL	4,2
GERMISTON	9,0
KRUGERSDORP	DROPS
LIBANON	NIL
PRETORIA	NIL
RUSTENBURG	9,0
SPRINGS|BRAKPAN	NIL
WATERVAL	3,7
ZUURBEKOM	2,0

Reservoir	Top Water Level	Capacity	Level 06h00
SIGNAL HILL	3.975	11.692	
YEOVILLE 1	7.925	17.729	
YEOVILLE 2	7.850	46.547	5.70
YEOVILLE 3	12.195	45.460	
BRIXTON	5.028	21.820	3.90
BEREA 1	2.666	4.091	
BEREA 2	6.070	22.730	4.10
PARKTOWN 1	4.241	23.639	
PARKTOWN 2	7.620	49.551	
HURST HILL1	4.267	23.184	3.10
HURST HILL 2	4.572	25.003	3.00
MEADOWLANDS	3.962	23.639	
CROWN GARDENS	3.885	22.730	
DOORNKOP	5.104	47.733	
AEROTON	7.900	45.460	
NATURENA	7.650	6.500	5.80
ALAN MANOR	7.650	7.650	
EAGLES NEST	7.650	23.000	
POWER PARK	7.650	46.000	
GLENVISTA	7.650	23.000	
CORRIEMOOR	7.650	18.600	
FAIRLAND	7.650	6.500	

A4

PUMPING

PUMPING	BLOEMENDAL	LIBANON	ROODEPOORT	ZUURBEKOM
WATER PUMPED	50.060	53.900		24.460
DELIVERY HEAD	1350.000	820.000	1010.000	
SUCTION HEAD	260.000	58.000	690.000	
AVERAGE LOAD	1312.000	612.500	358.330	719.160
MAXIMUM DEMAND	21384.000	780.000	960.000	755.440
LOAD FACTOR	6.000	78.000	37.000	95.000
HIGHEST MAXIMUM DEM.	21384.000	780.000	960.000	755.440
PROGR. LOAD FACTOR	3.000	4484.000	576.000	96.000
TOTAL KILOWATT HOURS	31488.000	14700.000	8599.920	17259.840
PUMP No. 1: TIME				24.400
Ml/d				
PUMP No. 2: TIME			9.000	
Ml/d			44.000	
PUMP No. 3: TIME		23.190		
Ml/d		53.900		
PUMP No. 4: TIME	23.750			
Ml/d				
PUMP No. 5: TIME				
Ml/d				

ZUURBEKOM WELL PUMPS

WELL	ML PUMPED	WELL LEVELS
1		
5		
7	1.020	
8	8.640	
TOTAL	9.660	
MAKEUP WATER	13.382	
TOTAL ZU	23.040	

ZUURBEKOM

CHLORINE :	FREE :	0.64	RESIDUAL :	0.63
TEMPERATURE :	MAX. :	29.0	MIN. :	10.5

GENERAL

	PREV. MORNING	06H00 MORNING
BAROMETRIC PRESSURE :	831	833
TEMPERATURE :	16.9	14.9
HUMIDITY :	92	89

A5

VEREENIGING PUMPED TO :		
ZWARTKOPJES	585.63	
EIKENHOF	170.74	
SASOLBURG		
VANDERBIJLPARK		
ISCOR	( )	
TOTAL	756.37	
ZUIKERBOSCH PUMPED TO :		
ZWARTKOPJES	180.80	
PALMIET	561.40	
EIKENHOF	437.10	
MAPLETON	327.40	
SASOLBURG		
ZWARTKOPJES PUMPED TO :		PUMPING HEADS :
BRAKPAN		
EAST RAND	55.09	21.27
FOREST HILL	193.31	29.98
MEREDALE	74.80	29.63
MEYERS HILL	81.93	29.98
YEOVILLE	59.77	31.67
ZWARTKOPJES PUMPED	464.90	
PALMIET PUMPED INC. MEYERSDAL	600.46	
EIKENHOF PUMPED	620.19	
MAPLETON PUMPED	326.80	

WATER LOSS

RESERVOIR LOSS
PIPELINE LOSS
TOTAL LOSS

WATER BALANCE

STATION	
VEREENIGING	756.37
ZUIKERBOSCH	1506.70
ZUURBEKOM	9.66
TOTAL PUMPED	2272.73
CHANGE IN ADD	
CONTENT DED.	198.60
WATER LOSS	
ASS. CONS.	2074.13
PROGR. ASS. CONS.	

WD930RP

DAILY DISTRIBUTION LOG - CENTRAL DEPOT

24 HOURS BEGINNING 06:00 ON THE 01 01 1996

PAGE NO: 5
96-05-13 09:25:47

DAILY ORDERS (Ml/d)

STATIONS	DATE	TIME	ZK	EH	PM	MP	SA+ZK	TOTAL
VEREENIGING PUMPING STATION	1011996	06	500	200			100	800
	1011996	16	500	200			100	800
	2011996	06	500	200			100	800
ZUIKERBOSCH PUMPING STATION	1011996	06	180	400	565	300		1445
	1011996	16	180	300	565	300		1345
	2011996	06	180	300	565	300		1345

PUMPING INSTRUCTIONS

ZUIKERBOSCH TO	:	
ZWARTKOPJES	:	
EIKENHOF	:	____
PALMIET	:	____
MAPELTON	:	____
TOTAL	:	____
VEREENIGING TO	:	____
ZWARTKOPJES	:	
EIKENHOF	:	____
TOTAL	:	____
ZUURBEKOM	:	____
GRAND TOTAL	:	____

PUMPING ORDERS

STATION	06h00	CHANGE
VEREENIGING	800	
ZUIKERBOSCH	1345	
ZUURBEKOM	10	
TOTAL ORDERED	2155	

PUMPING INSTRUCTIONS

ZUIKERBOSCH TO	:	
ZWARTKOPJES	:	
EIKENHOF	:	____
PALMIET	:	____
MAPELTON	:	____
TOTAL	:	____
VEREENIGING TO	:	____
ZWARTKOPJES	:	
EIKENHOF	:	____
TOTAL	:	____
ZUURBEKOM	:	____
GRAND TOTAL	:	____

A6

Example Monthly User Consumption Report

WD313R *** R A N D W A T E R *** PAGE NO : 1

C O N S U M E R D R A W O F F R E P O R T 1996/05/14 12:13

PIPELINE	CONS.	NAME:	199511	199512	199601	199602	199603	199604
N1	1	RAND WATER	0	0	0	0	0	0
N1	1	RAND WATER	124	139	159	160	100	195
N1	1	RAND WATER	343520	356340	368670	354980	616010	383190
N1	1	RAND WATER	1	3	3	4	3	5
N1	101	SPOORNET	0	0	0	0	0	0
N1	101	SPOORNET	0	0	0	0	0	0
N1	101	SPOORNET	0	0	0	0	0	0
N1	366	KINROSS MINES LTD	57020	62850	73010	69490	74680	72410
N1	366	KINROSS MINES LTD	40678	22989	35926	42104	48663	45364
N1	453	POPLAR MINES (PTY) LTD	0	0	0	0	0	0
N1	511	JOHANNESBURG MUNICIPALITY	0	0	0	0	0	0
N1	515	KINROSS DORPSRAAD	36770	37270	39350	40120	42000	40330
N1	518	BETHAL MUNICIPALITY	188520	200090	210560	199110	203630	200390
N1	519	TRICHARDT MUNISIPALITEIT	29000	27210	29090	35020	30000	29330
N1	522	DEVON HEALTH COMMITTEE	20147	20244	20609	18635	20303	8583
N1	522	DEVON HEALTH COMMITTEE	0	0	0	0	0	0
N1	526	STADSRAAD VAN EVANDER	5	5	6	7	6	5
N1	532	LEANDRA MUNISIPALITEIT	88140	93990	96380	88810	113790	65590
N1	532	LEANDRA MUNISIPALITEIT	3109	3062	4064	3120	3233	3057
N1	556	EASTERN SERVICES COUNCIL	9308	8120	7284	8009	29045	15905
N1	763	RACHMANN E H	36	50	35	21	26	23
N1	774	MULLER B	203	183	230	191	182	179
N1	781	TALJAARD J A	0	0	0	0	0	0
N1	786	DEPT OF PUBLIC WORKS	5302	5336	4593	5983	5168	6366
N1	804	CRONJE MR T F	0	0	0	0	0	0
N1	806	BARNARD T A	100	111	181	189	138	148
N1	809	LOMBARD J G	0	0	0	0	0	0
N1	834	SWARTS C J	99	130	92	86	97	93
N1	837	DU TOIT F P	13	15	58	52	90	67
N1	840	JACOBS A P	87	31	15	16	13	14
N1	884	TALJAARD J A	0	0	0	0	0	0
N1	884	TALJAARD J A	0	0	0	0	0	0
N1	884	TALJAARD J A	70	35	52	38	34	34
N1	884	TALJAARD J A	0	0	0	0	0	0
N1	925	STRYDOM J P	55	110	83	60	45	75
N1	927	MEYERS S	6	8	8	6	8	9
N1	977	KRUGER A S M MEV	21	0	0	0	0	0
N1	1637	ENGELBRECHT D J	123	138	111	94	135	151
N1	1645	WILLEMSE C H J	14	16	26	14	16	14
N1	1717	COMBRINK M R H	424	332	464	379	737	420
N1	1717	COMBRINK M R H	47	70	127	44	63	31
N1	1722	VAN WYK MR G E	0	0	0	0	0	0
N1	1747	BRITS A J	0	0	0	0	0	0
N1	1851	WYBORN MR A H	0	0	0	0	0	0
N1	1871	DU PLESSIS MR F J	169	123	55	71	0	0
N1	1936	BREEDT L VAN W	0	0	0	0	0	0
N1	1968	KRUGELL MR J C	120	74	149	69	72	69
N1	1978	W & E BOERDERY (EDMS) BPK	0	0	0	0	0	0
N1	2007	ROARTY MR J C	33	32	27	10	14	14
N1	2129	LOUW MR J J	77	70	76	60	47	186
N1	2141	ROARTY MR A	61	97	87	61	41	44
N1	2198	LANGENSTRASS ALAN	0	0	0	0	0	0
N1	2200	GOLDEN LAY FARMS LIMITED	1089	1094	1212	88	0	364
N1	2409	LABOR CONSTRUCTION	0	0	0	0	0	0
N1	2429	R U C MINING CONTRACTING COMPA	0	0	0	0	0	0
N1	2429	R U C MINING CONTRACTING COMPA	0	0	0	0	0	0

A8

*** R A N D W A T E R ***

C O N S U M E R D R A W O F F R E P O R T

PIPELINE	CONS.	NAME:	199511	199512	199601	199602	199603	199604
N1	2429	R U C MINING CONTRACTING COMPA	0	0	0	0	0	0
N1	2429	R U C MINING CONTRACTING COMPA	0	0	0	0	0	0
N1	2481	ZIETSMAN DR P H	88	0	0	0	0	0
N1	2493	MULLER MR A P	0	0	0	0	0	0
N1	2547	CRISICK PROPERTIES LTD	42	55	57	29	23	33
N1	2575	STEYN MR H S	0	0	0	0	0	0
N1	2577	T P A ROADS	0	0	0	0	0	0
N1	2577	T P A ROADS	0	0	0	0	0	0
N1	2577	T P A ROADS	0	0	0	0	0	0
N1	2591	VENTER MR D J	0	0	0	0	0	0
N1	2601	BOART EXPLORATION SERVICES	0	0	0	0	0	0
N1	2604	LOMBARD MR J C	14	17	15	12	16	21
N1	2608	COETZER MR J G	0	0	0	0	0	0
N1	2615	L E T CONSTRUCTION (PTY) LTD	0	0	0	0	0	0
N1	2642	POTGIETER P P J MR	0	0	0	0	0	0
N1	2648	TAYLOR MR G O	0	0	0	0	0	0
N1	2723	ERWEE W J	25	23	43	44	37	29
N1	2771	ENGELBRECHT J.H	30	69	45	3	6	17
N1	2778	SNYMAN J.L	0	0	0	0	0	0
N1	2779	P K FARMING ENTERPRISES (PTY)	0	0	0	0	0	0
N1	2779	P K FARMING ENTERPRISES (PTY)	4	5	5	3	3	0
N1	2779	P K FARMING ENTERPRISES (PTY)	1683	1904	1585	1677	1897	1647
N1	2779	P K FARMING ENTERPRISES (PTY)	0	0	0	0	0	0
N1	2780	LAMPBRECHT A.M	0	0	0	0	0	0
N1	2795	HENNIE STOETERY EDMS BPK	0	0	0	0	0	0
N1	2807	MAREE J J	0	0	0	0	0	0
N1	2809	FOURIE D H	0	0	0	0	0	0
N1	2816	SCHURZ B	0	0	0	0	0	0
N1	2816	SCHURZ B	0	0	0	0	0	0
N1	2833	NOLTE D.G.H.	11	0	0	0	0	0
N1	2835	MEYER J.G	624	518	726	475	657	461
N1	2836	MEYER F.A.	260	192	298	159	318	405
N1	2851	TAYLOR G O	0	0	0	0	0	0
N1	2864	HEIN E M	58	42	59	34	37	47
N1	2899	HATTINGH G F S	0	0	0	0	0	0
N1	2907	W A FOURIE	697	533	685	652	691	801
N1	2913	NELL J H	11	7	11	8	15	11
N1	2933	LESLIE NOODWATERSKEMA	1144	883	1067	886	1581	990
N1	2941	VENTER W A	0	0	0	0	0	0
N1	2948	VAN DER WALT J P	55	33	35	24	30	26
N1	2952	VISSER F A J	26	47	45	30	38	39
N1	2953	VAN DER WOUDEN D	0	0	0	0	0	228
N1	2956	SMANT J	716	780	824	433	240	219
N1	2957	PISTORIUS J F W T	29	17	50	5	1	1
N1	2982	MOUTON P L	7	8	7	6	3	4
N1	3023	SMITH P	25	23	24	23	23	16
N1	3031	TALJAARD J A	141	69	139	50	300	208
N1	3031	TALJAARD J A	217	157	241	21	2	8
N1	3032	BEUKES A J	12	0	0	0	0	28
N1	3044	PEDEN S W	99	96	98	79	63	170
N1	3093	BOTHMA H N DE C	0	113	104	113	107	100
TOTAL	N1		830509	845908	898955	871867	1194507	878164

A10

WD313R

*** R A N D W A T E R ***

C O N S U M E R D R A W O F F R E P O R T

PAGE NO : 3
1996/05/14 12:15

PIPELINE	CONS.	NAME:	199511	199512	199601	199602	199603	199604
N2	361	LESLIE BRACKEN GOLD MINES LTD	55660	57850	62570	58620	64040	62260
N2	361	LESLIE BRACKEN GOLD MINES LTD	20	0	50	0	0	0
N2	784	KRUGER J C	142	127	102	91	103	136
N2	826	PRETORIUS I J	0	0	0	0	0	0
N2	833	LAMPRECHT P M	0	0	0	0	0	0
N2	833	LAMPRECHT P M	0	0	0	0	0	0
N2	2010	R U C MINING CONTRACTING CO (P	0	0	0	0	0	0
N2	2218	KANHYM ESTATES	0	0	0	0	0	0
N2	2390	VIVIERS H J	0	0	0	0	0	0
N2	2913	NELL J H	0	0	0	0	0	0
TOTAL	N2		55822	57977	62722	58711	64143	62396

WD313R

*** R A N D W A T E R ***

C O N S U M E R D R A W O F F R E P O R T

PAGE NO : 4
1996/05/14 12:15

PIPELINE	CONS.	NAME:	199511	199512	199601	199602	199603	199604
N3	365	WINKELHAAK MINES LTD	111190	103870	113120	109070	112580	105460
N3	526	STADSRAAD VAN EVANDER	0	0	0	0	0	0
N3	526	STADSRAAD VAN EVANDER	0	0	0	0	0	0
N3	526	STADSRAAD VAN EVANDER	62810	60430	69410	61980	67150	24940
N3	850	ADLEM J	0	0	0	0	0	0
N3	884	TALJAARD J A	0	0	0	0	0	0
N3	3031	TALJAARD J A	0	0	0	0	0	0
N3	3031	TALJAARD J A	21	17	26	9	31	30
TOTAL	N3		174021	164317	182556	171059	179761	130430

WD313R

*** R A N D W A T E R ***

C O N S U M E R D R A W O F F R E P O R T

PAGE NO : 5
1996/05/14 12:15

PIPELINE	CONS.	NAME:	199511	199512	199601	199602	199603	199604
N4	1	RAND WATER	0	0	0	0	0	0
N4	535	DORPSRAAD EMBALENHLE	279500	102080	501360	285090	304050	296430
N4	1747	BRITS A J	0	0	0	0	0	0
N4	1873	ZEKVELD T A	0	0	0	0	0	0
TOTAL	N4		279500	102080	501360	285090	304050	296430

WD313R

*** R A N D W A T E R ***

C O N S U M E R D R A W O F F R E P O R T

PAGE NO : 6
1996/05/14 12:15

PIPELINE	CONS.	NAME:	199511	199512	199601	199602	199603	199604
N5	365	WINKELHAAK MINES LTD	8913	8778	8174	8138	9216	7139
N5	533	SECUNDA MUNISIPALITEIT	359810	356280	384130	370810	414890	391050
N5	2218	KANHYM ESTATES	0	0	0	0	0	0
N5	2609	SECUNDA COLLIERIES MINE EXPANS	9120	13380	10580	16240	20920	23650
TOTAL	N5		377843	378438	402884	395188	445026	421839

APPENDIX B: SOFTWARE USERS MANUAL

Reservoir System Pumping Optimisation

User Manual

Version 1.02
March 1998

Neil Manson
Water Systems Research Group
Private Bag 3
University of the Witwatersrand
2050
n-manson@civen.civil.wits.ac.za

1. INTRODUCTION

Reservoir System Pumping Optimisation, or *RSPO* for short, is a tool developed for operators of municipal reservoir and pump station systems. It allows the operator to analyse data measured from their system, simulate the system, and calculate an optimum pumping policy.

This software was developed as part of a project entitled "Development of a Model to be used for the Economic Optimisation of the Pumping and Design Policies of Reservoir Systems". This project was funded by the Water Research Commission and conducted by B. Barta and N. Manson of the Water Systems Research Group, University of the Witwatersrand. This manual describes how to use the software, but details of the calculations used are given in Water Research Commission report No 757/1/98.

1.1. Reservoir Systems

A reservoir system will generally consist of a number of pump stations. There may, in fact, be no pump stations, in which case it will be a gravity fed system which will not require optimisation of the pumping. There will also usually be one or more reservoirs in the system, and a number of users whose requirements the system is designed to meet. *RSPO* models this generalised system by lumping all the pumping into a single pump station, all the reservoirs into a single reservoir, and all the user demands into a single demand. Figure 1 shows a schematic of the model used by *RSPO*.

If the actual system contains more than one pump station or reservoir, RSPO automatically sums the reservoir contents and pumping rates as the data is imported. The reason for this lumping is that it greatly simplifies the analysis and simulation, without affecting the optimum pumping policy that is eventually calculated.

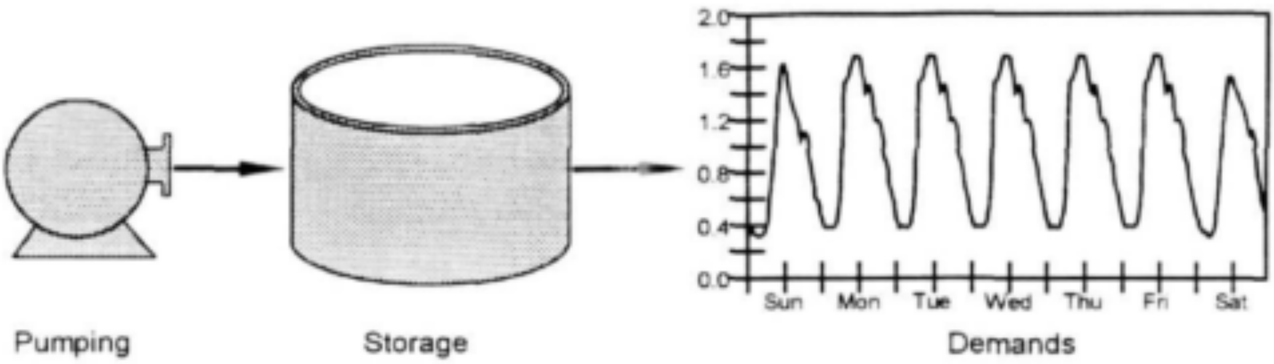

Figure 1 *RSPO* System Model

1.2. Operation or Pumping Policies

Each pump or pump station in a system such as that described will be able to pump water to the reservoir at certain rates. A simple pump station with only a single pump may have only two settings, namely *'on'* and *'off'*. A more complex system may have a number of pumps which can all run simultaneously, or may have variable speed pumps that can be set to run at any speed within a certain range. The flow rate at which the pump station will deliver water to the reservoir for any particular setting will depend on the types and number of pumps running, their operating speeds, the difference in height and distance between the pump station and the reservoir, and a number of other factors.

The operation or pumping policy for a reservoir system is the set of rules by which the operator decides on which pump setting to use at any particular time to meet the demands of the users. There will be certain limitations on this policy, for example, the reservoir must not be allowed to run dry, or to overflow. There are also costs involved in pumping water to the reservoir, for example, the cost of energy to drive the pumps, the cost of running the pump at less than its best efficiency, the cost of changing from one setting to another, and others. The function of this program is to assist the operator in finding the policy that allows the users demands to be met in the least costly manner possible. This is what is meant by the Optimum Operation Policy.

1.3. Optimisation of the Pumping Policy

How does one find an optimum policy? The optimum policy will be very dependent on the system limitations, the available pump settings and the various costs involved. The policy will also be dependent on the demands from the user, and this is where the major difficulty arises in trying to find an optimum policy. We can never know what the users demands will be in the future. What *RSPO* does is to assume that the future demands will be statistically similar to the demands in the past. Based on this assumption, historical demand data can be analysed and used to calculate an optimum pumping policy. As long as the future demands do remain statistically similar to the historical demands, the policy calculated will be optimum. But if anything changes, the policy will no longer be optimum, and it will be necessary to recalculate a new optimum policy. For this reason, it is recommended that operators recalculate the operation policy fairly often to ensure that it remains optimum. In this case, 'fairly often' means at least once a year, but preferably whenever changes are know to occur in the system. Changes in the system would include changes in the costs, in the demands or in the system configuration.

The main operations of the program are shown in Figure 2 below. These operations create a new system on the computer, import the historical demand data from the actual system and then analyse the data, simulate the reservoir and optimise the pumping policy. RSPO can then be used at the start of each time step to calculate the optimum pump setting. Each of these steps are discussed in detail in the following sections.

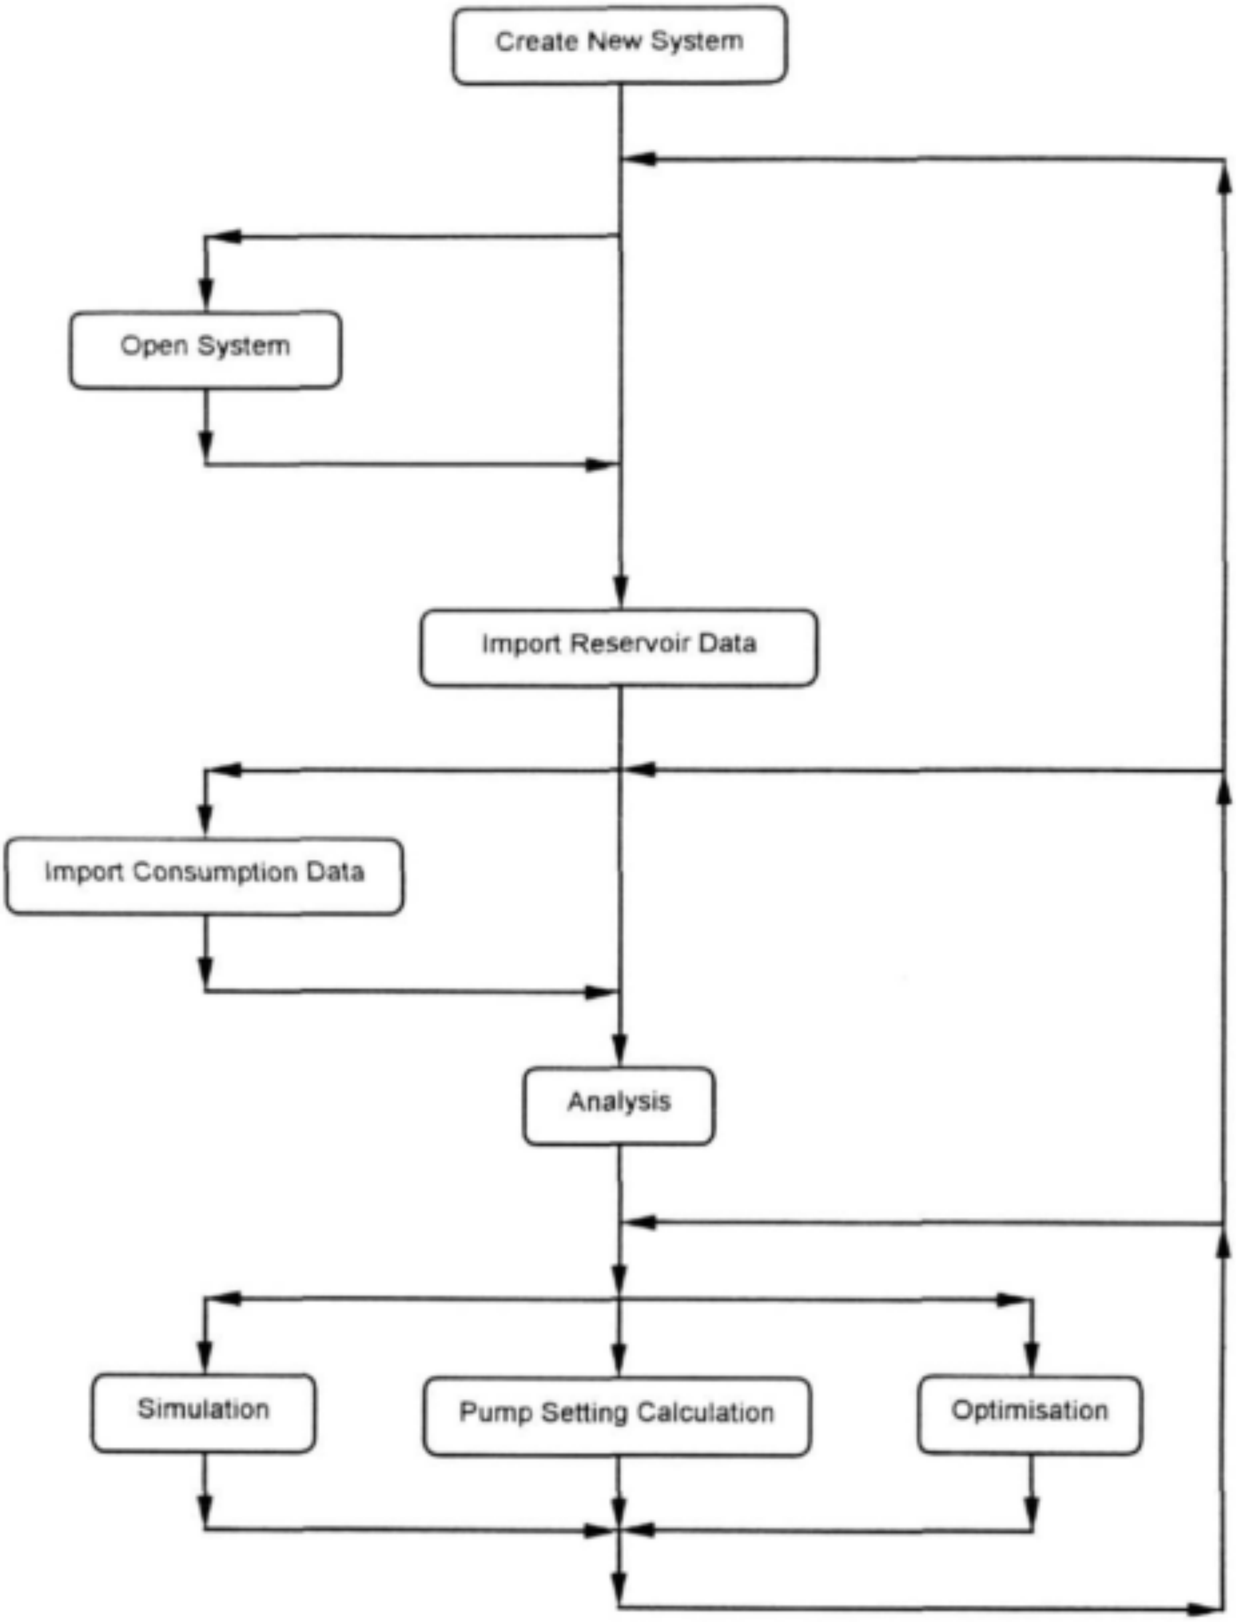

Figure 2 Sequence of Main Operations

2. GETTING STARTED

The first part of this manual is written as a tutorial for you to follow. It is recommended that you work through the tutorial sections, and use the example files supplied, as this is the quickest way to see what the program can do, and how to use it.

2.1. System Requirements

RSPO was written for the Microsoft® Windows™ operating system, version 3.1. It will also run under Microsoft® Windows™ 3.11 and Microsoft® Windows95™. Any computer that is capable of running any of these operating systems will be able to run *RSPO*. At least a 486 processor is recommended as *RSPO* performs some intensive mathematical calculations that can take some time to complete. The faster the processors speed, the quicker these calculations will be completed.

2.2. Installing the Software

To install *Reservoir System Pumping Optimisation* you need to insert the program disk into your 3½-inch 'Stiffy' drive. This will usually be the 'A:' drive, but it may be the 'B:' drive on your computer.

Follow the steps for the operating system you are using as given below:

<u>Microsoft® Windows™ 3.1x</u>

1. From *Program Manager*'s menu, select **File** and then **Run**.
2. Type in 'A:\SETUP' and press ENTER.
3. The *Setup* program will then run and the Welcome Screen will be displayed.
4. Click on the NEXT button or press ENTER.
5. Select the drive and sub-directory in which you want *RSPO* to be installed. You can also type in a new sub-directory. Then click on the NEXT button or press ENTER. To accept the default directory of 'C:\RSPO', just click on the NEXT button or press ENTER.
6. You will then see a list of optional program components. If you do not want to install any of these components, click on them or press the SPACEBAR to deselect them. When only the options you want are selected click on the NEXT button or press ENTER.
7. A confirmation screen will be displayed. Check that the options are what you want. If the are not, click on the BACK button to change the options. If the options are correct, click on the INSTALL button or press ENTER.
8. *Setup* will now copy some files to your hard drive and unpack the files. It will also install a new program group in *Program Manager*. When *Setup* is complete, click on the DONE button or press ENTER to exit from *Setup*. To start *RSPO*, double click on the *Res. Sys, Pumping Optimisation* program group in *Program Manager*, and then double click on the *Reservoir System Pumping Optimisation* item, which appears as 

<u>Microsoft® Windows95™</u>

1. Click on **Start**, and select **<u>R</u>un**.
2. Type in 'A:\SETUP' and press ENTER.
3. The *Setup* program will then run and the Welcome Screen will be displayed.
4. Click on the NEXT button or press ENTER.
5. Select the drive and sub-directory in which you want *RSPO* to be installed. You can also type in a new sub-directory. Then click on the NEXT button or press ENTER.
6. You will then see a list of optional program components. If you do not want to install any of these components, click on them or press the SPACEBAR to deselect them. When only the options you want are selected click on the NEXT button or press ENTER.
7. A confirmation screen will be displayed. Check that the options are what you want. If the are not, click on the BACK button to change the options. If the options are correct, click on the INSTALL button or press ENTER.
8. *Setup* will now copy some files to your hard drive and unpack the files. It will also install an *RSPO* item in your start menu.
9. To run *RSPO*, click on **Start**, select **Res. Sys. Pumping Optimisation** and then **Reservoir System Pumping Optimisation**.

2.3. How to use this Manual

If this is the first time you have used the program *Reservoir System Pumping Optimisation*, it is recommended that you read through the first two chapters of this manual, **SETTING UP A SYSTEM** and **PROCESSING SYSTEM DATA**. These chapters explain, in a step-by-step manner how to use the program. You can follow along using the example data supplied with the program. The remaining two chapters, **MENU OPTIONS** and **RSPO FILES**, give reference information on all the available commands and the files used by *RSPO*.

2.4. Manual Conventions

This User Manual uses a number of conventions to assist in explaining the use of the Program *RSPO*.

`Monospace`	Monospaced text in a border will be used to indicate the contents of text files used by *RSPO*.
Italics	Italic text will be used for program and data file names.
Bold	Bold text will be used to indicate menu options that the user needs to select.
SMALL CAPS	Small capital letters will be used to name any of the keys on the keyboard that the user needs to press, or for the names of buttons on the screen that the user should click on.
Double Underline	Double underlined text will be used to indicate cross references to other sections and chapters in this manual.

Now that you have installed the software, and know how to use this manual, let us proceed with optimising a system.

2.5. Disclaimer

This software and the associated data and files are supplied "as is". The Author, the Water Research Commission, the Water Systems Research Group and the University of the Witwatersrand cannot and do not guarantee that any functions contained in the Software will meet your requirements, or that its operations will be error free. The entire risk as to the Software performance or quality, or both, is solely with the user and not the Authors. You assume responsibility for the installation, use, and results obtained from the Software.

The Author makes no warranty, either implied or expressed, including with-out limitation any warranty with respect to this Software documented here, its quality, performance, or fitness for a particular purpose. In no event shall the Authors be liable to you for damages, whether direct or indirect, incidental, special, or consequential arising out the use of or any defect in the Software, even if the Authors have been advised of the possibility of such damages, or for any claim by any other party.

All other warranties of any kind, either express or implied, including but not limited to the implied warranties of merchantability and fitness for a particular purpose, are expressly excluded.

3. SETTING UP A SYSTEM

3.1. Creating a New System

When you start *RSPO* you will see the following main screen.

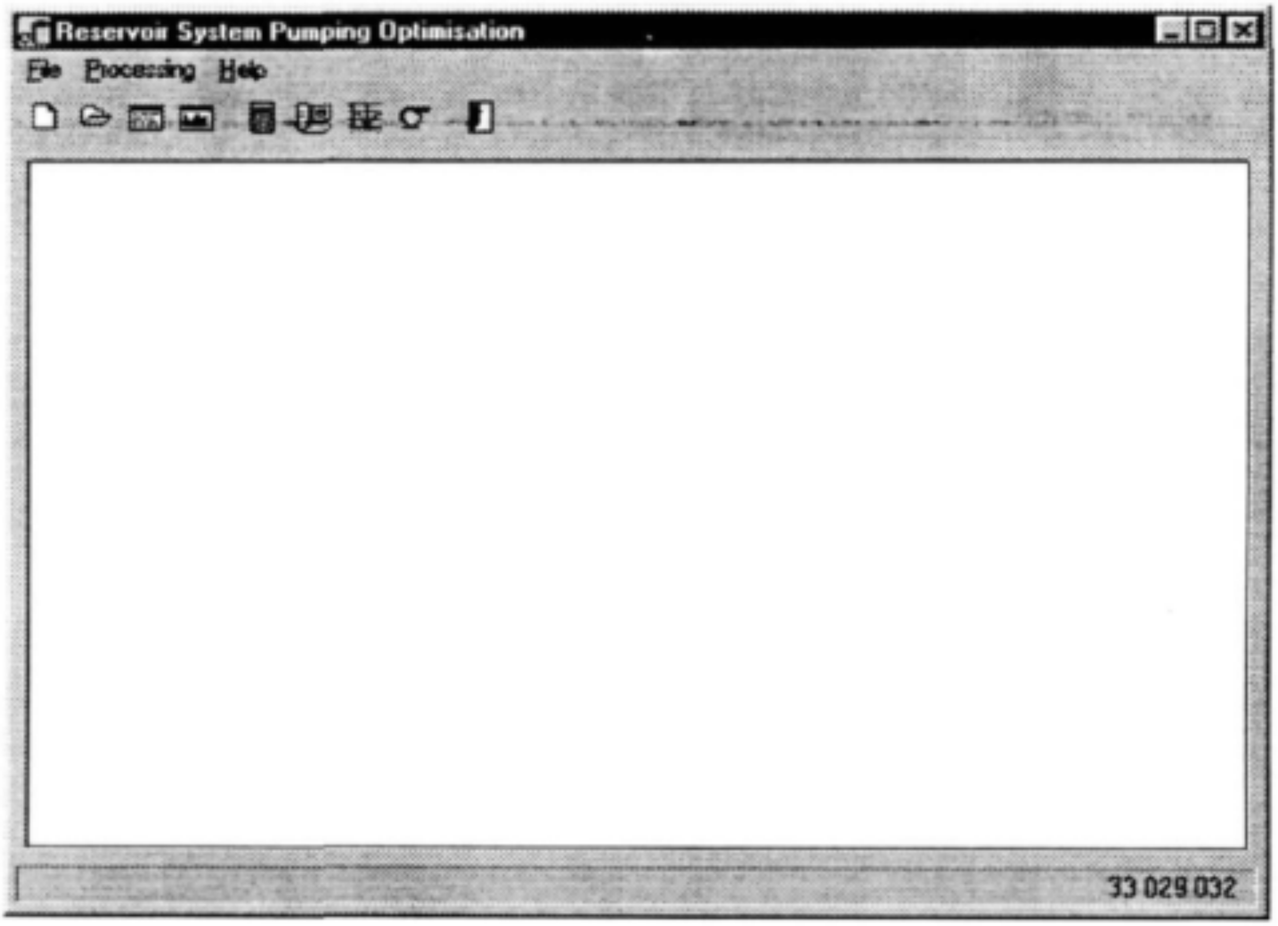

Figure 3 *RSPO* Main Screen

Before you can start analysing a system, you need to create a new system on your computer. To do this, choose **File|New** from the menus or click on the  toolbar button. You will then see the Create New System dialog, as shown below:

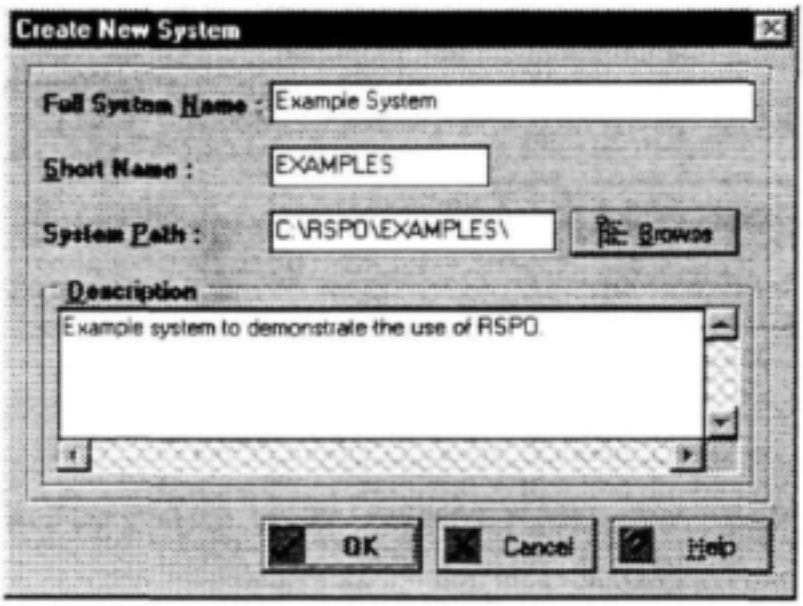

Figure 4 Create New System Dialog

To create the new system, type in a name for the system in the Full System Name box. This is the name that will be used to refer to this particular system in all the output report files. The example that will be used throughout this manual will be called '*Example System*'. When you have typed in the name, press TAB to move to the next box. You will see that *RSPO* will automatically fill in the Short Name and System Path boxes, based on the Full Name. The Short Name is a maximum of 8 characters long, and is used to name some of the files that are generated by *RSPO*. The System Path is the path to the sub-directory in which all the files for this particular system will be kept. This manual assumes that the program was installed in the default directory of '*C:/RSPO*'. If you have installed RSPO in a different directory, your system may look slightly different to this manual. The short name for this *Main Example System* will be '*EXAMPLES*', and the path is '*C:/RSPO/EXAMPLES*'.

If you want to keep the values assigned by *RSPO* then you can just press TAB three times to get to the Description box, or you can change the Short Name and System Path to anything you like. You can also click on the BROWSE button to select a new System Path. In the Description box you can type any other information that you would like to associate with the system. For example, type "Example system to demonstrate the use of RSPO."

Once all the boxes contain the correct values, press ENTER, or click on the OK button. *RSPO* will then create a new sub-directory for your system files and store a *System Initialisation File* in that sub-directory. You will also notice that the title bar of the Main Screen now shows the Full Name of the system, so that you can always be aware of which system you are working on. For details on all the different file types and their default names and extensions, please see the chapter entitled **RSPO Files** at the end of this manual.

3.2. Importing Data

Once the new system has been created, the next step is to import data to analyse. The *System Run Data* consists of the reservoir content, the pumping rate and the demand for each time step in the range of data. This data is imported from a standard ASCII text file. The data in the file must be in the format specified in the chapter **RSPO Files** at the end of this manual. Be careful that the data in the text file is correct and that it contains no gaps, as no checking is done on data imported from a text file.

To import data from a text file, select **File|Import|Text Data** from the menus. A File Open dialog box will appear. This is a standard dialog that you will see every time you need to open a new file. An example is shown below:

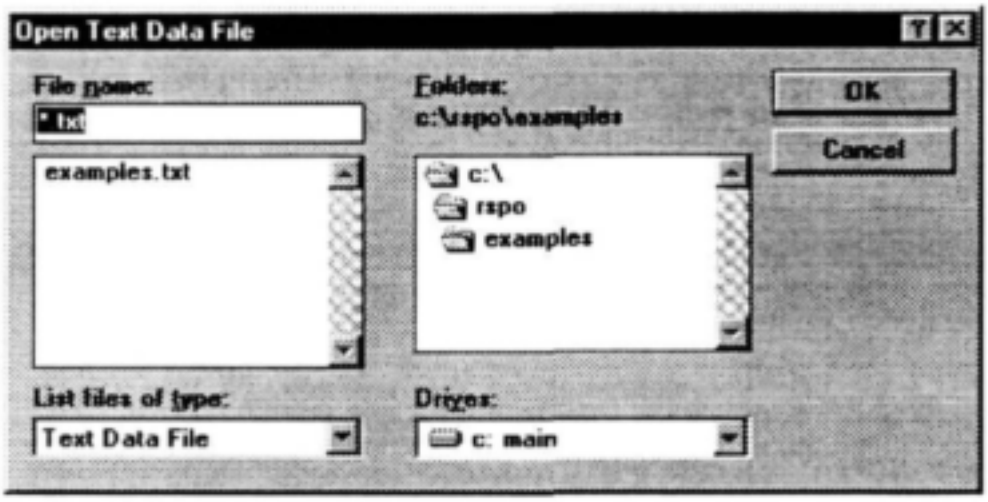

Figure 5 Open File Dialog

A set of example data is provided in the file *EXAMPLES.TXT*. Select this file and click on the OK button. The Import Text Data dialog will now appear, as shown below

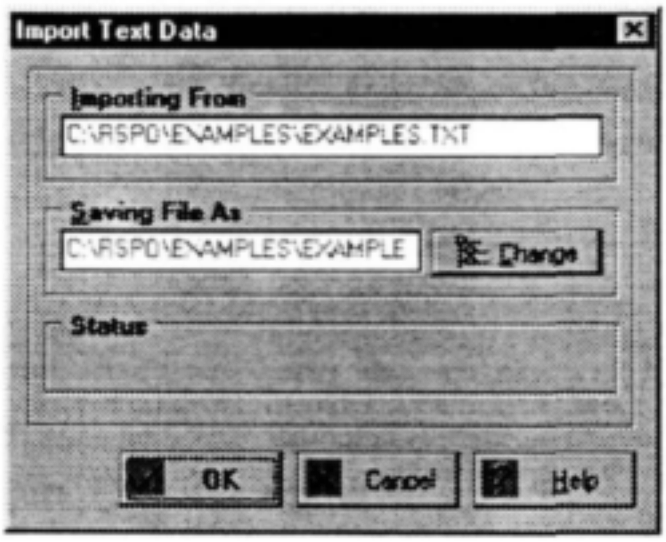

Figure 6 Import Text Dialog

Press ENTER or click on OK to start the import. You can also change the name of the file to which the data will be saved by clicking on the CHANGE button. *RSPO* will now import the text data. As it is doing so, the date and time of the current data record will be displayed in the Status panel so you can see the progress of the import.

3.3. Graphing the Data

Once you have imported the data, it is usually a good idea to graph it as this allows you to get a general impression of what the data looks like, and helps you to find any abnormalities. *RSPO* provides a tool to graph the imported system data. From the menus, select **File|Graph**, or press the toolbar button. Select the **.RUN* file containing the data you would like to graph, in this case, *EXAMPLES.RUN*. A Graph Dialog, such as the one in the figure below, will then be displayed.

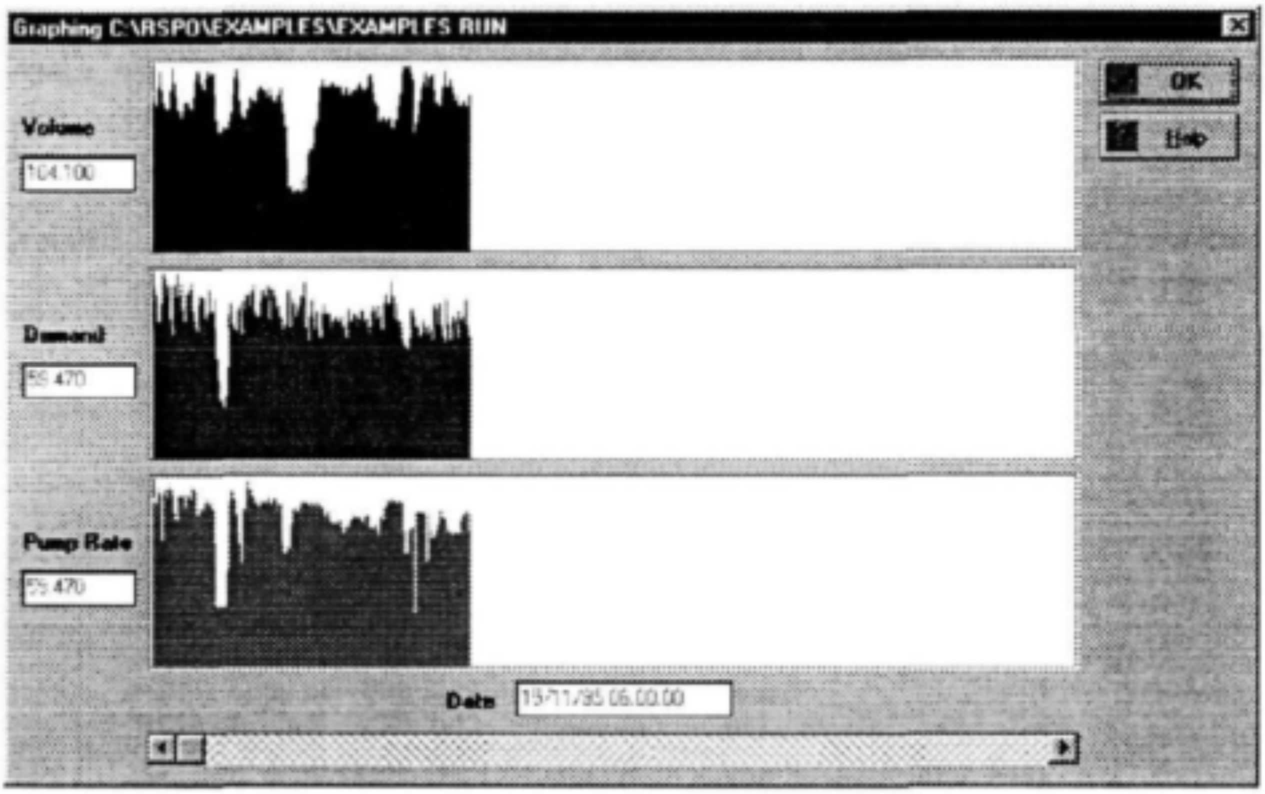

Figure 7 Graph Dialog

You can use the scroll bar at the bottom of the dialog to move the black line (shown here on the extreme left) through the data. The date and time, volume, demand and pump rate values for the data point indicated by the black line are also shown in the text boxes. When you are done viewing your data, you can press ENTER or click on OK.

3.4. Exporting Data

It is often useful to be able to export the system data to a standard ASCII text file. This can then be used in a graphing package, for example. This is especially useful for data that has been synthetically generated by RSPO. For example, to export the data that has just been imported, proceed as follows. Choose **File|Export|System**

Data from the menu. An Open File dialog will be displayed. Select the *System Run File* that you would like to export, i.e. *'EXAMPLES.RUN'* and click on OK. The Export Data to ASCII dialog will then appear, as shown below.

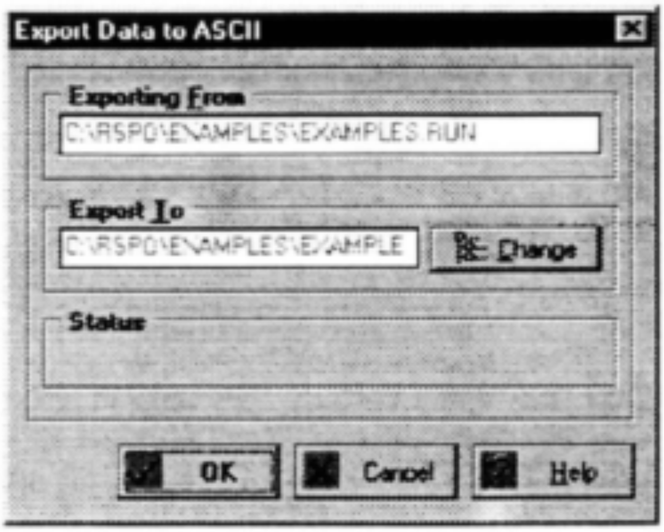

Figure 8 Export Data to ASCII Dialog

If you want to change the name of the file in which the exported data will be saved, click on the CHANGE button. Once you are satisfied with the file names, click on OK. The status panel will display the date and time of the data point that is currently being exported. When this process is complete. *RSPO* will write out a *Run Text File* (*.RTX), and display the file.

These reports show the report name, the date and time generated, the file name of the exported file, and then list the values for all the available data points.

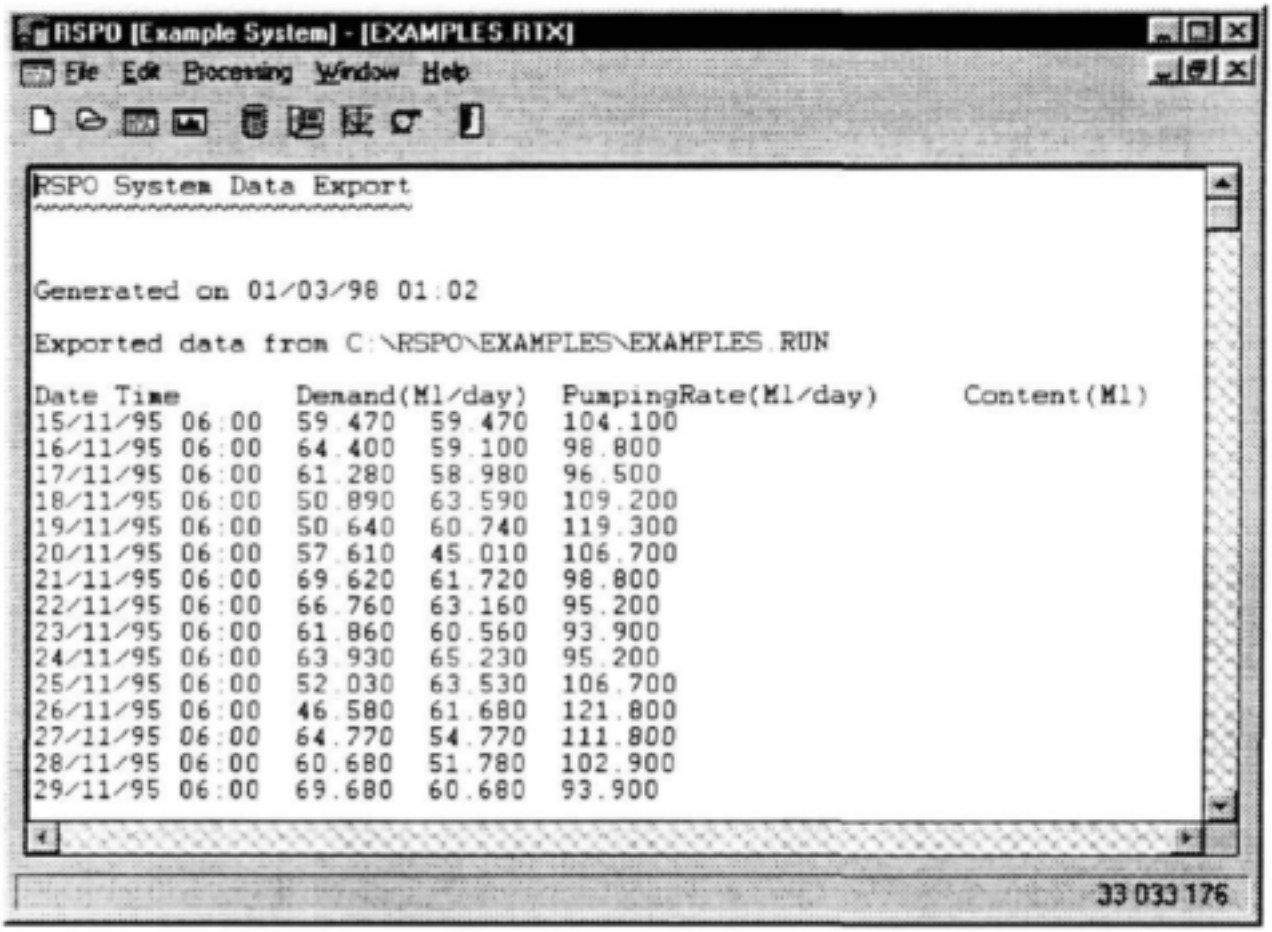


```
RSPO System Data Export


Generated on 01/03/98 01:02

Exported data from C:\RSPO\EXAMPLES\EXAMPLES.RUN

Date Time         Demand(Ml/day)  PumpingRate(Ml/day)      Content(Ml)
15/11/95 06:00    59.470   59.470   104.100
16/11/95 06:00    64.400   59.100   98.800
17/11/95 06:00    61.280   58.980   96.500
18/11/95 06:00    50.890   63.590   109.200
19/11/95 06:00    50.640   60.740   119.300
20/11/95 06:00    57.610   45.010   106.700
21/11/95 06:00    69.620   61.720   98.800
22/11/95 06:00    66.760   63.160   95.200
23/11/95 06:00    61.860   60.560   93.900
24/11/95 06:00    63.930   65.230   95.200
25/11/95 06:00    52.030   63.530   106.700
26/11/95 06:00    46.580   61.680   121.800
27/11/95 06:00    64.770   54.770   111.800
28/11/95 06:00    60.680   51.780   102.900
29/11/95 06:00    69.680   60.680   93.900
```

Figure 9 ***Run Text File*** **Display**

4. PROCESSING SYSTEM DATA

Once you have imported your data, you are in a position to begin processing them. The first step is to analyse the data as this calculates the statistics necessary for both the simulation and optimisation options.

4.1. Analysing the system Data

To analyse a set of data, choose **Processing|Analysis** from the menu, or click on the toolbar button, and an Open File dialog will appear. Select the *System Run File* that you want to analyse and click on OK. You will then see the following Data Analysis Dialog.

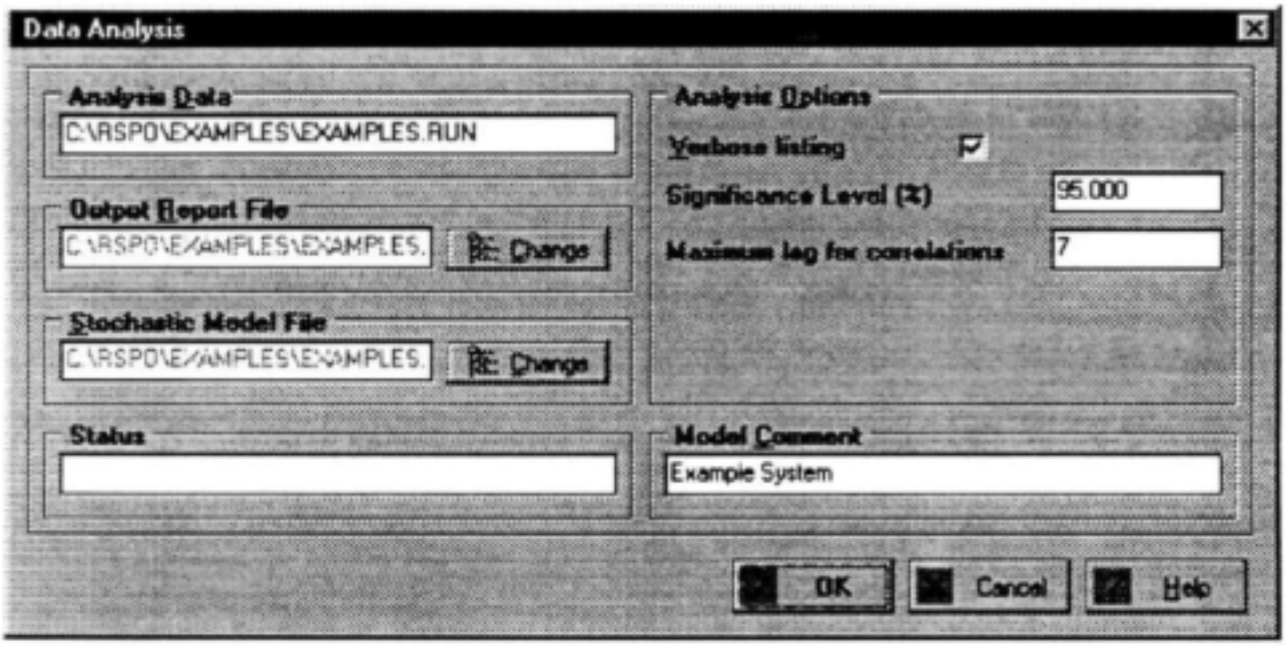

Figure 10 Data Analysis Dialog

The first edit box on this screen shows the *System Run File* that will be analysed. *RSPO* automatically assigns names to the *Analysis Report File* and the *Stochastic Model File*. You can change the names of these files by pressing either of the CHANGE buttons. It is recommended that you keep the names the same as each other, except for the extensions, so that you know which report is associated with which model.

The Verbose Listing option can be turned on if you want the report to contain the intermediate results of some of the calculations. This option only affects the output report, and not the actual calculations. The significance level is used in the calculation of the periodic component and in fitting the probability distribution. It is recommended that you do not set the significance level to a value of less than about 90%.

If you set the Maximum Lag for Correlations to a value greater than zero, *RSPO* will calculate the autocorrelation values for your data, up to the value set. These autocorrelations do not affect the stochastic model, but are written in the report for your information.

You can enter any information you like in the Model Comment edit box, up to 255 characters. This is useful to remind yourself of the exact settings used in the calculation of the model.

When you have entered all the required values, click on OK to begin analysing the data. You may have to wait quite a while, especially if you have large amounts of data or a relatively slow computer, so now might be a good time to get a cup of coffee. While the analysis is running, information will be displayed in the Status edit box to let you know what is being done. When the analysis is complete, *RSPO* will write out an *Analysis Report File*, and display it. Each section of this report file is discussed below:

```
RSPO Analysis Report
~~~~~~~~~~~~~~~~~~~~

Example System

Generated on 01/03/98 01:02

Input Data File:        C:\RSPO\EXAMPLES\EXAMPLES.RUN
Output Stochastic Model:                 C:\RSPO\EXAMPLES\EXAMPLES.MDL
```

The first section of the report shows the report name, the full system name, the date and time the report was generated. It also shows the names of the input and output files that were used and generated by the analysis.

```
Start Date & Time       15/11/95 06:00
End Date & Time         12/05/96 06:00

Data Points             180
Minimum Demand                          19.16 Ml/day
Maximum Demand                          69.68 Ml/day
Minimum Pumping Rate                    19.65 Ml/day
Maximum Pumping Rate                    66.76 Ml/day
```

This next section of the report shows some basic information about the data available for analysis.

```
m                       -0.0337752708416917
c                       56.2718311847758
Linear R Squared        0.0478160553849888
a                       54.6400060516249
b                       0.99954905429196
Log R Squared           0.0161850032044853

Linear trend of the form y = mx + c removed, where:
m                       -0.0337752708416917
c                       56.2718311847758
LinR2                   0.0478160553849888
```

This section of the report shows the calculation of the secular trend. *RSPO* calculates both a linear and a logarithmic trend, and removes the one that fits the data most accurately. In this case, a better fit is indicated by value of R Squared that is closer to one. If the Verbose Listing option is not checked, only the last three lines of this section are included in the report.

The following periodic components had a significance of
greater than 95% and were removed:

Annual Cycle: No Cycle Calculated

Weekly Cycle: Average = 0.0185984126984373

Period	Alpha	Beta	Significance
7.00	-1.49565921055988	-2.83107663492425	0.999636137683088
3.50	-2.03839456339651	-0.16131969373324	0.963720186333785

Daily Cycle: No Cycle Calculated

The next section shows the calculation of the periodic trend components in the demand data. An annual trend will only be calculated if there is more than a years data available, and a daily trend can only be calculated if the time step is less that 12 hours. Any trends that have a significance higher than that set on the Data Analysis dialog are shown in the report, and removed from the demand data. For details on how the periodic component is calculated, see Water Research Commission report No 757/1/98.

Estimated Parameters for Residual Stochastic Component

Mean	0.00138188586646455
StdDev	7.35387825011152
Skew	-2.22349057666834
Kurt	11.227225833007

(Estimated by method of moments)

The estimated parameters for the residual component are calculated by the method of moments from the demand data after the secular and periodic trend components have been removed.

Serial autocorrelations

Lag	Correlation Coefficient
0	1
1	0.701062386158384
2	0.548878474945079
3	0.463176962038485
4	0.316486163596267
5	0.183619980720542
6	0.0944963105141409
7	0.0714724566914905

Serial correlation coefficients are only calculated if the Minimum Lag setting on the Data Analysis dialog is greater than zero. These give the values of the autocorrelation between points that are separated by "Lag" points.

Test statistic	0.1461
Estimated Autoregressive Parameters	
Phi(1,1)	0.701062386158384
First guess at AR parameters	
Phi(1,1)	0.701062386158384
Value:	710.886922634639
Autoregressive Model of order 1 fitted, with the following best fit autoregressive parameters	
Phi(1,1)	0.585310014373619
Iterations:	34
Value:	700.867418176444

An autoregressive model is used to simulate the autocorrelation in the demand data. The Autoregressive parameters are first estimated using the recursive estimation formula described in Water Research Commission report No 757/1/98. The estimates are then used as an initial guess for the optimisation routine. The change in the output value from the first guess to the best fit gives an indication of how accurate the first guess was. If there is very little change, then the first guess was a good one. Once the best-fit values for the autoregressive parameters have been found, the autoregressive component is removed from the demand data. This leaves only the Independent Random component

Estimated Parameters for Independent Random Component

Mean	-0.000235858848420868		
StdDev	5.33644556519468		
Skew	-0.585781382189186		
Kurt	5.08183373853231		
(Estimated by method of moments)			
Bins	12		
Elements/Bin	18		
Min Value	-363.3598		
Bin No	Bin End	Observed	Expected
1	-21.0232	0.0000	0.0073
2	-5.8425	18.0000	24.6174
3	-3.3189	18.0000	23.4373
4	-2.0049	18.0000	15.5839
5	-1.0178	18.0000	12.7437
6	-0.0054	18.0000	13.5415
7	1.2370	18.0000	16.5700
8	2.3099	18.0000	13.6404
9	3.8773	18.0000	17.7868
10	5.8944	18.0000	17.8317
11	13.2105	18.0000	23.0428
12	355.5471	0.0000	1.1972
Total Values		180	180.0000
Chi Squared	10.8799		
Iterations	127		
Min Chi Squared	3.2747		

A Normal probability distribution was fitted to
the independent random component, with the following best fit parameters:

Mean	0.10274512216477
StdDev	4.45537476662543

The Chi Squared Goodness-of-Fit test gives a significance of
99.3274285495777%

The last section of the analysis fits a normal distribution to the independent random component. The mean, standard deviation, skewness and kurtosis are calculated by the method of moments. Although this method is easy to calculate, it does not necessarily give the best fit. A Chi Squared (χ^2) test is performed, and the parameters are adjusted until the best fit is obtained. If the verbose listing option is checked, the details of the χ^2 test are given in the report. These 'best fit' parameters are then used in the stochastic model. The significance of the fit is also given, and should be as close to one as possible.

When this analysis is complete, the results can be used to run simulations of the system, or to optimise the system, as described in the following sections.

4.2. Running Simulations

There are two main reasons for running simulations of the performance of a reservoir system. The one is to see the effect of different operating systems on the total running cost of the system, and the other is to use the stochastic model fitted during the analysis phase to generate synthetic demand data. Synthetic demand data is data that is generated by the computer, but is statistically similar to the historical measured demand data. Synthetic demand data is useful for testing the reservoirs system for much longer periods than the available historical demand data.

To run a simulation, select **Processing|Simulation** from the menu, or click on the toolbar button. You will then see the following System Simulation Dialog.

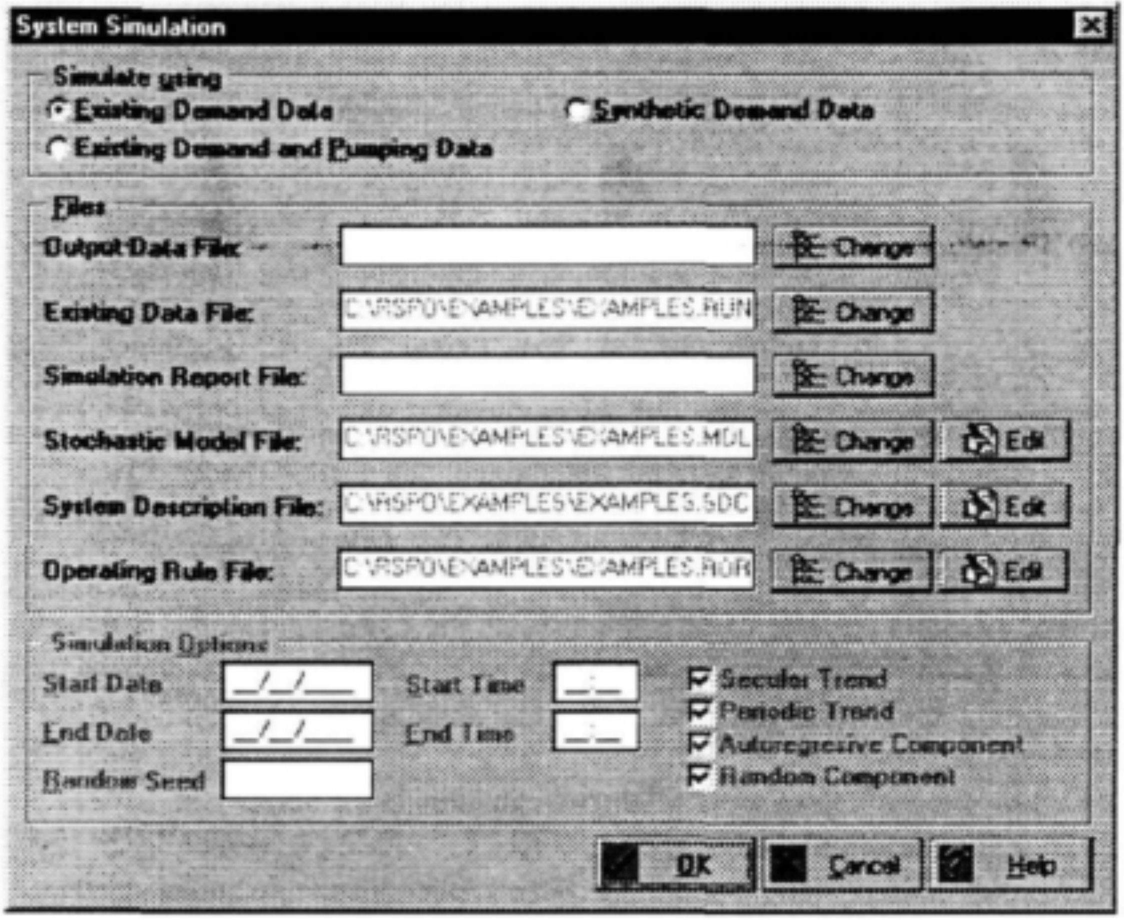

Figure 11 System Simulation Dialog

In the first box you can select whether to simulate the system using one of the following options:

- Existing Demand Data - Simulations using this option will take the demand data from an existing *System Run File* and use the operating rule stored in an existing *Reservoir Operating Rule File* to calculate the pumping rate for each time step. This option also requires an existing *Stochastic Model File* in order to predict the next time steps demand.
- Existing Demand and Pumping Data - Simulations using this option take both the demand and pumping rates from an existing *System Run File*. This option is used to calculate the operating costs for the historical demand data as a reference.
- Synthetic Demand Data - This option generates the demand data from the *Stochastic Model File*, and calculates the pumping rates from the *Reservoir Operating Rule File*. This option can be used to generate *System Run Files* that are much longer than the available historic demand data.

All simulations will require you to enter the names of an *Output System Run File* and a *Simulation Report File*. By default, these files will have the same name, except for the file extension, but you can change this if you need to. The *Output System Run*

File will contain the Demand, Pumping Rate and Reservoir Volume data that will be generated by the simulation. The *Simulation Report File* will contain information on the results of the simulation. This file is discussed in detail later in this section.

The names of all the other files are automatically generated by *RSPO*, but you can change any of them if you need to by clicking on the appropriate CHANGE button. If the option you select requires a *Stochastic Model File*, you must use a file generated by using the Analysis option of this program. You can manually edit the *Stochastic Model File*, but this is not recommended unless you are sure of what you are doing.

All the simulation options also require a *System Description File*. A typical *System Description File* is shown below. The file consists of a number of sections, each with a section heading enclosed in square brackets "[...]". Each section then contains a number of variable descriptions. Each description has a variable name, followed by an equal sign, "=", followed by the variable's value. For example, in the [Reservoir] section, there is a line that says "Maximum Volume=124.354". The variable name is "Maximum Volume", and describes the maximum available volume in this reservoir. The value after the equal sign is 124.354 which tells the program that this reservoir can contain a maximum of 124.354 Ml of water.

In the *System Description File*, the semicolon character ";" is used as a comment character. Any information after the semicolon is ignored and does not need to be in the file. It is just there for your information.

Each section of the *System Description File* will be described separately below:

```
[Header]
System=C:\RSPO\EXAMPLES\EXAMPLES.RPO
Comment=Example Reservoir System to Demonstrate the use of RSPO
Time=01:02
Date=01/03/98
```

The [Header] section describes which system the *System Description File* applies to. The System variable must contain the full path of the *System Initialisation File*. The Comment variable can contain any information you like, up to 255 characters. The Time and Date variables are useful to keep track of when the file was last edited. If you edit the file by clicking on the EDIT button in the System Simulation Dialog, *RSPO* will automatically update these variables, but if you edit the file any other way, you will have to update these values manually.

```
[Reservoir]
Start Volume=104.10  ; Simulation initial volume
Max Volume=124.354; Maximum possible volume in the reservoir
High Volume=117.700            ; High limit volume
Low Volume=80.000  ; Low Limit volume
Min Volume=0          ; Absolute minimum volume, usually zero

; Each of the following penalties must be given in Rands per Ml
Max Penalty=1000               ; Penalty for exceeding the maximum
                               ; volume. Should include the cost of
                               ; lost water.
High Penalty=0        ; Penalty for exceeding the high limit
                               ; volume
Low Penalty=25                 ; Penalty for not meeting the low limit
                               ; volume
Min Penalty=1000               ; Penalty for running the reservoir dry

; Storage costs
Storage Cost Slope=10          ; Basic cost of storage, in Rands per Ml
Storage Cost Offset=10         ; Reservoir fixed costs in Rands
```

The [Reservoir] section describes the constraints and costs that are applicable to the reservoir in the system. The Min Volume and Max Volume variables describe the absolute minimum and maximum volumes allowable in the reservoir. All the volumes must be in megalitres (Ml). The High Volume and Low Volume variables are warning limits on the volume in the reservoir. The low volume warning is often used to ensure that the volume in the reservoir is always sufficient to ensure that an emergency demand can be met. The Start Volume variable describes the volume that is assumed to be in the reservoir at the start of a simulation. This must be within Max Volume and Min Volume limits.

The next four variables allow you to impose a financial penalty if the limits are not met. The penalties must be in Rands per Ml over or under the limit. If the Max Volume limit is exceeded, both the Max Penalty and the High Penalty will be applied. Similarly if the Min Volume limit is not met.

The storage costs allow you to apply a straight line cost function to the cost of storing water in the reservoir. The cost at each time step is calculated as Cost = Storage Cost Slope * Reservoir Volume + Storage Cost Offset. A positive slope will cause the optimum operation policy to keep the reservoir as empty as possible, while a negative slope will cause the optimum operation policy to keep the reservoir as full as possible.

```
[Pumping]
; Number of possible pump settings, excluding all pumps off
Settings=3

; Coefficients of equation giving flow (Ml/d) as a function of
; reservoir content (Ml). Content is the variable x
Flow1=0.00002*x^2-0.04322*x+39.39856
Flow2=0.00003*x^2-0.05935*x+62.01738
Flow3=-0.00001*x^2-0.04711*x+88.58809

; Coefficients of equation giving power consumption (kW) as a function of
; Flow (Ml/d). ; Flow is the variable y
Power1=0.00020*y^3+0.00929*y^2+4.29534*y+242.22940
Power2=0.00007*y^3+0.00641*y^2+4.44414*y+470.52090
Power3=0.00000*y^3-0.00057*y^2+4.95610*y+1016.76150
```

The [Pumping] section describes the pumping that is available in the system. Settings gives the number of different settings that are available. This will include all the combinations of all the pumps that can be run simultaneously. If there is a variable speed pump available in the system, it must be approximated by a number of individual settings. The setting that has all pumps turned off is always assumed, and therefore does not need to be included in the *System Description File*.

For each of the settings there must be a Flow? variable. This gives an equation that describes the flow rate as a function of x, the reservoir content. These flow rates must be in increasing order, that is, Flow1 must deliver a smaller flow rate than Flow2 over the entire range of reservoir contents. The reservoir content, x, must be in Ml and the flow rate must be given in Ml/d.

The Power? variables are similar to the Flow? variables. For each of the settings, there must be a Power? equation that can be used to calculate the power consumption for a given flow rate. The flow rates, y, must be in Ml/d, and the power must be in kW.

```
[Energy]
; Number of Power tariff zones
; The zones MUST total 24 hrs, and must have NO overlaps and NO gaps!
Zones=2

; Zone start and end times
Zone1 Start=06:00
Zone1 End=18:00
Zone2 Start=18:00
Zone2 End=06:00

; Zone power cost in cents per kWh
Zone1 Cost=5.94
Zone2 Cost=4.75
```

The energy section describes the cost of energy that will be used for pumping that is available to the system. Provision is made for multiple tariff zones. For each of the zones given by the Zones variable, there must be a start time and an end time, as well as a cost, in cents per kWh. The zones must be continuous, and total exactly 24 hours. If they do not, the results of the simulation will be unpredictable and inaccurate.

```
[Switching]
; Cost of switching from one pump setting to another, in Rands.
; There should be up to (Settings+1)^2 switches. Default value for
; switches not listed is R0,00
From0to0=0
From0to1=25
From0to2=50
From0to3=75
From1to2=25
From1to3=50
From2to3=25
```

The [Switching] section describes the cost of switching from one pump setting to another. *It is not necessary to list every single switch cost*, but any costs not listed are assumed to be zero. The costs can be for switching pumps on and for switching them off, and they may also be negative.

You can change the *System Description File* by clicking on the CHANGE button, and you can edit any of these settings by clicking on the EDIT button.

The *Reservoir Operation Rule File* contains the operation policy that is used to calculate the pumping rate at any particular stage of the simulation. This file is not needed if the Existing Demand and Pumping Data Simulation option is chosen. A typical *Reservoir Operation Rule File* is shown below.

```
[Header]
System=C:\RSPO\EXAMPLES\EXAMPLES.RPO
Comment=Example Reservoir Operating Rule
Generated=01:01:00  01/09/97

[Change Levels]
Level1=50
Level2=0
Level3=-25
DontChange=0
```

The [Header] section is similar to the header section of the *System Description File*, and describes which system this operation policy belongs to. If the *Reservoir Operation Rule File* was generated by the Optimisation function of this program, the Generated variable will show the date and time the file was created.

The [Change Levels] section describes the actual pumping policy. During the simulation, the pump setting is calculated for a particular time step at the beginning of that step by calculating the expected reservoir volume. The expected reservoir volume is the current volume less the best estimate of the demand during the time step. Each of the Level? variables indicate at which reservoir level to change from one pump setting to the next. For example, Level1=50 indicates that, if the expected content is less than 50, the pump setting should be setting 1. If the expected content is less than Level2, that is 0Ml, then the pump setting should be Setting 2. If the expected contend is higher than the highest Level value, then all the pumps will be turned off. There should be as many Level? variables as the value of the Settings variable in the *System Description File*.

The DontChange value is a variable that governs whether or not to change the setting if the new setting is sufficiently close to the old setting. For example, if the current setting is setting 0 and the new expected content is 48 Ml, then the new setting should be setting 1. But if the DontChange value is 10% then the DontChange range will be 10% of the difference between Level1 and Level2, that is 5 Ml. The new expected contend is within 5 Ml of the change level so no change in setting will occur. This variable helps to reduce the cost of running a reservoir system when the cost of switching from one setting to another is relatively high.

The *Operating Rule File* can be changed or edited by clicking on the appropriate CHANGE or EDIT buttons.

If you choose the Synthetic Demand Data option then you will have to enter the Simulation Options. The Start Date, Start Time, End Date and End Time options describe the period over which you would like to simulate the reservoir option. The Random Seed variable is useful if you wish to run different simulations with the same synthetic demand data. In that case you would enter the same number for the Random Seed for each simulation. The Random Seed can be any number between 0 and 65535. If you don't need to do this, just leave this field blank. The other four check boxes allow you to include or exclude specific components of the stochastic model. It is recommended that all four of these be checked.

To follow along with this manual, make sure the Existing Demand Data option is selected. Change the *Output Data File* to 'SIMULATE.RUN', and you will see that the *Simulation Report File* is automatically updated to 'SIMULATE.SMR'. Click on the OK button to start the simulation. Simulations will usually run very quickly, unless you are simulating a very long period. When the simulation is complete, the *Simulation Report File* will be displayed, as discussed below:

```
RSPO Simulation Report
~~~~~~~~~~~~~~~~~~~~~~


Example System

Generated on 01/03/87 01:02


Simulation run using existing Demand data

Output System Run File:       C:\RSPO\EXAMPLES\SIMULATE.RUN
Demand data taken from:       C:\RSPO\EXAMPLES\EXAMPLES.RUN
Stochastic Model File:        C:\RSPO\EXAMPLES\EXAMPLES.MDL
System Description File:      C:\RSPO\EXAMPLES\EXAMPLES.SDC
Reservoir Operation Rule:     C:\RSPO\EXAMPLES\EXAMPLES.ROR
```

This first section gives details of the type of report, the system for which the simulation was run, and the names of all the files that were used or created by the simulation.

```
Simulation Results Summary
~~~~~~~~~~~~~~~~~~~~~~~~~~

Simulation started at    15/11/95 06:00
and completed at         12/05/96 06:00
Number of data points         180
Time step                1d 00h 00m 00s
Exceeded Max Volume           0
Exceeded High Volume          0
Below Low Volume              176
Below Min Volume              0
```

The next section of the *Simulation Report File* summarises the results of the simulation. It gives details of the period of the simulation, the number of data points, and the time interval between points. It also gives a count of the number of times each of the limits were not met.

```
Reservoir Costs
Storage Cost                R89,378.04
Total Max Volume penalty             R0.00
Total High Volume penalty            R0.00
Total Low Volume penalty         R141,797.15
Total Min Volume penalty             R0.00
                           ----------------
Total Reservoir Costs            R231,175.20
```

The next section of the simulation report details the costs that were incurred for the reservoir. These costs are calculated from the Penalty and Storage Cost functions in the [Reservoir] section of the *System Description File*. The total of the reservoir costs is also shown.

```
Pumping Costs
Power Consumed                  123,290.40 kW
Power Cost                 R158,156.92
Switching Cost              R1,475.00
                           ----------------
Total Pumping Costs              R159,631.92

                           ================
Total System Costs         R390,807.12
                           ================
```

The last section gives the costs incurred by the pumping facility of the system. These costs are calculated from the [Pumping] section of the *System Description File*. If the Existing Demand and Pumping Data option was chosen when running the simulation, there will be no pumping costs in the simulation report, as it is not possible to calculate them. The final line gives the total system costs, which are the sum of the Reservoir and the Pumping Costs.

The main objective of the simulation function is to compare the cost impact of various different operation policies. The next section describes the Optimisation function, that can be used to calculate the optimum operation policy.

4.3. Optimising the System

The main function of the *RSPO* program is to find the optimum operating policy. To do this, you need to either select **Processing|Optimisation** from the menu, or click on the toolbar button. The System Optimisation Dialog will be displayed, as shown below:

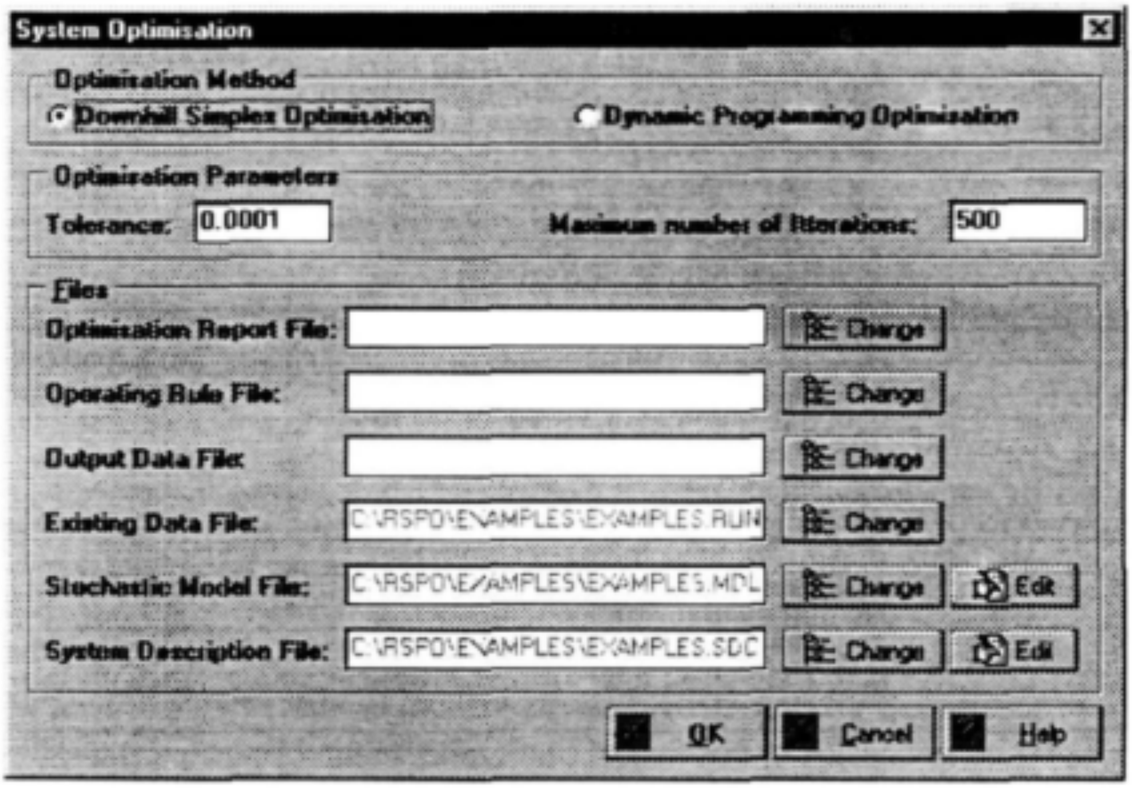

Figure 12 System Optimisation Dialog

The first choice you need to make is between the two optimisation methods, namely Downhill Simplex Optimisation or Dynamic programming Optimisation. The Downhill Simplex method is the method that will calculate the Optimum Operation Policy, and save it to a *Reservoir Operation Rule File* for later use. The Dynamic programming method doesn't find an optimum policy, but rather finds the best possible decision for each individual time step in the data being optimised. This is available as a comparison to the Optimum policy option.

To calculate the optimum policy, select the Downhill Simplex Optimisation method. Next, set the Optimisation Parameters. The tolerance controls the accuracy of the solution found. The smaller the tolerance, the more accurate the solution, but the longer it will take to find. The Maximum iterations controls the maximum length of time that the program will continue trying to find the optimum policy. If you are not sure what values to use for these parameters, just leave them with the default values. If you are going to try a number of optimisations, it is recommended that you set the tolerance to a fairly large value (0.1) to start with. When you have found an approximate solution that you are happy with, set the tolerance back to a small number, and run the optimisation again to find an accurate solution.

You then need to select the files that you want to use in the optimisation. Enter a name for the *Optimisation Report File* by pressing the first CHANGE button. This name will automatically be used for the Operating Rule and Output Data files. These can be changed if you wish. If you select an existing *Reservoir Operating Rule File* it will be overwritten with the new, optimum policy found by this optimisation. The policy in the old *Reservoir Operating Rule File* will also be used as a starting guess for the optimisation procedure.

The biggest problem with this Downhill Simplex Optimisation method, as with most other optimisation methods, is that it can only find a local optimum solution. For this reason, it is a good idea to set up a number of operation policies that are as different as possible from one another, and then run optimisations using each of them as the first guess. If you do this, then you can be pretty sure that the best result will be the best possible result.

Ensure that the Downhill Simplex Optimisation option is selected, that the tolerance is 0.0001, and that the Maximum number of iterations is 500. Change the *Optimisation Report File* name to 'OPTIMUM.OPT' and click on the OK button to start the optimisation. A status box will be displayed in the bottom left of the System Optimisation Dialog to let you know how the optimisation is progressing. When the optimisation completes, the *Optimisation Report File* will be displayed, as discussed below. The optimisation also produces a *System Run File* which contains the same reservoir volume, demand and pump rate data as would be produced by running a simulation with the optimum operation policy.

```
RSPO Optimisation Report
~~~~~~~~~~~~~~~~~~~~~~~~

Example System
Optimised using the Constrained Downhill Simplex Optimisation method.

Generated on 01/03/98 01:02

Demand Data taken from        C:\RSPO\EXAMPLES\EXAMPLES.RUN
Stochastic Model File:        C:\RSPO\EXAMPLES\EXAMPLES.MDL
System Description File:      C:\RSPO\EXAMPLES\EXAMPLES.SDC
Output System Run File:       C:\RSPO\EXAMPLES\OPTIMUM.RUN
Optimised Rule File:          C:\RSPO\EXAMPLES\OPTIMUM.ROR
```

The first section of the report gives the report name, the system name, the optimisation method used, the date and time that the optimisation was run, and the file names of all the files that were used or created by the optimisation

```
Optimisation started at        01/03/98 01:02:47
and completed at       01/03/98 01:02:51
Number of iterations   25
```

The next section gives the duration of the optimisation and the number of iterations required to find the optimum policy.

```
Optimum Change Levels:
Level Value          Plus           Minus
1      78.0421594192    0.3555077029    -0.6945316757
2      37.6128234757    0.2417619458    -0.3447371967
3    -115.4355822447    5.4840238712    -2.5882990277
DontCh   0.0000000000   0.0000000000    0.0000000000
```

The next section gives details of the Optimum Operation Policy. For each change level, the optimum value found is given, with a positive and negative range. The range gives you an idea of the accuracy of the solution. The optimum DontChange value, and its range, is also given.

```
Exceeded Max Volume        0
Exceeded High Volume       1
Below Low Volume           106
Below Min Volume           0

Reservoir Costs
Storage Cost               R140,443.89
Total Max Volume penalty            R0.00
Total High Volume penalty           R0.00
Total Low Volume penalty        R28,417.97
Total Min Volume penalty            R0.00
                           ----------------
Total Reservoir Costs          R168,861.86

Pumping Costs
Power Consumed                  124,289.21 kW
Power Cost                 R159,438.19
Switching Cost                R725.00
                           ----------------
Total Pumping Costs            R160,163.19

                           ================
Total System Costs         R329,025.06
                           ================
```

The last section of the *Optimisation Report* gives a detailed running cost breakdown, as you would see if you ran a simulation using the Optimum Operation Policy. Once you have found the optimum operation policy, you can use it to run your reservoir system in the most economic way possible, as discussed in the next section.

If you select the Dynamic Programming Optimisation Method you will see that the optimisation parameters change to "Number of States". Because of the way that dynamic programming works, the volume in the reservoir must be approximated by a number of integral states rather than a continuous range. The higher the number of states the finer the resolution of the result, but the longer it will take to solve. There is also a restriction on the maximum number of states due to the memory of the computer. The maximum number of states is 8192 divided by the number of time steps in the data that you are optimising. If you enter a number of states that is larger than this maximum, the number of states is automatically reset to the maximum.

For the Dynamic Programming Optimisation Method, you will also have to set up all the files as for the Downhill Simplex Method, except that no *Reservoir Operation Rule File* is needed. When you start the optimisation, the status box will be displayed, showing how much of the optimisation is complete. When it is complete, the *Optimisation Report File* is displayed and the optimum run data is saved in the *System Run File*. The *Optimisation Report File* for Dynamic Programming Optimisation shows the same information as for the Downhill Simplex Optimisation, except that no optimum policy is given.

The Downhill Simplex and Dynamic Programming Methods are two different ways of finding the cheapest way to operate a reservoir system for a given set of demand data and cost information. The Dynamic Programming method will find the best possible way, but does not give any information on how to operate the system in the future. The Downhill Simplex Method can sometimes give results that are not the global best policy, and the type of policy is fixed, but it does give information on how to run the system in the future. Once an optimum policy is found, it can continue to be used until the future demand is no longer statistically similar to the historic demand, or until the system costs change.

4.4. Calculating Pump Settings

Once you have calculated an optimum operation policy using the Downhill Simplex Optimisation option, you can use the new policy to run your reservoir system on a regular basis. At the start of each time step, usually a day, you can calculate at what setting your pumps should operate for that time step as follows. Select **Processing|Pump Setting** from the menus, or click on the toolbar button, and the Calculate Pump Setting Dialog will be displayed, as shown below:

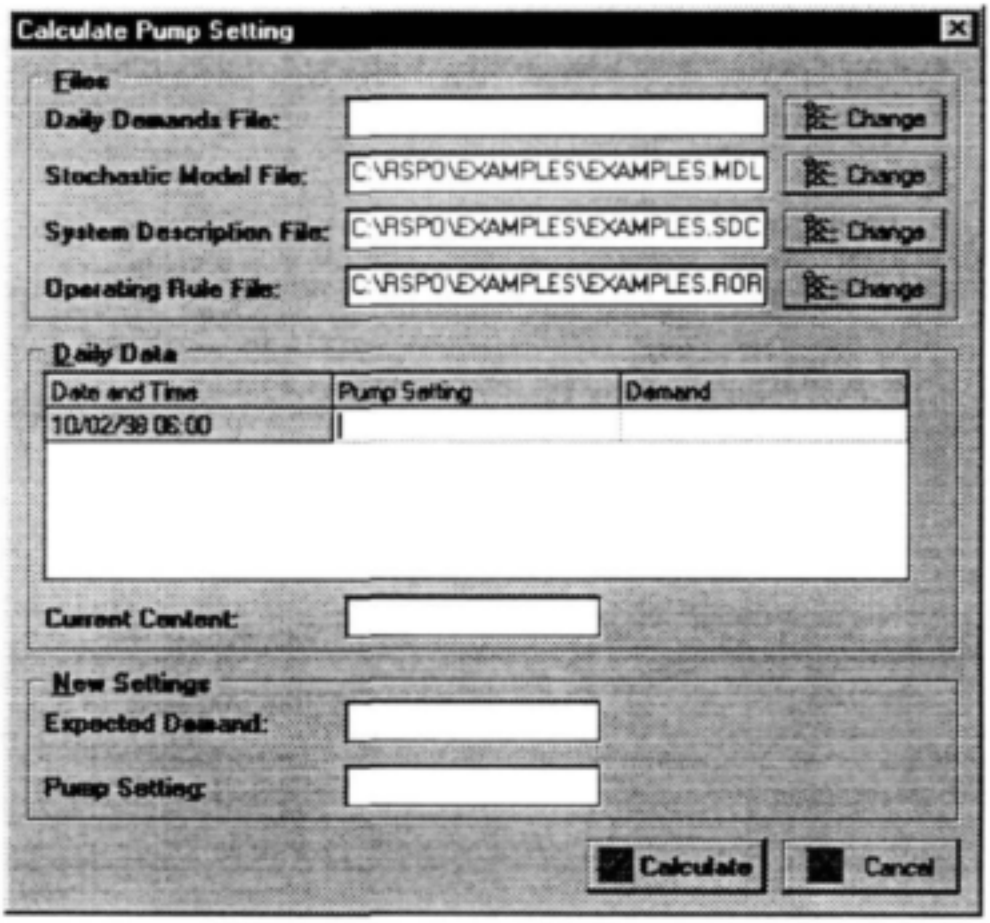

Figure 13 Calculate Pump Setting Dialog

The number of rows in the Daily Data table will depend on the Stochastic Model being used to predict the expected demand. For each row in the table, fill in the pump setting and the demand for the time step that ends at the date and time shown in the left-hand column. Also fill in the current volume in the reservoir in the box labelled Current Content. When these values are correct, click on the CALCULATE button. If this is the first time you have used this option, you will see a Save As dialog, in which you can enter the name of the file which will save all the daily demand information. This file saves information you have already entered, so that you will never need to enter the same information again.

When you have selected a *Daily Demands File*, *RSPO* will calculate the demand expected for the current time step, and the optimum pump setting to meet that demand. You would normally do this at the start each day to calculate the pump setting for that day. If you do this regularly for a number of months, the information built up in the *Daily Demands File* will be useful for future analysis and optimisation.

The preceding sections described, in a step-by-step manner, how to use the *Reservoir System Pumping Optimisation* program. The remaining sections give reference information on all the available commands and all the file types used by *RSPO*.

5. MENU OPTIONS

5.1. File|New

Creates a new reservoir system on your computer. This is done by creating a new sub-directory, which contains a *System Initialisation File*.

Dialog Box Options

Full System Name:

The full name of the reservoir system that you are going to model.

Short Name:

The name that will be used for the system files. It must be not more than eight characters long. *RSPO* will automatically truncate the full name to eight characters, but you can change this, by typing in anything you wish.

System Path:

The path to the new sub-directory. By default this is a sub-directory of the directory which contains the RSPO.EXE file, but can be changed by typing in a new path or clicking the BROWSE button.

Description:

An area where you can enter any text information regarding the reservoir system you wish.

Toolbar button: 

See also: Creating a New System

5.2. File|Open

Displays a standard Open File dialog box that allows you to select an existing *RSPO System Initialisation File* (*.RPO). The system described in the *System Initialisation File* becomes the current system, and the full name of the new system is shown in the title bar of the program window.

Dialog Box Options

File Name

Select or type the name of the document you want to open. This box lists documents with the filename extension selected in the List Files Of Type box. To see a list of files with a particular extension, type an asterisk (*), a period, and the three-character extension, and then press ENTER.

List Files Of Type

Select the type of file you want to see in the File Name list. The types available will depend on the particular function that is being executed. For example, if the **File|Open** menu was selected, then the only file type available will be RSPO System Files (*.RPO).

Drives

Select the drive that contains the file you want to open.

Directories

Select the directory that contains the file you want to open.

Toolbar button:

5.3. File|View

Allows you to view any text file that is less than 32kb in length. A standard Open File dialog is displayed, allowing you to select the file that you would like to view. Multiple selections are allowed. For a description of the Open File dialog box options, see **File|Open**. When you press OK, a document window will be opened within the program window, displaying the contents of the file you selected. The **Edit** and **Window** menu options will also become available. For details of these commands, see the appropriate sections later in this chapter.

Toolbar button:

5.4. File|Graph

Displays graphs of the reservoir volume, demand and pumping rate in a *System Run File*. A standard Open File dialog will be displayed, allowing you to select the *System Run File* (*.RUN) that you would like to view. For a description of the Open File dialog box options, see **File|Open**. When you select a file, the Graphing dialog will be opened with the name of the file that you selected in the title bar.

The scroll bar at the bottom of the dialog allows you to scroll left and right through the data displayed. The Date text box shows the date of the current data point, which is also indicated by the black line on the graphs. The volume, demand and pump rate for the current data point are also shown in the text boxes on the left of the dialog. When you have finished looking at the data, click on OK to close the dialog.

Toolbar button:

See also: **Graphing the Data**

5.5. File|Import|Text Data

Imports data from a standard ASCII, tab delimited, text file. A standard Open File dialog is displayed, allowing you to select the text file from which you want to import data. For a description of the Open File dialog box options, see **File|Open**. When you have selected the file you want to import, the Import Text Data dialog will be displayed. The formatting of the text file is described under **Input Files**.

Dialog Box Options

Importing From

The file name and path from which the data is being imported

Saving File As

The name and path of the file that will contain the imported data. By default, this will be the same as the text file from which the data is being imported, but it can be changed by pressing the CHANGE button.

Status

During the importing of data, the date and time of the current data point will be displayed in the status box.

See also: **Importing Text Data**

5.6. File|Export|System Data

Exports system run data to a standard ASCII, tab delimited, text file. A standard Open File dialog will be displayed, allowing you to select the *System Run File* (*.RUN) that you would like to export. For a description of the Open File dialog box options, see **File|Open**. When you select a file, the Export Data to ASCII dialog will be opened.

Dialog Box Options

Exporting From

The file name and path from which the data is being exported

Export To

The name and path of the file that will contain the exported data. By default, this will be the same as the *System Run File* from which the data is being exported, but it can be changed by pressing the CHANGE button.

Status

During the exporting of data, the date and time of the current data point will be displayed in the status box.

See Also: **Exporting Data**

5.7. File|Exit

Closes the *Reservoir System Pumping Optimisation* program. These is no need to save files before closing as they are always saved when they are created. You can also close *RSPO* by pressing ALT-F4.

Toolbar button:

5.8. File|1 (2, 3, 4)

Use the numbers and names at the bottom of the **File** menu to quickly open one of the four most recently used systems.

See also: **File|Open**

5.9. Edit|Copy Ctrl+C

Copies selected text to the Clipboard. This command is available only if there is at least one document window open.

Text that you copy to the Clipboard replaces the previous contents. If no text is selected, the Clipboard will not be changed.

To quickly copy information to the Clipboard, press CTRL-C.

5.10. Processing|Analysis

Analyses the demand data in a *System Run File* and calculates the secular trend, the periodic trend, the serial correlations and fits a linear autoregressive stochastic model. A standard Open File dialog will be displayed, allowing you to select the *System Run File* (*.RUN) that you would like to analyse. For a description of the Open File dialog box options, see **File|Open**. When you select a file, the Data Analysis dialog will be opened.

Dialog Box Options

Analysis Data

The name of the *System Run File* containing the demand data to be analysed.

Output Report File

The name of the file in which the *Analysis Report* will be saved. The default name is the same as the *System Run File* being analysed, but this can be changed by pressing the CHANGE button. The information in this file is described under **Report Files**.

Stochastic Model File

The name of the file in which the *Stochastic Model* will be saved. The default name is the same as the *System Run File* being analysed, but this can be changed by pressing the CHANGE button. The information in this file is described under **Configuration Files**.

Status

Text indicating the progress of the analysis will be displayed in this box once the analysis has been started.

Verbose Listing

Check this option if you would like the analysis report to contain some of the intermediate calculation results as well as the main results.

Significance Level (%)

The significance level is used in the calculation of the periodic component and in fitting the probability distribution. It must be a number between 0% and 100%, but a setting of less than about 90% is not recommended.

Maximum lag for correlations

Indicates the maximum number of time steps between serial correlations. If this is zero, serial correlations are not calculated.

Model Comment

Enter any other information you would like recorded in the *Stochastic Model File* in this text box. You can save up to 255 characters.

Toolbar button:

See also: **Analysing the system Data**

5.11. Processing|Simulation

Runs a simulation of a reservoir system, and calculates the system running costs.

Dialog Box Options

Simulate using

Select one of the following options to determine what demand and pumping rate data are used in the simulation process:

Existing Demand Data: Reads the demand data from the *Existing Data* file, but calculates the pump rates from the *Operating Rule* file.

Existing Demand and Pumping Data: Reads both the demand and pumping rate data from the *Existing Data* file.

Synthetic Demand Data: Uses the *Stochastic Model* to generate synthetic demand data, and uses the *Operating Rule* file to calculate the pumping rates.

Output Data File

The *System Run File* in which to save the volume, demand and pumping rate results of the simulation. Click on the CHANGE button to select a file name.

Existing Data File

The *System Run File* from which the demand and pumping data will be read, depending on the Simulate using option selected. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button.

Simulation Report File

The name of the file in which the running cost results of the simulation will be saved. The default file name is the same as the *Output Data File*, but this can be changed by clicking on the CHANGE button.

Stochastic Model File

The name of a *Stochastic Model File*. This file is not required if the Simulate using Existing Demand and Pumping data is selected. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button. This file can also be edited by clicking on the EDIT button, but this is not recommended unless you are sure of the effect your changes will have.

System Description File

The name of a *System Description File*. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button. This file can also be edited by clicking on the EDIT button.

Operating Rule File

The name of a *Reservoir Operating Rule File*. This file is not required if the Simulate using Existing Demand and Pumping data is selected. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button. This file can also be edited by clicking on the EDIT button.

Start Date

The date from which to start a simulation. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

Start Time

The time from which to start a simulation. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

End Date

The date at which to end a simulation. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

End Time

The time at which to end a simulation. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

Random Seed

An integer between 0 and 65535 that is used to start the random number generator. This is useful if you would like to run multiple simulations using the same synthetic demand data. To do this, use the same Random Seed for each simulation. If no seed is entered, a seed is calculated from the system clock. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

Secular Trend

Check this option if you would like the synthetic demand data to be calculated using the secular trend from the stochastic model. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

Periodic Trend

Check this option if you would like the synthetic demand data to be calculated using the periodic trend from the stochastic model. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

Autoregressive Component

Check this option if you would like the synthetic demand data to be calculated using the autoregressive component from the stochastic model. This option requires that the Random Component option also be checked. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

Random Component

Check this option if you would like the synthetic demand data to be calculated using the random component from the stochastic model. Unselecting this option will automatically unselect the Autoregressive Component option. This setting is only available if the Simulate using Synthetic Demand Data option is selected.

Toolbar button:

See also: **Running Simulations**, **RSPO Files**

5.12. Processing|Optimisation

Calculates the optimum operating rule, or finds the best set of operation decisions for given set of data, depending on the Optimisation Method selected.

Dialog Box Options

Optimisation Method

Select one of the following options to determine what demand and pumping rate data are used in the simulation process:

Downhill Simplex Optimisation: Calculates the optimum reservoir operating rule for the given system and demand data.

Dynamic Programming Optimisation: Calculates the best series of pump settings for the given system and demand data.

Tolerance

This governs the accuracy of the optimum solution found. The lower the tolerance, the better the accuracy of the solution, but the longer it will take to find a solution. This setting is only available if the Downhill Simplex Optimisation method is selected.

Maximum number of Iterations

This limits the length of time that the program will spend looking for an optimum solution. If an optimisation stops because the maximum number of iterations has been reached, the solution found may not be very accurate. This setting is only available if the Downhill Simplex Optimisation method is selected.

Number of States

The number of integral states into which to approximate the volume of the reservoir. The larger the number of states, the more accurate the solution, but the longer it will take to find a solution. This setting is only available if the Dynamic Programming Optimisation method is selected.

Optimisation Report File

The name of the file in which the results of the Optimisation will be saved. Click on the CHANGE button to change the file name.

Operating Rule File

The name of a *Reservoir Operating Rule File* in which to save the calculated optimum operating rule. If the file selected exists it will be used as a starting guess for the optimisation process, and it will be overwritten with the new operating rule. This file is not required if the Dynamic Programming Optimisation method is selected. The default file name is the same as the *Optimisation Report File* name, but can be changed by clicking on the CHANGE button.

Output Data File

The *System Run File* in which to save the volume, demand and pumping rate results of the optimisation. The default file name is the same as the *Optimisation Report File* name, but can be changed by clicking on the CHANGE button.

Existing Data File

The *System Run File* from which the demand data used in the optimisation will be read. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button.

Stochastic Model File

The name of a *Stochastic Model File*. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button. This file can also be edited by clicking on the EDIT button, but this is not recommended unless you are sure of the effect your changes will have.

System Description File

The name of a *System Description File*. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button. This file can also be edited by clicking on the EDIT button.

Toolbar button:

See also: **Optimising the System**, **RSPO Files**

5.13. Processing|Pump Setting

Calculates the expected demand and the optimum pump setting based on the selected *System Description file*, *Stochastic Model file* and *Operation Rule file*.

Dialog Box Options

Daily Demands File

The name of a *Daily Demands File*. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button.

Stochastic Model File

The name of a *Stochastic Model File*. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button.

System Description File

The name of a *System Description File*. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button.

Operating Rule File

The name of a *reservoir Operation Rule File*. The default file name is the same as the system short name, but can be changed by clicking on the CHANGE button.

Daily Data

This table should contain the demand and pump setting for each time step required for the prediction of the demand. The Time and Date in the left-hand column specifies the end of each time step for which data is required.

Current Content

The volume of the reservoir, in megalitres, at the start of the current time step.

Expected Demand

The predicted demand for the current time step.

Pump Setting

The pump setting calculated from the expected demand and the selected *Reservoir Operation Rule* file.

Toolbar button: 

See also: **Calculating Pump Settings**, **RSPO Files**

5.14. Window|Cascade

Arranges all the open document windows overlapping each other so that the title bars are visible and the current window in on top. This command is available only if there is at least one document window open.

5.15. Window|Tile Horizontally

Arranges all the open documents one above the other so that they can all be seen without overlapping. The current window will be the uppermost window. This command is available only if there is at least one document window open.

5.16. Window|Tile Vertically

Arranges all the open documents next to each other so that they can all be seen without overlapping. The current window will be the leftmost window. This command is available only if there is at least one document window open.

5.17. Window|Close

Closes the current document window. This command is available only if there is at least one document window open.

5.18. Window|Close All

Closes all the open document windows. This command is available only if there is at least one document window open.

5.19. Window|1 (2, 3, ...)

Use the list of open document windows at the bottom of the **Window** menu to select a document which you want to be the current document. A check mark (✓) next to one of the document titles indicates which is the current document. This command is available only if there is at least one document window open.

5.20. Help|About RSPO

Displays information about your copy of *Reservoir System Pumping Optimisation*, including the version number; the copyright, and the authors e-mail address. Unfortunately, a full on-line help system was not implemented for this release of the software, so no other help functions are available.

6. RSPO FILES

The following sections describe all the files used or generated by *RSPO*. Each file type description starts with the name of the file type, and then has the file extension in brackets. This is then followed by a description of what the file contains, where it is used, and any formatting details.

6.1. Input Files

Text Data File: (*.TXT)

This file contains the System Run data in a standard tab delimited ASCII format. The *Text Data File* can have any number of heading rows containing any type of information except for the tab character (ASCII #9). These lines are ignored. The first tab character indicates the first line of the actual data. If the text in the same line as this first tab character is not a date, it is assumed that this is a row of column headings, and it is also ignored. From this point on, until the end of the file, each row must contain a date and time, the demand, the pumping rate and the reservoir volume, in that order, and each separated by a tab character. The demand and pumping rates must be in units of Ml/day, and the reservoir volume must be in units of Ml. The date and time must be the same as the short date and short time strings specified in the Microsoft® Windows™ operating system. The rows of data must be sorted in chronological order, and have no gaps. This is the same format as the exported data that will be discussed under *Run Text Files* late in this section.

6.2. Data Files

System Run Files: (*.RUN)

These are *RSPO* data files containing Date and Time, Reservoir Content, Pump Rate and Demand data for each time period in a specific range. They could be historical data imported from an actual system or data generated by the simulation or optimisation functions of the program.

Daily Demands Files: (*.DDD)

These files contain Date and Time, Pump Setting and Demand for each time step that the Calculate Pump Setting dialog was used to calculate the pump setting.

6.3. Configuration Files

System Initialisation File: (*.RPO)

This is the main file that describes each reservoir system on your computer. It contains information such as Full Name and Description of the system.

System Description File: (*.SDC)

This is the file that contains the technical description of the reservoir and pumping system. It contains information about the size of the reservoir, the limits on the reservoir contents, the available pump settings and their capabilities, as well as all the cost functions of the system. For a full description, see the chapter on **Running Simulations**.

Stochastic Model Files: (*.MDL)

These are text files giving details of the stochastic model that is used to generated synthetic demand data and to predict the next time steps demand. The calculations used to create this file, and that use these results are described in detail in Water Research Commission report No 757/1/98.

Reservoir Operation Rule Files: (*.ROR)

These file describe the policy that is used to decide what the new pump setting should be at the start of each time step. A detailed description of this file type is given in the chapter on **Running Simulations**.

6.4. Report Files

Analysis Report: (*.ANR)

The *Analysis Report* gives the results of all the calculations performed during the analysis of system data. A full description of these results is given in the chapter on **Analysing the System Data**.

Simulation Reports: (*.SMR)

These reports give the results of running a simulation of a reservoir system. They include information on the period of the simulation, the running costs of the simulation and a count of any constraint violations. They also indicate what simulation options were chosen and the input and output files used. A typical *Simulation Report* is shown in the section **Running Simulations**.

Optimisation Reports: (*.OPT)

There are two type of *Optimisation Reports*, depending on the optimisation method chosen. Both types are shown in the section **Optimising the System**.

If the Downhill Simplex method was chosen, the *Optimisation Report* gives details of the files used in the optimisation and the resultant optimum operating policy. The accuracy of the optimum policy is also indicated.

If the Dynamic Programming method was chosen, the *Optimisation Report* also gives information on the files used for the optimisation, and it details the running costs in the same way as a *Simulation Report*.

6.5. Exported Text Files

Run Text Files: (*.RTX)

These are ASCII text files containing tab delimited data, exported from one of the *System Run Files*. Their format is the same as that discussed for the *Text Data File*.