Genetic Algorithms for Water Resource Assessments

P.G. van Rooyen[1], S.J. van Vuuren[2]

[1]WRP Consulting Engineers (Pty) Ltd, P.G. van Rooyen
[2]University of Pretoria, S.J van Vuuren

Abstract

The operating costs (in particular pumping costs) of the various water resource systems in South Africa amounts to hundreds of millions of Rands per year and much effort is invested by water resource planners to implement cost saving operating rules that do not jeopardize the long term assurance of supply. The current method of optimisation is through the application of scenario analysis as part of an iterative approach to determine the most feasible rule. A conceptual description of the application of Genetic Algorithms for operating rule optimization for the Water Resources Yield Model (WRYM), which is one of the main water resource analysis models applied in South Africa, are provided. The paper presents information from a Water Resource Commission funded project as described in the reports WRC Report No. 1388/1/05 and WRC Report No. 1144/1/01.

Keywords: *Optimisation, water resource operating rules, system analysis, genetic algorithm.*

1 Introduction

1.1 Background

The location of minable minerals in South Africa has historically stimulated urban and industrial developments in areas with limited water resources. Over the years the bulk water supply systems were expanded extensively to store and convey water from water rich river system to meet the growing water requirement of urban and industrial development. Furthermore, the need for energy to support the urbanized population's activities, both electrical and liquid fuel, resulted in the development of coal fired power stations and coal-to-liquid fuel conversion plants which are situated near the coal fields in areas where the water resources has already been fully utilized.

This urbanisation pattern has resulted in the development of bulk water supply infrastructure of which a significant portion involve transferring water from low lying areas through pumping conduits requiring substantial amounts of electrical energy. The associated operating costs (pumping costs) amounts of hundreds of millions of Rands per year with the result that significant efforts are invested by water resource planners to ensure cost saving rules and measures are implemented.

The method of optimisation that is currently carried out is through a process of scenario simulation, where different rules are analysed and evaluated using network simulation models. This process requires a high level of human intervention in defining the scenarios and processing the results in an iterative cycle. Typically the selection of the scenario of rules is bases on the experience of the annalist and the knowledge of the particular system being optimised. Although this process has provided acceptable results in the past, the question that always lingers is whether or not a more optimum-operating rule exists that has not been considered by the annalist.

In the literature survey that was carried out as part of Phase 1 of the WRC Project, Report Number 114/1/01 (**van Vuuren, S.J. et.al., 2001**), it was shown that Genetic Algorithms can be applied for the optimisation of the water resource system operating rules. All the examples found in the literature, however, is based on water resource modelling techniques and analysis procedures that are different from those used in South Africa. The research challenge therefore lied in developing a conceptual genetic algorithm that is applicable to the currently applied analysis technology and uses the existing water resources models as point of departure.

1.2 Overview of Genetic Algorithms

An algorithm is any procedure that takes in data and modifies it according to a step-by-step set of instructions. Every program ever written is an example of an algorithm. Genetic Algorithms are programs that stimulate the logic of Darwinian selection, if one understands how populations accumulate, differences over time due to the environmental conditions acting as a selective breeding mechanism, then you understand GAs. Put another way, understanding a GA means understanding the simple, iterative process that underpin evolutionary change. The issue, of course, is how best to get that "selection pressure" translated into a program procedure and applied to your problem.

Holland, 1975, the pioneer of Genetic Algorithms defines the evolutionary nature of the algorithm as follows:

- Start with an initial population that is randomly generated, but contains the variability parameters characteristic of the population.

- The fitness of each individual in the population is assessed according to a fitness function.

- The probability of each individual to survive is proportional to its fitness.

- The individuals of the next generation are selected on probability and through a genetic transformation process of crossover and mutation, ensuring that the solution is not localized within the solution environment.

GAs are suited to solve problems that are not vulnerable to attack by brute force methods because the sheer number of potential solutions defies the possibility of testing them all. Such problems are typically multi-constrained, that is the solution must be a balance of conflicting or synergistic properties. When considering a problem with multiple dependencies you are normally forced to admit the possibility of isomeric solutions i.e. solutions that give the same result using different processing routes. So for some problems there is no such thing as the "best solution" but instead you are looking for members of a fuzzy set of solutions that can be defined as "good enough".

2 Development framework

The first step in the process of developing the conceptual GA model was to define a framework that could serve as a guide for the formulation of the conceptual model. This framework was defined by considering the current modeling environment and modeling systems applied in South Africa as well as taking pointers and suggestions from the literature.

2.1 Current modeling environment

Introducing a new procedure, the Genetic Algorithm (GA), into the water resource analysis environment, where a large pool of knowledge is founded on present technologies, requires careful consideration. The aim is to implement the GA so that a high degree of the current knowledge and proven methodologies are maintained, while also allowing for flexibility for accommodating future analysis requirements.

To this end, relevant guidelines are listed below:

- The GA should be developed for the existing water resource models that are used in practice for the management of most of the countries water resources. These models are the Water Resources Yield Model (WRYM) and the Water Resources Planning Model (WRPM). This will ensure the existing pool of knowledge (methods and expertise) can benefit by adding optimisation functionality to the current models that are already known and applied.

- Cognisance should be taken of the analysis techniques that are used in practice in South Africa. Currently an important result that quantifies the supply capability of a water resource system is the Firm Yield that can be abstracted. Two types of Firm Yield values are usually determined; (1) the Historical Firm Yield, which is based on analyses of the historical sequence and (2) the Stochastic Firm Yield, which also defines the reliability at which the yield is available (**DWAF, 1987**). The aim with the optimisation of an operating rule for a system is, in most cases, to maximise the firm yield in order to "stretch" the resource as far as possible. Minimising the operating costs, in particular pumping costs, is normally a secondary objective and usually opposes the objective of maximising the firm yield.

- In order to build confidence in the application of GA technology it is deemed essential to introduce (develop and verify) the GA technology in phases. The intention is to prove the value of a GA by using basic yield analysis and relatively simple systems as a first step. The GA can then thereafter be extended to include probabilistic projection analyse of complex systems, involving changing system characteristics where demands are growing, new reservoirs and transfer conduits are implemented over a planning horizon.

- It is proposed that the GA first be developed to only optimise the operating rules of a system, and as a further phase, be extended to include the optimisation of the capacities (sizes) of water resource infrastructure such as conveyance capacity of water conduits and the storage capacity of reservoirs.

- Although the time it will take the GA to compute a solution is an important factor to consider, the focus should initially be on ensuring the viability of the technique and the validity of the results. Once the integrity of the GA has been established and proven, resource can be employed to reduce the computational requirements.

- Within the context of the above guideline, the applied water resource analysis techniques allow for different types of analysis which required different computational requirements. For example, historical analysis, where only one hydrological sequence is analysed is computationally significantly less intensive compared to stochastic analyses where up to a thousand sequences are analyses in some cases. The conceptual design of the GA should allow for both types of analysis techniques.

- The basic design of the GA should be to maintain the flexibility and generic structure that is provided by the network models. This will involve making the definition of the GA as generic as possible so that it can be applied to most of the water resource systems and defined through model input data files.

- The requirement of the National Water Act (Act 36 of 1998) to implement Licensing will change the emphasis of water resource analysis in South Africa from yield determining assessments of large systems to system analysis dealing with multiple users that are dispersed over the catchments of the river systems. Water resource analysis for Licensing will broadly focus on two aspects, (a) accounting for the interaction users have on one another in the system and (b) determining how the reconciliation of water balances will be achieved in a fair and equitable manner. The GA design will have to take these requirements into account.

2.2 Pointers from the literature

Information from past research, as presented in the literature, indicates that the formulation of a GA consists of four main elements, (a) the encoding scheme, (b) the GA operators, (c) the process (water resource system) that is optimised, and (d) the objective function as translated into the GA fitness function. These processes are described briefly in the following sections.

2.2.1 Encoding scheme and initialisation process

Encoding involves the process where the "genes" of the GA is converted into the input parameters of the process to be optimised. In our case the parameters are the variables defining the operating rules and the process to optimise, in the water resource system. An important component of this encoding process is to ensure that the constraints or limitations on the parameter are adhered to when assigning their values.

From the literature it has been noted that some form of random process is used to generate the initial set of genes (**Goldberg,1989; Olivera, et.al , 1997 and Ndiritu, 2003**). Thereafter the genes are converted into parameter values and the feasibility of the values verified. If a parameter set is infeasible, that gene is rejected and the process is repeated until a full set of initial members (also called the Initial Population) is produced. The constraints and limitations that are applicable to the operating rule parameters of the WRYM are discussed in subsequent sections**.**

From the literature it was pointed out that there are advantages to define an encoding scheme that groups decision variables together, as apposed to dealing with each parameter individually (**Olivera, et.al,1997**). In doing this, one may also be able to capture some of the constraints and limitation presented above. A further important requirement specified in the literature is that the selection of the initial population should be such that the parameter values must be uniformly distributed over the search space (**Olivera, et.al,1997**).

The components making up the genetic algorithm are briefly described in the next sections.

2.2.2 Fitness function

The fitness function consists of two components (a) the objective function representing the yield, reliability and pumping costs, and (b) the fitness scheme which involves the calculation of a numerical value for each gene that reflect how good a gene has performed relative to the others.
The objective function can include multi objectives and typically a weighting factor is used to represent the importance of different objectives.
The fitness values of each gene are used to select which genes become parents for the determination of the next generation of children genes.

2.2.3 Pairing

Pairing is the process where pairs of genes are selected based their fitness (the genes with high fitness values have a higher probability of being selected) for crossover. This selection process is also referred to as "Tournament Selection" in some literature sources.

2.2.4 Crossover

Working with real-values chromosomes, as apposed to binary chromosomes, was defined in the literature to be a more clear-cut way for the design of the GA. Given that real-value chromosomes will be use in the proposed development, two means of crossover were suggested by (**Olivera, et.al. 1997**), namely, uniform and quadratic. Uniform crossover basically involves the copying of parameter values from one of the parents to a child. The selection of which parents parameter values is copied in based on a binary random variable. For the quadratic crossover a weighting system is applied where a new values is calculate for the child chromosomes based on both the parent's parameter values. The weights are such that the parent with the best fitness has the largest influence (contributes the most) on the child's parameter value. As was the case with the initial population, the feasibility of the children's chromosomes or genes has to be tested. Adjustments are then made to the parameters of the infeasible cases where needed.

2.2.5 Mutation
This is the process where, usually a small number, of child chromosomes are changed "arbitrarily" in order to create some maverick chromosome to assess the fitness of a diverse set of cases in the search space. Mutated chromosomes also have to be tested for feasibility or the method of mutation should be such that the feasibility is preserved.

3 Defining water resource system operating rules

In order to formulate the conceptual GA model it is necessary to first identify and briefly describe the variables (model input data) that define the operating rule of water resource systems. The methods of affecting the operating rules differ in modeling system and the focus here is on the variables applicable in the Water Resources Yield Model (WRYM). It is not practical to provide an in depth description of WRYM in this paper and the reader is referred to the reports *Water Resources Yield Model: Procedural Manual* (**DWAF, 2005**) and Water Resources Yield Model: User Guide (**DWAF, 1999**) for details on the models functionality, input data definitions and simulation capabilities.

There are four basic variables applied in the WRYM defining the operating rule of a water resource system, namely:
- Number of operating zones in a reservoir.
- Bottom elevation level of reservoir zones.
- Penalty value of the water in reservoir zones.
- Channel penalties.

These are described in the following sections.

3.1 Reservoir operating zones
One of the mechanisms of defining the operating rules in the WRYM is to divide the reservoir into horizontal zones with each zone representing a portion of the storage volume of the reservoir. In the WRYM each zone's dimension is defined by the bottom elevation of the zone and the volume of water is deduced from the elevation-volume relationship. It is proposed that each reservoir be allowed to have a fixed number of zones (say ten zones). The number of zones does not influence the behavior of the system's operation directly and therefore would not serve as a meaningful GA optimisation input variable.

3.2 Bottom elevation level of reservoir zones
The difference between the bottom elevations of two zones, that are "stacked" on top of each other, defines (indirectly through the elevation-volume characteristics) the volume of storage that is contained in a zone. It is important to note that, with the WRYM, different zone levels can be defined for each of the 12 months of the year. This flexibility allows for the design of operating rules that accommodate seasonal variability. However, in most of the systems that are modelled in South Africa, constant zone levels are used throughout the year.

3.3 Penalty value of the water in reservoir zones
WRYM currently provides for the definition of one penalty value per reservoir zone for all the months of the year. This implies that the sequence of drawdown among the zones in different reservoirs is the same throughout a particular scenario analysis. It would be possible to change the input data structure of WRYM to accommodate different zone penalties for the 12 months of the year. However, it is proposed to only consider such an enhancement once the GA has been proven viable with the current structure.

There are certain constraints that have to be adhered to when assigning values to the zone penalties. In most cases, where the "Rule Curve" (see WRYM Users Guide for definition) is set to be equal to the full supply level of the reservoir, the penalty values that are assigned to the zones must be in ascending order, following the stacked sequence of the zones from top to bottom.

It is important to note that the penalty values assigned to each reservoir zone has two functions. Firstly it represents the "unit constraint" of drawing water from the reservoir zone and secondly the, same penalty value serve as a "unit gain" of storing water in the particular zone. This second functionality is to enable transferring water from one reservoir to another. The model also has a second mechanism for effecting transfer between reservoirs by means of a two arc channel, as explained in the next section.

3.4 Channel penalties
A channel is the modelling element that represents a conduit (links) between reservoirs or nodes. The penalty structure of channels can be divided into two basic types and the distinction between the types lies in the number and direction of the arcs that makes up a channel.

For this discussion, **Type (a)** is a channel that is only defined by one arc and the (flow) direction of the arc is the same as that of the channel. In this case the penalty value represents a "unit constraint" that has to be overcome before water will flows in the channel. For flow to occur in the channel water must be draw or pull through at a "unit gain" that is greater than the "unit constraint" of the channel.

Type (b) channel consists of two arcs, one arc in the same direction as the channel (also referred to the forward arc) and the second arc, in the opposite direction of the channel or backwards arc. The operation of a **Type (b)** channel is to "pull" water through the channel at a selected penalty value. The "pull" action is achieved by the arc that is in the opposite direction of the channel direction and that the minimum flow constrained placed on the forward arc is equal to required transfer rate.

In summary the function of a **Type (a)** channel in a water resource system is to serve as a conduit through which water can be drawn by means of either a reservoir zone penalty value (unit gain) or a **Type (b)** channel, in other words a **Type (a)** channel cannot "pull" water through itself. However, a **Type (b)** channel is capable to drawing water through itself with a "pulling force" that is equal to the penalty value that is assigned to arc that is in the opposite direction of the channel.

During the operating rule optimisation process the channel penalty values are changed in relation to the penalties assigned to the reservoir zones to implement a particular operating rule. In general the type of channel (**Type (a)** or **Type (b)**) would remain the same for all the scenarios (test cases) with the only change being the penalty values of the channels.

4 GA definition

The processes that make up the GA and their interactions with one another are best illustrated by the flow diagram shown in **Figure 1**. The indicated processes and the linkages between them were compiled using similar algorithmic flow diagrams and GA descriptions found in the literature.

4.1 Step 1: Initial gene pool generation

The first process of the GA involves generating initial sets of genes that defines the operating rule values for the intended water resource system elements using a random selection method that results in parameter values that are uniformly distributed over the search space. In order to be able to duplicate results across different computer systems it is required to implement a "software" random number generator (similar to the subroutines that are currently applied in the WRYM) as apposed to a computer dependant random number generator. These random number generators make use of seed values that will be changed during the evaluation and testing phase to illustrate that convergence is achievable with different initial gene populations. Once it is shown that convergence is achieved, the seed values will be fixed in the production version of the model.

In order to only allow penalty values that are not in conflict with the penalties of the water resource elements that are excluded from the optimisation process, it is necessary to limit the possible selection of penalty values for optimisation within a pre-defined range. This range of penalty values will be fixed for a particular system configuration, however, the range has to be defined in relation to the penalty values assigned to the water resource elements that are not optimised. Both the penalty value search range and the penalties for the water resource elements that are not optimised have to be defined by the model user.

4.2 Step 2: Test the feasibility of the gene sets

This procedure evaluates each gene set to identify if they represent feasible penalty definitions by complying with the following constraints:

- Penalty values should be in ascending order from the top zone to the bottom zone for each reservoir.
- Penalty values for **Type (b)** channels have to be such that the penalty for the reverse arc must always be greater than the penalty for the forward arc. (This relative order of the penalty values is used by the model to define the directions of the arcs.)

4.3 Step 3: Calculate the fitness of the gene sets

This step involves three basic processes, first each gene set is converted into penalty values that defines an operating rule scenario. The water resource system is then analysed iteratively, with the operating rule scenario in place, to determining the firm yield of the system. (The iterations referred to here are different runs with the same inflow sequence, historical or stochastic, where the target abstraction of the yield channel is changed between the runs to find the Firm Yield value.) Other simulations results such as the flow through pumping channels will be extracted from the simulation output (this will be for the run where the firm yield is imposed as the target draft) and summarised in a form required by the fitness function. Additional simulation result, such as the volume of spillage from the system, may also be used as a component of the fitness function. The model will be developed in such a way that the user can select which channel flows should be included in the fitness function and allowances will be made to assign weights to each of the selected channel flows and the firm yield result.

4.4 Step 4: Tournament selection

This procedure uses the result of the fitness function of a particular gene set, relative to the sum of the fitness function results of all the gene sets, to select which of the parent gene sets are to form pairs for crossover. The selection procedure is carried out by a roulette-wheel parent-selection scheme (**Goldberg, 1989**) which assures a higher probability of selecting the parents that has the best fitness function. The operation of the roulette-wheel procedure requires the generation of random numbers, and for the same reasons as presented in **Step 1**, a software random number generator will be used for this purpose.

Olivera et.al, 1997, performed tests on two selection schemes, window and ranking (**Davis, 1991**), and found that the ranking scheme was preferred. It is therefore proposed to also apply the ranking scheme in the GA model. In most of the GA applications presented in the literature an additional feature is applied in the selection of the child genes call elitism. This involves using the genes of the best solution, unchanged, as one of the child genes. This ensures that the best solution of one generation is kept unchanged to compete with the next new generation of children genes.

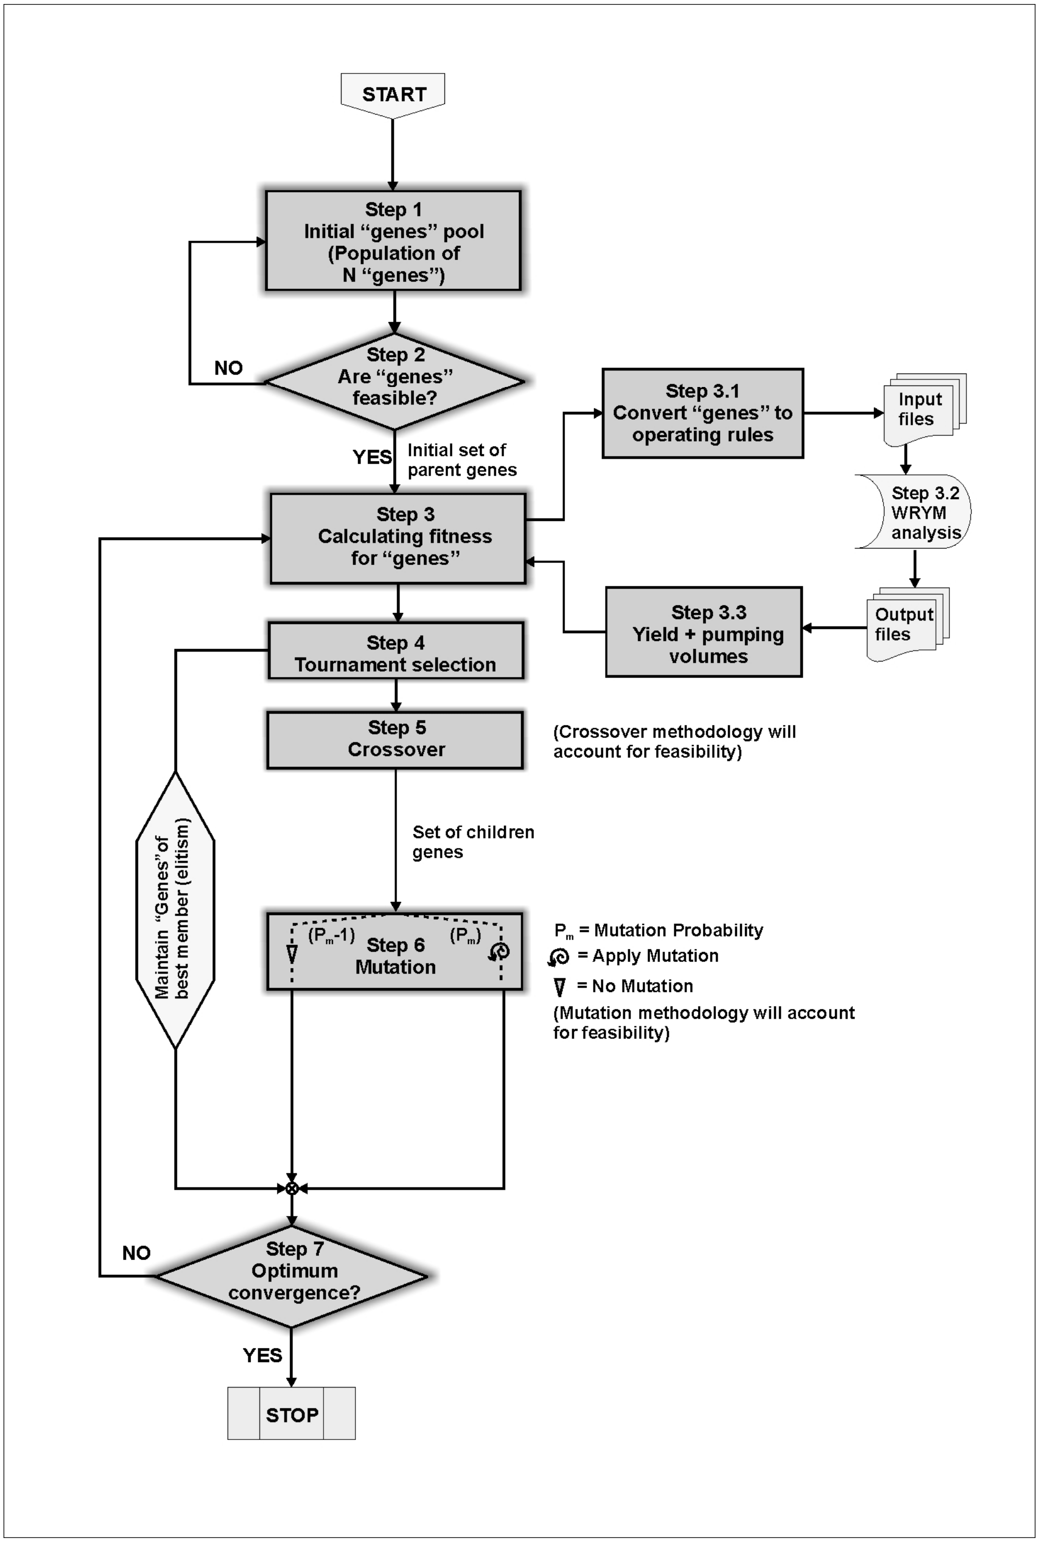

Figure 1: Flow diagram of Genetic Algorithm

4.5 Step 5: Crossover operation

Crossover is the GA operator that creates a new generation of gene sets, two genes from each parent pair. In the classic strain representation of a gene (strain is a string of zeros and ones), crossover involves exchanging (swapping) parts of the two parent genes to form the new children genes. Given that real value chromosomes will be used for the GA the procedure would be somewhat different. In the context of the selection of the reservoir zone and channel penalty values, it is proposed to use either uniform or quadratic methods of crossover to produce the new generation as suggested by **Oliveira, et.al, 1997**.

In order to maintain the constraints the suggestion of **Olivera, et.al, 1997** and others are that parameters should be grouped rather than working with each parameter individually. In the context of the WRYM this would mean that the penalty values for the zones in a reservoir and for the arcs in a channel would respectively be treated as groups. Another group would be the zone levels of each reservoir. The uniform crossover method would then select one of the groups of parameters of a parent as the parameters for the child. The selection process to pick which parent's parameter group is transferred to the child will be based on the results of a random generator that produces either a one or a zero value. The selection is then as follows, if the random generator's result is one, the parameter group from the first of the two parents is used and vice verse if the random generator's value is zero. The generation and selection process is repeated for all the parameter groups that are part of the optimisation process.

4.6 Step 6: Mutation

Mutation is the process where maverick genes are generated to ensure the search space is explored randomly to avoid the optimisation process only achieving local optimum solutions. From the literature it is show that the percentage genes of each generation on which mutation should be applied ranges from 0.05 to 0.20.

In order to maintain the feasibility of the operating rules after mutation it will be required to only allow mutated parameter values, penalties and zone levels, to change within defined limits. A zone level and a zone penalty value could only be changed to a value the lies between the values of the zones above and below. The value will be calculated by using a uniform distribution selection method.

4.7 Step 7: Optimum convergence

The final step is to check whether an optimum answer has been found. This will be based on a criteria that will indicate by how the rate of improvement from the one iteration to the next has changed and if the rate is large the process will loop back to **Step 3**, otherwise the iterative search will stop.

5 Proposed analysis procedure

Within the context of the existing structure of the WRYM, **Figure 2** gives a schematic presentation of how the GA can be implemented.

As indicated, the Genetic Algorithm Subroutine will interact with the Yield Search Subroutine by imposing the operating rule parameters onto the system and it will receive back the yield results and relevant pumping volumes. The purpose of the Yield Search Subroutine is to control the Network Simulation Subroutine through an iterative search process to find the yield, historical firm yield or stochastic yield at a given reliability. The Yield Search routine will output relevant simulation results to the existing output files (*SUM.OUT and *PLT.OUT). The step-by-step results from the GA will present output to a new file, summarising the process and behaviour of the GA for evaluation purposes.

What is apparent from the discussions presented above is that the elements (channels and reservoir zones) that should form part of the optimisation process are dependant on the characteristics of the particular water resource system and have to be defined as an input to the GA algorithm. For example, the dummy dams (a modeling dam representing a cluster of farm dams) would not form part of the optimisation process and should therefore be excluded from the list of elements to be optimized. The data structure of the GA should therefore allow the system annalist to select the water resource elements for optimization. This selection based on their understanding of the constraints and flexibilities that exist in the water resource system. The definition of the optimisation definition will be achieved through a separate input file that will contain the parameter for the GA Subroutine.

6 Conclusions and recommendations

From the literature it was identified that genetic algorithms have been used successfully in optimising the operation of water resource systems. The conceptual GA for the WRYM, described in the paper, indicate that the operating mechanisms of the WRYM are suitable for the application of GA optimisation.

It is recommended that the research be taken further to develop and test a genetic algorithm for use with the WRYM.

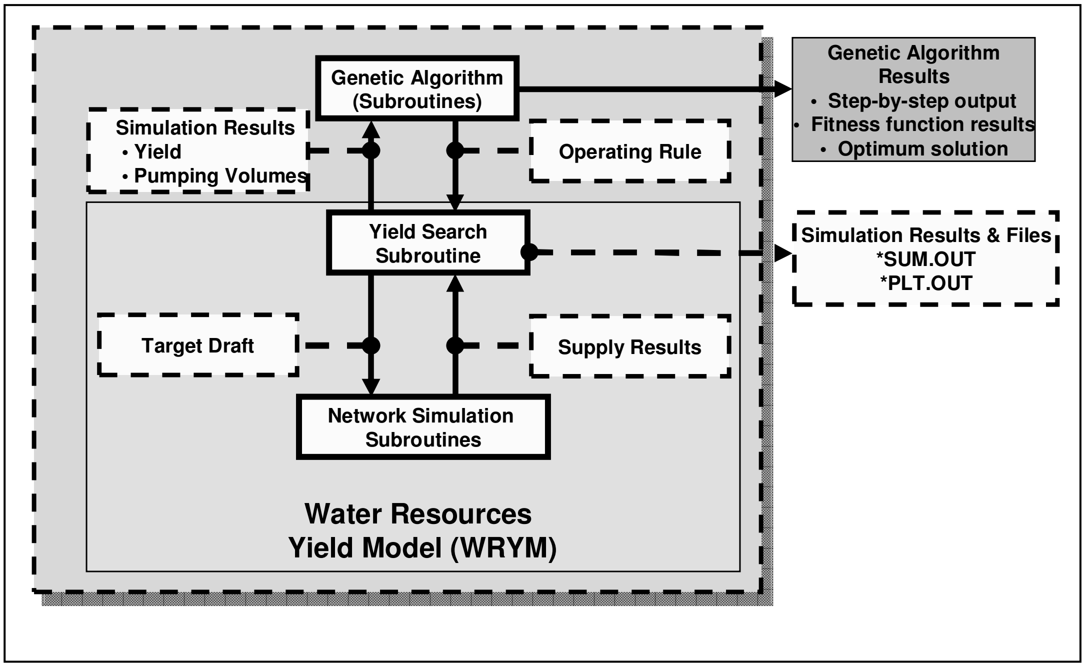

Figure 2: Schematic representation of the GA linked to the WRYM

7 Acknowledgements

The research presented in this paper was undertaken as a project funded by the Water Research Commission, whose support is greatly acknowledged.

References

Davis, L. 1991. Handbook of genetic algorithms, Van Nostrand Reinhold, New York.

DWAF 1987. Department of Water Affairs and Forestry. Report No. PC000/00/6986. Yield analysis – terms and procedures. Compiled by BKS Water Resources Associates as part of the Vaal River System Analysis, Pretoria, SA.

DWAF. 1999. Water Resources Yield Model ~User Guide – Release 4.1.1. Report No. P000/001/98 compiled by BKS (Pty) for the Department of Water of Water Affairs and Forestry, Pretoria, SA.

DWAF. 2005. Water Resource Yield Model. Procedural Manual. Internal report of the Department of Water Affairs and Forestry, compiled by WRP Consulting Engineers (Pty) Ltd.

Goldberg, D.E. 1989, Genetic Algorithm in Search, Optimization, and Machine Learning, 2nd ed., Addison-Wesley, Reading. Mass.

Holland, J.H. 1975. Adaptation in Natural and Artificial Systems. University of Michigan Press. 2nd Edition 1992.

Ndiritu, J.G. 2003. Reservoir system optimisation using penalty approach and multi-population genetic algorithm., Water SA, Vol. 29, 273-280, July 2003.

Olivera, R. and Loucks, D.P. 1997. Operating rules for multi-reservoir systems, Water Resources Research, Vol: 33, 839-852.

Van Vuuren, S.J., van Rooyen, P.G., Fouchè, W.L., Haarhof, J. 2001. WRC Report No. 1144/1/01, Water Research Commission, Pretoria, South Africa.

Van Vuuren, S.J., van Rooyen, P.G., van Zyl, J.E., van Dijk, M. 2001. WRC Report No. 1388/1/05, Water Research Commission, Pretoria, South Africa.