Reservoir system optimisation using a penalty approach and a multi-population genetic algorithm

JG Ndiritu
School of Civil and Environmental Engineering, University of the Witwatersrand, Private Bag 3, WITS 2050, South Africa

Abstract

A multi-population genetic algorithm (GA) was used to optimise a system of two reservoirs that supplies monthly varying demands and environmental flow requirements. Optimisation aimed at minimising the penalty resulting from non-supply of water and the occurrence of low reservoir storage states that would limit non-consumptive utilisation of water in the reservoirs. Four cases were analysed viz. Case I: Reservoir capacities and demands were fixed and the operating rules were optimised; Case II: Demands were fixed and the reservoir capacities and operating rules were optimised; Case III: Reservoir capacities were fixed and the demands and operating rules were optimised; and Case IV: Reservoir capacities, demands and operating rules were optimised. The genetic algorithm obtained reasonable solutions for all cases. A detailed analysis of Case IV obtained several high-performance solutions of varied sizes and supply capabilities. This analysis revealed specific limitations of supply reliability and the expected storage states of one of the reservoirs. The analysis also obtained the ranges within which the optimal monthly operating rules for the system are expected.

Keywords: genetic algorithm, penalties, operating rules, demands, reservoir capacities

Introduction

The reservoir system optimisation problem is often quite complex. At the planning and design stage, decisions on system configuration and component sizes have to be made. Once this has been done, the operating rules that will maximise system performance need to be formulated. Ideally, the two problems should be dealt with together although system yield and sizing analysis is often taken as separate from the system operation problem. The comprehensive analysis of reservoir management and operations models by Yeh (1985) and Loucks et al. (1981) and others provides ample evidence of this. Yield analysis methods that incorporate operation scheduling are, however, in practical use. For example, the water resource yield model that is widely used in South Africa applies the storage state balancing space rule approach (Basson et al., 1994). The space rule operating policy ensures that all reservoirs in the system are drawn down together in a manner that allows all the water stored in the reservoirs and all water entering the system to be available to meet the target draft without a shortfall until all the reservoirs have failed simultaneously. The reservoirs are thus operated in a manner that utilises the overall storage capacity to the maximum and prevents any unnecessary water spillages. Johnson et al., (1991) have used this approach in reservoir operation analysis. The associated water resource planning model also widely used in South Africa refines these rules using an iterative network flow programming approach (Basson et al., 1994). This approach defines the system as a network configuration of arcs and nodes with each arc having a lower and an upper flow bound and also an associated cost per unit of flow. The optimisation aims at minimising the total cost of flow in the network and the operating rules are derived from the optimal solution. Hsu and Cheng (2002) have used the network flow approach to optimise a water resource system in Taiwan. This paper is aimed at demonstrating the application of the genetic algorithm (GA) method to a practical reservoir system optimisation problem including capacity and yield analysis and also system operation. The hydrology is based on a real system but hypothetical demands and environmental flow requirements are applied. The GA is a good candidate for reservoir system optimisation as it possesses some unique advantages over many classical optimisation methods. These are discussed in the next section. Although the GA has been researched and applied fairly extensively, it is only recently that applications to reservoir operation have been reported (Oliviera and Loucks, 1997, Wardlaw and Sharif, 1999, Sharif and Wardlaw, 2000). As Van Vuuren (2002) indicates, the potential of the GA has not been fully utilised in the South African water industry.

A water resource system could be designed to maximise yield or to minimise the penalty caused by non-supply of water and/or non-utilisation of reservoirs due to low storage levels. The application of the GA to maximise system yield, subject to probability constraints of supply and reservoir storage states for the same system is reported elsewhere (Ndiritu, 2002) while the penalty approach is applied here. Penalties should ideally be obtained from socio-environmental-economic analysis of the costs of non-supply and non-availability of water in reservoirs. As Basson et al. (1994) indicate, data availability in general is extremely limited while many intangible social and political factors also often come into play. Relative penalties agreed upon by the water management body and all the stakeholders could be a reasonable alternative. The penalties applied here are hypothetical but are not considered unrealistic.

The hydrology of the system analysed is based on the Elands River catchment in South Africa up to Mkombo Dam which supplies domestic and industrial needs. The other dam in the catchment, Rust de Winter, supplies irrigation water.

The genetic algorithm

Because detailed descriptions of the basics of the GA are widely available, only a brief explanation of how the GA works is given

☎+2711 7177134; fax: +2711 3391762; e-mail: JNdiritu@civil.wits.ac.za
Received 12 December 2002; accepted in revised form 23 April 2003.

here. The GA is a population-based optimisation method based on the principle of survival of the fittest. An initial population of possible solutions to the problem being optimised is generated randomly in a coded form. The representation of the population as a code enables the exchange of the valuable 'genetic' material within individual members. To ensure that this happens, the selection of the members to create the next generation is carried out in a manner that favours the better performing ones. Crossover, the exchange of parts of the codes of a specific proportion of the selected parents to obtain the 'children', then follows. The proportion of the selected parents to involve in crossover is the probability of crossover. The generated children are then subjected to mutation – the change of a small proportion (probability of mutation) of the genetic code to encourage search into regions beyond the parent population. The processes of selection, crossover and mutation (termed as a generation) are repeated until the set termination criterion is met. The termination criterion could be the convergence to an optimum (indicated by no further improvement with successive generations) or a set maximum number of generations or system simulations. Numerous modifications to the basic GA have been made to suit specific applications and to improve GA performance.

The GA code applied here was written by the author in Fortran 90. It used binary coding , tournament selection and multiple point crossover. A description of this GA follows. The value of each decision variable *i* say v_i was represented as a randomly generated binary substring of a specified bit length (*l*). The actual value v_i was computed by linear interpolation between the lower bound Vl_i and the upper bound Vu_i of the search range. For example, if a code 0110 was generated, then $v_i = Vli+[Vu_i-Vli]\ [0(2^3) + 1(2^2) + 1(2^1) + 0(2^0)]/(2^4 - 1) = Vli + [Vu_i - Vli][6/15]$. A point in the *n*-dimensional search space consisted of the values of the *n* decision variables and was represented as a string (chromosome) composed of the binary substrings of the individual variables. The performance of each chromosome was obtained by first decoding the chromosome to obtain the variables v_i and using these to simulate the problem and thereby obtain the objective function. Tournament selection was implemented as follows: Two sets of chromosomes were randomly selected from the population with each consisting of a specified number of chromosomes called the tournament size. The best performing chromosome from each set was obtained. The process was repeated $c \times p$ times where *c* is the probability of crossover and *p* the population size. Each selected pair was used to generate a 'child' of the new population in the crossover step. To implement multi-point crossover, locations corresponding to variable substring boundaries of the chromosomes were generated randomly. The bit material (genetic code) at these locations was exchanged to obtain a child. Each child replaced a randomly chosen member of the initial population. Mutation was implemented by changing the bit values (replacing 0 with 1 and 1 with 0) of randomly selected positions and chromosomes. The elitist (best performing) member of the parent population was maintained as a member of the new population.

The GA does not require linearisation like the linear programming approaches. It also does not suffer the 'curse of dimensionality' of dynamic programming. The GA uses the actual objective function values and not gradients of the response surface as gradient search methods do. It can therefore easily handle rough and discontinuous response surfaces. Many GA users also consider it an advantage that the GA searches with a population of solutions and not just from a single point like local search methods. The traditional GA, however, often fails to locate global optima where these are known and methods to improve its robustness are often applied. The GA used in this study applied independent sub-population searches followed by shuffling and then repeated independent sub-population searches as a measure against local optimum traps. This approach had been developed in previous work (Ndiritu and Daniell, 1999; 2001). The use of binary coding discretises the search space and search is consequently limited to the resulting multi-dimensional grid. This leads to the failure of the GA to fine-tune to the optimum which invariably lies outside this grid. To deal with this, the GA is sometimes combined with a local search method to fine-tune after an initial GA search. The GA used here applied a gradual search range reduction if successive generations located the decision variable within a small portion of the search range. To maintain robustness of search with range reduction, a search range shifting (hill climbing) routine was included. To use fine-tuning and range shifting, two search ranges were set for each decision variable; an initial range that may be varied as the optimisation proceeded, and an ultimate search range that was fixed. Ndiritu and Daniell (1999, 2001) give detailed accounts of these improvements. Figure 1 is a flow chart of the improved multi-population GA. The GA is termed as a multi-population GA because it uses several sub-populations and not a single one. Though the GA is easily modifiable, some modifications, like the ones made here, increase the number of optimisation parameters of the GA itself that have to be suitably selected. The common trial-and-error approach was used to obtain these and the following parameter values were adopted: A crossover probability of 1.0, a mutation probability of 0.05, a tournament size equal to half the sub-population size, the number of crossover positions equal to the number of decision variables, and 12 subpopulations with 48 members each. The range reduction and shifting procedures were carried out after every 5 generations. The maximum number of generations for all sub-populations was set at 500 while the maximum number of system simulations was set at 100 000. These limits were chosen subjectively. All the optimisations, however, converged before these limits were reached. Termination through convergence was applied when the average of the best objective function values of the second last 5 generations exceeded 99% of the average of the best objective function values of the last 5 generations.

System simulation and penalty computation

The Elands River system up to Mkombo Dam is located to the north-east of Pretoria, South Africa. The upper dam, Rust de Winter, has a catchment area of 1 145 km^2 and a mean annual runoff (MAR) of 19.8 mm. The incremental area to Mkombo is 2 578 km^2 and the MAR from this area is 3.9 mm. The MAR from the whole catchment to Mkombo is 32.8 Mm^3/a. The live storage capacity of Rust de Winter and Mkombo Dams is 26.9 and 204.6 Mm^3 respectively. The Department of Water Affairs and Forestry (DWAF) provided 77 years of simulated monthly runoffs and point rainfalls at the two sites. DWAF also provided the monthly average Symon's pan evaporations and area-volume data for the reservoir sites. Second-order polynomials were used to model the area-capacity relationships. The monthly distributions of demand from the two reservoirs were obtained from a previous study (DWAF, 1989).

Figure 2 shows the system configuration and how the penalties were computed. The water supply and environmental penalties were factored in direct proportion to the volume not supplied and a factor (in boldface) which indicates the relative cost of non-supply of a unit volume of water. The penalties were chosen subjectively but reflect the approach applied in the South African water resource yield model and the water resource planning model.

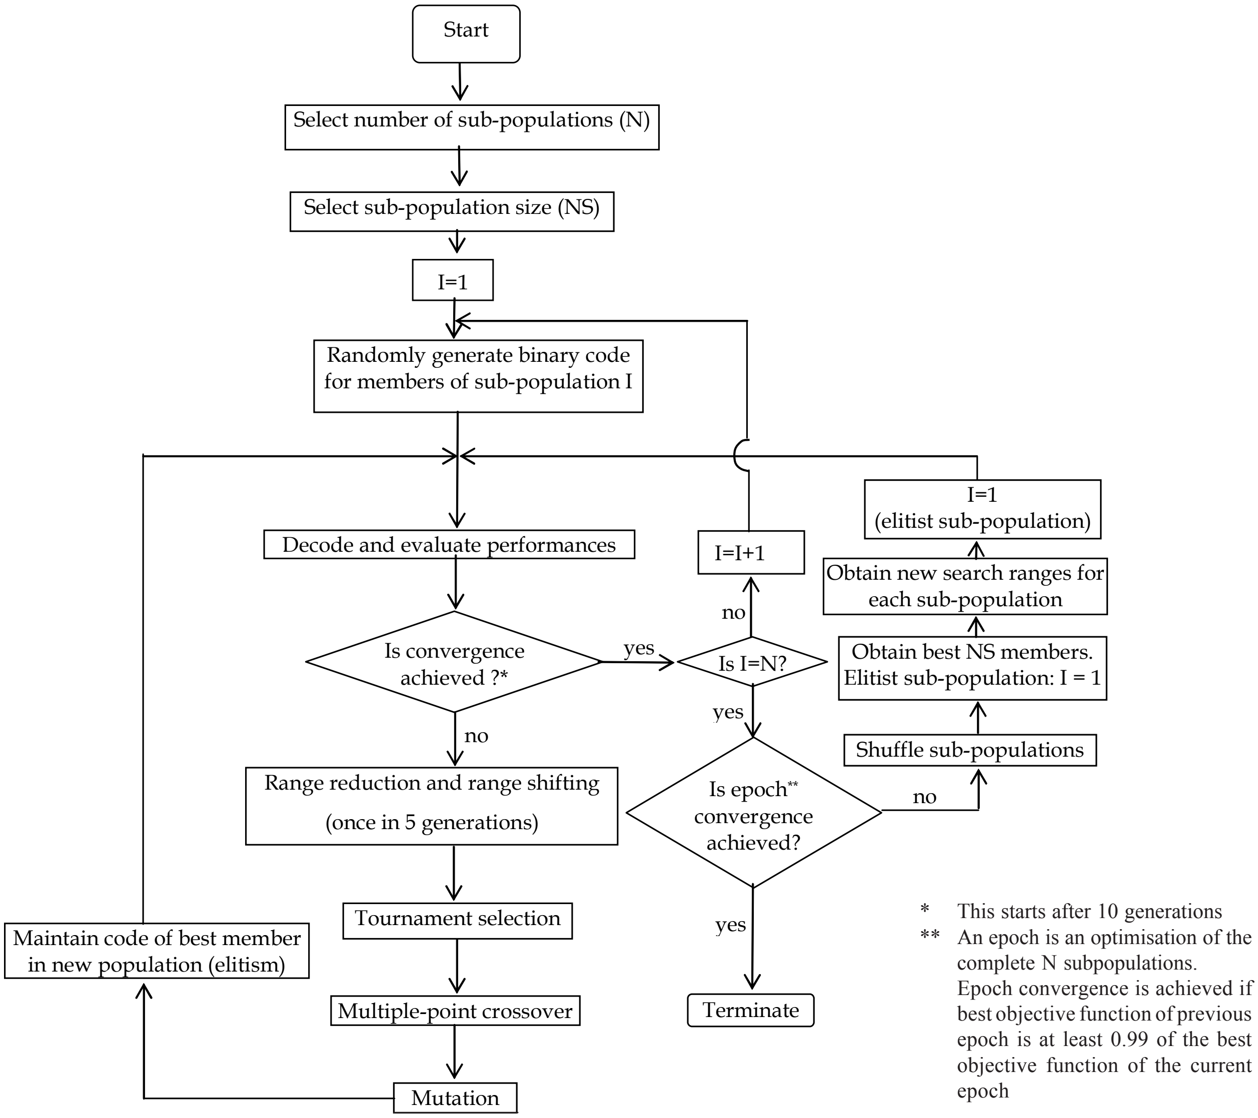

Figure 1
Flow chart of multi-population genetic algorithm

Rust de Winter Dam was taken to supply irrigation (as in the real system) and lower penalties were assigned to this demand than Mkombo Dam demands which are domestic and industrial. More severe supply restrictions were also allowed for Rust de Winter Dam. The storage state penalties were factored in proportion to live storage capacity. The penalties as given in Fig. 2 were then scaled against the maximum possible penalties – zero supply, zero environmental flow and no water in reservoir live storage. The water supply and environmental penalties were grouped together since both have the same dimensions while the storage state penalties were scaled separately. The two scaled penalties were then combined and weighted to give a single overall penalty that took values only in the range 0 to 1. The water supply and environmental flow penalties were subjectively weighted at twice the storage state penalties. The objective function then took the form:
maximise 1-overall scaled penalty.

System simulation was carried out on a monthly basis using mass balance assuming the reservoirs were initially half full. The simulation included reservoir operating rules allowing four levels of supply restrictions, thus giving three rule curves. These were allowed to vary on a monthly basis. Defining these curves therefore required 72 values to be obtained from optimisation. The following four scenarios that represent possible problems that a water resource analyst may encounter, were investigated.

Case I The reservoir capacities and demands were fixed and the operating rules that maximise system performance were optimised.

Case II The demands were fixed and the capacities and operating rules that maximise system performance were optimised.

Case III The capacities were fixed and the operating rules and demands that give the best performance were optimised.

Case IV The capacities, demands and the operating rules that maximise system performance were optimised.

The number of parameters to optimise were therefore 72, 74, 74 and 76 for Cases I, II, III and IV respectively. For Cases I and II, the demands were set to high values of 90% of the firm yield obtained in a previous study of the system (DWAF, 1989). The search ranges for the capacities were 0 to 50 and 0 to 250 Mm^3 for Rust de Winter and Mkombo Dams respectively for Cases II and IV. The search

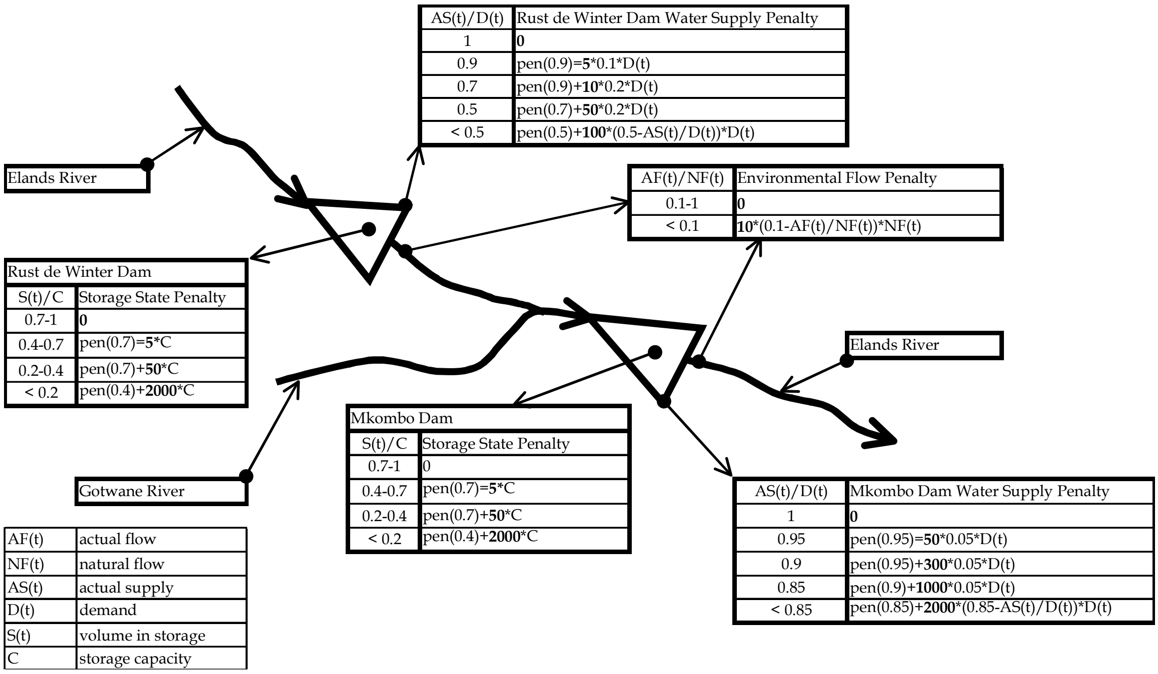

Figure 2
System configuration and penalty computation

TABLE 1
Optimised capacity, drafts and supply

Parameter	Case I	Case II	Case III	Case IV
Capacity Rd W (Mm^3)	26.943*	22.056	26.943*	4.366
Capacity Mkombo (M^3)	204.592*	195.784	204.592*	211.039
Target draft RdW (Mm^3/a)	13.600*	13.600*	0.502	0.491
Target draft Mkombo (Mm^3/a)	8.000*	8.000*	7.376	7.962
Actual supply RdW (Mm^3/a)	10.390	10.075	0.500	0.485
Actual supply Mkombo (Mm^3/a)	7.453	7.504	7.185	7.802
% of supply/demand (RdW)	76.39	74.08	99.56	98.84
% of supply/demand (Mkombo)	93.16	93.79	97.41	97.99
Objective function	0.982	0.985	0.991	0.996
Number of function evaluations	12672	12672	12672	12672

* these parameters were set to the specified values - were not optimised.

ranges for the demands for Cases II and IV were 0 to 13.6 and 0 to 8.0 Mm^3/a for Rust de Winter and Mkombo Dams respectively. The upper limits are the 90% firm yield values from DWAF (1989). The fixed reservoir capacities for Cases I and III were set to the active storages of the existing reservoirs. The 72 operating rule curve values could vary from 0.01 to 0.99 with the constraint that the width of each storage restriction zone was at least 0.1 of the reservoir capacity. This constraint was implemented by factoring the objective function value by 0.5 for any violation. The environmental flow requirements were assumed to be 10% of the monthly natural flow. None of the 4 cases attempted to model the real system as it exists especially with regards to demands and non-consumptive water utilisation. All the cases were therefore theoretical but were considered realistic enough to indicate how well the GA would handle real reservoir system optimisation problems.

Results and discussion

Tables 1 and 2 present the optimised results for the four cases and Fig. 3 the corresponding monthly operating rule curves. A single optimisation took about 6 min on a 1 000 MHZ Pentium IV processor. As expected, Case IV obtained the highest objective function value because more parameters were available to optimise. Fixing the demand (Case II) imposed a more severe constraint than fixing the capacity (Case III). Case III therefore gave a higher objective function than Case II. The exceedance probabilities (Table 2) indicated that there was a difficulty meeting the Mkombo supply at restrictions lower than 10% even for Cases III and IV where the target draft was optimised. This could be attributed to high evaporation losses due to the large surface area and/or the environmental flows which were not subject to optimisation. Cases I and II revealed high exceedance probabilities of storage states

TABLE 2
Exceedance probabilities of optimised solutions

Supply/ demand	Case I	Case II	Case III	Case IV
Rust de Winter				
1	0.592	0.6916	0.0162	0.0617
0.9	0.592	0.6916	0.0162	0.0617
0.7	0.4307	0.4394	0	0.0141
0.5	0.2197	0.2489	0	0.0054
Mkombo				
1	0.8074	0.7716	0.4459	0.369
0.95	0.8074	0.7716	0.4459	0.369
0.9	0.4524	0.4188	0.0271	0.0162
0.85	0.0736	0.026	0	0
Storage state	**Case I**	**Case II**	**Case III**	**Case IV**
Rust de Winter				
0.7	0.618	0.6342	0.0152	0.0509
0.4	0.4535	0.4675	0	0.0162
0.2	0.29	0.3139	0	0.0097
Mkombo				
0.7	0.882	0.8604	0.4448	0.395
0.4	0.5703	0.5249	0.053	0.0335
0.2	0.2392	0.2197	0	0
Environmental flows				
Rust de Winter	0.1101	0.1156	0.0794	0.0477
Mkombo	0.2434	0.1168	0.2732	0.1442

lower than 0.2 for both reservoirs.

In light of the uncertainties in practical water resources management, the provision of alternative solutions of satisfactory performance is likely to be more worthwhile than the single solution that gives the best objective function value. The ability of the GA to provide these has been observed and highlighted in water supply and system distribution applications but not prominently in water resource systems analysis. The system analysed here is relatively simple, it therefore does not test or demonstrate this ability thoroughly but is nonetheless useful. Case IV was selected for this purpose.

The average objective function value of the best solutions from the 24 subpopulations of case IV (2 epochs of 12 subpopulations each) was 0.992. The standard deviation of the objective function values was 0.003 indicating that the performance of all the 24 solutions was practically identical. The relationships among capacities and drafts from the 24 solutions are illustrated in Fig. 4. Curve fitting equations are not included as they would not serve any specific purpose. Many alternative combinations of capacities and demands (target drafts) can be obtained from this figure. Figure 5 presents the exceedance probabilities of given supply levels and storage states obtained from the 24 solutions. The inability of Mkombo Dam to supply 95% or 100% of the demand at a reliability greater than 0.8 is seen in Fig. 5b. Figure 5d shows the difficulty of Mkombo Dam to maintain a storage state greater than 70% for more than 25% of the time. The average rule curves with bands of ±1 standard deviation for the 24 solutions of Case IV are given in Fig. 6. These constitute the probable range of the optimal operating rules for the system. Figure 6 reveals the absence of notable seasonal variations of the operating rules. This implies that operating rules based on the seasonal instead of monthly time interval may be adequate. This would simplify the optimisation problem considerably. The rule curves for the third level of supply restriction in Figs. 3 and 5 seem unrealistically low for actual reservoir operation. This problem could be avoided by setting the lower search limit of the rule curve to a reasonably higher value than 0.01 as used here. By providing larger weights to the storage state penalties in relation to the weights for water supply and environmental flow, the unrealistically low rule curves could be avoided.

It is worthwhile remembering that although the analysis here used real hydrological data from the Elands River catchment, the demands, the environmental flow requirements and the penalty structure were hypothetical. The analysis is therefore not an indicator of the performance of the existing system.

Conclusions

A multi-population genetic algorithm (GA) has been applied to optimise a system of two reservoirs that supplies monthly varying demands and environmental flows. Optimisation aimed at minimising the penalty resulting from non-supply of water and the occurrence of low reservoir storage states that would inhibit non-consumptive water utilisation. The problem used real hydrological data from the Elands River catchment in South Africa but the demands, the environmental flow requirements and penalty structure were hypothetical. The GA obtained reasonable least-penalty solutions for the four cases analysed. A more detailed analysis of one of the cases demonstrated the ability of the GA to provide several high-performance solutions of varied sizes and supply capabilities. The expected range of monthly operating rule curves for optimal system performance was obtained. The analysis also revealed the inability of one of the reservoirs to supply demands at low restriction levels and high reliability, and to maintain high storage states. This is an indicator that GA optimisation could be used to systematically identify the critical areas of multi-reservoir system performance.

The system applied is relatively simple and so did not incorporate a search of alternative system configurations. This aspect could be tested using a more complex system that could also apply stochastically generated sequences in place of the single one applied here. Analysis of complex systems with stochastically generated data could be expensive computationally. Improving the efficiency and effectiveness of the GA, and the use of other population-based optimisation techniques (e.g. SCE-UA, Duan et al., 1992) are thus reasonable areas of further research.

Acknowledgements

The author acknowledges the provision of hydrological and other data by the Department of Water Affairs and Forestry (DWAF), Hydrology Section.

References

BASSON MS, ALLEN RB, PEGRAM GGS and VAN ROOYEN JA (1994) *Probabilistic Management of Water Resource and Hydropower Systems*. Water Resources Publications, Colorado, 424 pp.

DEPARTMENT OF WATER AFFAIRS AND FORESTRY (DWAF) (1989) Elands River Water Resources Investigation.

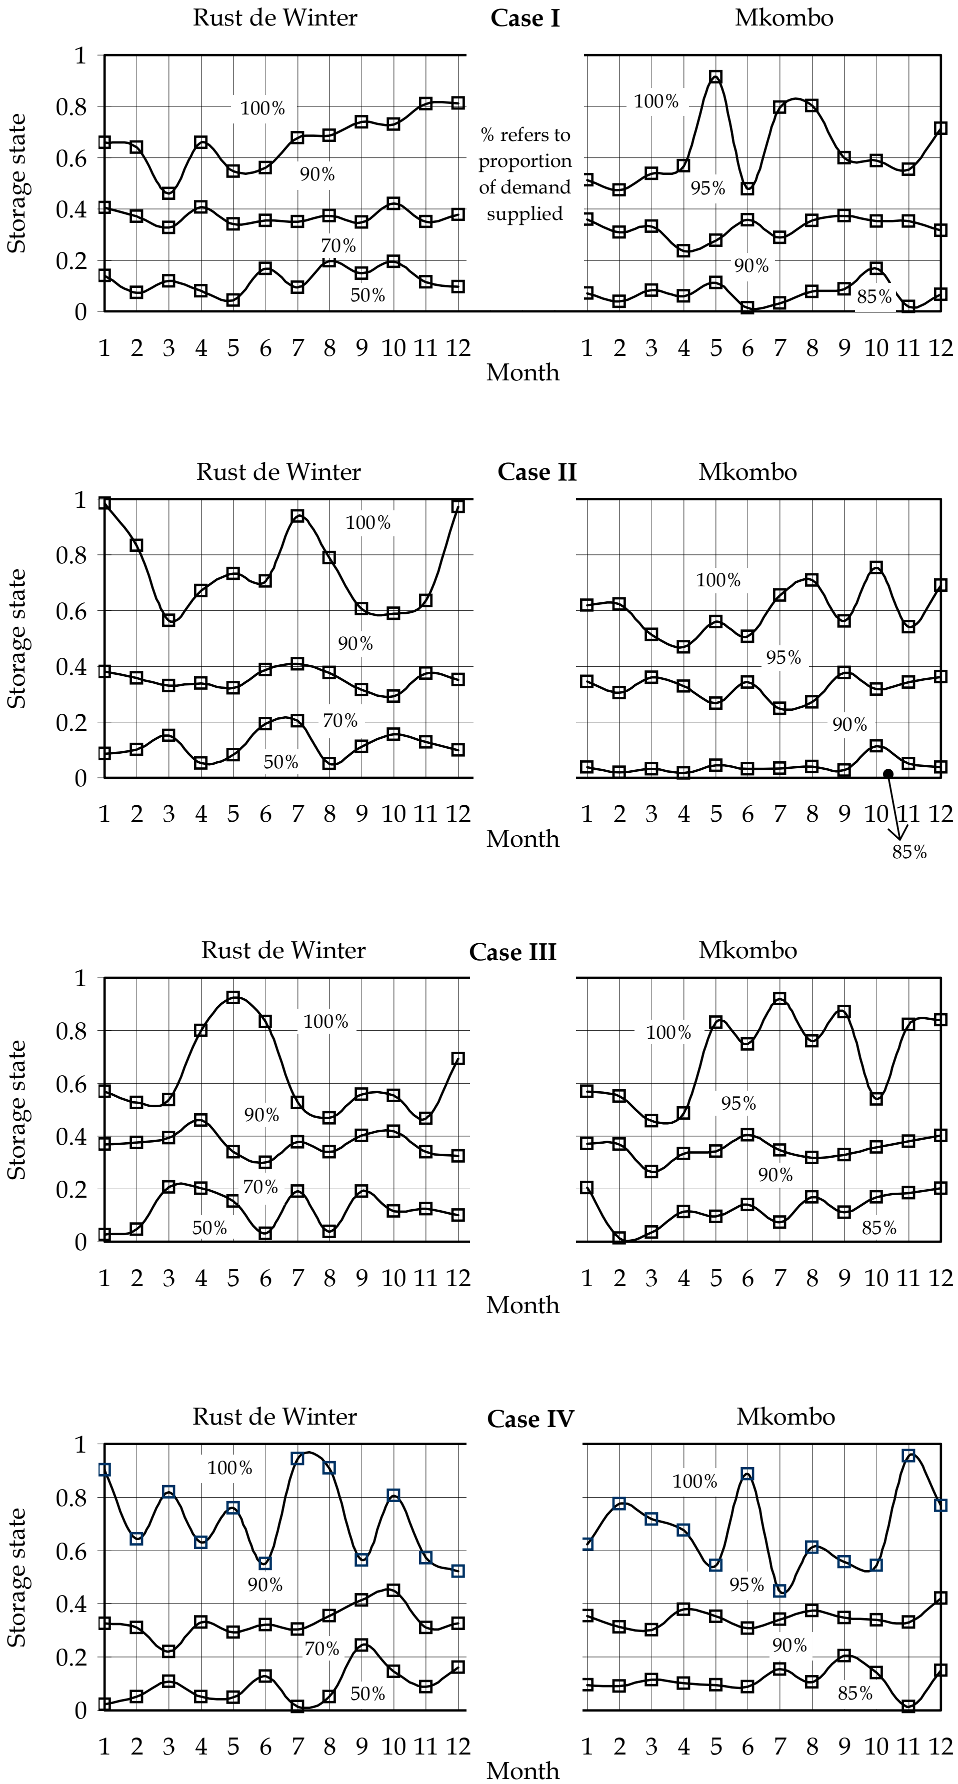

Figure 3
Operating rule curves of optimised solutions

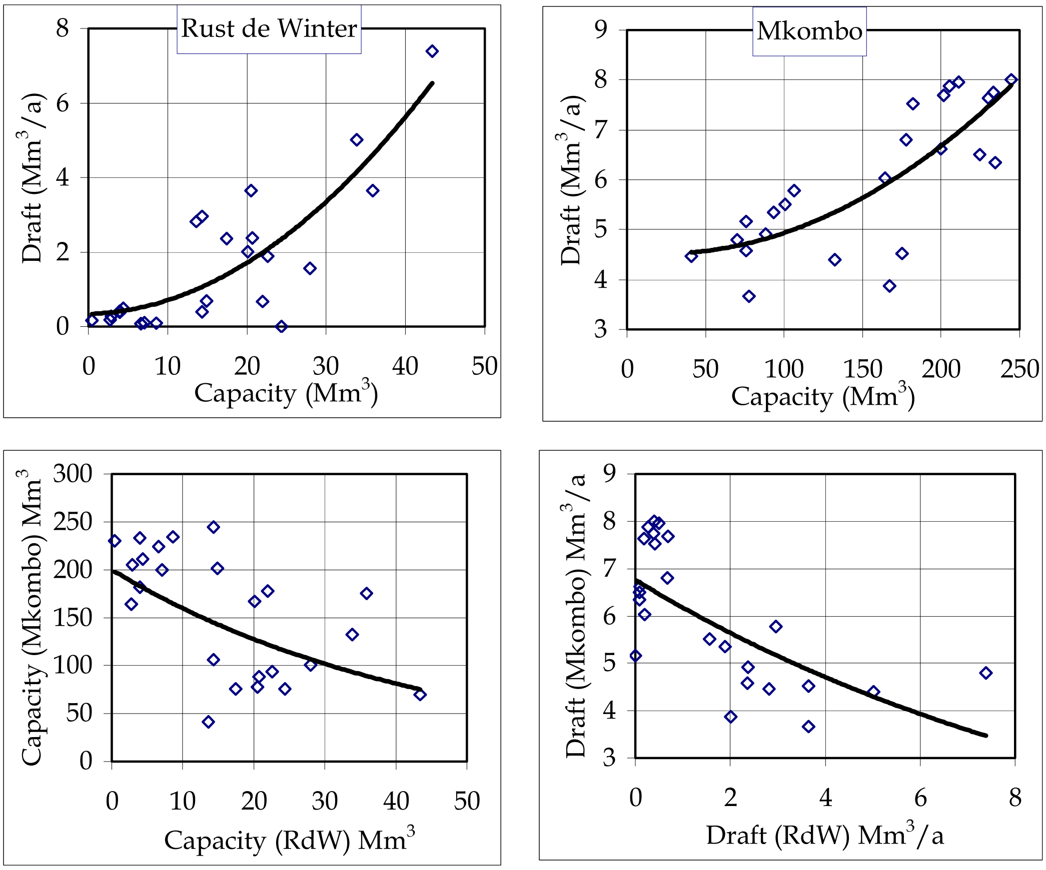

Figure 4
Relationships of optimised demands and capacities for Case IV

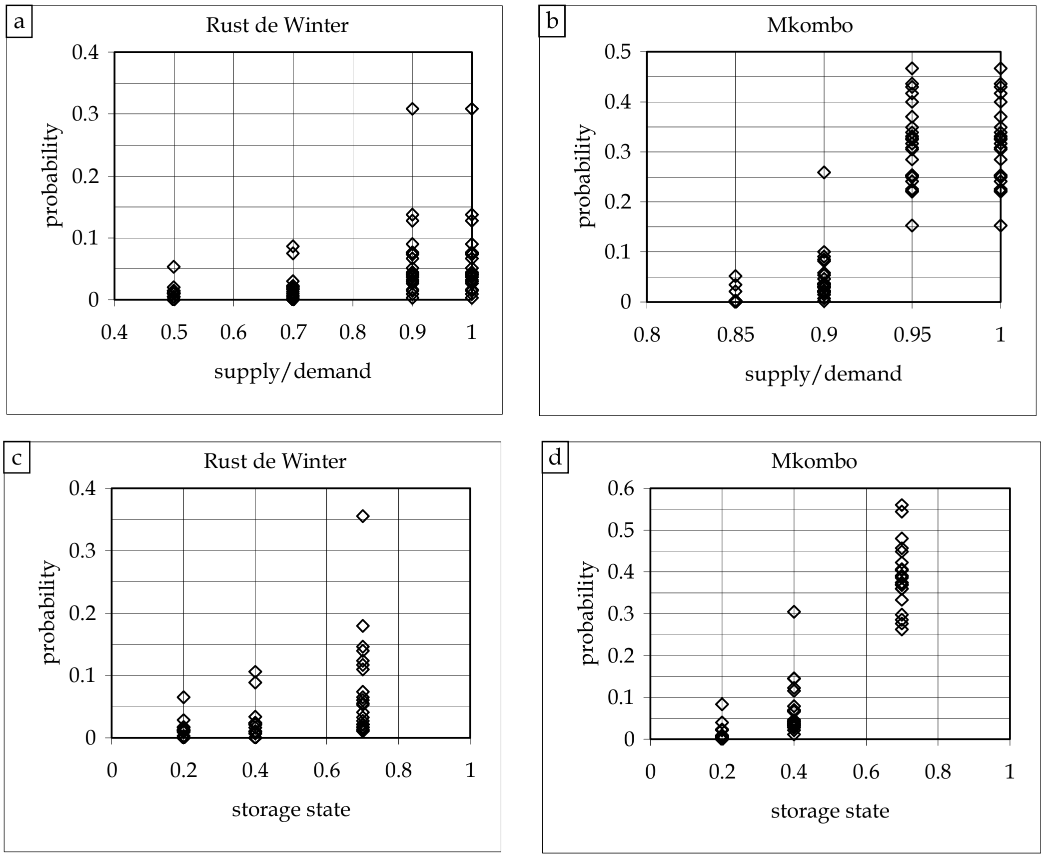

Figure 5
Exceedance probabilities of optimised solutions for Case IV

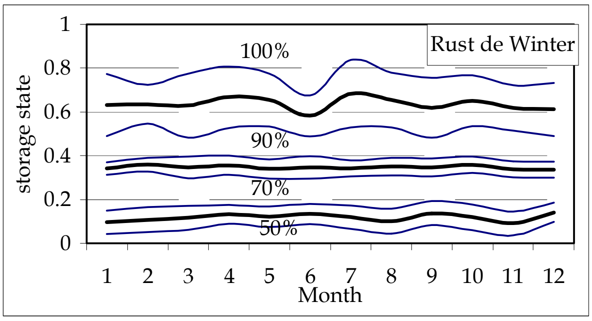

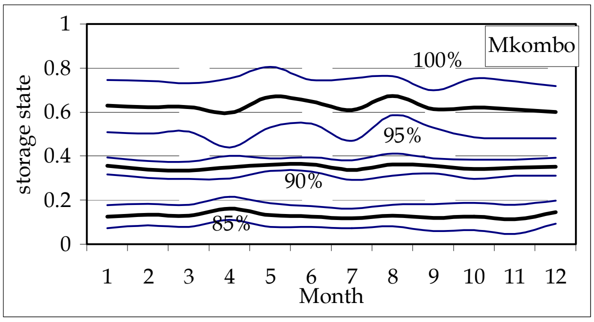

Figure 6
Average operating rule curves with ±1 standard deviation bands for Case IV

DUAN Q, SOROOSHIAN S and GUPTA V (1992) Effective and efficient global optimization for conceptual rainfall-runoff models. *Water Resour. Res.* **28** (4) 1015-1031.

HSU NS and CHENG KW (2002) Network flow optimization model for basin-scale water supply planning. *J. Water Resour. Plann. Manage.* ASCE. **128** (2) 102-112.

JOHNSON SA, STEDINGER JR and STASCHUS K (1991) Heuristic operating policies for reservoir system simulation. *Water Resour. Res.* **27** (5) 673-685.

LOUCKS DP, STEDINGER JR and HAITH D (1981) *Water Resource Systems Planning and Analysis*. Prentice-Hall Englewood Cliffs, N. J.

NDIRITU JG and DANIELL TM (1999) Assessing model calibration adequacy via global optimization. *Water SA* **25** (3) 317-326.

NDIRITU JG and DANIELL TM (2001) An improved genetic algorithm for rainfall-runoff model calibration and function optimization. *Math. and Comp. Modell.* **33** (6-7) 695-706.

NDIRITU JG (2002) Optimizing water resource system design and operation using genetic algorithms. *Proc. Int. Conf, Water the Life Blood of Mankind*. Arusha Tanzania, 11-13 Dec.

OLIVIERA R and LOUCKS DP (1997) Operating rules for multireservoir systems. *Water Resour. Res.* **33** (4) 839-852.

SHARIF M and WARDLAW R (2000) Multireservoir system optimization using genetic algorithms: Case study. *J. Comp. Civil Eng.* ASCE. **14** (4) 255-263.

VAN VUUREN SJ (2002) Application of genetic algorithms – Determination of the optimal pipe diameters. *Water SA* **28** (2) 217-226.

WARDLAW R and SHARIF M (1999) Evaluation of genetic algorithms for optimal reservoir system operation. *J. Water Resour. Plann. Manage* ASCE. **125** (1) 25-33.

YEH WW-G (1985) Reservoir management and operation models: A state-of-the-art review. *Water Resour. Res.* **21** (12) 1797-1818.