Multiobjective Genetic Algorithms for Pump Scheduling in Water Supply

by
Dragan A. Savic, Godfrey A. Walters and Martin Schwab[1]
School of Engineering
University of Exeter
Exeter, EX4 4QF

Introduction

Seeking cost reduction and energy savings in water supply by improving the operation of pumps is both an obvious as well as a very efficient possibility to consider. Without making changes to the basic elements of a water supply system, remarkable reductions in operation costs can be achieved by optimising the pump scheduling (PS) problem. The UK Water Research Centre (WRC) [1985] estimates the scope for savings for an average system to amount to ten percent of the current pumping cost. Considering that in the UK the overall energy cost for pumping in water supply is about £70 Million p.a., the estimated savings are very high.

The pump scheduling problem can be formulated as a cost optimisation problem which aims to minimise marginal costs of supplying water, whilst keeping within physical and operational constraints (e.g. maintain sufficient water within the system's reservoirs to meet the required time varying consumer demands) [WRC, 1985]. There are different causes for costs, most of which can be regarded as energy costs. Important characteristics of the costs related to the energy consumption are the electricity tariff structure and the civil engineering components of the system, determining the relative efficiencies of the available pumpsets, the head through which they pump and marginal treatment costs. To achieve the aforementioned aim of pump scheduling, an algorithm should be devised to determine which of the pumps available within the system, should be used during which interval of the optimisation period.

Besides energy costs, pump maintenance is another important cost factor in water supply. Although the cost for wear on the pumps caused by the switching of a pump cannot be quantified easily, it can be assumed that it increases as the number of "pump switches" increases. According to Lansey et al. [1994], a pump switch is defined as "turning on a pump that was not operating in the previous period". Thus the aim is to limit pump maintenance costs by limiting the number of pump switches.

The issue of optimised pump scheduling has been subject to increased research in recent years. A comprehensive review of optimisation approaches to pump scheduling in given by Ormsbee and Lansey [1994]. They considered methods using linear, dynamic and non-linear programming as well as specialised forms of these methods (e.g. mixed-integer linear programming). Their classification is made

[1] ERASMUS student, University of Stuttgart, Germany

according to the programming method implemented, the type of the system that is to be optimised (e.g. single or multiple tanks), the type of the hydraulic model used, the type of the demand model used and the resulting control policy. An attempt to develop and use Genetic Algorithms (GAs) for pump scheduling has been reported by Mäckle et al. [1995]. They used a binary GA with linear ranking to optimise the pump schedule for a one-reservoir, 4-pump system over a 24-hour period. They concluded that GAs and other evolutionary programming methods can offer a lot of new possibilities for solving pump scheduling problems, in particular by allowing different constraints and cost factors to be considered more easily than with other methods.

This paper presents several improvements of the single objective GA and the results of the further investigation into the use of multiobjective GAs for solving the pump scheduling problem. A multiobjective approach used in this work deals with both the energy cost and the pump switching criterion, at the same time. The performance of the algorithm is tested for different demand profiles and additional requirements and compared to that of the single objective GA.

Pump Scheduling Problem

The pump scheduling problem presented by Mäckle et al. [1995] is used in this work. The system considered consists of one water distribution reservoir which is supplied by four pumps through a single water main. The optimisation period is set to one day as historic patterns of the water demand of an average day are commonly used for pump scheduling. The time interval over which the electricity tariff structure is repeated was also modelled to be 24 hours.

Table 1. Pumping capacities of the fixed speed pumps

Pump	Amount of water pumped in one hour [m^3]	Amount of electricity used in one hour [kWh]
Pump 1	10	12
Pump 2	30	30
Pump 3	50	44
Pump 4	100	80

As in many electricity supply systems, the model incorporates a cheap night and a more expensive day tariff. For the problem used in this paper the day-time tariff cost is set to be twice that charged during the night. The period for which the higher day-tariff applies is from 7 a.m. until 8 p.m. The optimisation period was divided into intervals of one hour, i.e. the pumps can be either switched on or off during a particular hour of the day.

There are four fixed speed pumps in the system. As each of the pumps can be run during any time interval, there are $2^4 = 16$ possible combinations of the pumps during each hour of the day. The (constant) pumping capacities of the pumps are listed in Table 1.

With hydraulic problems like the pumping of water, one has to consider friction losses and other head influences to avoid oversimplification. Details about head influences can be found in [Hicks, 1971]. Friction influences are water flow rate, pipe diameter, pipe length, interior condition of the pipe and characteristics of the fluid. The relative friction loss increases with an increase in the amount of water pumped. For this problem it was assumed that the friction loss for flow rates smaller than 100 m^3/hour is small and can therefore be neglected. The values listed in Table 2 are for the pump combinations with a flow rate greater than 100 m^3/hour.

Table 2. Influences of the pumping head on the electricity consumption

Pumps switched on	Amount of water pumped in one hour [m^3]	Amount of electricity used in one hour without head influence [kWh]	Amount of electricity used in one hour with head influence [kWh]
4	100	80	80
1+4	110	92	92
1+3	130	110	111
1+3+4	140	122	124
1+2	150	124	127
1+2+4	160	136	141
1+2+3	180	154	165
1+2+3+4	160	166	182

The maximum storage volume in the reservoir was set to 800 m^3 [Mäckle et al., 1995]. If the reservoir is filled to this maximum, this volume can be taken from the reservoir and after taking it out, there is a minimum amount of water left in the reservoir for safety reasons (e.g. fire fighting). The water source is assumed not to have any restrictions. The 24-hour consumer demand profile was adopted from the problem presented by Mäckle et al. [1995].

Genetic Algorithm

Single Objective Approach

In the pump scheduling case, each pump during a certain time interval is represented by one bit of the string. This means that if the bit's value is zero, the pump it represents is switched off during this time interval. If the value is one, the pump is switched on. A string consisting of 24×4 bits describes the problem completely. Thus the overall number of possible solutions to the pump scheduling problem is 2^{96}.

The fitness function for the single objective approach developed by Mäckle et al. [1995] aggregates the cost for energy and consumption and penalties for violation of the constraints of the system. The electricity cost C_e is given as:

$$C_e = T_n \sum_{i=1}^{7} C\left[PC(i)\right] + T_d \sum_{i=8}^{20} C\left[PC(i)\right] + T_n \sum_{i=21}^{24} C\left[PC(i)\right]$$

where, i is the time interval of the day, T_n is the night time tariff, T_d is the day time tariff, $PC(i)$ is the pump combination of the time interval i and $C[PC(i)]$ is the electricity consumed in time interval i. Other cost factors can be easily incorporated into the algorithm.

Three system constraints are considered in an approach similar to that of Mäckle et al. [1995].The following constraint violations are incorporated into the current algorithm as penalties:

(1) minimum reservoir level,

(2) the initial water level should be reached in the reservoir at the end of the optimisation period, and

(3) maximum reservoir level (not included directly).

The resulting cost function in terms of energy consumption and constraint violations is:

$$Cost = C_e + \sum Penalty$$

The fitness of individual chromosomes is calculated as the inverse of the overall cost.

Multiobjective Approach

In contrast to single-criterion models, which proceed to the identification of a so-called "optimal" solution, multicriterion methods provide a choice of trade-off solutions from which a decision maker can select a suitable one to implement. Hence, the aim of the multiobjective (MO) approach is to find a spread of good, trade-off solutions with respect to all objectives. In the case of pump scheduling, the two objectives are the minimisation of energy costs and the minimisation of pump switches. As it is desirable to deal with two objectives separately, i.e. without aggregating them into a single fitness function, the fitness of the individuals has to be determined according to both objectives individually.

The energy cost function used for a single objective case is used for the MO approach as well. The fitness of a chromosome according to the pump switch objective is defined as the inverse value of the number of pump switches a particular pump schedule requires. As a result, a solution requiring fewer switches is considered better according to this objective than a solution with more pump switches.

The multiobjective approach used in this work is based on the concept of pareto optimal ranking as suggested by Goldberg [1989]. The method consists of three steps:

(1) finding the non-dominated individuals in the current sub-population,

(2) assigning a rank to the individual of this pareto-optimal set, and

(3) removing this group of non-dominated individuals from the current sub-population and thus creating a new sub-population.

The aforementioned process is repeated until all individuals of a generation have been assigned to a group. The first set found is ranked one. The rank assigned to the individuals of the following pareto-optimal sets found in this generation is increased by one for every repetition. Finally, all individuals are assigned a fitness score according to their rank.

In addition to providing a set of trade-off solutions, the MO algorithm has to ensure that there are no infeasible members of the final population, i.e. in the set of solutions that is presented to the decision maker (pareto-optimal set). Because of the fact that penalties assigned to an infeasible individual only affect the cost objective, there is a chance of having an infeasible solution in the final pareto-optimal set. To prevent this from occurring a method had to be devised to deny equal probability of reproduction to the infeasible solutions that are in the pareto-optimal set. This was achieved by treating feasibility as an additional objective but with the highest priority. In addition, this objective is used only to distinguish between the top two groups (ranked 1st and 2nd).

The Improved Genetic Algorithm

To increase the exploitation features of the GA the option of combining it with a local search method is implemented in this paper. Two different local search strategies based on two different definitions of the neighbourhood of a binary string representing a pump schedule, were investigated. The following are the two neighbourhood definitions:

(a) the neighbourhood is defined as the set of binary strings differing from the given string in exactly one bit, and

(b) the neighbourhood consists of strings which contain two bits of different values, representing the same pump at different times and are identical in all the other bits.

Regardless of which definition is used, if the new string obtained by searching the neighbourhood represents a better solution, the search proceeds with the new string. The term “better solution” implies different decisions for the multiobjective (MO) and single objective (SO) approach. In the SO approach the solution newly created by the local search operator is kept if its costs have improved compared to the costs of the chromosome that was altered in the last step of the local search, i.e. the starting solution. Otherwise, the starting solution is kept. In the MO approach the new solution is kept if the two following requirements are met:

(1) neither the cost nor the switches have deteriorated compared to the starting solution, and

(2) the individual has not become infeasible by the action of the local search operator.

In addition to the local search methods, the progressive penalty assignment was applied in this work. A penalty function that assigns comparatively low penalties at the beginning of the optimisation process and increases penalties during the course of optimisation, was used:

$$Penalty = y + \frac{m}{e^{-s\cdot(g+x)} + 1}$$

where, m is the maximum penalty chosen, g is the number of generations produced at the time the function is called, and y, s and x are the parameters of the penalty function.

Results

The two local search methods were applied to both the SO and MO approaches. When applied to the SO approach, the first local search (neighbourhood case "a") performed better than the second local search method (neighbourhood "b"). The reason for that may be in the fact that the second method neither changes the amount of water that is pumped nor shifts the pumping work from one pump to another. However, when applied with the MO approach, both local search methods performed equally well. In both cases the GA combined with a local search found good solutions after fewer evaluations of the objective function.

If there is no detailed knowledge of the problem, and the field of application of an optimisation algorithm is not very limited, the algorithm must be able to cope with changes to the model. Furthermore, finding an optimal solution quickly and with a comparatively small computational effort is often required. If the changes to the problem, compared to a previous problem, are not radical, it might be sensible to start the optimisation from solutions found before. In this work, changes in the amount and the profile of the consumer demand and alternatives of the initial water level were considered to be important in this context. Again, both the SO and MO approaches were tested using the aforementioned changes.

Altering the Demand Profile

The solutions used as an initial population for the test run were produced by running the optimisation algorithm with the original demand profile for the same number of generations and also the same parameter configuration as the runs with the alternative demand profiles. The first alternative demand profile tested differs from the original one only in the distribution of the demand over the day. Hence, the sums of the hourly demands are the same for both profiles and amount to 2570 m^3 in 24 hours. The second alternative demand profile has higher peaks around lunch time and during the late hours of the day. These have been chosen according to the recommendations given by the Water Research Centre [1985]. The overall amount of water consumed is 2670 m^3 and therefore higher than the overall demand of the original profile.

The use of the SO approach with the first alternative demand profile and with the initial population seeded with the final solution from the original run resulted in an improvement over the run which started with a randomly created initial population. Changing the distribution of the demand over the day leads in this case to lower costs for pumping the same amount of water during 24 hours. As expected, the best solution for the second alternative demand profile is slightly more expensive than the one found for the original demand profile as more water has to be pumped. However, starting from a population seeded with the original solution leads to an improved performance of the GA. The improvement achieved was less pronounced than in the case of the first demand profile because the changes to the system were considerable.

The solutions found using the MO approach and with the initial population seeded with solutions obtained in previous runs, were always well distributed in the cost/switch plane. On average, the number of trade-off solutions in the pareto-optimal set was larger than in the case of solutions obtained with a

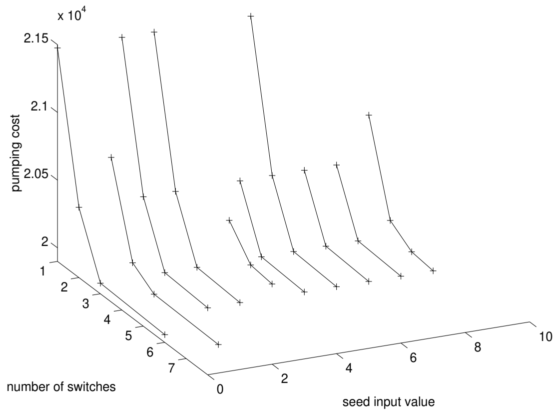

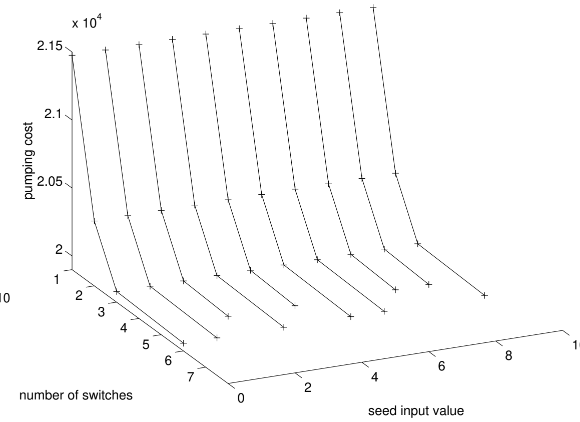

randomly created initial population (Figure 1).

Figure 1. Performance of the MO approach with the alternative demand profile and: randomly created population (left); an initial population obtained from previous run (right)

Altering the Initial Water Level in the Reservoir

The initial water level in the reservoir at the beginning of an optimisation period is also the water level that is to be achieved again at the end of the period. Thus, changing the initial water level does not alter the total amount of water that has to be pumped, but the times at which the water must be pumped are changed. Two different cases were considered: (a) lower initial level (75%), and (b) higher initial level (125%).

Getting the initial population for the SO optimisation run by seeding the algorithm with results of similar runs enabled the search to find good solutions for altered initial water level considerably faster than in the case of randomly generated initial population. In the case of the higher initial water level (125%) the algorithm with the seeded initial population even converged on a better final solution (Figure 2). Similar results, in terms of cost, were found for the MO approach as well.

Conclusions

In this work the approach to pump scheduling for water supply using the simple, single objective (SO) GA, implemented by Mäckle et al. [1995], was investigated and several improvements were introduced. These improvements include:

(a) a hybridisation of the GA search methodology with the local (neighbourhood) search, and

(b) the multiobjective treatment of the pump scheduling problem.

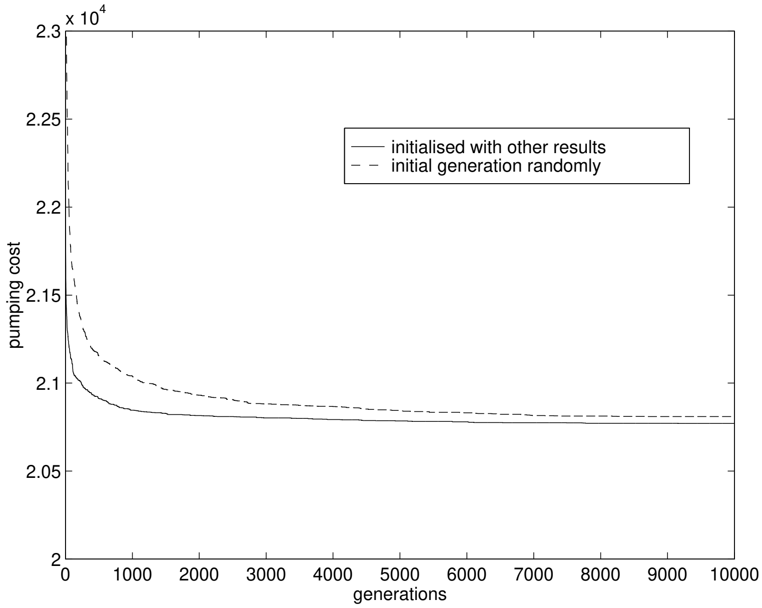

Figure 2. Course of evolution for a SO solution with random or previous-run initialisation

Two different methods of defining the neighbourhood were investigated and implemented with both the SO and MO approaches. In both implementations, the improvement in GA performance was observed after hybridisation.

In addition to energy costs, the MO approach implemented in this work considered pump switching as an additional objective. The algorithm was tested and improved by the following measures:

(a) progressive assignment of penalties for constraint violation, and

(b) introduction of feasibility of solutions as an additional objective (with the highest priority).

Based on tests performed, the MO approach achieved very good solutions in terms of both objectives and generally provides a good set of alternative solutions within the pareto-optimal group.

Robustness of the two GA approaches was tested by changing the amount and the profile of the consumer demand and by altering the initial water level in the reservoir. It was found that even in the case when changes were large (25% change in the initial reservoir level) the seeding of the initial population with the solutions from previous runs brought improvements in both efficiency (speed) and quality of the solutions found.

Acknowledgements

This work was supported in part by the UK Science and Engineering Research Council, grant GR/J09796

References

Goldberg, D.E., 1989, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley.

Hicks, et al., 1971, *Pump Applications Engineering*, McGraw-Hill, New York.

Lansey, K.E. and K. Awumah, 1994, Optimal Pump Operations Considering Pump Switches, *Journal of Water Resources Planning and Management, ASCE*, 120(1), 17-35.

Mäckle, G., D.A. Savic and G.A. Walters, (1995), Application of Genetic Algorithms to Pump Scheduling for Water Supply, *Genetic Algorithms in Engineering Systems: Innovations and Applications*, *GALESIA '95*, IEE Conference Publication No. 414, Sheffield, UK, pp. 400-405.

Ormsbee, L.E. and K.E. Lansey, 1994, Optimal Control of Water Supply Pumping Systems, *Journal of Water Resources Planning and Management, ASCE*, 120(2), 237-252.

Water Research Centre, (1985), *Pump Scheduling in Water Supply*, Swindon, Wilshire, U.K.