Water Network Rehabilitation with a Structured Messy Genetic Algorithm

By D. Halhal [1], G.A. Walters [2], D Ouazar [3] and D.A. Savic [4]

Abstract

The importance of water distribution network rehabilitation, replacement and expansion is discussed. The problem of choosing the best possible set of network improvements to make with a limited budget is presented as a large optimisation problem to which conventional optimisation techniques are poorly suited. A multi-objective approach is described, using capital cost and benefit as dual objectives, enabling a range of non-inferior solutions of varying cost to be derived. A Structured Messy Genetic Algorithm is developed, incorporating some of the principles of the Messy Genetic Algorithm, such as strings which increase in length during the evolution of designs. The algorithm is shown to be an effective tool for the current optimisation problem, being particularly suited both to the multi-objective approach and to problems which involve the selection of small sets of variables from large numbers of possibilities. Two examples are included which demonstrate the features of the method and show that the algorithm performs much better than a standard Genetic Algorithm for a large network.

[1] Engineer, Water and Electricity Distribution Co.(RAID), 5 Rue Okba Ibn Naffiy, BP 286, Tangier, Morocco

[2] Senior Lecturer, School of Engineering, University of Exeter, Exeter, EX4 4QF, UK

[3] Head of Hydraulic Systems Analysis Laboratory, Mohammadia School of Engineers (EMI), BP765, Agdal, Rabat, Morocco

[4] Lecturer, School of Engineering,University of Exeter, Exeter, EX4 4QF, UK

Key Words: Water distibution, genetic algorithms, multi-objective, optimisation, rehabilitation, networks, systems, design.

Introduction

As water distribution systems age, they experience problems of deteriorating infrastructure, water loss and service disruption. They gradually lose carrying capacity, eventually failing to fulfil their specified function of distributing potable water continuously and within prescribed pressure limits. Consumer demands increase with time, and low pressure and high head loss problems become frequent, giving rise to excess energy expenditure and causing consumer discontent. The frequency of leaks and breaks increases, resulting not only in loss of water, but also in contamination through cracked pipes. More money becomes necessary to re-establish or maintain adequate levels of service, but funds are usually limited, allowing only partial solutions to the problem. Hence the decision process for rehabilitation, replacement, and/or expansion of existing systems to meet current and future demands forms a major area of interest.

Improvement in a system's performance can be achieved through replacing, rehabilitating, duplicating or repairing some of the pipes or other components (pumps, tanks, etc.) in the network, and also by adding completely new components. It is likely that funding will be available to modify or add only a small number of components in a network at any time. The decision problem is to choose which components to add or improve (and how to improve them) to maximise the benefits resulting from the changes to the system whilst minimising the costs.

The benefits from system rehabilitation are subject to rapidly diminishing returns with respect to cost, which is usually limited by funding for which other schemes compete. When maximising benefit subject to a cost limit, a single-criterion optimisation algorithm will seek for the solution which yields the highest benefit, neglecting other cheaper solutions yielding lower benefits which might be preferable given the competition for funds. The problem can be modelled more usefully as a multi-criteria optimisation problem, for which a range of improvement schemes of differing costs and benefits is obtained, up to a maximum cost corresponding to the whole of the available fund. The present paper describes such an approach, which permits a judicious choice of solution along the trade-off curve between benefit and cost, according to the part of the total fund releasable and the benefit resulting from the rehabilitation scheme.

The paper also demonstrates the effectiveness of combining a technique the authors call "Structured Messy Genetic Algorithms" (SMGA) with multi-objective optimisation to form a model capable of handling real network improvement problems, in particular the selection of solutions involving small sets of components from within large water distribution systems. The development of the SMGA is presented in this paper for the multi-objective problem, using the concept of Pareto-optimality (Goldberg 1989). The algorithms is tested with and without a technique known as "fitness sharing", and compared in performance with a standard Genetic Algorithm (Goldberg 1989) using two examples.

Previous work

A number of published works have dealt with pipe rehabilitation and replacement based on historical statistics and costs. Shamir and Howard (1979) proposed a model for making pipe replacement decisions based on pipe breakage history and the cost of repairing and replacing pipes. The optimal time for replacing existing pipes is then obtained. Kettler and Goulter (1985) identified a relationship between breakage rate and pipe diameter for cast-iron and asbestos-cement pipes, which shows a decreasing pipe failure rate with increasing diameter, and a correlation between the number of pipe failures and pipe age. This result indicates an increasing number of breaks per year over time. Walski (1986) developed a framework for making water system rehabilitation decisions and gave suggestions about rehabilitation alternatives for some specific problems. He developed (Walski 1987) cost data for pipe cleaning and duplicating and made a comparison of different alternatives. O'Day (1982) stated that rehabilitation decisions must be based on accurate data for breaks and leaks geographically gathered and analysed but not on general tendencies, rules of thumb or ages of pipes which are not sound criteria for this purpose. Ramos (1985) proposed a methodology for evaluating direct and indirect benefits from pipe replacement to establish a benefit/cost ratio representing a criterion for justifying replacement of pipes of small size. Andreou & Marks (1986) presented a methodology for modelling breaks in large diameter cast-iron water pipes depending on stages of deterioration of the pipe.

Several authors have developed algorithms for optimal rehabilitation decisions. Woodburn et al (1987) presented a model for determining the minimum cost for the rehabilitation, replacement, and/or expansion of an existing water distribution system based on a combination of non-linear optimisation and hydraulic simulation procedures. The decision variables were the length and diameter of pipe to be replaced in each arc and the length of old pipe to be relined in each arc. The model allows replacement and/or rehabilitation of a portion of a pipe, which may not be realistic. Kim and Mays (1994) tackled the problem using integer variables to represent the decision for the replacement and rehabilitation instead of using pipe lengths as decision variables. Murphy & Simpson (1992) used a standard GA to find the optimal solution of the Gessler (1985) network. The method chooses the optimal combination among the eight alternative decisions possible for each of the eight decision variables (pipes). Murphy et al (1994) solved the “Any town” network (Walski et al 1987), introducing pumps and reservoirs, as well as pipes, as decision variables.

The above models for optimal rehabilitation are all based on the single objective of minimising cost subject to performance constraints, such as minimum pressure for consumers. They cannot readily be applied to the present problem of maximising the benefit subject to limits on funding. Neither has the use of any of the above models for large network rehabilitation problems been reported.

Problem formulation

The purpose of the approach presented herein is to find the optimal way to invest some or all of a limited amount of money to enhance the performance of an existing water distribution system. In this model, performance may be enhanced by:

- increasing the hydraulic capacity of the network by cleaning, relining, duplicating or replacing existing pipes,
- increasing the physical integrity of the network by replacing structurally weak pipes,
- increasing system flexibility by adding additional pipe links
- improving water quality by removing or relining old (cast iron) pipes

For simplicity, in this paper the only network components that are considered are the pipes, although in a more general formulation pumps, tanks, valves and reservoirs could all be incorporated without invalidating the algorithms.

This type of problem can be formulated in general as:

$$\begin{aligned} &\text{Maximise} \quad f(i) = \text{Benefit}(i) \\ \text{and} \quad &\text{Minimise} \quad F(i) = \text{Cost}(i) \\ &\text{subject to:} \quad \text{Cost}(i) \leq \text{Fund} \end{aligned} \tag{1}$$

where Benefit(i) is the benefit resulting from adopting solution i, Cost(i) is the cost of the solution i and Fund is the maximum budget available for the improvement of the system.

The Objective functions

Benefit

In general, the benefit attributed to a solution is a measure of the improvement in the level of service experienced by consumers combined with a measure of the decrease in costs associated with system operation and maintenance. The simplified model used here for demonstrating the algorithms developed in this paper combines four factors, with weights chosen by the user. The total benefit of a solution i is therefore expressed as the weighted sum of four benefits:

$$\text{Benefit}(i) = w_h \,\text{Benefit\_hyd}(i) + w_p \,\text{Benefit\_phy}(i) + w_f \,\text{Benefit\_flx}(i) + w_q \,\text{Benefit\_qual}(i) \tag{2}$$

where Benefit_hyd(i) is the pressure benefit derived from better hydraulics

Benefit_phy(i) is the maintenance benefit derived from better physical condition of the pipes

Benefit_flx(i) is the operational benefit derived from greater network flexibility

Benefit_qual(i) is the quality benefit derived from pipe replacement

and w_h w_p w_f and w_q are the respective weightings

The four components of the benefit are explained in more detail below.

Hydraulic Benefit

The benefit from improvement of hydraulic performance is quantified as the difference between the pressure deficiencies in the initial network before improvement (Deficiency(0))and in the solution found (Deficiency(i)).

$$\text{Benefit\_hyd}(i) = \text{Deficiency}(0)-\text{Deficiency}(i) \tag{3}$$

The deficiency is calculated as the sum of the pressure shortfalls or excesses, Δj, at consumer nodes (j), weighted by the nodal demands (Qj):

$$\text{Deficiency} = \alpha \sum_{j\in jpm} (\Delta j *Qj) + \beta \sum_{j\in jpx} (\Delta j *Qj) \tag{4}$$

where *jpm* is the set of nodes with pressure below the minimum

jpx is the set of nodes with pressure above the maximum

α, β are weights to allow different emphasis on pressure shortfalls or excesses.

To determine the nodal pressures and hence the deficiency of a solution, the steady-state hydraulic network solver, EPANET (Rossman 1993), is used.

Physical Integrity Benefit

The improvement in structural condition of the network pipes resulting from a solution i reduces future repair costs, with the benefit calculated as follows:

$$\text{Benefit\_phys(i)} = \sum_{j=1}^{Nrnw} Cost_{rep}(j) \qquad (5)$$

where *Nrnw* is the number of renewed pipes

$Cost_{rep}(j)$ is the break repair cost of the pipe *j*, equal to the present value of the breakage repair costs for the existing pipe over 10 years. (New pipes are assumed to have no annual repair cost during their first 10 years as they are usually under warranty for this period.)

The repair cost is given by the following formula (Walski, 1986)

$$Cost_{rep}(j) = \sum_{t=tp}^{t=tr} \frac{J(t)C_b(j)}{(1+r)^{t-tp}} \qquad (6)$$

where $C_b(j)$ is the repair cost of a break for pipe *j*

r is the interest rate

tp is the present year

tr is year *tp*+10

J(t) is the break rate in year *t*, which is approximated by

$$J(t) = Jo(1+b)^t \qquad (7)$$

where *Jo* is the break rate in year 0 (break/km/yr.)

b is the break rate growth coefficient

t is the time (yr.)

Flexibility Benefit

Each new pipe in parallel with an old one (duplication) increases the number of paths taken by water and therefore the flexibility of the system. It gives rise to a benefit quantified as follows:

$$\text{Benefit\_flx}(i) = \sum_{j=1}^{Ndp} \text{Diam}(j) \tag{8}$$

where *Ndp* is the number of new parallel pipes

Diam(j) is the diameter of the new pipe.

Quality Benefit

A low Hazen Williams (HW) factor for a pipe is usually a sign of corrosion, tuberculation or scaling, all of which create sites for the development of micro-organisms and/or discoloured water. Therefore, for each old pipe lined or replaced having a HW coefficient less than a specified value, a benefit for water quality improvement is assigned to the solution and is calculated as follows:

$$\text{Benefit\_qual}(i) = \sum_{j=1}^{Nrl} \text{Length}(j)) \tag{9}$$

where *Nrl* is the number of renewed and/or lined pipes

Length(j) is the length of the renewed or lined pipe j.

The above represents a very simplified approach for assessing the benefits of a solution, and judgement is required in choosing weightings that reflect the relative importance of the four main areas of benefit for the network under consideration.

4.2 Cost

The cost of a new pipeline is taken as the supply cost plus the cost of installation. For the purpose of this study, the costs used are those applicable to Morocco, from which country the examples are taken. They apply to asbestos cement pipes with a service pressure rating of 12 bars, this being the most commonly used pipe in the country. The cost of replacing an existing pipe is taken as slightly higher than that for a new pipeline, and the costs associated with cleaning and lining are assumed, as these techniques have not as yet been used in Morocco. These costs together with available pipe sizes are given in Table 1.

The Structured Messy Genetic Algorithm Approach

The algorithm used for the optimal rehabilitation problem is called a Structured Messy Genetic Algorithm (SMGA) and was derived by the authors from the Genetic Algorithm (GA) first proposed by Holland (1975). GAs are stochastic numerical search procedures based on the mechanics of natural selection and evolution. They have their roots in the biological processes of survival and adaptation. Excellent descriptions of the GA and subsequent developments can be found in Goldberg 1989, Davis 1991 and Michalewicz 1992. A brief description of the main GA features is given here for continuity.

Genetic Algorithms

Most GAs encode the design parameters for a problem onto a data string. An individual solution to the problem becomes a string of numbers, analogous to a "chromosome" in nature. The basic iteration involves a population of individuals (solutions) evolving through many generations, with new individuals replacing previous generation members. New individuals result from "reproduction", the first phase of which involves a stochastic selection process to choose sets of "parents" from the population which are on average fitter (better) than the population mean. The fitness is determined by evaluating the objective function (e.g. cost or benefit) from the encoded design parameters of the individual. Using stochastic reproduction mechanisms such as "crossover", the chromosomes from the parents are combined to form new individuals (offspring).

Mutation can be applied with a low probability to randomly change elements of the data strings, thereby introducing new values of design parameters.

Structured Messy Genetic Algorithm (SMGA)

The SMGA, introduced herein, has some features in common with the Messy GA (Goldberg et al 1989, Goldberg et al 1990) such as the coding scheme and variable length strings. It is based on progressively building up the complexity of the individuals in successive populations of solutions, in a structured manner, partially imitating the natural evolution of complex life-forms from single cell organisms.

SMGA starts with a complete enumeration of all single digit strings (elements equivalent to single cell organisms) to form an initial population. In other words all possible single variable decisions are evaluated. The best elements are retained in an initial population to provide subsequent populations with elemental building blocks of high fitness. In the basic iteration, elements from the initial population are added onto the current population members to form a new population with longer strings, the process being known as concatenation.

The population then undergoes an optimisation process using a conventional GA for a number of generations until meeting a termination criterion, whereupon a new population is generated by concatenation. This sequence continues until no improvement can be achieved after two successive concatenation steps, or the population strings reach a predetermined length, defined in accordance with the available fund. The flowchart of Figure 1 outlines the SMGA algorithm.

Water distribution systems are usually of large size, while the available fund is often very limited and permits only a small number of arcs or other elements to be upgraded, renovated or strengthened. SMGA takes advantage of the problem structure and can therefore easily handle large systems, while other optimisation techniques, including SGA, are either poorly adapted to the kind of problem or face computer memory difficulties.

SMGA encodes only those decision variables which are active. Only the relevant arc numbers and the rehabilitation decision for each of those arcs are stored in the string, which increases in length through the evolution process. The total search space is thereby explored with strings of small maximum length, (corresponding to the maximum number of pipes that can be financially considered), thus avoiding very serious computer memory problems for systems of large size. In contrast, the SGA stores data for each network element, whether active or not, giving rise to very long strings for large networks. In addition, the search space is for SMGA is reduced with respect to the SGA which considers all combinations of arcs as candidate solutions. This speeds up the process of finding the optimal or near optimal solutions, thus reducing the CPU time consumption. In fact, when only p arcs among q can be financially considered, and each can take n alternative solutions, the search space contains a number of solutions equal to:

$$n^p \frac{q!}{p!(q-p)!}$$

while it contains n^q possible solutions when the total number q of arcs is considered as in standard GAs.

Applying Genetic Algorithms to Multi-objective optimisation

Like most optimisation problems, water distribution design has multiple objectives, such as minimising cost and maximising benefit. It is possible to use a single-criterion formulation by using the other criteria as constraints, (e.g. maximising benefit subject to a cost limit), or by incorporating multiple criteria into a single objective function, using weighting factors. Neither of these approaches is entirely satisfactory.

In multi-objective optimisation, the different problem criteria are evaluated separately and each solution is assessed according to these terms. The result is not a single point but a set of points known as a Pareto optimal set. Each point of the Pareto optimal set is not dominated by any other point, i.e. in going from one point to another, it is not possible to improve on one criterion without making at least one of the other criteria worse.

By maintaining and continually improving a population of solutions, a GA can search for many non-dominated solutions at the same time, which makes it a very attractive tool for solving multi-objective optimisation problems. It is, however, necessary to control the distribution of solutions within each population to prevent convergence onto a single solution. In the present application this is done by maintaining sub-populations of varying benefit for defined ranges of solution cost (classes).

Although the multi-objective problem can be handled with a SGA, the SMGA has some significant advantages for the present problem. With a standard GA, evolution starts with all variables taking randomly assigned values. It is very unlikely that this will result in an initially even distribution of solutions over all cost classes, but using techniques such as Pareto optimal ranking and fitness sharing (both described below) an even distribution of solutions will eventually evolve. However, when the fund available for rehabilitation is limited, as in the present application, and the system's size is large, only solutions of low cost are of interest. A standard GA will create many generations of high cost solutions before producing a set for which the cost is less than the available fund. This delays both the formation of an adequate range of sub-populations and the subsequent improvement of their elements, thus affecting severely the SGA's efficiency and performance.

SMGA, in contrast, approaches the problem in a progressive way, starting with short strings of one digit length and increasing them during the process until reaching their maximum length. It generates solutions over the required range of costs from early in the process, and then focuses on improving them. This not only overcomes the initial delay experienced by SGA in finding an adequate range of feasible solutions, but also speeds up the process of forming the final set of non-dominated solutions.

Although SMGA can be used for single objective optimisation, e.g. to maximise benefit subject to a cost limit, its efficiency is impaired by the early generation of high benefit but high cost elemental "building blocks", which are of limited use when more complex solutions evolve. In contrast, the multi-objective approach provides "building blocks" with a range of costs and benefits, ideal for subsequent generation of more complex solutions. Hence it can be seen that SMGA and the multi-objective approach are mutually beneficial, and together can form an efficient model capable of solving large system improvement problems.

To deal with multi-objective problems, the following extensions to GAs were incorporated and developed.

Pareto optimality

Pareto optimal ranking, suggested by Goldberg (1989) is a method of assigning fitness to individual solutions which takes into account each of the different optimization criteria. The method involves finding the set of individuals in the current population that are non-dominated. These individuals are assigned the highest rank and removed from further contention. The next set of non-dominated individuals is determined from the remaining population and assigned the next highest rank. This process continues until the entire population is suitably ranked. A fitness value is then assigned to each individual based on the rank to which it belongs.

To illustrate Pareto optimality ranking, assume a population of 10 individuals, plotted according to cost and benefit on Fig 2, has to be ranked. The curve Ω is the trade-off curve between the benefit and the cost. The points connected by this curve form a set of non-dominated solutions because there are no points better than these on all criteria. These points, indexed 1 in Fig 3, are good solutions although none is best for both criteria. The set {A1, B1, C1, D1} is of points not dominated by any other point. All these points are ranked one and then removed from further contention. The next set of non-dominated solutions {A2, B2} is assigned a rank of two and then excluded. The rank three is assigned to the next non-dominated set {A3, B3, C3}, and finally, the set {A4}, containing one solution, is assigned rank four. The population is then completely ranked, and the fitness can be assigned to each element as a function of its ranking. For the present work, fitness is given by:

$$fit(i) = \frac{1}{rank(i)} \qquad (10)$$

where *fit(i)* is the fitness of individual *i*, and *rank(i)* is the ranking of individual *i*.

6.2 - Fitness sharing

Fitness sharing was introduced by Goldberg and Richardson (1987), and analysed in detail by Deb (1989). Its basic idea is to force individuals to share available resources by dividing the population into sub-populations of similar individuals. Its goal in multi-modal function optimisation, is to distribute the population over a number of different peaks in the search space, with each peak receiving a fraction of the population in proportion to the magnitude of that peak. In a multi-objective optimisation, fitness sharing methods are useful in stabilising the multiple sub-populations that arise along the Pareto-optimal front, thereby preventing excessive competition between distant population members (Goldberg 1989).

A fitness sharing scheme is used here to maintain stable clusters of solutions over a range of cost classes, a technique known in GA terminology as niche formation and speciation. The fund is divided into a number of cost intervals (5 to 10 is used here). Each individual having a cost lying within the bounds of an interval is assigned to the interval, thus forming sub-populations or classes of solutions having a cost within specified limits. Fig. 3 illustrates the formation of classes.

Class i is said to be full if it contains a number of solutions greater than or equal to $\frac{N_{pop}}{Nclass}$, where N_{pop} is the population size, and $Nclass$ is the total number of classes.

The shared fitness is determined as the fitness of an individual divided by the number of individuals belonging to its class.

$$shared\_fit(j) = \frac{fit(j)}{Ncl(i)}$$

where $shared\_fit(j)$ is the shared fitness of individual j, $fit(j)$ is its fitness determined by Pareto optimality ranking and $Ncl(i)$ is the number of individuals belonging to class i.

The number of classes or sub-populations is variable within the SMGA formulation. It starts with a small number (5 classes), as the cost of one digit string solutions is usually small compared to the available fund, and increases during the process until reaching a predetermined limit.

The SMGA parameters

Population size

Population size is a very sensitive parameter for the GA process because, with too small a population, the GA will converge rapidly to a local optimum, while, with a large one, there will be a long waiting time for an improvement in the process. A population size of 40 elements was found suitable for the current work.

Coding

An integer coding, (strings of integers corresponding to the design variable decisions), is used as it fits in well with the problem formulation. In this case, as the number of decision variables changes from population to population, the string contains two codes: one indicates the decision variable reference (the pipe number) while the other indicates the decision number. In practice, the string is split into two independent sub-strings each representing one code type; i.e., instead of representing the string as: [(2,1) (1,8) (5,2) (3,4) (6,2)], where left numbers correspond to the arcs, and right numbers correspond to decisions, it is represented as:

arc sub-string	[2 1 5 3 6]
decision sub-string	[1 8 2 4 2]

This representation is flexible and allows a genetic operator to be applied to one sub-string independently of the other.

Crossover

Crossover describes the process by which parents' coded data strings are combined to form new coded strings for their offspring. In its simplest form, a single point along the coded string is selected at random, and the code strings from two parents are broken at this point. The tail ends are switched and the strings recombined to form the codes for two offspring. In the SMGA, a two point crossover (the string being broken at two points with the middle section of code switched), was found more effective than one point crossover, and is applied to both arc and decision sub-strings. In this application, a single offspring is created from two parents. The newly formed arc sub-string has to be checked to see if any arcs (pipes) are duplicated, with duplicated arcs being replaced randomly by others.

Mutation

Occasional random alteration of digits protects the GA process against premature loss of potentially useful genetic material. Each digit of an offspring's string is altered with a low probability p_m , usually held constant throughout the evolution in a SGA. However, in the SMGA the mutation rate p_m is set as a function of the string length, resulting in a decreasing probability of mutation for each digit as string lengths increase. Mutation is applied to arc sub-strings as well as to decision sub-strings with a check to eliminate duplication in the former.

The Evolution Iteration

In the SMGA, each generation has a population of individuals with strings longer those of the previous generation. Two steps are performed to generate a new population. First a concatenation process creates individuals for the new population by adding one digit to the strings of each individual in the old population. The added digit is chosen either from the initial population of good single element solutions, or randomly if duplication of arcs occurs. The process is performed one or more times, depending on the string length specified for the new generation. The second step allows evolution of the new population with the string length held constant and using one of the two conventional GA sub-models outlined below. The evolution is allowed to occur for a restricted time. The concatenation step is then reapplied and the cycle repeated until no improvement is achieved after two concatenation steps or a maximum specified string length is reached.

The GA Sub-Models

The two GA sub-models which were tested within the SMGA were first a "Steady State" GA and second a "Traditional" GA. Both use a selection procedure to chose two individuals from the current population as parents from which an offspring is formed by crossover and mutation as described earlier. The selection procedure adopted is a stochastic process under which the probability of an individual's selection is proportional to its fitness. The offspring's cost and fitness are then evaluated.

The methods differ in the way in which offspring replace older individuals. In the Steady State GA, the offspring is placed in its appropriate cost class, replacing the worst individual if the class is over-full. If the class is not full, the worst individual of another full class is removed, thereby maintaining a constant size (steady-state) population. To keep adequate diversity in each class, duplicate solutions are not allowed. At the start, some classes will be over-full and others under-full, but after a few iterations all classes will have a balanced number of individuals. The Steady State GA terminates after a number of objective function evaluations, the number being a function of the string length.

In the Traditional GA approach, the offspring becomes a member of a new population which completely replaces the old one when enough offspring have been created. Pareto optimal ranking and fitness sharing techniques respectively deal with the creation of diverse sub-populations and the distribution of solutions among the various classes formed. The Traditional GA terminates after a number of generations, the number being a function of the string length. This GA is less time consuming than the steady state one for similar results, and so was adopted for the final model.

The SGA

To gauge the performance of the SMGA, a SGA was also developed and tested for comparison. An integer coding was adopted with the string length equal to the number of arcs (pipes) in the network, and the integer value in each string position defining the action to be taken for the corresponding arc. The selection and crossover operators were the same as those presented earlier. As the string length is constant, the mutation rate was set to a fixed value (0.05).

Examples

Example 1: A Small Looped Network

Figure 4 shows a small looped network consisting of 15 pipes, 9 nodes and 7 loops. A minimum pressure head of 20 m is sought at each node. There are 8 possible decisions for each pipe, as shown in Table 2. Tables 1, 3 and 4 give the costs, network data and available pipe sizes. The annual growth of pipe break frequency is assumed at 0.06, the available fund is assumed at 2000000, while the interest rate is taken to be 0.08. The coefficients used to weight and evaluate the benefits of solutions are given in Table 5.

First, the multi-objective problem was solved using the SMGA procedure. Two versions of the program were run, the first with and the second without fitness sharing. Each version was run three times for a total number of 5000 objective function (OF) evaluations. Each evaluation involved analysing the network using EPANET for a single set of demands.

The results do not show noticeable dominance of one version over the other. However, when using the fitness sharing method, the population is better distributed among the different classes. Solutions having a cost less than 1.1 x Fund were plotted according to Cost and Benefit on Fig 5, which shows the best solution for each generation for both versions.

The problem was also run using the SGA with fitness sharing. Comparison with the SMGA results demonstrates that SMGA managed to produce a slightly better set of non-dominated solutions, as shown in Fig 6. Table 6 gives a selection of the non-dominated solutions identified together with the costs, the corresponding benefits, pipeline decisions and the minimum nodal heads obtained. It should be noted that none of these solutions achieves the desired minimum nodal pressure of 20 m, with node 2 or node 5 being the most pressure critical. This is consistent with the available fund being less than required to solve all the problems in the network.

Example 2: a real system

The second example, figure 7, is a real water distribution network for a town of 50000 inhabitants in Morocco. It contains 115 nodes and 167 pipes, and is fed from one reservoir. The fund is fixed at 2500000, while the required minimum head, annual break growth and interest rate are the same as in example 1.

The total number of possible solutions for an unlimited string length equals 6.55×10^{150}. For a realistic maximum string length of 40 digits, the search space is reduced to 8.67×10^{74} possible solutions, which still remains huge. For example, a complete enumeration of this “reduced” space would consume 2.75×10^{62} centuries of CPU time, assuming that 1000 objective functions can be evaluated every second.

As a test, SMGA and SGA, both with fitness sharing, were run 3 times for 25000 objective function evaluations each time. Initially, SGA consumed the total number of evaluations without converging to a set with costs less than the available fund. Subsequently a high penalty function was introduced to drive solutions towards the economically feasible region. High cost solutions are penalised harshly and so have a low chance of contributing to future generations, while solutions having costs less than or equal to the fund are conserved. Even with this modified penalty, it took more than 7000 objective function evaluations for the SGA to converge to a feasible set, and sub-populations in only four out of eight cost classes were formed at the end of the process. In contrast, SMGA formed sub-populations in all cost classes from the beginning. The results for this large problem demonstrated an outstandingly superior performance for SMGA over SGA as shown in Fig.8.

Conclusions

As water distribution systems age, their pipes lose carrying capacity and become increasingly susceptible to breaks and cracks. Combining this with an ever increasing demand for water can give rise to loss of satisfactory performance and uneconomic operation. Rehabilitation, replacement and/or expansion of water distribution networks becomes very important and involves increased expenditure, which will usually greatly exceed any available funding.

The problem of deciding on the optimal improvements is a multi-criteria one, as the increase in benefit from the rehabilitation operation can be achieved only at the expense of an increase in cost. A multi-objective optimisation method was successfully developed based on a Structured Messy Genetic Algorithm (SMGA), introduced here by the authors, and was compared in approach and efficiency with a Standard Genetic Algorithm. Both use the concept of Pareto optimal selection, and have been developed to find the best way to invest judiciously some or all of an available budget on improving the carrying capacity of a water distribution system, the physical integrity of its pipes, the water quality and the system flexibility.

The methods have been tested using two examples, the first being a small hypothetical network and the second a moderately large real system. The results showed that the SMGA method is greatly superior to SGA for the large network.

A method based on the techniques of niche and species was developed, involving the division of a population into several cost classes. A fitness sharing scheme was also implemented to assess its influence on the GA process. It prevents overcrowding of a class by dividing the fitness of each individual in a class by the total number of individuals in that class. The results showed that the introduction of fitness sharing does not produce any significant improvement, but ensures a good distribution of the population among the classes.

The structure of the problem is such that only a small subset of the design variables (pipeline upgrading options) will be selected in economically feasible solutions due to funding constraints. The progressive building up of solutions from simple elements developed for SMGA, combined with the multi-objective approach, which keeps a range of good solutions with varied costs throughout the process, proved very effective. This was especially true for the largest network, for which SMGA greatly outperformed a more traditional SGA approach.

Acknowledgement

The authors gratefully acknowledge the support of the British Council, Binnie Black and Veatch and the U.K. Engineering and Physical Sciences Research Council (grant GR/J09796).

Appendix 1: References

Andreou, S. and Marks, D.H.(1986). “A New Methodology for Modelling Water Pipe Breaks”, *Water Forum ‘86’*, ASCE, 2, 1726-1733.

Davis, L. (1991). *Handbook of Genetic Algorithms*, Van Nostrand Reinhold, New York.

Deb, K.(1989). “Genetic Algorithms in multi-modal function optimisation”, Masters thesis, The Centre for Genetic Algorithms Report No. 89002, University of Alabama.

Gessler, J.(1985). “Pipe Network Optimisation by Enumeration”. Proc. Speciality Conference, Computer Applications for Water Resources, ASCE, Buffalo, New York, 572-581.

Goldberg, D.E. (1989). *Genetic Algorithms in Search, Optimisation and Machine Learning*, Addison Wesley.

Goldberg, D.E., Korb, B. and Deb, K.(1989). "Messy Genetic Algorithms: Motivation, Analysis, and First Results", *Complex Systems*, 3, 349-530.

Goldberg, D.E., Deb, K and Korb, B.(1990). "Messy Genetic Algorithms Revisited: Studies in Mixed Size and Scale", *Complex Systems*, 4, 415- 444.

Goldberg, D.E. and Richardson, J.(1987). "Genetic Algorithms with sharing for multi-modal function optimisation", *Proc. 2nd Inter. Conf. Genetic Algorithms*, 41-49.

Holland, J.H.(1975). *Adaptation in Natural and Artificial Systems*, Ann Arbor Science Press, Ann Arbor, Michigan.

Kettler, A.J. and Goulter, L.C.(1985). "An Analysis of Pipe Breakage in Urban Water Distribution Networks". Canadian Journal Civ. Eng., 12 (2), 286-293.

Kim, H. J. and Mays, L.W.(1994). "Optimal Rehabilitation Model for Water Distribution Systems". *Journal of Water Resources Planning and Management*, ASCE, 120 (5), 674-692.

Michalewicz, Z. (1992). *Genetic Algorithms + Data Structure = Evolution Programs*. Springer-Verlag.

Murphy, L.J. and Simpson, A.R.(1992) “Genetic Algorithms in Pipe Network Optimisation”, *Research Report N° R93*, Department of Civil and Environmental Engineering, University of Adelaide, Australia.

Murphy, L.J., Dandy, G.C. and Simpson, A.R.(1994). “Optimum Design and Operation of Pumped Water Distribution System”, *Proc. Conf. on Hydraulics in Civil Engineering*, Institution of Engineers, Brisbane, Australia.

O’Day, K.(1982). “Organising and Analysing Leak and Break Data for Making Water Main Replacement Decisions”, *Journal AWWA*, 74 (11), 589-596.

Ramos, W.L.(1985). “Benefit/Cost Analysis Procedure for Determining Water Main Replacement”. *AWWA National Conference*, 1-13.

Rossman, L.A.(1993). *Epanet Users Manual.* U.S. Environmental Protection Agency, Ohio, USA.

Shamir, U. and Howard, C.D.(1979). “An Analytic Approach to scheduling Pipe Replacement”. *Journal AWWA*, 71 (5), 248-58.

Walski, T. M.(1986). “Making Water System Rehabilitation Decisions”, *Water Forum ‘86’*, ASCE, 467-474, 1986.

Walski, T.M.(1987). (ed) *Water Supply System Rehabilitation*. Task Commitee on Water Supply Rehabilitation Systems, ASCE.

Walski, T.M, Brill, E.D., Gessler, J., Goulter, I.C., Jeppson, R.M., Lansey, K., Han-Lin Lee, Liebman, J.C., Mays, L., Morgan, D.R., and Ormsbee, L. (1987). "Battle of the Network Models: Epilogue". *J. Water Resources Planning and Management*, ASCE, 113 (2), 191-203.

Woodburn, J., Lansey, K. and Mays, L.W.(1987). "Model for the Optimal Rehabilitation and Replacement of Water Distribution System Components". *Proc. Nat. Conf. Hydraulic Eng*., ASCE, 606-611.

Appendix 2: Notation

b	=	break rate growth coefficient
Benefit_hyd	=	hydraulic benefits
Benefit_phy	=	physical condition benefits
Benefit_flx	=	flexibility benefit
Benefit_qual	=	quality benefit
$C_b(j)$	=	repair cost for break in pipe j
$Cost_{rep}$	=	break repair cost
diam(j)	=	diameter of new pipe j
F( )	=	function
f( )	=	function
fit(i)	=	fitness of solution i

i	=	integer index
J(t)	=	break rate in year t
j	=	integer index
jpm	=	set of nodes with pressures below minimum
jpx	=	set of nodes with pressures in excess of maximum
Length(j)	=	length of pipe j
Nclass	=	number of classes
Ncl(i)	=	number of individuals in class i
Ndp	=	number of new parallel pipes
Nrl	=	number of renewed or lined pipes
Nrnw	=	number of pipes renewed
Npop	=	population size
n	=	number of solutions
p	=	number of arcs considered
Qj	=	demand at node j
q	=	number of arcs in network
r	=	interest rate
rank(i)	=	ranking of solution (i)
shared_fit(j)	=	shared fitness of individual j
t	=	time
tp	=	present year
tr	=	present year +10
w_h	=	weighting of hydraulic benefits

w_p = weighting of physical condition benefits

w_f = weighting of flexibility benefits

w_q = weighting of quality benefits

α = weighting coefficient for pressure shortfalls

β = weighting coefficient for pressures excesses

Δj = pressure shortfall or excess at node j

Ω = cost-benefit curve

Table 1: Available Pipe Sizes and Alternative Solution Costs

Diam (mm)	New pipe	Clean & line	Replace existing pipe
80	100	85	110
100	175	100	190
150	220	150	240
200	320	220	350
300	550	300	600
400	780	410	850
500	980	500	1050
600	1350	630	1500

Table 2: Options for Each Decision Variable

Decision	Option for an existing pipe	Option for a new pipe
1	Leave as it is	Diam 80 mm
2	Clean and line	Diam 100 mm
3	Duplicate with Diam 150 mm	Diam 150 mm
4	Duplicate with Diam 200 mm	Diam 200 mm
5	Duplicate with Diam 300 mm	Diam 300 mm
6	Duplicate with Diam 400 mm	Diam 400 mm
7	Renew with the same Diam.	Diam 500 mm
8	Renew with next largest Diam.	Diam 600 mm

Table 3: Node Data

Node	Elevation (m)	Demand (l/s)
1	20	15
2	22	15
3	20	30
4	16	70
5	33	20
6	30	55
7	21	25
8	22	15
9	75	Reservoir

Table 4: Pipes Data

Pipe	From node	To node	Length (m)	Diam (mm)	H-W Coef	Break (br/km/yr)	Repair Cost
1	1	2	1300	80	130		
2	1	3	1500	80	130		
3	1	4	1500	80	130		
4	2	5	1700	150	70	0.5	6000
5	3	5	3100	200	100	1.1	8000
6	3	6	1200	150	120		
7	4	8	1000	300	110	0.1	15000
8	4	9	2100	150	100	0.25	6000
9	5	6	1500	150	120		
10	5	7	3600	200	80	1.7	8000
11	6	7	2000	300	100	0.1	15000
12	6	9	3000	200	110	1.02	8000
13	8	7	1000	200	120		
14	7	9	1500	150	90	0.4	6000
15	8	9	1000	200	90	1.5	8000

Table 5: Coefficients Used in the Benefit Calculations

Coefficient	label	value
Hydraulic benefit weighting	w_h	125
Physical integrity benefit weighting	w_p	400
Flexibility benefit weighting	w_f	50
Quality benefit weighting	w_q	0.01
Pressure deficiency weighting .	α	0.6
Pressure excess weighting	β	0.4

Table 6: Non-dominated Solutions.

Cost	Benefit	Minim. Pressure	Min. Press Node	Arc 1	Arc 2	Arc 3	Arc 4	Arc 6	Arc 13	Arc 14	Arc 15
1940000	1433.72	7.47	2	80	80	150	Leave	Leave	300*	Leave	400*
1911000	1428.96	10.28	5	150	80	150	Leave	Leave	Leave	300*	200*
1780000	1417.67	5.85	2	80	80	150	Leave	Leave	Leave	400	Leave
1755000	1414.81	3.82	2	80	80	150	Leave	Leave	Leave	300*	200*
1585000	1388.83	1.00	2	80	80	80	150*	Leave	Leave	300*	Leave
1435000	1369.52	-2.47	2	80	80	150	Leave	Leave	Leave	300*	Leave
1380000	1311.64	-7.06	2	80	80	150	Leave	Leave	150*	Leave	300*
1367500	1305.26	-15.00	2	80	80	100	Leave	Leave	Leave	300*	Leave
1255000	1260.33	-27.38	2	80	80	80	Leave	Leave	Leave	300*	Leave
1160000	1125.96	-13.54	2	80	80	150	Leave	Leave	Leave	Leave	300*
930000	1030.88	-30.02	2	80	80	150	Leave	Leave	Leave	Leave	200*
750000	848.15	-60.56	2	80	80	80	Leave	Leave	Leave	Leave	200*

* Add parallel pipe of diameter shown.