A general multi-objective hyper-heuristic for water distribution network design with discolouration risk

Kent McClymont, Ed Keedwell, Dragan Savić and Mark Randall-Smith

ABSTRACT

The optimisation of water distribution networks (WDNs) by evolutionary algorithms has gained much coverage in the literature since it was first proposed in the early 1990s. Despite being well studied, the problem and objectives continue to evolve as demands on water companies change. Motivated by the increased focus on reducing the risk of discolouration, this study examines a three objective version of the WDN design problem which takes into account cost, head excess and discolouration risk. Using this formulation, this paper presents a method for producing optimised network designs aimed at reducing discolouration risk in the network design phase and thus reducing the associated long-term maintenance and operational burdens of the system. This paper discusses the use of a discolouration risk model and, using this model, the optimisation of network design, specifically pipe diameters, to produce a range of high quality self-cleaning networks. The network designs are optimised using the Markov-chain hyper-heuristic (MCHH), a new multi-objective online selective hyper-heuristic. The MCHH is incorporated in to the known NSGA-II and SPEA2 and supplied with a range of heuristics tailored for use on the WDN design problem. The results demonstrate an improvement in performance obtained over the original algorithms.

Key words | discolouration, heuristic, hyper-heuristic, multi-objective, optimisation

Kent McClymont (corresponding author)
Ed Keedwell
Dragan Savić
College of Engineering,
Mathematics and Physical Sciences,
University of Exeter,
Exeter,
UK
E-mail: K.McClymont@exeter.ac.uk

Mark Randall-Smith
Mouchel,
Clyst Works,
Clyst Road,
Topsham,
Exeter EX3 0DB,
UK

INTRODUCTION

The UK water industry is tightly regulated by Ofwat, the UK regulatory body, with the performance of water companies closely monitored by a range of indicators; from water quality, customer service (e.g., sufficient pressure) to customer satisfaction. Recently, an emphasis has been placed on customer satisfaction in particular, which is partly measured by monitoring customer complaints and contacts. Motivated by these regulatory demands, water companies are now focusing efforts on reducing the frequency of water discolouration events (Cook 2007) prior to customer contacts occurring. Indeed, discolouration events (the visible discolouration of water at the tap) have been attributed to approximately 30% of all complaints received by water companies in the UK (Cook 2007).

It should be noted, however, that the drive to reduce the number of discolouration events is not isolated to the UK, but experienced in many countries, worldwide. Importantly,

the phenomenon does not appear to display regional variances, outside of temperature and material differences, with Boxall & Prince (2006) demonstrating the validity of UK models abroad. As such, it can be supposed that any advancement relating to discolouration modelling may have world wide application.

However, despite its apparent prominence in industry, discolouration risk is seldom included in studies that optimise the design of water distribution networks (WDNs). In this paper we investigate the use of discolouration risk as an objective in a multi-objective algorithm and propose a new hyper-heuristic called the Markov-chain hyper-heuristic (MCHH) for the multi-objective optimisation of WDNs. The proposed hyper-heuristic is able to adapt its behaviour during optimisation to take advantage of the best performing heuristics for a given problem and could incorporate expert-derived problem-specific heuristics (although fairly generic

heuristics are studied here). The use of discolouration risk and the MCHH together demonstrate a powerful new direction for WDN optimisation that can incorporate real-world business considerations and expertise in the optimisation process.

Paper outline

A key aim of this study is to examine the impact of pipe diameters on discolouration risk and demonstrate a method to reduce that risk in parallel with optimising network design and/or rehabilitation cost whilst meeting head requirements in a novel consideration of the WDN design and rehabilitation problems. The 'Background' section discusses various relevant discolouration risk models from the literature. One of these models, the Cohesive Transport Model (CTM), is used in this study as an objective in the proposed multi-objective WDN problem. The interaction between the problem's competing goals is investigated by using the traditional objective of minimising cost and discolouration risk. Discolouration risk is included to represent each design's impact on the propensity for pipes to accumulate material and thus generate discoloured water. However, a natural by-product of reducing discolouration risk may be the increase in head losses and thus reducing pressures within the system, so characteristics such as velocity and node head have also been considered as constraints in the optimisation process. There then follows a section which introduces the WDN design problem and discusses some relevant approaches from the literature.

In order to optimise the problem, the MCHH, a newly proposed multi-objective selective hyper-heuristic from the field of optimisation (McClymont & Keedwell 2011) is described and is applied to the WDN design and rehabilitation problem. A review of hyper-heuristics is also presented. The proposed hyper-heuristic is designed in an abstract fashion to facilitate simple integration into existing optimisers and allow for the tailored heuristics to be changed and adapted for each new problem to be solved. A range of low-level heuristics are used by the hyper-heuristic, each tailored for use in WDN design/rehabilitation problems. These heuristics include classic evolutionary algorithm (EA) operators such as mutation and crossover.

Designs are evolved using the well known NSGA-II (Deb *et al.* 2003) and SPEA2 (Zitzler *et al.* 2002) in addition to MCHH variants of both these algorithms. The problem is

modelled using EPANET (Rossman 2000) and a state-of-the-art model of stored material using shear stress (Boxall *et al.* 2001; Boxall & Saul 2005). The experiment is conducted on three benchmark datasets (*Two Loops*, *Hanoi* and *Anytown*) and three real-world networks, described in the Results section. Optimisation results are compared with hyper-heuristics from the literature as well as results presented in (Shie-Yui Liong & Atiquzzaman 2004) who contrast a shuffled complex evolutionary (SCE) approach to a variety of preceding evolutionary approaches on these datasets. Finally, the Conclusion section discusses the optimisation results which reveal a correlation between cost and discolouration potential in addition to the trade-off with head excess.

BACKGROUND

Modelling discolouration

Despite the varying approaches, only a handful of discolouration risk modelling software packages are available to water companies at present. The most notable of which are: the PODDS I-IV models developed by the University of Sheffield (Boxall *et al.* 2001; Boxall & Saul 2005); the DRM software developed by Mouchel (Dewis & Randall-Smith 2005); and prototype models founded on discolouration risk ranking (Vreeburg *et al.* 2005).

The CTM (Boxall *et al.* 2001; Boxall & Saul 2005) provides a method for calculating the volume of accumulated material within a network, which is measured as turbidity (expressed in Nephelometric Turbidity Units, NTU). The model is also used to calculate the volume of material mobilised given specific hydraulic events in the network. In theory, the impact of mobilised material is proportional to and limited by the material stored in the network and thus by reducing the potential material in the network the associated discolouration risk will also decrease. Taking that relationship into consideration, this paper demonstrates the process of evolving networks that have shear stress characteristics which encourage self-cleaning and, therefore, the prevention of discolouration events by reducing the material stored in the network. The notion of self-cleaning thresholds have been investigated a number of times by Boxall & Prince (2006)

who consider shear stress thresholds and Buchberger *et al.* (2008) who consider velocity thresholds.

The PODDS models are based upon research into the CTM (Boxall *et al.* 2001; Boxall & Saul 2005) which analyses discolouration risk using a shear stress approach. This study utilises a software implementation of the CTM, called the Discolouration Propensity Model (DPM) that builds upon Mouchel's DRM software and incorporates the CTM.

The implementation of the discolouration model used in the formulation of the WDN design problem is based upon the following equations. The most pivotal aspect of the model's design is the shear stress equation:

$$\tau = \rho g R_h S_0 \tag{1}$$

where τ is shear stress, ρ is water density, g is gravitational acceleration, R_h is hydraulic radius, and S_0 is hydraulic gradient. The CTM applies each link's maximum daily shear stress values (known as the daily conditioning shear stress; Cook 2007) in the calculation of the potential material stored in each link. Using the shear stress obtained above and the relationship

$$\tau'_s = \frac{c^b - c_{\max}}{k} \tag{2}$$

given in Boxall & Saul (2005), it is possible to derive the stored material C (measured as turbidity in NTU) for each pipe in a network, where: $C_{\max}$ is the maximum possible material stored in the pipe; τ'_s is the layer strength; and k and b are calibrating constants.

The daily conditioning shear stress values are calculated based on the hydraulic gradient values produced by an underlying hydraulic model – in this case the well known EPANET (Rossman 2000). Appropriate NTU values are then calculated given the proper calibration of the constants k and b. Using these NTU values, a risk score can be assigned to each pipe. The k and b values used in the model were set to 2 and 1, respectively; based upon calibration data given in Boxall & Saul (2005). To more accurately determine the k and b constants for a new network, a process of flushing and sampling is advised to obtain accurate turbidity measures for each specific network.

WDN design problem

WDNs are built by water companies in order to provide water services to the end users with the aim of satisfying their demand. A WDN is comprised of pipes, nodes (junctions and demand points), hydraulic devices (such as pumps) and sources (tanks and reservoirs) and constitutes the infrastructure that delivers water from the source (e.g., reservoir) to various locations where it is drawn from the network for consumption (e.g., residential housing or industrial sites).

Real world WDNs are, more often than not, large and complicated structures which are commonly interlinked with other neighbouring networks. Indeed, the set of national WDNs in the UK alone represents a significant infrastructure that requires constant operational management, maintenance and rehabilitation. In order to satisfy consumer demand, the networks must be constructed with a good layout that connects to all points of demand and designed to provide the best possible hydraulic conditions and operational requirements. In essence, the creation of a WDN can be broken down into three constituent parts: layout, design and operation.

Each of these parts represents complex individual problems that require significant expertise and resources to solve. Although, it should be noted that decisions made for one problem has an effect on the next; i.e., selecting a network layout will affect the possible design choices and thus overall cost of the network. However, attempting to optimise all three parts simultaneously would represent a significantly more complex task and would need to take into account a much larger range of decision maker considerations which are not normally modelled in the individual problems. For example, the design problem is primarily concerned with the sizing (diameters) of pipes in the network. Changing pipe sizes effects the hydraulic conditions in the network and hence the quality of the network based on its ability to serve the various demand points.

The problem is complex as the overall hydraulic conditions are affected by each pipe and so changes to one pipe will have a different effect on the overall conditions depending on the sizes of all the other pipes in the network. As such, each pipe cannot be designed in isolation, but rather as a combination of sizes for all pipes in the network. Furthermore, even for relatively small networks, the number of possible combinations of pipes is very large which makes enumeration of all the possible designs impossible within reasonable time. If, for example, there were six potential

sizes for each pipe in a network of just thirty pipes, there would be $6^{30} = 2.21 \times 10^{23}$ possible combinations – far more than is possible to compute within a reasonable time; the relationship between the number of pipes and the increase in search space means that WDN optimisation is known to be a NP-hard problem (Yates *et al.* 1984).

The non-linearity of the hydraulic equations also has an effect on the complexity of the search space, creating a multi-modal landscape. Multi-modal problems are particularly difficult as good network designs are separated by regions of less-good or infeasible network designs. As a result, early methods like linear programming (Schaake & Lai 1969) and hill-climbing algorithms are not as effective at solving these problems and often get stuck in local optima: i.e., the best network design in the local area of the search space but not optimal in the context of the total (global) search space.

WDN optimisation

The optimisation of WDNs by EAs has gained much coverage in the literature since first proposed in the early 1990s (Murphy & Simpson 1992; Simpson *et al.* 1994). Traditionally, the WDN design problem has been formulated as a single objective problem where the quality of the network is based solely on the economic impact of the design, i.e., given a fixed layout, the optimal network design is one which meets the hydraulic requirements with the least possible cost (Savic & Walters 1997). The hydraulic constraints are usually given as an acceptable range of node pressures and/or pipe velocities. Single-objective EAs, in particular, were shown to solve this problem effectively (Savic & Walters 1997).

However, since the original problem of minimising cost has been demonstrated as being efficiently solved by the early EA methods and, in combination with the emergence of the multi-objective approach, multi-objective formulations of the problem are now more commonly used. These multi-objective approaches are able to provide a trade-off surface for the many competing aspects of new WDN designs and take into account an assortment of different factors influencing WDN design; such as reliability, operational and maintenance costs and water quality (e.g. Formiga *et al.* 2003; Farmani *et al.* 2005). However, as the needs of water companies change and the measures of a

network's performance evolve, new formulations of the problem will continue to be proposed and thus any proposed method for solving the problem should be as flexible as possible to accommodate future variations of this traditional problem. Indeed, these multi-objective approaches consider many more objectives and illustrate the many additional factors that influence the design of a WDN, such as discolouration risk.

WDN problems have been solved with a wide and varied selection of optimisation techniques, from early approaches, such as (Savic & Walters 1997) that used traditional EAs, to more complex approaches with online learning (di Pierro *et al.* 2009) and, more recently, hybrid methods like a multi-algorithm genetically adaptive multiobjective (AMALGAM) method (Raad *et al.* 2010). Indeed, AMALGAM is a good example of the more recent trend to examine self-optimising search strategies such as memetic algorithms (Ong *et al.* 2006) and hyper-heuristics (Burke *et al.* 2003). Being an intentionally generalised method, hyper-heuristics in particular show great promise in providing an effective means of augmenting traditional and well known optimisation methods like NSGA-II (Deb *et al.* 2003) and SPEA2 (Zitzler *et al.* 2002) to further improve the efficiency of the optimisation search process whilst also ensuring a robust method that can more easily be applied to new forms of the WDN problem. The hyper-heuristic approach aims to achieve this improvement through the automated process of selecting and generating heuristics for hard computation problems (Burke *et al.* 2009), which enables the methods to adapt as well as removing much of the burden of parameter tuning from user.

Hyper-heuristics

Although a number of traditional meta-heuristic techniques have been applied to the WDN design problem, such as differential evolution (Vasan & Simonovic 2010; Suribabu 2010), these methods often rely on specific types of operations to work effectively. In contrast, hyper-heuristics are methods which are designed to automate the process of selecting heuristic operators (like mutation and crossover) which greatly improves the efficiency of the search and removes the need to strictly specify the type of operations used to create new solutions. Conceptually, hyper-heuristics

can be thought of solving the problem of heuristic selection by using 'heuristics to choose heuristics' (Cowling *et al.* 2000).

Indeed, traditional meta-heuristic methods, like EAs are essentially iterative processes that use fixed strategies for generating new solutions (i.e., network designs) to a problem that sequentially employ low-level heuristic operators, such as crossover and mutation, to enable the search of the space of possible solutions in order to find optimal solutions to the given problem (Deb 2001). In contrast, by supplying an online selective hyper-heuristic with a wider range of heuristic operators, for example a set of different mutation operators, the algorithm is able to learn and select the most appropriate heuristics for the current problem at run time and so improve the efficiency of the search process. In effect, the selective hyper-heuristic decides which heuristic to apply next at each generation in the search process, allowing for the possibility of applying mutation for a few generations before then applying crossover and *vice versa*. The hyper-heuristic approach has been demonstrated in the literature to be very effective at solving a range of real-world optimisation problems (Burke *et al.* 2010); especially problems from operations research.

A key benefit of the abstract nature of hyper-heuristic selection methods is that they are distinct from the selection and propagation of good solutions in a population. For example, in addition to creating new solutions with mutation and crossover heuristics, EAs use selection strategies to determine which solutions will be propagated from one generation to the next. These two parts, solution generation and solution selection, are distinct and are often contained within different program modules. As a result, the modular design of EAs makes them an ideal base algorithm that can be easily modified to include higher-level optimisation elements, such as selective hyper-heuristics. This combination of generic hyper-heuristics embedded in efficient EA with highly specialised low-level heuristic operators makes hyper-heuristic methods well placed for solving difficult problems like WDN design will minimal effort.

However, whilst hyper-heuristics have been shown to be very effective at solving a range of single-objective combinatorial optimisation problems (Burke *et al.* 2010), few examples of multi-objective approaches exist in the

literature. Indeed, much of the hyper-heuristic literature is focused on demonstrating the efficacy of new methods on single- or bi-objective (which has unique properties when compared with higher objectives) combinatorial problems, with few examples of higher order multi-objective methods such as Burke *et al.* (2005). Furthermore, those methods that do exist are often focused on the generation of low-level heuristics (known as generative hyper-heuristics) rather than methods to select heuristics during optimisation (selective hyper-heuristic), such as McClymont & Keedwell (2010).

METHOD

The novel formulation of the multi-objective WDN design problem and a novel online selective hyper-heuristic (MCHH) is presented in this section.

Multi-objective WDN design problem

As outlined earlier, the formulation of the WDN design problem used in this study incorporates the evaluation of a network's propensity to cause discolouration events. This was done by using discolouration risk as an objective which was calculated using the DPM software based on the CTM model (Boxall *et al.* 2001; Boxall & Saul 2005), where the EA is to minimise this potential for pipes to store material.

In order to fairly represent the complete effect of changing pipe diameters in the network, the discolouration risk objective was combined with the traditional objective of minimising cost in addition to minimising head excess. These three objectives are:

1. Cost of network infrastructure:

$$f_0 = \sum_{i=0}^{L} f_{\text{cost}}(\text{diameter}_i, \text{length}_i) \tag{3}$$

2. Sum of cumulative potential material after daily conditioning shear stress for all pipes in the network (L):

$$f_1 = \sum_{i=0}^{L} f_{\text{turbidity}}(\tau'_{s_i}, R_{\text{h}}, S_0) \tag{4}$$

3. Sum of the cumulative head excess (head over 35 m):

$$f_2 = \sum_{i=0}^{L} \max(0, \text{head}_i - \max \text{Head}) \qquad (5)$$

The $f_{\text{turbidity}}$ function returns the stored material in a pipe, given the shear stress (τ'$_s$), hydraulic radius (R_h) and hydraulic gradient (S_0). The maxHead was set at 35 m, where F_3 returns the excess in recorded head (head_i). The cumulative head excess was calculated by summing, for every hourly time step, the excess on each demand node above a given maxHead threshold value.

The f_{cost} function calculates the cost of the network design based on the length and diameter of the pipes in the network. The cost tables for the three benchmark networks are used from the original papers and are available online at http://centres.exeter.ac.uk/cws/. The total cost of the network infrastructure assets (pipes) was calculated for all pipes in the network. For Two Loops and Hanoi, it was assumed that the entire network was to be replaced or, perhaps more appropriately, newly constructed. This allowed the optimising algorithm to take into account the overall cheapest network, as opposed to the cheapest network rehabilitation relative to an initial design. In contrast, for Anytown and the three real-world networks, only specific pipes were able to be resized, basing the problem on the network rehabilitation scenario. For these networks, the first objective (cost) was only calculated for the new pipes with all other objectives remaining the same; i.e., cost was calculated for new pipes but discolouration and head calculated for all nodes in the network to reflect the hydraulic changes experienced across the whole network.

In addition to the three objectives outlined above, a node head constraint was placed on all solutions. If the network design resulted in any head deficit (a node with head below 30 m), then a penalty was added to each of the three objectives. The penalty was equal to the maximum value for each objective. In effect, this ensured that all feasible networks always dominated networks with head deficit and provided a strong selection pressure for feasible solutions and prevented the algorithms wasting evaluations searching infeasible regions of the search space. A second constraint was also added for excess velocity (over 2.5 m s^{-1}). Again, if the network contained pipes with excess velocity, the same penalty was added to each objective. If both constraints were broken, then two penalties would be added.

The Markov-chain Hyper Heuristic (MCHH)

Hyper-heuristics are designed to operate independently of any specific problem or set of heuristics. Instead, hyper-heuristics aim to optimise the selection heuristics for the current problem, effectively acting as 'heuristics to choose heuristics' (Cowling *et al.* 2000). The MCHH is an online approach that can be integrated into existing meta-heuristics, such as an EA. The MCHH uses meta-data collected during the search, such as the number of dominating solutions generated by a heuristic, to learn the most effective heuristics for the current problem. As shown in Figure 1 below, when incorporated in an EA, the MCHH is applied after the new generation of solutions have been created and evaluated.

The steps used in the MCHH are shown below in Figure 2, which outlines the MCHH incorporated in a $(\mu + \lambda)$ Evolution Strategy (Laumanns *et al.* 2000). The steps in 2.1 relate to the selective hyper-heuristic operations, operating on meta-data that quantifies the quality of each heuristic's performance, given in more detail below. In effect, the MCHH evaluates how effective each heuristic is at generating new solutions and adjusts the likelihood of
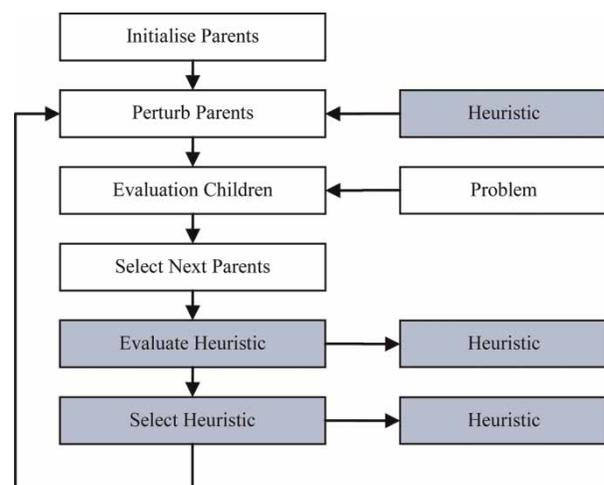

Figure 1 │ Schematic representation of an EA with embedded online selective hyper-heuristic processes (shaded).

```
1.      Initialise parent population (μ)
2.      Repeat:
   2.1.        If ε generations since last episode then
        2.1.1. Calculate performance (p) of current heuristic
        2.1.2. If p > γ then increase the weight from last heuristic to current heuristic by α
        2.1.3. Else if p < γ then decrease the weight from last heuristic to current heuristic by β
        2.1.4. Select next heuristic
   2.2.        Vary parents (μ) using current heuristic to generate children (λ)
   2.3.        Evaluate children (λ)
   2.4.        Select parents for the next generation (μ′) from union of parents (μ) and children (λ)
   2.5.        Update the archive (ω) with children (λ)
```

Figure 2 │ Pseudocode for the Markov chain hyper-heuristic (MCHH) algorithm incorporated in a $(\mu + \lambda)$ Evolution Strategy (a simple EA).

selecting that heuristic in future generations based on this score.

In Figure 2, the terms μ and λ represent the parent and child populations. p is the performance score of the current heuristic, given below, γ is a performance threshold that controls the adaption of the weights, and α and β are the reward and penalty scores applied to the weights, respectively.

Heuristic performance

The performance measure calculated in step 2.1.1 (Figure 2) of the MCHH is outlined briefly below. Although a large number of studies have been conducted on hyper-heuristics, the majority of these are focused on single-objective problems. As such, few frameworks exist which evaluate heuristic performance on multi-objective problems. Examples from the literature include objective specific learning methods, such as the TSRoulWheel selective hyper-heuristic (Burke *et al.* 2005). Whilst the TSRoulWheel method is shown to be effective in the original study, and to be reasonably good in an initial study of the MCHH (McClymont & Keedwell 2011), it relies on improvements in one objective resulting in improvements in another. This is because it evaluates each objective in turn, effectively formulating the problem as a set of single-objective sub-problems. However many real-world problems include objectives with varying degrees of correlation, such as the formulation of the WDN design problem used in this study, making the TSRoulWheel less amenable to this type of problem.

The MCHH employs a more general approach based on Pareto optimality, a fundamental technique used in the majority of multi-objective EAs. The Pareto dominance relationship is used to assign a quality performance measure to each heuristic, calculating the ratio of dominating solutions produced by the heuristic and indicating the chances of generating dominating solutions when applied in future iterations. The score is calculated, for each solution in the new child population, as the ratio of solutions it dominates in the parent population and then averages these ratios to produce a single score. This is shown in Equation (6). Theoretically, good heuristics (for moving towards the Pareto front) will have a high probability of generating dominating solutions.

$$p(h, \mu, \lambda) = \frac{\sum \forall_{a \in \lambda} \forall_{b \in \mu} dom(a, b)}{|\mu| \times |\lambda|}$$

$$\text{where } dom(a, b) = \begin{cases} 1, a \prec b \\ 0, a \nprec b \end{cases} \tag{6}$$

The function $p(h, \mu, \lambda)$ shown in Equation (6) returns the average ratio of parent solutions μ dominated by each child solution in λ produced by heuristic h. The terms a and b refer to an individual child and parent, respectively, whilst the function $dom(a, b)$ returns an integer (0 or 1) value indicating whether a dominates b.

Markov chain

As implied by the name, the MCHH uses a Markov chain to guide the selection of heuristics and applies online reinforcement learning to adapt the transition weights in the Markov chain. By using a Markov chain to control the selection of heuristics and adapting the transition weights from one heuristic to another, the MCHH is able to not only learn which heuristics are effective, but what sequence of heuristics are most effective. For example, a heuristic may be good in general but a combination of two other heuristics, when applied in a specific sequence, may perform even

better. The approach is designed to try and learn these transition sequences to further improve the optimisation process. Firstly, the MCHH constructs a fully connected Markov chain with one state for each heuristic, i.e., each state in the chain is connected to every other state and to itself (see Figure 3). The weight of each edge out of a state represents the probability of moving from the current state (the current heuristic) to the destination state (the heuristic to be applied subsequently), where all edges out of each state sum to one.

The MCHH traverses this Markov chain by stochastically selecting the next heuristic using roulette selection biased by the outbound edge weights. Each heuristic is applied ε times (set to two in the experiments below) before selecting the next heuristic. The next heuristic is then applied ε times before again selecting another heuristic, and so on. At the end of each episode (ϵ applications of a heuristic) the quality score is calculated and the weight of the last edge traversed by the MCHH, the edge used to move to the current heuristic, is updated. The transition weights in Markov chain can also be represented as a matrix of $m \times m$, shown in Figure 4, where m is the number of heuristics. The row represents the current state and the columns the potential state to move to.

After applying the heuristics for one episode of ε applications, the heuristic's quality is calculated. If the resulting score is greater than some threshold γ, the weight corresponding to the last transition (made to get to the current operator) is increased by α, otherwise, the weight is degraded by β. Once the weight has been adjusted, the sum of outflow edges from the previous state are normalised to 1 to maintain the fidelity of the matrix. After normalising, the effect of increasing or decreasing an edge in the Markov chain will decrease or increase the other edges, respectively.
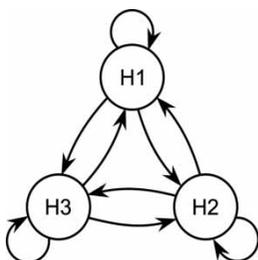

Figure 3 | Example Markov chain with three states representing three heuristics.

Figure 4 | Transition weight matrix for a Markov chain with three states representing three heuristics.

The repetition of this process should allow the matrix to converge on a set of probabilities for moving between individuals in the set of heuristics. This process identifies the good links between heuristics, with sequencing controlled by the edge direction, giving probabilistic information about combinations of heuristics.

EXPERIMENTAL SETUP

This section describes a novel application of online selective hyper-heuristics, such as the MCHH, to the multi-objective WDN design problem.

Experimental data

Experiments to demonstrate the benefits of considering discolouration risk during the optimisation run and examine the efficacy of the MCHH were conducted on three well-known benchmark networks (Two Loops, Hanoi and Anytown) in addition to three real-world networks. The Two Loops network consists of eight links which connects the six nodes and a reservoir. The Hanoi network consists of 34 links which connects the 32 nodes and a reservoir. The Anytown network consists of 1 reservoir, 1 pumping station, 2 tanks, 22 nodes, and 42 links. However, only a selection of links in the Anytown network were available for resizing, treating the problem as a rehabilitation test case. The tanks and pumps in the Anytown network were fixed to the initial setting. The input files and cost models are available at http://centres.exeter.ac.uk/cws/.

The three real-world networks are located in the South West of England. The smallest network contains 1 reservoir, 52 junctions and 68 pipes, the medium sized network

contains 1 reservoir, 1 tank, 81 junctions and 107 pipes, and the largest network consists of 2 reservoirs, 160 junctions and 213 pipes.

In addition to applying the MCHH, the experiment was used to study the effect of pipe diameter on network self-cleaning characteristics subject to node head constraints and as such so no other characteristics were altered during the optimisation process. The experiment used EPANET to simulate the hydraulic effects of pipe diameter changes over a 24 hour extended period. The DPM was then applied to calculate the potential material stored in each pipe given the revised hydraulic conditions. The hydraulics, potential for discolouration and cost calculation formed the basis of the objective function for the optimising algorithms.

Optimisers

The optimisation of the pipe diameters was completed using a NSGA-II and SPEA2, two modified variants of these algorithms which included the MCHH operations as well as Simple Random and TSRoulWheel. NSGA-II and SPEA2 were given populations of 10 and run for 2,000 generations for both the normal optimisation runs and the optimisation runs with the MCHH. The parameters for the MCHH were set as follows: $\gamma = 0.2$, $\alpha = 0.1$, $\beta = 0.1$ and $\varepsilon = 2$. A schematic view of the optimisation process is given in Figure 5.

As shown in Figure 5, both Cost and Potential are calculated using the EPANET output and Head Excesses and Violations directly from recorded Head values. The normal EA elements (i.e., NSGA-II operations) shown in solid boxes. Hyper-heuristic elements (heuristic selection) shown in dashed boxes and only included in the MCHH variants.
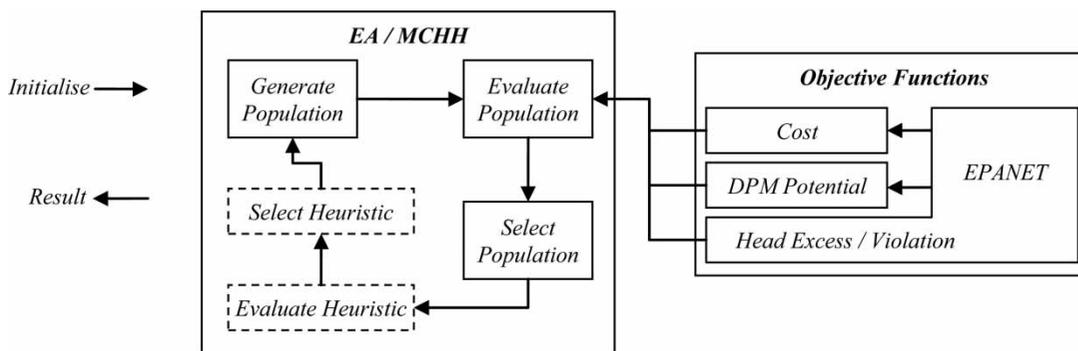
Hyper-heuristics for comparison – Simple Random

Perhaps the most basic selective hyper-heuristic is the 'Simple Random' (Burke *et al.* 2005) which randomly selects a heuristic at each generation. Although simple, even this basic approach can provide significant improvement in performance over meta-heuristics that employ a fixed heuristic. However, it is assumed that although Simple Random performs well on many problems, a more sophisticated strategy should be able to provide further improvement in performance. Simple Random has been included in this study to provide a benchmark of performance for the MCHH and TSRoulWheel (Burke *et al.* 2003). Simple Random was embedded in NSGA-II using the same approach as the MCHH; selecting random heuristics at each generation. The algorithm was also used as a control to demonstrate that the heuristics alone did not provide the significant improvement in performance observed in the MCHH variants.

Hyper-heuristics for comparison – TSRoulWheel

Although many hyper-heuristic algorithms have been proposed in the literature, there are few that can be applied to the online selection of heuristics on multi-objective problems. However, of the few hyper-heuristic methods that can be applied to multi-objective problems, the TSRoul-Wheel (Tabu Search Roulette Wheel) algorithm is perhaps the most notable and general multi-objective scoring method from the literature (Burke *et al.* 2003). In TSRoul-Wheel, each heuristic is rated on its performance on each objective, which informs the update of selection weights

Figure 5 │ Schematic view of the experimental design.

by reinforcement learning, i.e., for a problem with m objectives and h heuristics, a weight matrix of $m \times h$ is produced. The algorithm first selects an objective to use as a basis from which it then selects heuristics, based on each heuristic's weight for that specific objective. Whilst this method is shown to be the most effective examined in Burke *et al.* (2003), and to be a reasonably good method in this study, it relies on improvements in one objective resulting in improvements in another as it evaluates each objective in turn, formulating the problem as a set of single-objective sub-problems. The TSRoulWheel hyper-heuristic was implemented as directed in Burke *et al.* (2003) with a learning rate of 1.

Heuristics

Four extra heuristics were supplied to the MCHH variants of the EAs in addition to the normal single-point additive Gaussian integer mutation operator ($\sigma = 1$) and the uniform crossover operators used in both the original algorithms. Three alternative parameterisations of the Gaussian integer mutation ($\sigma = 0.01, 0.075, 0.75$) and a simple replacement heuristic was used that created a new random resampled solution rather than perturbing an existing solution.

Performance measure

The hypervolume indicator (Bader *et al.* 2008) (which was normalised to 1) was used to monitor the performance of

each of the six algorithms, calculated using random samples drawn from within the objective space. The hypervolume indicator gives a scalar representation of the ratio of objective space dominated by the population, illustrated in Figure 6.

The hypervolume was calculated by sampling in the range of possible values for each objective and calculating the ratio of sampled points that were dominated by the population (those in the shaded region). The hypervolume was normalised by dividing the number of sample points dominated by the population by the total number of sample points. Once a sample set had been generated it was kept and used for all hypervolume calculations on that problem for all algorithms and trials. Each of the four algorithms was run 20 times and the hypervolume results averaged to ensure a fair comparison of performance.

RESULTS

The results for the three benchmark and three real-world networks are given in Table 1 which shows the best result obtained by each algorithm on each objective. Each algorithm returned a Pareto front (population) of the most optimal solutions it was capable of locating. The Cost, Discolouration Risk and Cumulative Head Excess columns report, for each algorithm on each network, the best value obtained across the population. For example, the Hanoi results for NSGA-II represent the best cost, discolouration risk and head excess from two different solutions. The
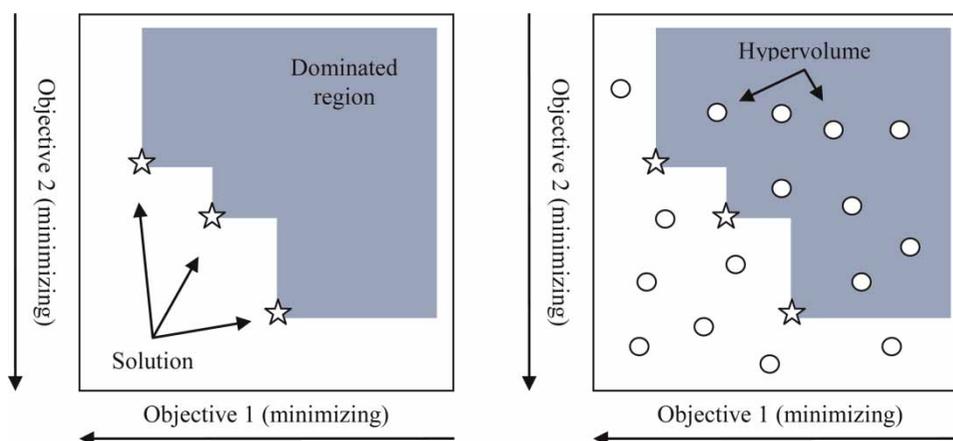

Figure 6 │ Illustration of the calculation of hypervolume using sampling.

Table 1 | Results from two meta-heuristics and two MCHH variants on the six test networks – best result shown in bold. Each algorithm is broken down into the best value obtained for each of the three objectives on each network; Cost, Discolouration Risk (Disc.) and Cumulative Head Excess (H.) Cost is given in USD (US$), Discolouration Risk in expressed in Nephelometric Turbidity Units (NTU) and Head Excess in meters (m)

	NSGA-II			SPEA2			MCHH (NSGA-II)			MCHH (SPEA2)		
	Cost	Disc.	H.	Cost	Disc.	H.	Cost	Disc.	H.	Cost	Disc.	H.
Two Loops	419,000	0	33	419,000	0	33	419,000	0	33	419,000	0	33
Hanoi	6.14×10^6	5.1×10^5	164	6.099×10^6	5.93×10^5	164	6.073×10^6	5.08×10^5	164	6.073×10^6	5.32×10^5	164
Anytown	1.34×10^7	2.12×10^6	223	1.65×10^7	2.21×10^6	218	1.34×10^7	2.022×10^6	209	1.46×10^7	2.01×10^6	218
Industry 1	3.42×10^7	2.42×10^7	263	3.56×10^7	2.38×10^7	246	3.22×10^7	2.33×10^7	214	3.29×10^7	2.54×10^7	214
Industry 2	5.21×10^7	7.76×10^7	198	5.25×10^7	7.82×10^7	215	5.012×10^7	7.043×10^7	191	5.17×10^7	7.26×10^7	191
Industry 3	9.89×10^7	1.74×10^8	378	9.73×10^7	1.75×10^8	389	9.71×10^7	1.42×10^8	348	9.65×10^7	1.44×10^8	360

numbers in bold indicate that the value was the best found for that object on that network across all the algorithms, e. g., NSGA-II and MCCH variant of NSGA-II both found the lowest cost solutions for Anytown.

The results given in Table 1 clearly shown an improvement in the best possible objective values obtained by the MCHH variants of NSGA-II and SPEA2. One of the two MCHH variants obtains the best objective value for all three objectives over all six problems, while the original NSGA-II and SPEA2 algorithms only locate best values on some of the objectives on the smaller, benchmark networks.

More detailed analysis of selected problems is given below in Figures 7, 8 and 9. The results discussed were obtained using the MCHH embedded in NSGA-II, and are used to illustrate the benefits of including discolouration risk modelling into the design of new and rehabilitation of existing WDNs. An analysis of the performance of the MCHH on these problems compared with two hyper-heuristics from the literature is given in the Conclusions section below.

The MCHH on WDN problems

As shown in the Hanoi results (Figure 7), the MCHH variants converge more quickly when compared with the original NSGA-II and SPEA2 algorithms. In addition to converging quickly (shown in the generational plot in Figure 7), both MCHH variants also consistently produce better final generation results – illustrated in the box plots in Figure 7 that show the distribution of results. Due to the short number of generations and limited number of evaluations, only the NSGA-II and SPEA2 variants of the MCHH located the known optimum of 419,000 for Two Loops. Simple Random and TSRoulWheel both fail to locate good solutions and so have poor (near to zero) hypervolume throughout the search.

Despite the poor performance on Hanoi, the performance of both Simple Random and TSRoulWheel improves significantly on the rehabilitation problems, as illustrated in the Anytown (Figure 8) and real-world network 2 (Figure 9) results. Indeed, TSRoulWheel performs very well on the Anytown problem and outperforms NSGA-II. While NSGA-II produces good results on the simple problems, such as Hanoi, as the complexity of the problem increases by the introduction of larger number of pipes, SPEA2 is shown to
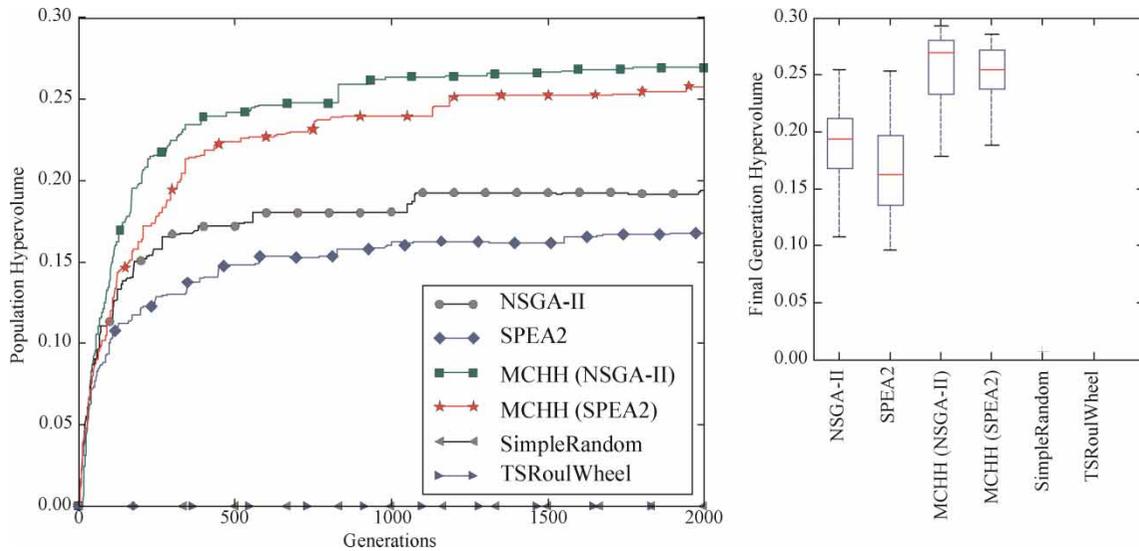
Figure 7 │ Two plots showing hypervolume results NSGA-II, SPEA2, both MCHH variants, Simple Random and TSRoulWheel on the Hanoi benchmark network. Left plot shows hypervolume over generations (averaged over all trial runs), right plot shows distribution of final generation hypervolume for all trial runs. Simple Random and TSRoulWheel are shown at the bottom of the figure along the *x*-axis.
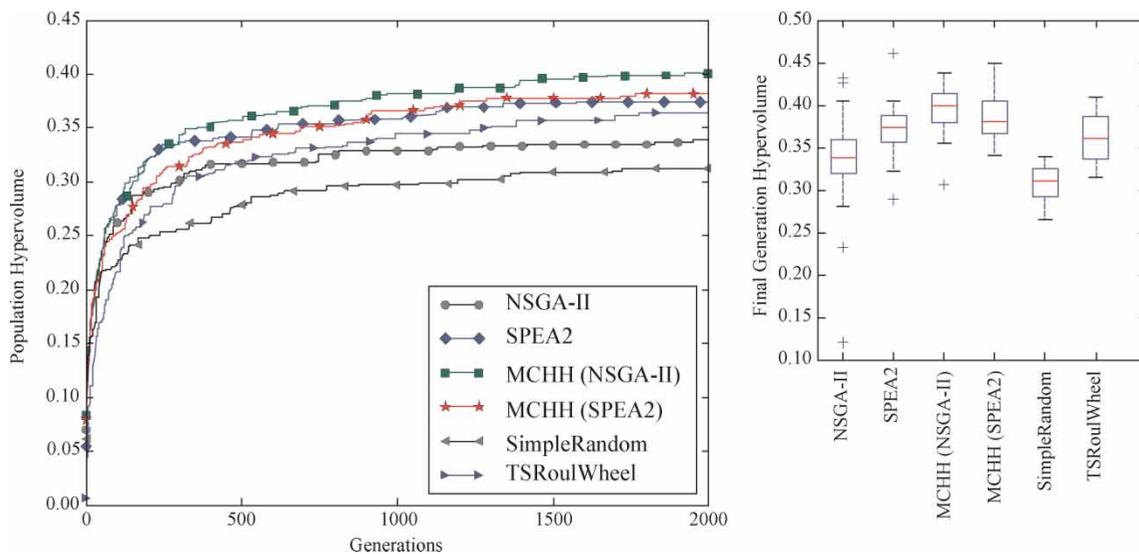

Figure 8 │ Two plots showing hypervolume results NSGA-II, SPEA2, both MCHH variants, Simple Random and TSRoulWheel on the Anytown benchmark network. Again, the left plot shows hypervolume over generations (averaged over all trial runs), right plot shows distribution of final generation hypervolume for all trial runs.

consistently outperform NSGA-II – again shown in Figures 8 and 9. Interestingly, this relationship is not mirrored in the MCHH variants, with the NSGA-II variant of the MCHH achieving the best results across all the problems. This might indicate a stronger tie between the variation and selection operations of SPEA2 compared with those of NSGA-II.

A common problem in optimisation is early convergence: converging on a poor solution quickly. In the experiments conducted in this study, the MCHH hyper-heuristic variants not only converged on a solution more quickly than the original algorithms but also located a better range of solutions, dominating a larger proportion
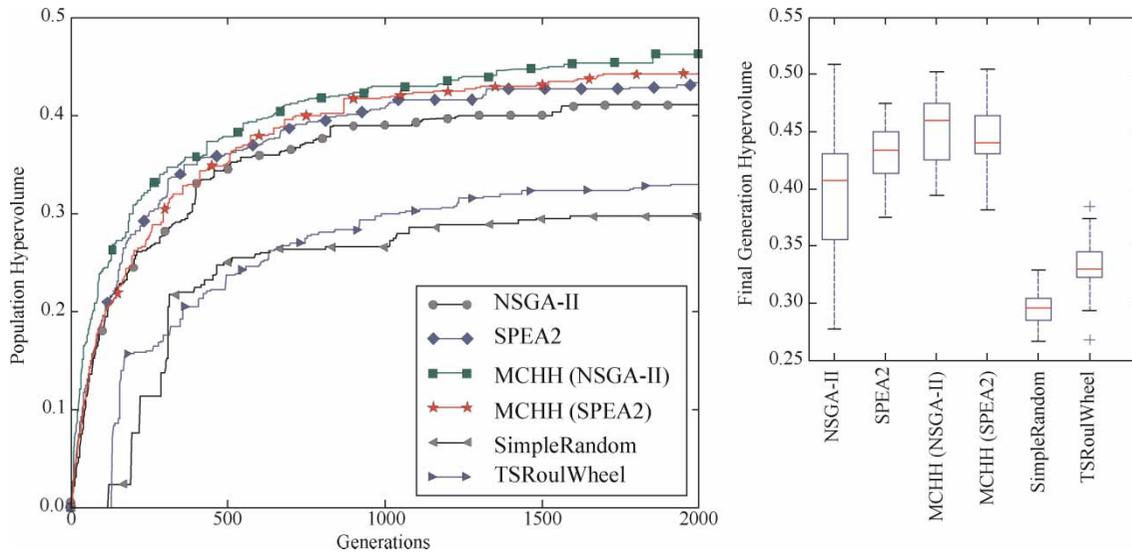
Figure 9 │ Two plots showing hypervolume results NSGA-II, SPEA2, both MCHH variants, Simple Random and TSRoulWheel on real-world network 2.

of objective space. Examining the final generation archives of SPEA2 and the final generation population of MCHH (NSGA-II) confirmed the greater spread of solutions, providing a better range of options to decision makers, as shown later in Figure 11.

All the results clearly demonstrate the MCHH variants' ability to utilise the additional heuristics which is shown by a better performance when compared with Simple Random and TSRoulWheel both of which were supplied with the additional heuristics. On the simpler problems, the MCHH performs much better than NSGA-II and SPEA2. However, it is also noticeable that as the problem complexity increases, the improvement in performance is less profound.

Heuristic weights

The improvement in performance of the MCHH over the traditional meta-heuristics like NSGA-II is directly a result of the MCHH's ability to select good heuristics during the search process. The online learning process enables the MCHH variants to utilise a wider set of low-level heuristics in a single optimisation run and is therefore better equipped to search for good solutions. However, as illustrated by the Simple Random hyper-heuristic, supplying an optimiser with a large set of heuristics does not guarantee it will use them efficiently or even improve the quality of the search.

As described in the Method section, the learning method and weight structure of the MCHH allows it to use these heuristics effectively by updating the probability of selecting each heuristic (their weights) during the search process.

An example of this is shown in Figure 10 that illustrates the overall weights assigned to each heuristic over generations from two individual trial runs on the Hanoi problem. As expected, the weight assignments vary between the two runs due to the different areas of the search space (the distribution of the population) being searched in both trial optimisation runs. Interestingly, both examples quickly downgrade the importance of the crossover operator, indicating the operator does not tend to produce dominating solutions compared with the other heuristics. The Resample heuristic, which samples new solutions, is also reduced in one of the runs (right pane, Figure 10) but interestingly still used in the other (left pane, Figure 10). The reduced weighting of the resample heuristic in the right pane results could be a reaction to the increased weight assigned to the mutation heuristic with the largest standard deviation. As both heuristics have strongly explorative behaviour, the MCHH is less likely to give strong weighting to both heuristics as this would result in a stronger stochastic element to the search which would further decrease the rate of convergence.

The MCHH always initialises all weights to be equal at the start of the search, however, as can be seen in Figure 10, the MCHH quickly biases the use of heuristics to those that
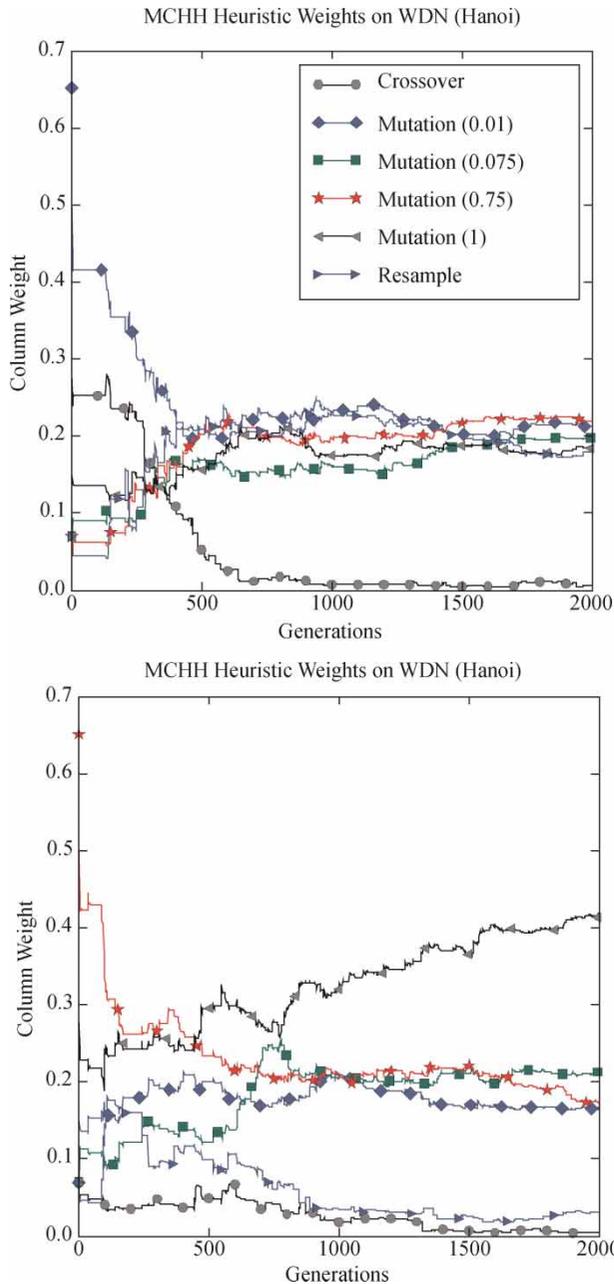
Figure 10 │ Two plots of normalised total weights (sum of all weights for moving to each heuristics) over generations assigned to each heuristic by the MCHH during the optimisation search process from two different trial runs on the Hanoi benchmark water network.

provide good results early on in the search. Indeed, while it appears as though the weights start at different points in Figure 10 due to the scale, the weights are changed within the first couple of generations. Clearly, the first few generations of the search have a strong impact of the later

weighting of the heuristics and contributes to the differences in the heuristic weightings between the different optimisation runs.

Discolouration risk in WDN problems

Both variants of the MCHH (NSGA-II and SPEA2) produced the best results across all the problems, with similar final generation results. As such, the NSGA-II variant of the MCHH results is discussed below as the trends found in the final populations occur in all sets of results.

Two Loops

The Two Loops network was used to examine the possibility of using discolouration propensity, in the form of stored turbidity, during the process of optimising WDN design. A set of solutions that satisfied the head constraints was found by optimising the network. Of these solutions, two were found that matched the known minimal cost of 419,000 and eliminated all discolouration potential – i.e., satisfied the self-cleaning threshold. These were the same as those found in Savic & Walters (1997). These results prove that, for a simple WDN, a solution may exist that is both self-cleaning and minimal for cost: the ideal solution.

A further set of solutions that satisfied head constraints with self-cleaning properties but incurred additional cost was identified. These improved upon the head excess of the minimal cost solutions and could be viable alternatives if leakage were a concern. Excessive pressure in a WDN leads to water loss through leakage from joints, fixtures and small breaks which, over a large number of networks, can result in significant costs for water companies (Al-Hemairi & Shakir 2006). Interestingly, the results show a degree of correlation between cost and discolouration (also shown in Hanoi results in Figure 11). It is hypothesised that this is not a perfect correlation as changes in pipe sizes close to sources will increase the possible shear stress downstream, increasing the cost and reducing discolouration.

Hanoi

The Hanoi benchmark was used to examine whether, for larger networks, it was possible to find solutions with
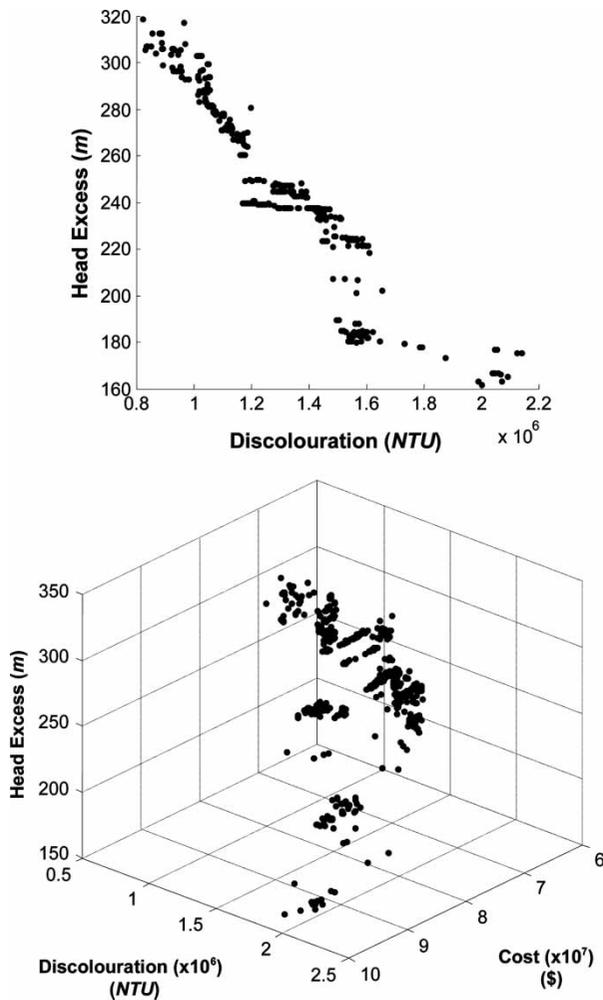
Figure 11 │ Two plots showing the combined Pareto front from all MCHH (NSGA-II) optimisation runs on the Hanoi benchmark problem for all three objectives – 2D view of Head Excess and a 3D view of the whole front.

known minimal cost and a self-cleaning feature. It was theorised that as the complexity of the network increases the likelihood of finding ideal solutions would decrease. As this study shows, this is indeed the case and it is not easy to find these solutions in networks larger than Two Loops, where a trade-off surface is more likely.

Figure 11 shows the results from the MCHH (NSGA-II) on Hanoi in three objectives. Whilst a Pareto front often resembles a curved line in two objective problems, it usually forms a plane in three objective problems, allowing for a trade-off between all three objectives. During the Hanoi experiment, the MCHH did locate a solution of known minimal cost, although it was not self-cleaning. Furthermore, it

can clearly be seen in Figure 11 that the correlation between cost and discolouration is retained, albeit with some noise. There is also a significant trade-off between the Head Excess and Discolouration with reduced head excess resulting in increased discolouration risk. In contrast, while there is a trade-off between cost and head excess for the more expensive networks (which suggests bottlenecks in the designs), the relationship between the two objectives becomes less strong as the size of the network decreases. This is shown by the curved edge of the front in the 3D plot in Figure 11 and the increase in the density (number of alternatives) of solutions as the cost scales down. At the lowest cost networks, the trade-off in solutions is found only between discolouration and head excess. If only head excess and cost or discolouration and cost were used as two objective problems, then the problem would be degenerate – i.e., the trade-off is lost close to the true Pareto optimal solutions.

The cheapest network satisfying the head-deficit constraint was found at US\$6.22 m, comparable with the SCE optimal solution found for EPANET2 (Rossman 2000). Interestingly, 16.7% of the pipes for this solution stored over 5 NTU of potential material (based upon recorded visible NTU). However a range of alternative solutions were also located that satisfied the constraints and for an increased cost of US\$68,422, this discolouration risk can be reduced to just 11.8% of pipes by adopting a slightly more expensive solution. By including the discolouration risk scores into the optimisation process, it is possible to consider alternative solutions that incur a marginally more expensive up-front cost but significantly reduce later maintenance costs by reducing the discolouration risk and allowing for more selective cleaning schemes to be introduced.

CONCLUSIONS

This paper presents a novel approach to pipe diameter design for WDNs and proposes the MCHH, a multi-objective online selective hyper-heuristic, which is applied to the problem. The experiment demonstrated that by incorporating discolouration propensity, calculated using the DPM, and excess head (thus taking some leakage concerns

into account) it is possible to lessen future discolouration mitigation burdens with a calculated increase in rehabilitation/construction costs. The method was tested on three benchmark networks (Two Loop, Hanoi and Anytown) and three real-world networks.

Two solutions with minimal cost and a self-cleaning feature were found for Two Loop. However, results for Hanoi highlighted the difficulties in finding minimal cost, self-cleaning networks for relatively larger networks. This study does show that a modest increase in cost can also attain a self-cleaning threshold in this larger network a result which may be more applicable to much larger, real WDNs.

As anticipated, the results from the application of this approach to the rehabilitation of real-world networks supports the theory that a trade-off curve of cost against discolouration potential becomes increasingly more varied and gradual as the complexity of the problem increases. The likelihood of finding minimal cost, self-cleaning solutions is also improbable for relatively large networks. Nevertheless, by applying this approach during rehabilitation planning, companies can also consider discolouration potential in addition to new demands and thus improve the self-cleaning properties of their networks for a marginal increase in cost.

In addition, this paper applied the MCHH to the same three objective formulation of the WDN design and rehabilitation problems. The MCHH was incorporated in the well-known NSGA-II and SPEA2 and supplied with a range of low-level heuristics tailored for use on WDN design and rehabilitation problems. Both the original optimisers and the MCHH variants were run on the problem. The results demonstrated an improvement in performance across all six benchmark (Two Loops, Hanoi and Anytown) and real-world networks used in the experiment and clearly demonstrated the improvement in performance that can be obtained over the original algorithms. In each case, the MCHH variants produced better results more quickly that the original algorithms, requiring minimal work to embed the MCHH.

Finally, this study conducted a comparison of the MCHH against two hyper-heuristics from the literature: Simple Random and TSRoulWheel. The results demonstrated the MCHH's ability to find a wider range of solutions across the networks as indicated by better hyper-volume results. Furthermore, the results illustrated the faster rate of convergence of the MCHH across all three of the real-world problems, allowing for shorter runs on each of the problems and providing significant reduction in the cost associated with optimising the networks.

ACKNOWLEDGEMENTS

REFERENCES

Al-Hemairi, H. A. & Shakir, R. H. 2006 Minimizing leakage rates in water distribution networks through optimal valves settings. In: *Proceedings of the World Environmental and Water Resources Congress 2007: Restoring our Natural Habitat*, ASCE, Omaha, Nebraska, USA, pp. 1–13.

Bader, J., Deb, K. & Zitzler, E. 2008 Faster hypervolume-based search using Monte Carlo sampling. *Multiple Criteria Decision Making for Sustainable Energy and Transportation Systems* **634**, 313–326.

Boxall, J. & Prince, R. 2006 Modelling discolouration in a Melbourne (Australia) potable water distribution system. *Journal of Water Supply Research Technology-Aqua* **55**, 207–219.

Boxall, J. B. & Saul, A. J. 2005 Modelling discolouration in potable water distribution systems. *Journal Environmental Engineering ASCE* **131**, 716–725.

Boxall, J., Saul, A. & Skipworth, P. 2001 A novel approach to modelling sediment movement in distribution mains based on particle characteristics. *Water Soft Systems* **1**, 263–273.

Buchberger, S. G., Blokker, M. & Vreeburg, J. 2008 Sizes for self-cleaning pipes in municipal water supply systems. In: *Proceedings of Water Distribution Systems Analysis 2008*, ASCE, Kruger National Park, South Africa, pp. 1–10.

Burke, E., Hart, E., Kendall, G., Newall, J., Ross, R. & Schulenburg, S. 2003 Hyper-heuristics: an emerging direction in modern search technology. In: *Handbook of Metaheuristics* (F. Glover & G. A. Kochenberger, eds). Kluwer, Amsterdam, pp. 457–474.

Burke, E. K., Hyde, M., Kendall, G., Ochoa, G., Ozcan, E. & Qu, R., 2010 Hyper-heuristics: A Survey of the State of the Art. Computer Science Technical Report No. NOTTCS-TR-SUB-0906241418–2747, School of Computer Science and Information Technology, University of Nottingham, UK.

Burke, E. K., Hyde, M., Kendall, G., Ochoa, G., Ozcan, E. & Woodward, J. 2009 A classification of hyper-heuristics approaches. In: *Handbook of Metaheuristics, International Series in Operations Research & Management Science* (M. Gendreau & J.-Y. Potvin, eds). Springer, Berlin, pp. 449–468.

Burke, E. K., Landa Silva, J. & Soubeiga, E. 2003 Hyperheuristic approaches for multiobjective optimisation. In: *Proceedings of the Fifth Metaheuristics International Conference* (MIC 2003), Kyoto, Japan, pp. 11.1–11.6.

Burke, E. K., Landa Silva, J. D. & Soubeiga, E. 2005 Multi-objective hyper-heuristic approaches for space allocation and timetabling. In: *Meta-heuristics: Progress as Real Problem Solvers, Selected Papers from the 5th Metaheuristics International Conference (MIC 2003)*, pp. 129–158, Springer, Berlin.

Cook, D. 2007 Field Investigation of Discolouration Material Accumulation Rates in Live Drinking Water Distribution Systems. Ph.D. Thesis, University of Sheffield, UK.

Cowling, P., Kendall, G. & Soubeiga, E., 2000 A hyperheuristic approach to scheduling a sales summit. *Lecture Notes in Computer Science* **2079**, 176–190, Springer, Berlin.

Deb, K. 2001 *Multi-Objective Optimization using Evolutionary Algorithms.* John Wiley & Sons, Chichester.

Deb, K., Mohan, M. & Mishra, S. 2003 A Fast Multi-objective Evolutionary Algorithm for Finding Well-Spread Pareto-Optimal Solutions. KanGAL Report No. 2003002, Indian Institute of Technology, Kanpur.

Dewis, N. & Randall-Smith, M. 2005 Discolouration risk modeling. *Water Asset Management International* **1.3**, 16–18.

di Pierro, F., Khu, S.-T., Savic, D. & Berardi, L. 2009 Efficient multi-objective optimal design of water distribution networks on a budget of simulations using hybrid algorithms. *Environmental Modelling & Software* **24**, 202–213.

Farmani, R., Savic, D. & Walters, G. 2005 Evolutionary multi-objective optimization in water distribution network design. *Engineering Optimization* **37**, 167–183.

Formiga, K. T., Chaudry, F. H., Cheung, P. B. & Reis, L. F. R. 2003 Optimal design of water distribution system by multiobjective evolutionary methods. In: *LNCS Evolutionary Multi-Criterion Optimization*, Vol. 2632 (C. M. Fonseca, P. J. Fleming, E. Zitzler, L. Thiele, K. Deb, eds). Springer-Verlag, Berlin.

Laumanns, M., Zitzler, E. & Thiele, L. 2000 A unified model for multi-objective evolutionary algorithms with elitism. In: *Congress on Evolutionary Computation (CEC 2000)*, La Jolla, CA, pp. 46–53.

McClymont, K. & Keedwell, E. C. 2010 Optimising multi-modal polynomial mutation operators for multi-objective problem classes. In: IEEE Congress on Evolutionary Computation (CEC), Barcelona.

McClymont, K. & Keedwell, E. C. 2011 Markov chain hyper-heuristic (MCHH): an online selective hyper-heuristic for

multi-objective continuous problems. In: *Proceedings of the 13th Annual Conference on Genetic and Evolutionary Computation*, Dublin, pp. 2003–2010.

Murphy, L. J. & Simpson, A. R. 1992 Pipe optimization using genetic algorithms. Research Report No. R93, Department of Civil Engineering, University of Adelaide, Australia.

Ong, Y. S., Lim, M. H., Zhu, N. & Wong, K. W. 2006 Classification of Adaptive Memetic Algorithms: A Comparative Study. *IEEE Transactions on Systems, Man and Cybernetics – Part B* **36**, 141–152.

Raad, D., Sinske, A. & van Vuuren, J. 2010 Multiobjective optimization for water distribution system design using a hyperheuristic. *Journal of Water Resources Planning and Management* **136**, 592–596.

Rossman, L. A. 2000 EPANET users manual. Technical Report No. EPA/600/R-00/057, United States Environment Protection Agency.

Savic, D. A. & Walters, G. A. 1997 Genetic algorithms for least-cost design of water distribution networks. *Journal of Water Resources Planning-ASCE* **123**, 67–77.

Schaake, J. & Lai, D. 1969 Linear programming and dynamic programming applications to water distribution network design. Report 116, MIT Press, Cambridge, MA.

Shie-Yui Liong, S. & Atiquzzaman, M. 2004 Optimal design of water distribution network using shuffled complex evolution. *Journal of the Institute of Engineering* **44**, 1.

Simpson, A. R., Dandy, G. C. & Murphy, L. 1994 Genetic algorithms techniques for pipe optimization. *Journal of Water Resources Planning and Management, ASCE* **120**, 423–443.

Suribabu, C. R. 2010 Differential evolution algorithm for optimal design of water distribution networks. *Journal of Hydroinformatics* **12**, 66–82.

Vasan, A. & Simonovic, S. P. 2010 Optimization of water distribution network design using differential evolution. *Journal of Water Resources Planning and Management* **136**, 279–287.

Vreeburg, J. H. G., Schaap, P. & van Dijk, J. C. 2005 Measuring discolouration risk: resuspension potential method. *Proceedings of the EWA Leading Edge Conference*, Prague.

Yates, D. E., Templeman, A. B. & Boffey, T. B. 1984 The computational complexity of the problem of determining least capital cost designs for water supply networks. *Engineering Optimization* **7**, 142–155.

Zitzler, E., Laumanns, M. & Thiele, L. 2002 SPEA2: Improving the strength pareto evolutionary algorithm for multiobjective optimization. In: *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001)* (K. C. Giannakoglou *et al.*, eds). International Center for Numerical Methods in Engineering (CIMNE), pp. 95–100.