WaterOpt: A Method for Checking Near-Feasibility of Continuous Water Supply

Bharadwaj Amrutur*‡, M. S. Mohan Kumar†‡, K. R. Sheetal Kumar†‡,
Lovelesh Patel*, Rajesh Sundaresan*‡, Nidhin K. Vaidhiyan*
*Department of Electrical Communication Engineering
†Department of Civil Engineering
‡Robert Bosch Centre for Cyber Physical Systems
Indian Institute of Science, Bangalore 560012, India

***Abstract*—This paper proposes a systematic approach to explore the search space of control parameters on an existing water network in a locality within Mysore, India, to see if continuous water supply can be provided to meet a nominal diurnal demand pattern. The control parameters are the initial states of the reservoirs at some reference hour and hourly maximum flow rates across control valves. The system state should return to its initial state after twenty four hours to follow a diurnal rhythm without overflow or underflow of reservoirs. The number of configurational changes should also be small. The paper proposes an optimisation framework that relaxes the constraints via convex penalty functions, performs approximate gradient descent to minimise the net penalty, and thus identifies a suitable configuration. The proposed framework is ideal for integration into a cyber physical system that can actuate the controls based on real-time demand and flow information.**

I. Introduction

The Karnataka Urban Water Supply and Drainage Board (henceforth Board) in the state of Karnataka, India, approached us to seek suggestions on how to convert the Devanoor Command Area (henceforth Devanoor) of Mysore, Karnataka, India, into a region where uninterrupted water supply[1] could be provided with minimal additional infrastructure. Devanoor spans about 2km x 2km with a population of about 160,000. Water is pumped into a mass distribution tank from which it flows via gravity to ten other overhead service reservoirs located at different parts of Devanoor (see Fig. 1). Water is then distributed from these reservoirs to the residential consumers, once again via gravity. The terrain is undulated. A significant fraction of the population (about 1/4) resides in a sublocality called Rajiv Nagar I Phase[2] (henceforth Rajiv Nagar I) that is on higher ground. Moreover the reservoirs are of insufficient capacity to meet demand during peak hours. At least five of the reservoirs are undersized. As a consequence, the Board and their execution partner (Jamshedpur Utilities and Services Company (JUSCO)) have found it difficult to meet, simultaneously, peak demand in this high population sublocality and yet equitable distribution across all of Devanoor. Ad hoc attempts to stock up the Rajiv Nagar I reservoir, in order to meet the peak hour demand, had led to disruption of water supply in the downstream localities.

This paper presents our systematic approach to explore the search space of control parameters in an automated fashion in order to decide if it is at all possible to provide uninterrupted water supply to meet a nominal diurnal demand pattern [2] in Devanoor. The framework we provide is ideal for integration into a cyber physical system that can actuate the network based on real-time demand and flow information.

The Devanoor water network has several flow control valves (FCV) that can control the maximum flow rates across the pipes. We take the control parameters to be the hourly maximum flow rates through these valves and the initial volumes in each reservoir. Our contributions are the following.

1) We demonstrate that there are settings that can achieve nearly continuous water supply in Devanoor to support a nominal diurnal demand pattern.

2) We show this is achievable with changes in settings restricted to a small subset valves, keeping the other valves fixed at constant maximum flow rates. Moreover, the number of changes in settings is small.

3) We show that the system is brought back to a state close its initial setting, within a desired target error, after twenty four hours. This readies the system for a diurnal rhythm.

A few remarks about our approach are in order. First, we take a somewhat naïve computational approach that sits on top of a hydraulic modeling tool (EPANET [3]). Our approach uses simple mass balance ideas for estimating inflows and outflows to and from the reservoirs, uses penalty functions for either overflow or underflow in the reservoirs, and employs gradient descent to explore the search space of control parameters. The advantage of our naïve approach is that it is scalable and generalisable to wider area networks. We next attempt to keep the number of changes to the valve settings to a low number by a suitable choice of penalty functions that *encourages sparsity* in the number of changes. We then enforce a twenty-four hour periodicity by introducing another penalty function for a deviation from a return to the initial condition after twenty four hours.

We now discuss some related works. A dynamic inversion based controller approach for control valve throttling to ensure target flows to all reservoirs in different zones of an undulat-

[1] This was part of an initiative to provide continuous water supply [1].

[2] Rajiv Nagar I is one of the two highly populated sublocatities in Devanoor.

978-1-5090-1161-2/16/$31.00©2016 IEEE

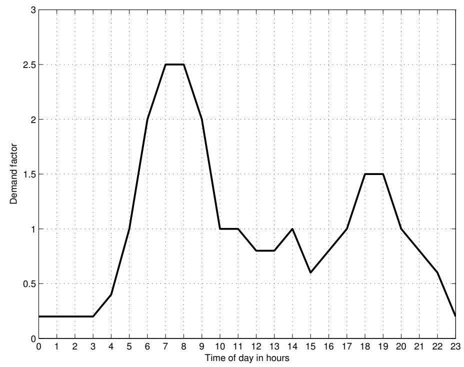

Fig. 2. Diurnal demand pattern

TABLE I
DEMAND, IN LITRES PER SECOND, AND CAPACITY, IN UNITS OF 100,000 LITRES, IN EACH METERED AREA.

S.No.	Metered area	Demand (ℓps)	Capacity (100,000ℓ)
1	Rajiv Nagar I	79.8	15.0
2	Tippu Park	37.4	10.0
3	N. R. Mohalla I	57.0	10.0
4	N. R. Mohalla II	14.4	2.5
5	Veeranagere	28.1	5.0
6	Rajendra Nagar I	39.6	15.0
7	Rajendra Nagar II	34.0	10.0
8	Rajiv Nagar II	19.6	10.0
9	Kesare	14.4	5.0
10	Bademakan	9.9	5.0

ing terrain was proposed in [4]. Another optimization based control to achieve supply management goals is discussed in [5]. This is applied to water transfer operations to the city of Sintra (Portugal) using a supervisory control system. An optimal valve control algorithm for minimization of leakage rates was presented in [6] and its performance was studied on an example network. This works differs from the above in that the search for a good configuration is fully automated.

The rest of this paper is organised as follows. In section II, we provide a very brief description of the Devanoor water network and highlight the key issues related to continuous water supply. In section IV, we describe the components of our method. In section III, we discuss our computational methodology in detail. In section V, we conclude with a short discussion of the outcomes of our research.

II. Key issues in the Devanoor area

Devanoor is divided into ten metered areas, each with a reservoir. See Fig. 1. Water from a treatment plant is pumped into a mass distribution tank ((MDT) either directly or via an intermediate mass balancing reservoir). This water then flows via gravity to each of the ten reservoirs of the metered areas.

The goal is to meet consumption of 180 litres (ℓ) per capita per day, including losses, as suggested by [2]. The diurnal demand pattern[3] [2] is given in Fig. 2. Taking the population in the various metered areas and the daily per capita consumption into account, we arrive at the average demand rate per second and report it in column 3 of Table I. Column 4 of that table provides capacities of the corresponding reservoirs. The supply from the MDT, by itself, is not enough to serve all ten metered areas during peak hours since the feeder mains are of inadequate sizes. Further, it can be gleaned from Table I that the reservoirs do not have enough capacity, by themselves, to meet demand during the peak hours[4]. Finally, as mentioned in the introduction, the terrain is undulated, but the supply lines are arranged in a cascade. Supply to higher ground reservoirs (for example, Rajiv Nagar I) affects the supply to the downstream sublocaties that are part of the cascade (Rajiv Nagar I, Tippu Park, N. R. Mohalla I and II, and Veeranagere). A combination of (1) supply via the feeder lines, and (2) storage in the reservoirs is therefore needed to meet peak demand and provide continuous water supply, if at all feasible.

There are additional complexities. The metered areas are not completely isolated. The feeder to the MDT also feeds a neighbouring area. So continuous water supply to Devanoor could affect supply to neighbouring areas. These issues are beyond the scope of this paper.

III. Modeling and methodology

Let time be slotted into hours for convenience. We will use t to denote the hours; $t = 0$ will denote a reference time, say midnight, in which case $t = T := 24$ refers to the following midnight. Our optimisation will run for one period which is a duration of T stages. T is the specified periodicity of the demand pattern. Let there be a total of R reservoirs including the main distribution tank. Reservoir r has capacity c_r. Each reservoir has its own set of parameters (for example staging height) all of which are important for calculating inflows and outflows based on hydraulic models. See the subsection III-A.

Let $v_r(t)$ denote the volume of water in the reservoir r at time t. The following are constraints for the time evolution of the state defined by $(v_r(t), r = 1, \ldots, R)$.

C1: No reservoir goes dry: $v_r(t) \geq 0$ for each r and t.
C2: No reservoir overflows: $v_r(t) \leq c_r$ for each r and t.
C3: Each reservoir returns to its initial state after T stages: $v_r(T) = v_r(0)$ for each r.

We have the following design variables at our disposal.

- The *initial volumes* in each reservoir

$$\mathbf{v}(0) := (v_r(0),\ r = 1, \ldots, R).$$

  We arbitrarily choose midnight as our reference zero time.
- The *maximum flow rate* at each valve and at each hour

$$\mathbf{a} := (a_j(t),\ j = 1, \ldots, J,\ t = 1, \ldots, T),$$

  where we have assumed a total of J valves.

[3]For simplicity, we take a standard demand pattern although in reality there is variation due to demographics, seasonality, emergency settings, etc. The area under the curve in Fig. 2 is 1 unit, which refers to (180/24)ℓ per capita per hour = 7.5ℓ per capita per hour.

[4]Consider Rajiv Nagar I and N. R. Mohalla II. The demand rate during 06:00 - 09:00 hrs is more than twice the daily average rate: $(79.8 \times 2)\ell$ps in Rajiv Nagar I and $(14.4 \times 2)\ell$ps in N. R. Mohalla II. The total demands works out to be 17.2 and 3.1 lakh litres (1 lakh =100,000), respectively, both of which exceed the respective capacities of 15.0 and 2.5 lakh litres.

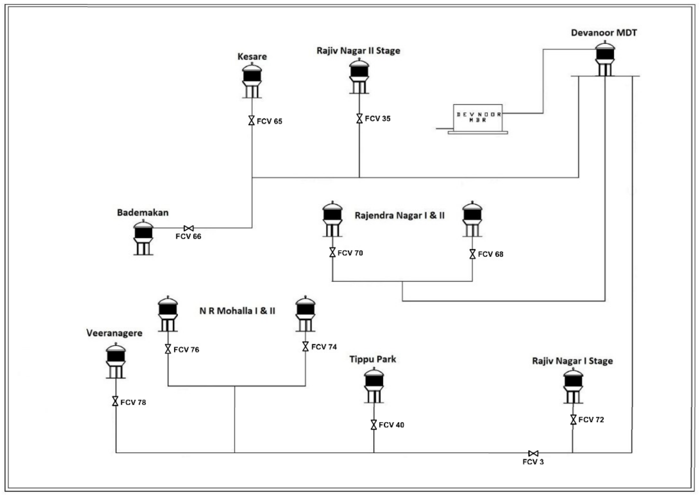

Fig. 1. The Devanoor water network with ten metered areas, each with its reservoir. FCV stands for flow control valve.

We consistently use r to index a reservoir, j to index a valve, and t to index time. There are R reservoirs, J valves, and T stages in a period. We use $\mathbf{a}(t)$ to denote the maximum flow rates at each valve, but at the indicated time t.

A configuration $\mathbf{x}$ is a particular setting of all the control variables, that is, the initial states and the maximum flow rate on all valves across the entire duration of a period. In symbols, $\mathbf{x} = (\mathbf{v}(0), \mathbf{a})$.

When we fix a configuration and follow the time evolution of the states, a certain trajectory ensues. This trajectory is, naturally, a function of the configuration. We make the following simplifying assumptions. Our cross validation, about which we remark in section V, supports these assumptions.

A1: The transience in flows due to the hourly changes in the maximum flow rate settings on the valves has minimal impact on the flow rates. The new flow rates are assumed to be instantaneously attained.

A2: The flow rates remain constant during the course of each hour. In reality of course reservoir water levels change during the course of the hour and may affect flow rates. We assume this change is negligible.

A3: The inflow rate into the mass distribution tank is constant and sufficient to meet Devanoor's daily average demand.

The water balance equation at time t is then

$$v_r(t+1) = v_r(t) + q_r(t) - u_r(t), \; r = 1, \ldots, R, \quad (1)$$

where $q_r(t)$ is the inflow into reservoir r and $u_r(t)$ is the usage from reservoir r during the tth hour. The quantity $u_r(t)$ comes from the diurnal demand pattern of Fig. 2. We next discuss the all important modeling of the inflows $q_r(t)$.

A. Modeling of the flows

The flow rates at any instant of time during the course of an hour, and therefore $q_r(t)$ for a particular t, are complex functions of the valve settings for that t, the water levels in the reservoirs during that period, the water network itself, its myriad parameters which include pipe dimensions, pipe material, nature of junctions, etc. A serious modeling of this will take us deep into the physics of hydraulic modeling, which we shall "outsource". We shall use a software toolkit called EPANET[5] as an oracle that shall tell us the physically realised flow rates for any configuration. Indeed, we call EPANET on an hourly basis with the settings and initial state for that hour as input. These settings are derived from the configuration. From the toolkit's output, we compute the flow rates into each reservoir, $q_r(t), \; r = 1, \ldots, R$, during that hour. Thus one may write $q_r(t) = \Phi_r(\mathbf{v}(t), \mathbf{a}(t))$ for suitable $\Phi_r, \; r = 1, \ldots, R$. One then views (1) as the driving equation for the system dynamics with the control variables included in $q_r(t)$. We thus have the state evolution equation:

$$v_r(t+1) = v_r(t) + \Phi_r(\mathbf{v}(t), \mathbf{a}(t)) - u_r(t), \quad r = 1, \ldots, R, \quad t = 0, \ldots, T-1. \quad (2)$$

We next describe the use of this simplified oracle-based model.

B. An overview of the computational methodology

We are now in a position to describe our methodology. Instead of meeting the indicated constraints C1 - C3 at all times, we relax them and allow for deviations with a cost. Let us first discuss the relaxations.

C1: When constraint C1 is violated, the reservoir has no water to meet the demand and the flows may be different. Our relaxation allows us to assume that the reservoir continues to meet demand, as if there is a secondary source of replenishment for this reservoir, at a cost.

C2: When constraint C2 is violated, the reservoir overflows. Our relaxation allows us to assume that the reservoir does not lose the overflowed water, at a cost.

C3: When constraint C3 is violated, the system ends up in a state different from the initial value, thereby affecting the

[5]EPANET is a public-domain, water distribution system modeling software toolkit developed by the United States Environmental Protection Agency's (EPA) Water Supply and Water Resources Division.

diurnal pattern. Our relaxation permits this at a cost. In reality, when this constraint is violated, an external agency should add or remove water to bring the state back to its initial state.

Additionally, we impose a new cost for frequent changes to the maximum flow rate settings on the valves. The total cost is then a function of the configuration. We next choose a configuration that minimises the overall cost. This tries to keep the violations to C1 - C2 and the number of changes to the settings to a minimum. For the minimising configuration, we then evaluate the trajectory and recognise that it meets the hard constraints C1 and C2, and gets us close to meeting C3.

IV. Components of our method

We now describe the components of our method and the many design choices motivated by common sense.

A. *Choice of cost functions*

Given a particular configuration $\mathbf{x} = (\mathbf{v}(0), \mathbf{a})$ which specifies the initial state of the reservoirs and the valve controls on each valve and each time period, given a particular diurnal demand pattern $\mathbf{u} = (u_r(t), r = 1, \ldots, R, t = 1, \ldots, T)$, the state evolution is given by (2). For this configuration, the following incremental costs are incurred.

(a) The cost of violating the constraint C1 and C2 at time period t on reservoir r is $f_{12}(v_r(t); c_r)$, where we take

$$f_{12}(\nu; c) = w_{12}(c) \cdot \left| \frac{\nu}{c} - \frac{1}{2} \right|^{\alpha}, \ \alpha \geq 1;$$

ν is a candidate volume and c is a candidate reservoir capacity. This is a convex function of ν (for each fixed c) that penalises any deviation of the reservoir state away from half its capacity. The weight $w_{12}(c)$ provides a weighting based on the capacity of the reservoir. There are of course many other penalty functions. One could also have a separate weight function for each reservoir, depending on the criticality of the reservoir. But there is enough flexibility made available by the choice of α that the above appears to be sufficient for our purposes.

(b) The cost of violating C3, of not returning to the initial state at the end of a period, on reservoir r is given by $f_3(v_r(T), v_r(0); c_r)$, where we take

$$f_3(\nu', \nu; c) = w_3(c) \cdot \left| \frac{\nu' - \nu}{c} \right|^{\beta}, \ \beta \geq 1;$$

ν is a candidate initial volume, ν' is a candidate final volume, and c is a candidate reservoir capacity. The weighting function $w_3(c)$ provides a capacity dependent weight. This (jointly convex in (ν, ν') for fixed c) penalty function could also depend on the reservoir with perhaps critical reservoirs having a higher weight. For Devanoor, we do not require this generality. The weighting function $w_3(c)$ also tells us how this cost compares with the cost of violating the constraints C1 and C2.

(c) Let $a_{\max,j}$ be the maximum flow rate across valve j. The cost of changing this setting on a particular valve j at time $t+1$ is $f_4(a_j(t+1), a_j(t); a_{\max,j})$, where we take

$$f_4(a', a; a_{\max}) = w_4(a_{\max}) \cdot \left| \frac{a' - a}{a_{\max}} \right|,$$

where, on a candidate valve, a is a maximum flow rate setting in some slot, a' is the changed setting in the following slot, and $a_{\max}$ is the maximum possible flow rate setting. Again, $w_4(a_{\max})$ is some weighting function that depends on maximum flow rate $a_{\max}$ across the valve. This penalty function has a sharp transition at $a' = a$ that encourages the algorithm to choose $a' = a$ as much as possible, yet keeping the penalty function convex. This is reminiscent of 'lasso' regression [7] in machine learning. Sharper penalty functions are possible, but at the expense of convexity.

The overall cost of a configuration is taken to be

$$\begin{aligned} F(\mathbf{x}) \ &= \ \sum_{t=0}^{T-1} \sum_{r=1}^{R} f_{12}(v_r(t); c_r) + \sum_{r=1}^{R} f_3(v_r(T), v_r(0); c_r) \\ &\quad + \ \sum_{t=0}^{T-1} \sum_{j=1}^{J} f_4(a_j(t+1), a_j(t); a_{\max,j}), \end{aligned}$$

where $v_r(t)$ are obtained from the state evolution (2). In the last term, $a_j(T)$ is defined to be $a_j(0)$.

While each of the component functions above are convex in their arguments (for fixed c or $a_{\max}$ as appropriate), the overall function may not be a convex function of the configuration. This is because of the possible, in greater generality, nonconvex dependence of Φ_r on the configuration. Our choice of convex penalty functions is to ensure that any nonconvexity, if it arises, is due to Φ_r. See also the discussion in Section V.

B. *The search for a good configuration*

We explore the search space via a simple gradient descent method. We first estimate a gradient and choose its negative as the direction of a move. We then perform a line search to identify the best step size. Finally, we make a noisy update, roughly as per the chosen direction and step size. These steps are carried out iteratively. We shall use n to denote the iteration index. Asympotically, the update matches the suggested step size and direction. Details follow.

Step 1: Estimation of the gradient. The configuration is made of variables for initial reservoir states (R variables) and variables for the maximum flow rate on each valve in each stage of a period (JT variables). Let $K = R + JT$ be the total number of variables which we index by k. The kth component of the gradient is estimated as

$$(\nabla F(\mathbf{x}))_k = \frac{F(\ldots, x_k + \varepsilon, \ldots) - F(\ldots, x_k - \varepsilon, \ldots)}{2\varepsilon},$$

where ε is a small constant.

Consider the computation of $F(\ldots, x_k + \varepsilon, \ldots)$. This is the cost of a new configuration $(\ldots, x_k + \varepsilon, \ldots)$ that is different from x in only one variable. We calculate the new state trajectory, via several calls to the EPANET toolkit for hourly flow rate computation, and then compute the new cost function associated with this changed configuration. A similar computation is done for $F(\ldots, x_k - \varepsilon, \ldots)$. The gradient is then estimated as per the formula. Some computational savings can be obtained by reducing the number of calls to the EPANET toolkit if ε is small and we are willing to ignore

TABLE II
CONSTANT FLOW RATE SETTINGS ON VALVES.

Valves/Reservoirs	Max-flow rate (ℓps)	Demand (ℓps)
Tippu Park	37.53	37.4
N. R. Mohalla II	14.42	14.4
Veeranagere	30.25	28.1
Rajendra Nagar I	39.46	39.6
Rajendra Nagar II	34.02	34.0
Rajiv Nagar II	19.50	19.6
Kesare	14.40	14.4

the effect of small changes to the reservoir water levels, but we will not elaborate these ideas here.

Step 2: Line search. We now perform an inexact line search to determine the length of the move along the direction opposite to the gradient $\mathbf{d} := \nabla F(\mathbf{x})$. An exact line search would choose a multiplier $\lambda^\star$ such that

$$\lambda^\star = \arg\min_{\lambda > 0} F(\mathbf{x} - \lambda \mathbf{d}).$$

An inexact line search replaces this with an iterative procedure: start with an aggressive value, say 1; then decrease in a geometric sequence until $F(\mathbf{x} - \lambda\mathbf{d}) < F(\mathbf{x}) - c\lambda|\mathbf{d}|$, where c is a design constant. Let L be the number of decreases in this geometric sequence. Call the finally chosen λ of this step 2 as λ_n, where n is the iteration index.

Step 3: Noisy gradient descent. To avoid traps at local minima, we perturb the move with a random Gaussian noise vector. Each component of this noise has variance $\beta_n = 1/(n+1)^\gamma$ where $0.5 < \gamma \leq 1$. This choice ensures that $\sum_n \beta_n = \infty$ and $\sum_n \beta_n^2 < \infty$. The former ensures sufficient exploration capability and the latter ensures that the added noise is of finite variance and can be averaged out. This choice is typical in stochastic approximation algorithms [8].

With $x^{(0)}$ as the initial iterate, the sequence of iterates is

$$\mathbf{x}^{(n+1)} = \mathbf{x}^{(n)} - \lambda_n \nabla F(\mathbf{x}^{(n)}) + \beta_n \mathbf{z}_n, \quad n \geq 0$$

where $(\mathbf{z}_n, n \geq 0)$, with each $\mathbf{z}_n \in \mathbb{R}^K$, is independent and identically distributed (iid) across n. The components of $\mathbf{z}_n$ themselves are iid with the standard normal distribution.

TABLE III
MAXIMUM FLOW RATE SETTINGS ON VALVES.

Time	Maximum flow rate settings (ℓps)		
	Rajiv Nagar I	N. R. Mohalla I	Bademakan
00:00	49.89	57.05	11.54
01:00	-	9.02	-
02:00	-	57.05	-
03:00	79.86	-	-
06:00	109.30	-	-
07:00	-	89.18	80.64
08:00	-	-	11.54
09:00	-	57.02	-
10:00	79.86	-	-

C. Implementation

Our implementation of the above algorithm is in the C programming language. The framework makes frequent calls to the EPANET toolkit for the hydraulic modeling[6]. While our description of the algorithm assumed synchronous, in particular hourly, demands and maximum flow rate settings on valves, the actual implementation handles asynchronous arrival of demand information and asynchronous changes to the valve settings. This is to enable interfacing with a real-time system. The asynchronous handling of tasks (changes in valve settings, changes in demands, etc.) is executed via a job scheduler that uses a time-based priority-queue data structure.

V. RESULTS AND DISCUSSION

We chose the weight functions with $w_{12}(c) = c, w_3(c) = 24c$. This could be justified as follows. Larger the capacity of the tank, the greater the weight. Also, the factor 24 ensures that the penalty for not returning to the initial state is treated on par with the penalty for the reservoirs overflowing or underflowing during the course of the day. We chose $\alpha = \beta = 2$ to fix the cost functions f_{12} and f_3. We further chose the weight $w_4(a_{j,\max}) = 500$, a constant. We arrived at these after several trials and a large weight for w_4 appeared essential to realise a sparse solution, but a systematic exploration of the space of these parameters remains to be done.

The resulting solution configuration under the chosen parameters is summarised in Tables II-III. Table II indicates the maximum flow rate settings on valves (leading to the indicated reservoirs) that could do with just one constant setting. As pointed out earlier, some reservoirs do not have enough capacity to meet the sustained demand during peak hours. Enough inflow must be maintained in the pipes leading to these reservoirs, and to enable this, the maximum flow rates must be carefully modulated across time on some valves. The valves and the corresponding reservoirs that need such a modulation are indicated in Table III. As can be seen from the table, only three valves need to be controlled across time. Two of the valves need four changes[7]. The valve controlling the flow into the Bademakan reservoir requires only two changes.

To cross-validate our solution, we use the EPANET multi-species extension toolkit which provides an extended twenty four hour simulation of the proposed configuration. Our proposed configuration in Tables II-III provides a near-feasible solution. The top subplot in Fig. 3 shows the fractional occupancy of the tanks (with reference to capacity). This roughly follows the diurnal usage pattern, but with a lead. None of the tanks overflow for the proposed configuration. None goes below 20% occupancy. Thus constraints C1 and C2 are met at all times of the day on all reservoirs. All reservoirs return close to their initial states, with the maximum difference normalised by the reservoir capacity being $\pm 2\%$.

We can at best control the maximum flow rate through a valve. The actual flow rate may be lower than the configured

[6] $2KT + L$ calls/iteration without the optimisations indicated in Step 1.

[7] Treat 57.05ℓps for the post-midnight hour at the N. R. Mohalla II valve as 57.02ℓps, the setting prior to the midnight hour.

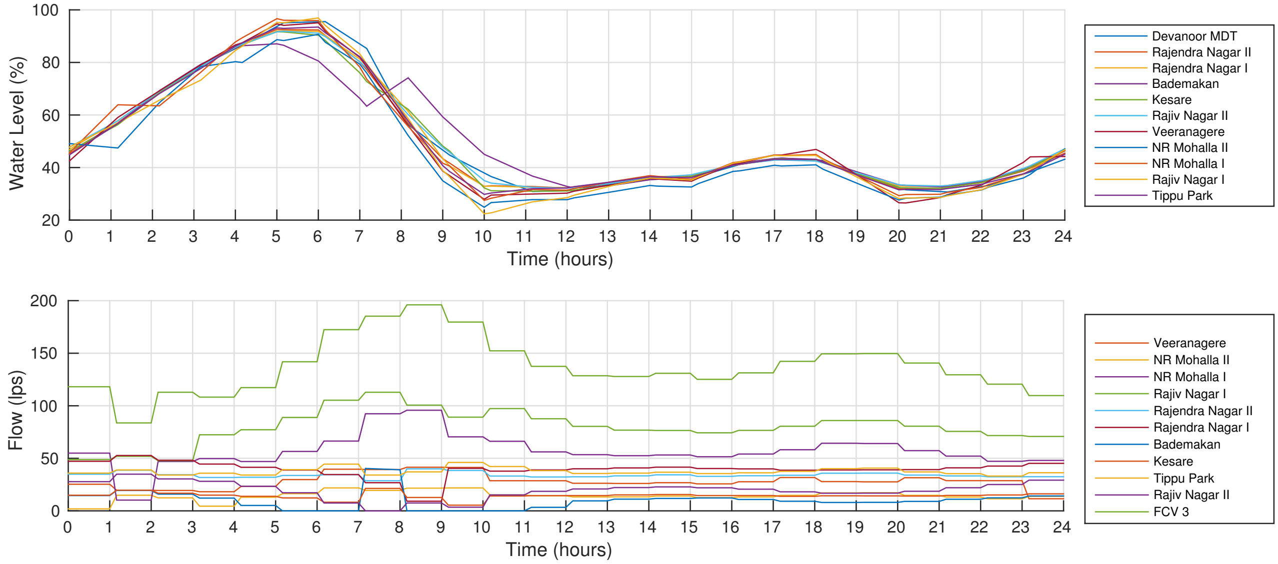

Fig. 3. The top figure indicates the water levels in the reservoirs, including the mass distribution tank, over time. The bottom figure provides the flow rates in the valves over time. There are 11 valves; ten of these lead to reservoirs, and the eleventh is FCV 3 given in Fig. 1.

maximum[8]. It may also fluctuate during the period. The bottom subplot in Fig. 3 indicates the realised flow rates at a ten-minute time resolution. Some pipes have zero flow rate for certain periods[9]. But our framework could be extended to include an additional cost function to reduce such occurrences.

We close this paper with some additional remarks.

1. The main advantage of our methodology is that we obtain a configuration using a fully automated procedure. Our simulation ran for 134,000 iterations on $K = 275$ variables and took four hours with no optimisations on a standard desktop machine[10]. With optimisations and distributed computation, we believe our method can be scaled to the level of a city.

2. One way to implement our proposal is via automated actuation of the FCVs. Our current proposal is of an open loop nature. As a logical next step, we plan to forecast demand [9] and close the loop to respond online to real demands.

3. Our method searched the space of "maximum flow rates" and used the EPANET tool as an oracle that identified the water network behaviour for a particular configuration. One could potentially explore the space of flow rates across valves. This is a convex optimisation problem with perhaps faster convergence. However, we may arrive at configurations, desired flow rates, that are physically infeasible. Instead, the search in the space of "maximum flow rate" controls and the use of EPANET ensured that the flows were always feasible.

[8] The maximum flow rate in all but two of the valves in Table II is just over the average demand. Since the reservoirs return close to the original state, the realised flow rates must be lower. Note that Rajendra Nagar I and Rajiv Nagar II have smaller configured flow rates than the daily average. The corresponding reservoirs will suffer a small deficit at the end of each period. As another example of realised flow rates being lower than the configured maximum, note that the searched maximum setting for the Bademakan valve in Table III is 80.64ℓps between 07:00 - 08:00 hrs. But the realised flow rate is approximately 40ℓps. Interestingly, the permissible limit, taking the diameter and the maximum velocity into account, is about 44ℓps. The maximum setting is well above this permissible maximum because our implementation ignored the per valve maximum flow rate setting. This can be easily fixed.

[9] Bademakan encounters zero flow rate twice, during 05:00 - 07:00 hrs and during 08:00 - 11:00 hrs. Rajiv Nagar II encounters zero flow rate during 07:00 - 08:00 hrs to help fill the Bademakan reservoir and prepare for the impending zero inflow into that reservoir from 08:00 - 11:00 hrs.

[10] Intel Core i5-4440 CPU at 3.10GHz $\times$ 4, 7.5 GB RAM, base system openSUSE 13.2 (Harlequin) (x86_64) 64 bit.

Acknowledgments

This work was supported by the Department of Electronics and Information Technology, Government of India, and the Robert Bosch Centre for Cyber Physical Systems (IISc).

References

[1] *Guidance Notes for Continuous Water Supply (24-7 Supply).* Ministry of Urban Development (Government of India), 2009.

[2] *Manual on Water Supply and Treatment (CPHEEO).* Ministry of Urban Development (Government of India), 1999.

[3] *EPANET 2,* U.S. Environmental Protection Agency, National Risk Management Research Laboratory, Office of Research and Development, Cincinnati.

[4] U. Manohar and M. S. Mohan Kumar, "Modeling equitable distribution of water: Dynamic inversion-based controller approach," *Journal of Water Resources Planning and Management*, 2013.

[5] G. Cembrano, G. Wells, J. Quevedo, R. Pérez, and R. Argelaguet, "Optimal control of a water distribution network in a supervisory control system," *Control engineering practice*, vol. 8, no. 10, pp. 1177–1188, 2000.

[6] P. W. Jowitt and C. Xu, "Optimal valve control in water-distribution networks," *Journal of Water Resources Planning and Management*, vol. 116, no. 4, pp. 455–472, 1990.

[7] T. Hastie, R. Tibshirani, J. Friedman, and J. Franklin, "The elements of statistical learning: data mining, inference and prediction," *The Mathematical Intelligencer*, vol. 27, no. 2, pp. 83–85, 2005.

[8] V. S. Borkar, "Stochastic approximation," *Cambridge Books*, 2008.

[9] M. Herrera, L. Torgo, J. Izquierdo, and R. Pérez-García, "Predictive models for forecasting hourly urban water demand," *Journal of hydrology*, vol. 387, no. 1, pp. 141–150, 2010.