

A Lagrangian Decomposition Approach for the Pump Scheduling Problem in Water Networks

Bissan Ghaddar, Joe Naoum-Sawaya*, Akihiro Kishimoto, Nicole Taheri, Bradley Eck

IBM Research, IBM Technology Campus Building 3, Mulhuddart, Dublin 15, Ireland.

{bghaddar@ie.ibm.com, joenaoum@ie.ibm.com, akihirok@ie.ibm.com, nicole.taheri@ie.ibm.com, bradley.eck@ie.ibm.com}

Abstract

Dynamic pricing has become a common form of electricity tariff, where the price of electricity varies in real time based on the realized electricity supply and demand. Hence, optimizing industrial operations to benefit from periods with low electricity prices is vital to maximizing the benefits of dynamic pricing. In the case of water networks, energy consumed by pumping is a substantial cost for water utilities, and optimizing pump schedules to accommodate for the changing price of energy while ensuring a continuous supply of water is essential. In this paper, a Mixed-Integer Non-linear Programming (MINLP) formulation of the optimal pump scheduling problem is presented. Due to the non-linearities, the typical size of water networks, and the discretization of the planning horizon, the problem is not solvable within reasonable time using standard optimization software. We present a Lagrangian decomposition approach that exploits the structure of the problem leading to smaller problems that are solved independently. The Lagrangian decomposition is coupled with a simulation-based, improved limited discrepancy search algorithm that is capable of finding high quality feasible solutions. The proposed approach finds solutions with gap guarantees. These solutions are compared to those found by a Mixed-Integer Linear Programming approach, which uses a piecewise-linearization of the non-linear constraints to find a globally optimal solution of the relaxation. Numerical testing is conducted on two real water networks and the results illustrate the significant costs savings due to optimizing pump schedules.

Keywords Large scale optimization, Lagrangian Relaxation, Pump Scheduling, Integer Programming.

*Corresponding Author

1 Introduction

With the demand for water expected to double by the year 2035 [26], water utilities are facing unprecedented challenges to maintain a reliable infrastructure with sustainable costs. The expected population growth and the shrinking supply of fresh water are among the main reasons for the urgent need to optimize the cost effectiveness of water resources. Energy costs are a major operating expense for water utilities; approximately 4% of all electricity used in the United States is attributed to the operation of potable water and waste water networks [7]. Besides the economic benefits, optimizing the operation of water networks could avoid unnecessary use of resources and minimize the ecological impact caused by pollution and greenhouse gas emissions.

Over the last few decades, practitioners and researchers have addressed several problems in water networks, mainly in reliability [32], network expansion [5, 33], pipe sizing [10], and network operations [28]. The problem of optimal pump scheduling has recently gained attention due to the increasingly complex electricity tariff schemes: energy utilities are providing incentives by offering cheaper electricity at low demand periods. For instance, the typical day-night pricing approach offers a night-time electricity price that is 50% cheaper than the day price. As renewable energy sources (i.e., wind and solar power) are being integrated into energy supply networks, electricity providers are beginning to offer dynamic electricity pricing, which varies in real time based on the realized supply and demand. Hence, forecasting electricity prices and optimizing pump schedules accordingly is becoming essential to capitalize on dynamic pricing and reduce operational costs.

Several approaches have been proposed for the optimal pump scheduling problem. Zessler and Shamir [37] proposed an iterative dynamic program that finds an optimal pump schedule for a 24-hour period given forecasted demands, the initial and final levels in the tanks, and the forecasted energy cost over the day. Brion and Mays [6] described a methodology that combines a non-linear optimization model with a hydraulic simulation model to implicitly enforce the conservation of energy and flow equations. Jowitt and Germanopoulos [20] presented a linear programming model where the constraints are linear approximations of the non-linear network equations; the parameters of the linear approximations are determined using simulation. Yu et al. [36] presented an optimization algorithm that uses the generalized reduced gradient method coupled with sensitivity analysis to find optimal schedules; this algorithm needs a feasible starting point and iterates so that all the interim points are feasible. Lansey and Awumah [22] introduced simplified hydraulic and cost functions for the pumps and implemented a dynamic programming optimization

algorithm. Sherali and Smith [31] use the Reformulation-Linearization Technique to construct a tight linear programming relaxation of the pump scheduling problem and then uses a branch-and-bound algorithm to obtain global optimal solutions. Sakarya and Mays [30] include water quality constraints and use a simulation optimization approach that combines a non-linear optimization solver with a water network simulator. McCormick and Powell [24] present a stochastic dynamic program that includes maximum demand charges by enforcing constraints and penalty costs on power use, which provides solutions for the case where a single maximum demand charge is considered. McCormick and Powell [25] presented a two-stage simulated annealing heuristic that produces high quality solutions in relatively short computational time. A two-stage heuristic approach was also proposed by Ulanicki et al. [35], where the first stage optimizes the tank trajectories and the second stage attempts to find a schedule that tracks the optimal trajectories. Another metaheuristic based on ant colony optimization was presented in López-Ibáñez et al. [23]. Giacomello et al. [15] also proposed a heuristic that is based on linear programming coupled with a greedy search algorithm and demonstrated the effectiveness of the method on water networks from the literature. Most recently, Gleixner et al. [16] proposed mixed-integer programming models coupled with problem-specific presolving for pump scheduling and demonstrated the effectiveness of the proposed approach on water networks from the industry.

The main challenge for the pump scheduling problem is the non-linearities in the hydraulic and energy conservation equations that define the way water flows through a network. These translate to non-linearities in the constraints and the objective of the optimization problem, which makes the problem difficult to solve, and there is no guarantee that any solution will be globally optimal. The problem is further complicated by multiple time periods (typically of 30 minutes each), which lead to a very large problem that is often challenging to solve within reasonable computational time. This paper proposes a Lagrangian decomposition [13, 14] approach to exploit the structure of the problem and alleviate the computational burden. The Lagrangian decomposition relaxes the constraint that links the time periods, decoupling the problem into smaller subproblems that are solved individually. This Lagrangian decomposition is implemented using a cutting plane approach. Similar approaches have been successfully applied to solve several challenging problems, such as [2, 8, 11, 14, 13, 17]. As the Lagrangian subproblem solutions are not likely to be feasible for the original problem, a neighborhood search based on Improved Limited Discrepancy Search (ILDS) [21] is proposed to explore the search space for high quality feasible solutions. Our proposed search algorithm couples ILDS with the EPANET [12] water network simulator,

which is used to verify the feasibility of solutions. ILDS provides an upper bound on the optimal objective function value, while the Lagrangian relaxation provides a lower bound. Thus, the proposed approach provides solutions of guaranteed quality. The proposed Lagrangian decomposition approach is further evaluated by comparing its performance against a Mixed-Integer Linear Programming (MILP) formulation; this MILP is a relaxation of the MINLP formulation that uses a piecewise-linearization of the non-linear constraints, and is solved using CPLEX [34]. Our results show that the proposed approach finds solutions that are close to optimal in only a fraction of the computational time of CPLEX. Furthermore, we show that dynamic pricing electricity tariff leads to a significant reduction in energy costs when compared to the day/night pricing. Finally, the numerical results show that the optimized pump schedules lead to a reduction in average tank levels thus reducing both pumping energy and pumping cost.

The contributions of the paper are threefold: (1) to the best of our knowledge, the proposed approach is the first to exploit the structure of the pump scheduling problem using Lagrangian decomposition in order to alleviate the computational burden, (2) the application of the simulation-based ILDS is novel and is shown to find solutions that are close to optimal in relatively short computational time, (3) we evaluate the effect of various dynamic pricing schemes that are currently being used in practice.

The rest of the paper is organized as follows. The problem formulation and the application of Lagrangian decomposition approach are detailed in Section 2. Computational results are presented in Section 3. Finally, Section 4 concludes and illustrates future research directions.

2 Problem Formulation

2.1 Notation

A water network consists of a set of nodes connected by pipes (or links). The nodes consist of three subsets: reservoirs (or water sources), tanks used to store water, and junctions that connect nodes and may have a demand. Pumps are a subset of the pipes, which are used to increase the pressure within the network by generating sufficient water flow to satisfy the demand. Without loss of generality, we assume that a pipe may not contain more than one pump and the pumps have a constant speed. We also assume that the planning horizon is divided into equally-sized time periods, and the water demand must be satisfied at each time period. Below is the terminology that we will use throughout the paper.

Parameters:

C_t : Electricity cost in time period t

γ : Specific weight of water

η_{ij} : Efficiency of pump installed on pipe (i, j)

Q_{ij}^U : Maximum rate of flow through pipe (i, j)

$D_{j,t}$: Demand at junction j in period t

E_i : Elevation of node i

P_j^L : Minimum water level in tank j

P_j^U : Maximum water level in tank j

A_j : Surface area of tank j

Δ_T : Length of each time period

Sets:

N : is the set of pipes

P : is the set of pipes that contain pumps

R : is the set of reservoirs

K : is the set of tanks

J : is the set of junctions

T : is the set of time periods

To formulate the pump scheduling problem, we define the following decision variables:

$q_{ij,t}$: Flow rate of link (i, j) in time period t ,

$h_{ij,t}$: Headloss in link (i, j) in time period t ,

$p_{i,t}$: Pressure at node i in time period t ,

$$s_{ij,t} = \begin{cases} 1 & \text{if the pump of link } (i, j) \text{ is on in period } t, \\ 0 & \text{otherwise.} \end{cases}$$

The following section presents the MINLP formulation for the pump scheduling problem and the proposed Lagrangian decomposition approach.

2.2 Mixed-Integer Non-linear Programming Formulation

The Mixed-Integer Non-linear Programming formulation of the pump scheduling problem is as follows:

$$[\text{OP}]: \min \sum_t^T \sum_{(i,j) \in P} \frac{\gamma C_t \Delta T}{\eta_{ij}} h_{ij,t} q_{ij,t} \quad (1)$$

$$\text{s.t. } q_{ij,t} \leq Q_{ij}^U \quad \forall (i,j) \in N/P, t \in T \quad (2)$$

$$q_{ij,t} \leq s_{ij,t} Q_{ij}^U \quad \forall (i,j) \in P, t \in T \quad (3)$$

$$\sum_i q_{ij,t} - \sum_k q_{jk,t} = D_{j,t} \quad \forall j \in J \quad (4)$$

$$q_{ij,t} ((p_{i,t} + E_i) - (p_{j,t} + E_j) - h_{ij,t}) \geq 0 \quad \forall (i,j) \in N/P, \forall t \in T \quad (5)$$

$$(p_{i,t} + E_i) - (p_{j,t} + E_j) - h_{ij,t} \leq 0 \quad \forall (i,j) \in N/P, \forall t \in T \quad (6)$$

$$s_{ij,t} ((p_{i,t} + E_i) - (p_{j,t} + E_j) - h_{ij,t}) = 0 \quad \forall (i,j) \in P, \forall t \in T \quad (7)$$

$$h_{ij,t} = (a_{ij} q_{ij,t}^2 + b_{ij} q_{ij,t} + c_{ij}) \quad \forall (i,j) \in N, \forall t \in T \quad (8)$$

$$P_{j,t}^L \leq p_{j,t} \leq P_{j,t}^U \quad \forall j \in K, t \in T \quad (9)$$

$$p_{j,t} = p_{j,t-1} + \frac{\sum_i q_{ij,t-1} - \sum_k q_{jk,t-1}}{\Delta T A_j} \quad \forall j \in K, \forall t \in T \quad (10)$$

$$q_{ij,t} \geq 0, s_{ij,t} \in \{0, 1\}. \quad (11)$$

The objective function (1) minimizes the power consumption. Constraints (2) enforce an upper bound on the rate of water flow in each of the pipes, and constraints (3) ensure that a pump must be on for the water to flow in the corresponding pipe. Constraints (4) conserve the flow of water, and constraints (5)-(8) are conservation of energy constraints based on the hydraulic properties of water networks. As shown in equations (8), the pipe headloss equation is assumed to follow a quadratic function where a_{ij} and b_{ij} are the resistance coefficients and c_{ij} is the shutoff head (see [9, 15] for details regarding the quadratic approximation of the pipe headloss function). Constraints (9) set the minimum and the maximum tank levels, and the constraints (10) are the tank mass balance equation. Finally, constraints (11) denote the binary and non-negativity conditions.

Problem [OP] is a non-convex MINLP that is computationally expensive to solve using a standard optimization software, with no guarantee of global optimality. In the following section, we describe a Lagrangian decomposition approach that exploits the structure of the problem.

The problem is decomposed by its time periods and smaller problems are solved iteratively in a cutting plane approach.

2.3 Lagrangian Decomposition

Problem [OP] exhibits a special structure that is suitable for Lagrangian relaxation [13], i.e., constraints (10) are the only constraints that link the time periods, and thus relaxing these constraints decomposes the problem into smaller problems that can be solved individually. Applying a Lagrangian relaxation to Constraints (10) with a multiplier μ , we get the following subproblem:

$$\begin{aligned} \text{[FSP]: } \min \quad & \sum_t \sum_{(i,j) \in P} \frac{\gamma C_t \Delta_T}{\eta_{ij}} h_{ij,t} q_{ij,t} + \sum_j \sum_t \mu_{jt} \left(p_{j,t} - p_{j,t-1} - \frac{\sum_i q_{ij,t-1} - \sum_k q_{jk,t-1}}{\Delta_T A_j} \right) \\ \text{s.t. } \quad & (2) - (9), (11), \end{aligned}$$

which decomposes into T subproblems, each of the form:

$$\text{[SP]}_t: \min \sum_{(i,j) \in P} \frac{\gamma C_t \Delta_T}{\eta_{ij}} h_{ij,t} q_{ij,t} + \sum_{j \in K} \left(\mu_{jt} p_{j,t} - \mu_{j,t+1} p_{j,t} - \mu_{j,t+1} \frac{\sum_i q_{ij,t} - \sum_k q_{jk,t}}{\Delta_T A_j} \right) \quad (12)$$

$$\text{s.t. } q_{ij,t} \leq Q_{ij}^U \quad \forall (i,j) \in N/P, \quad (13)$$

$$q_{ij,t} \leq s_{ij,t} Q_{ij}^U \quad \forall (i,j) \in P, \quad (14)$$

$$\sum_i q_{ij,t} - \sum_k q_{jk,t} = D_{j,t} \quad \forall j \in J \quad (15)$$

$$q_{ij,t} ((p_{i,t} + E_i) - (p_{j,t} + E_j) - h_{ij,t}) \geq 0 \quad \forall (i,j) \in N/P, \quad (16)$$

$$(p_{i,t} + E_i) - (p_{j,t} + E_j) - h_{ij,t} \leq 0 \quad \forall (i,j) \in N/P, \quad (17)$$

$$s_{ij,t} ((p_{i,t} + E_i) - (p_{j,t} + E_j) - h_{ij,t}) = 0 \quad \forall (i,j) \in P, \quad (18)$$

$$h_{ij,t} = (a_{ij} q_{ij,t}^2 + b_{ij} q_{ij,t} + c_{ij}) \quad \forall (i,j) \in N, \quad (19)$$

$$P_{j,t}^L \leq p_{j,t} \leq P_{j,t}^U \quad \forall j \in K \quad (20)$$

$$q_{ij,t} \geq 0, \quad s_{ij,t} \in \{0, 1\}. \quad (21)$$

Each subproblem [SP]_t is solved independently, and the solution of each [SP]_t is a vector $[q_t^h, p_t^h, h_t^h, s_t^h]_{h=1, \dots, \bar{H}_t}$ that minimizes the cost and satisfies constraints (12)-(21) at time t , where \bar{H}_t is the set of indices

for feasible solutions to $[\text{SP}]_t$. Hence, $[\text{SP}]_t$ can be written as:

$$Z[\text{SP}]_t = \min_{h=1, \dots, \bar{H}_t} \left\{ \sum_{(i,j) \in P} \frac{\gamma C_t \Delta_T}{\eta_{ij}} h_{ij,t}^h q_{ij,t}^h + \sum_{j \in K} \left(\mu_{jt} p_{j,t}^h - \mu_{jt+1} p_{j,t}^h - \mu_{jt+1} \frac{\sum_i q_{ij,t}^h - \sum_k q_{jk,t}^h}{\Delta_T A_j} \right) \right\}.$$

For a given value of the Lagrangian variables μ , it follows that $\sum_{t \in T} Z[\text{SP}]_t$ gives a lower bound on the optimal objective function of $[\text{OP}]$. The Lagrangian bound is the best lower bound that is given by the Lagrangian problem

$$\begin{aligned} & \max_{\mu} \sum_{t \in T} Z[\text{SP}]_t \\ & = \max_{\mu} \sum_{t \in T} \min_{h=1, \dots, \bar{H}_t} \left\{ \sum_{(i,j) \in P} \frac{\gamma C_t \Delta_T}{\eta_{ij}} h_{ij,t}^h q_{ij,t}^h + \sum_{j \in K} \left(\mu_{jt} p_{j,t}^h - \mu_{jt+1} p_{j,t}^h - \mu_{jt+1} \frac{\sum_i q_{ij,t}^h - \sum_k q_{jk,t}^h}{\Delta_T A_j} \right) \right\}. \end{aligned}$$

The Lagrangian problem is a nondifferentiable optimization problem that can be solved by a cutting plane method. Thus, defining

$$\theta_t := \min_{h=1, \dots, \bar{H}_t} \left\{ \sum_{(i,j) \in P} \frac{\gamma C_t \Delta_T}{\eta_{ij}} h_{ij,t}^h q_{ij,t}^h + \sum_{j \in K} \left(\mu_{jt} p_{j,t}^h - \mu_{jt+1} p_{j,t}^h - \mu_{jt+1} \frac{\sum_i q_{ij,t}^h - \sum_k q_{jk,t}^h}{\Delta_T A_j} \right) \right\},$$

the Lagrangian master problem is:

$$\begin{aligned} [\text{FMP}]: \quad & \max \sum_{t \in T} \theta_t & (22) \\ \text{s.t.} \quad & \theta_t \leq \sum_{(i,j) \in P} \frac{\gamma C_t \Delta_T}{\eta_{ij}} h_{ij,t}^h q_{ij,t}^h \\ & + \sum_{j \in K} \left(\mu_{jt} p_{j,t}^h - \mu_{jt+1} p_{j,t}^h - \mu_{jt+1} \frac{\sum_i q_{ij,t}^h - \sum_k q_{jk,t}^h}{\Delta_T A_j} \right) \quad \forall h \in \bar{H}_t, \quad \forall t \in T. & (23) \end{aligned}$$

Since \bar{H}_t is not known beforehand, problem $[\text{FMP}]$ is solved iteratively starting with a relaxation that ignores the unknown cuts (23). At each iteration, the relaxation of $[\text{FMP}]$ is solved and the resulting optimal values of the Lagrangian multipliers μ are used to solve the subproblems $[\text{SP}]_t$, thus leading to new inequalities of the form (23) that tighten the relaxation of $[\text{FMP}]$. Due to the relaxation, the optimal solutions $[\text{SP}]_t$, $\forall t \in T$ provide a lower bound on the optimal Lagrangian bound (optimal solution of $[\text{FMP}]$), and at every iteration the optimal solution of the relaxation of $[\text{FMP}]$ provides an upper bound. The algorithm is stopped when the upper and lower bounds are

within the desired ϵ gap. By relaxation, the Lagrangian bound is a lower bound on the optimal objective function value of [OP]. In Section 2.3.2, we describe a simulation-based heuristic that uses the Lagrangian solution to find a feasible solution to [OP] and hence an upper bound on the optimal objective function value of [OP].

In the following section, we discuss additional operational constraints that are imposed for practical implementation.

2.3.1 Additional Operational Constraints

Our computational experiments showed that the optimal pump schedules often contain frequent pump switching that would be impractical to carry out. Moreover, frequent activation and deactivation of pumps in short time periods reduces the lifetime of pumps and adds little benefit in reducing the energy cost. Hence, operators often require that pump schedules conform to a set of operational constraints that make them applicable in practice. For example, the following operational constraints would be typical:

- A pump should remain ON for a minimum amount of time τ_1 .
- A pump should remain OFF for a minimum amount of time τ_2 .
- A pump cannot be switched ON more than a maximum number of times max_starts during the planning horizon.

The above conditions can be enforced by including new binary variables and constraints into problem (1)-(11). The following binary decision variables:

$$y_{ij,t} = \begin{cases} 1 & \text{if the pump on link } (i,j) \text{ is turned ON at period } t, \\ 0 & \text{otherwise,} \end{cases}$$

$$z_{ij,t} = \begin{cases} 1 & \text{if the pump on link } (i,j) \text{ is turned OFF at period } t, \\ 0 & \text{otherwise,} \end{cases}$$

and the following constraints:

$$-s_{ij,t-1} + s_{ij,t} - y_{ij,t} \leq 0, \quad \forall (i, j) \in P, \forall t > 0, t \in T, \quad (24)$$

$$-s_{ij,t'} + y_{ij,t} \leq 0, \quad \forall (i, j) \in P, \forall t \in T, \forall t' \in T : t \leq t' \leq t + \tau_1, \quad (25)$$

$$s_{ij,t-1} - s_{ij,t} - z_{ij,t} \leq 0, \quad \forall (i, j) \in P, \forall t > 0, t \in T, \quad (26)$$

$$s_{ij,t'} + z_{ij,t} \leq 1, \quad \forall (i, j) \in P, \forall t \in T, \forall t' \in T : t \leq t' \leq t + \tau_2, \quad (27)$$

$$\sum_t y_{ij,t} \leq \text{max-starts}, \quad \forall (i, j) \in P, \quad (28)$$

ensure that the operational constraints are met. Constraints (24) force $y_{ij,t}$ to take a value of 1 if the pump on link (i, j) is turned ON at period t , and constraints (25) force the pump to remain ON for at least τ_1 time periods after it is switched ON. Constraints (26) force $z_{ij,t}$ to take a value of 1 if the pump on link (i, j) is turned OFF at period t , and constraints (27) force the pump to remain OFF for at least τ_2 time periods after it is switched OFF. Finally, constraints (28) sets an upper bound on the maximum number of times each pump can be switched ON during the planning horizon.

We note that constraints (24)-(28) link the time periods together and incorporating them in the Lagrangian decomposition approach makes the subproblem non-decomposable. These operational constraints could be relaxed in a similar way to the relaxation of constraints (10), but this may lead to a large number of Lagrangian multipliers. Instead, we elect to fully relax constraints (24)-(28) and not include them in the computation of the Lagrangian bound; however, these constraints are enforced in the simulation-based heuristic that is described next. We note that the Lagrangian bound remains a valid lower bound.

2.3.2 Simulation-based Improved Limited Discrepancy Search

Due to the relaxation, the Lagrangian decomposition algorithm often yields solutions that are not feasible for the original constraints. For our problem, the operational constraints are not included in the Lagrangian subproblems, making it even more unlikely that the Lagrangian solution is feasible. Search algorithms and metaheuristics are often used to transform an infeasible solution into a feasible one. Limited Discrepancy Search [18] and its improved version, the Improved Limited Discrepancy Search (ILDS) [21], were introduced as solution approaches for constraint satisfaction problems, where a subset of a search tree is explored for a given amount of time. To

the best of our knowledge, for the first time, ILDS is combined with a branch-and-bound algorithm and applied to solve a constraint optimization problem. ILDS starts from a given solution and explores branches of a search tree in order to find a feasible solution that satisfies the constraints, and potentially with a better objective function value. In the context of our approach, ILDS takes the solution of the Lagrangian decomposition and finds a feasible solution; the search tree is a set of nodes corresponding to the variables, and a branch of a node corresponds to assigning the variable a value of either 0 or 1.

Let $x_{k,t}$ be a variable in the domain $D = \{0, 1\}$, which represents the ON/OFF status of pump k ($1 \leq k \leq |P|$) at time period t ($1 \leq t \leq |T|$). The variables are arranged in the search tree by increasing time, and then by the variables that increase in order of pump identifier. A depth first search is then conducted, where at each node the partial assignment (i.e., values of the variables that are currently selected by the search) is checked for feasibility of the operational constraints. When values are assigned to all the variables at time t , EPANET hydraulic simulator [12] is used to verify if the current schedule satisfies the demand and to calculate the power consumption. If the partial assignment is infeasible or if the power consumption exceeds the current lowest known power consumption, then the search backtracks without assigning any values to the remaining variables. Here, the ‘lowest power consumption’ denotes the power consumption associated with the lowest cost pump schedule that satisfies the operational constraints and meets the demand for all the periods.

Given that the starting point is the solution of the Lagrangian decomposition, the ILDS approach is expected to find a feasible point after a small modification to the starting point. Let j denote the number of variables that have different values than the starting point. Starting with $j = 1$, the ILDS approach performs a depth first search by allowing re-assignments of the values of j variables. Upon completing the search, j is incremented and the search is repeated. If the algorithm finds a better solution, it immediately stops the current search and resets the starting point to the new better solution and j to 1. The search is then restarted. The algorithm terminates when a predetermined time limit is reached or when the search space is exhausted with $j = |P||T|$.

The computational results presented in Section 3 show that the proposed Lagrangian decomposition approach coupled with the simulation-based ILDS yields solutions with tight optimality gaps. To further assess the computational performance, in the following section we present a mixed-integer linear programming formulation of the pumps scheduling problem which can be

solved using CPLEX.

2.4 Mixed-Integer Linear Programming Formulation using Piecewise-Linearization

To assess the performance of the proposed Lagrangian decomposition approach, we consider the mixed-integer linear programming (MILP) formulation of the pump scheduling problem. The MILP formulation is obtained by using a piecewise-linearization of the non-linear equations. This leads to a MILP that can be solved with well-established solvers, such as CPLEX [34].

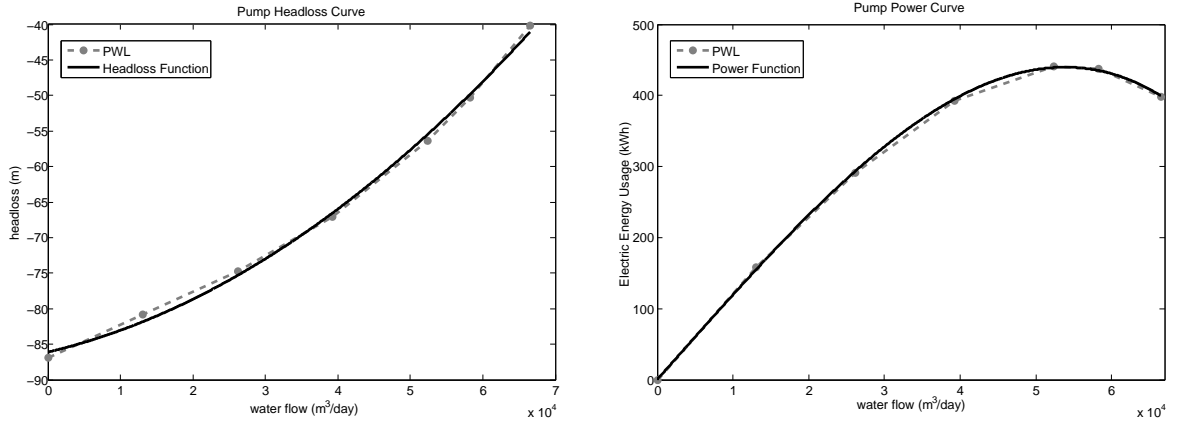


Figure 1: Example of Piecewise-Linear Functions Used in Formulation

The non-linear energy conservation equations (5)–(8) can be approximated by:

$$(p_{i,t} + E_i) - (p_{j,t} + E_j) = \hat{H}_{ij}(q_{ij,t}) \quad \forall (i, j) \in N, \quad \forall t \in T, \quad (29)$$

where $\hat{H}_{ij} : \mathbb{R} \rightarrow \mathbb{R}$ is a piecewise-linearization of the headloss function for a given pipe. Similarly, the non-linear objective function (1) is approximated with a piecewise-linearization. A sample of the piecewise-linearizations of the non-linear headloss and objective functions are shown in Figure 1.

Although, the formulation obtained by using piecewise-linearizations can be solved with commercial MILP solvers, the main computational challenge is due to the large number of additional binary variables that are used to model the piecewise-linear functions (there is one binary variable for each time period, for each pipe, for each piece of the piecewise-linear function). In order to reduce the number of binaries, we implement a pre-processing technique that uses the bounds on the flow rate values to eliminate unnecessary parts of the piecewise-linear functions.

2.4.1 Problem Pre-processing

For each pipe, a piecewise-linear function with κ pieces is used to describe the headloss function. The headloss function $\hat{H}_{ij}(\cdot)$ relates the difference between the pressures in its two adjacent nodes to the flow rate. Because of the elevation and pressure bounds of the connected nodes, it is often the case that the actual headloss will lie within a range that is contained by fewer than κ pieces. That is, due to the bounds on the pressures and since $\hat{H}_{ij}(\cdot)$ is an increasing function of $q_{ij,t}$, the feasible range of the flow can be explicitly calculated using the headloss equation (29), and the binary variables of the piecewise-linear pieces outside of this feasible flow range can then be eliminated. The solution time of the piecewise-linear formulation was greatly reduced when the pre-processing algorithm was implemented.

3 Numerical Testing




Figure 2: IND1 Water Network




Figure 3: Poormond Water Network

This section presents results for numerical experiments that were conducted on two networks. We call the first network IND1 (Figure 2). It consists of one reservoir, two pumps, and two tanks (4 nodes and 4 pipes) and is based on a real network provided by one of our industry partners. We call the second network Poormond (Figure 3). Poormond is a larger network that is also based on a real water distribution system in the United Kingdom. It consists of one reservoir, seven pumps, and five tanks (47 nodes and 44 pipes). The Poormond network is adapted from the Richmond water distribution network that is used in the literature [15]. The Richmond network available online [27] was modified to the Poormond network considered here by

1. controlling pumps by time rather than tank level;

2. removing Tank E to avoid use of an implied control valve;
3. lowering the elevation of seven nodes having zero demand by 10m to avoid errors due to negative pressures;
4. using the electricity price patterns of this paper.

The scheduling horizon is 24 hours, and the decisions to turn a pump ON or OFF are made at 30-minute intervals. Although a scheduling decision can be made every 30 minutes, operational constraints are required in practice to avoid frequent pump switching that shortens the life of the pumps. For example, in our implementation, pumps are required to stay on for at least 1 hour (2 periods) when switched ON. We also require that each pump is switched ON a maximum of 6 times during the 24-hour planning horizon. The final water level in the tanks is required to be at least the starting level, to ensure terminal conditions. These conditions are based on discussions with our industry partners and the approach that they currently use for the operation of their pumps.

The computations were performed on a Red Hat Linux blade server with a 3.5Ghz processor. The Lagrangian decomposition was implemented in AMPL [1] and BONMIN [3, 4], and IPOPT [19] was used as a non-linear optimization solver. The simulation-based ILDS that is described in Section 2.3.2 was implemented in C++ and EPANET [12] is used as a hydraulic and water network simulator. After the first feasible solution is found by ILDS, the time limit for ILDS is set to 200 seconds for IND1 and 500 seconds for Poormond. Finally, the MILP problem is implemented in CPLEX 12.4 OPL.

3.1 Computational Testing

This section provides computational results comparing the proposed Lagrangian decomposition against the MILP formulation that is solved with CPLEX. A time limit of 5 hours was imposed on CPLEX. Table 1 presents the results corresponding to the IND1 and Poormond networks. For each network, 5 different sets of energy prices are used. The prices are based on the real energy prices that occurred on May 21-25, 2013 and can be obtained from [29]. The following results are reported in Table 1.

Price : Energy Prices.

LB : Lagrangian bound which provides a lower bound on the optimal solution of [OP].

Iter : Number of iterations performed by the cutting plane algorithm.

CPU_{LR} : Total computational time in seconds for the cutting plane algorithm.

UB : Objective function value of the solution that is found by ILDS which provides an upper bound on the optimal solution of [OP].

GAP^a : % Gap between the Lagrangian upper and lower bounds

$$GAP^a = 100 \times \frac{UB-LB}{UB}.$$

CPU_H : Total computational time in seconds for ILDS.

CPU_T : Total computational time for the Lagrangian decomposition and ILDS, $CPU_T = CPU_{LR} + CPU_H$.

obj : Objective function value for the solution that is found by CPLEX.

GAP^b : CPLEX % Gap.

CPU : CPLEX total computational time.

Price	Lagrangian Decomposition							CPLEX		
	LB	Iter.	CPU_{LR} (sec.)	UB	GAP^a	CPU_H (sec.)	CPU_T (sec.)	Obj	GAP^b	CPU
IND1										
May 21	513.6	12	349	522.7	2%	203	552	530.8	2%	>5 hours
May 22	485.0	11	209	503.2	4%	201	410	508.4	3%	>5 hours
May 23	502.2	11	208	521.5	4%	203	411	537.7	4%	>5 hours
May 24	540.9	12	92	568.1	5%	203	295	568.4	3%	>5 hours
May 25	478.6	11	99	502.1	5%	201	300	490.5	2%	>5 hours
Average			191.4		4%	202.2	393.6		3%	>5 hours
Poormond										
May 21	123.3	25	2888	130.7	6%	555	3443	-	-	>5 hours
May 22	120.2	28	4134	139.4	14%	1312	5446	-	-	>5 hours
May 23	125.7	26	4408	140.2	10%	586	4994	-	-	>5 hours
May 24	138.7	28	4411	151.8	9%	577	4988	-	-	>5 hours
May 25	105.5	31	4739	130.3	19%	2289	7028	-	-	>5 hours
Average			4116		12%	1063.8	5179.8		-	>5 hours

-: No feasible solution was found within the 5 hours time limit.

Table 1: Computational Results evaluating the Lagrangian Decomposition Approach.

The results show that the proposed Lagrangian decomposition approach is capable of finding better solutions than CPLEX in significantly less computational time. Furthermore, CPLEX is only capable of finding solutions for the small IND1 network while for the Poormond network CPLEX could not find a feasible solution within the 5 hours time limit. Even for IND1, CPLEX

stopped at the 5 hours time limit with solutions that are within 3% gap on average. We note that for IND1, the feasible solutions that were found by the Lagrangian decomposition approach had better objective function value except for IND1 on May 25. For the Poormond network, the proposed Lagrangian decomposition found solutions with gaps ranging between 6% and 19% while CPLEX failed to find solutions for all the tested instances.

In the following section, we evaluate two electricity tariffs, day/night pricing and dynamic pricing, in terms of energy costs in pump scheduling.

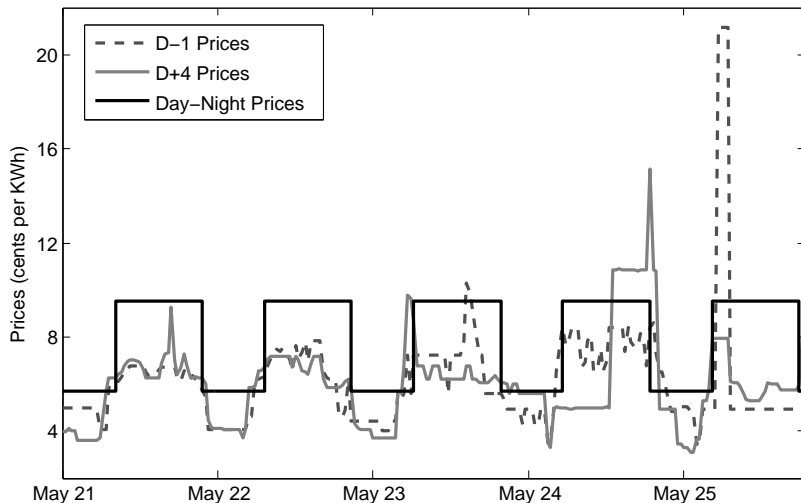


Figure 4: Dynamic and Day/Night Energy Prices.

3.2 Dynamic Pricing vs. Day/Night Pricing

This section compares two electricity tariff schemes available in Ireland based on the results that are obtained using the proposed Lagrangian decomposition approach. The Day/Night pricing scheme offers a 50% discount on the standard electricity rate during night hours (8:00am-11:00pm in winter and 9:00am-12:00am in summer). There is also a dynamic pricing scheme, where prices change every 30 minutes. Single Electricity Market (SEM) is the wholesale electricity market operating in the Republic of Ireland, which provides the forecasted and the realized electricity prices on their website [29]. We refer to the prices forecasted one day-ahead as “D-1 prices”. The actual prices that are realized are settled 4 days afterwards and hence we refer to the actual prices as “D+4 prices”. Hence, the pump schedules are optimized assuming D-1 prices for the energy cost, while the actual realized cost is computed based on the D+4 prices. Figure 4, illustrates the

different prices for the period May 21–25 2013, which are used for testing in this paper.

In this section, we evaluate the two pricing schemes in terms of impact on the optimal pump schedules and the resulting costs. First, we compute the optimal pump schedules and associated costs using the day/night pricing scheme. We compare this to the cost of using the forecasted D-1 prices to find optimal schedules, and computing the actual realized cost with D+4 prices. A summary of the results is given in Table 2. Note that the Poormond network contains a single set of demands and hence the daily costs are the same when assuming the day/night pricing. In contrast, IND1 was provided by our industry partner and contains different demands for each day, so the daily cost changes when using day/night pricing. The results indicate significant savings when dynamic pricing is used. Dynamic pricing leads to an average energy cost reduction of 24% for the IND1 network and 34% for the Poormond network. Thus, there is strong incentive for water utilities to adopt dynamic pricing, which further motivates the need for optimization methods that consider the changing electricity tariff and can find low-cost pump schedules.

Name	Cost in €				
	Day/Night	Forecasted D-1	Savings	Realized D+4	Savings
IND1					
May 21	699.5	522.7	25%	488.7	30%
May 22	684.8	503.2	27%	518.0	24%
May 23	693.9	521.5	25%	516.0	26%
May 24	675.3	568.1	16%	608.4	18%
May 25	688.3	502.1	27%	513.1	25%
Average			24%		25%
Poormond					
May 21	185.1	130.7	29%	114.6	38%
May 22	185.1	139.4	25%	124.0	33%
May 23	185.1	140.2	24%	116.7	37%
May 24	185.1	151.8	18%	136.0	27%
May 25	185.1	130.3	30%	119.3	36%
Average			25%		34%

Table 2: Costs with Dynamic Pricing vs. Day/Night Pricing

3.3 Pump Schedules

In this section, we illustrate the pump operation and the simulated tank levels for the pump schedules that were obtained using the proposed Lagrangian decomposition approach. A summary of the pump schedules is given in Table 3. Figure 5 also illustrates the computed pump schedules for May 21. As expected, we do not see a significant difference in the total number of hours that the pumps are ON, which is due to the fact that the system operator of IND1 network has put significant effort into matching the pumping schedules with the demands.

IND1												
	Current Schedule						Optimized Schedule					
	May 21	May 22	May 23	May 24	May 25	Average	May 21	May 22	May 23	May 24	May 25	Average
Pump 1	5.0	4.2	5.0	4.8	5.3	4.9	7.3	6.5	6.2	7.7	6.8	6.9
Pump 2	14.3	14.7	14.3	14.2	14.8	14.5	11.8	12.5	12.5	11.7	12.8	12.3
Poormond												
Pump 1	3.5	3.5	3.5	3.5	3.5	3.5	3	4	3.5	4	4	3.7
Pump 2	18	18	18	18	18	18.0	16	16	14	17	16	15.8
Pump 3	7	7	7	7	7	7.0	6	6	6	5.5	6	5.9
Pump 4	17	17	17	17	17	17.0	18	18	18	18	18	18.0
Pump 5	17	17	17	17	17	17.0	13.5	12	12.5	11.5	13.5	12.6
Pump 6	12.5	12.5	12.5	12.5	12.5	12.5	13	13.5	13.5	13.5	13	13.3
Pump 7	13.5	13.5	13.5	13.5	13.5	13.5	10	12	13.5	11	11	11.5

Table 3: Number of Hours where a Pump is ON.

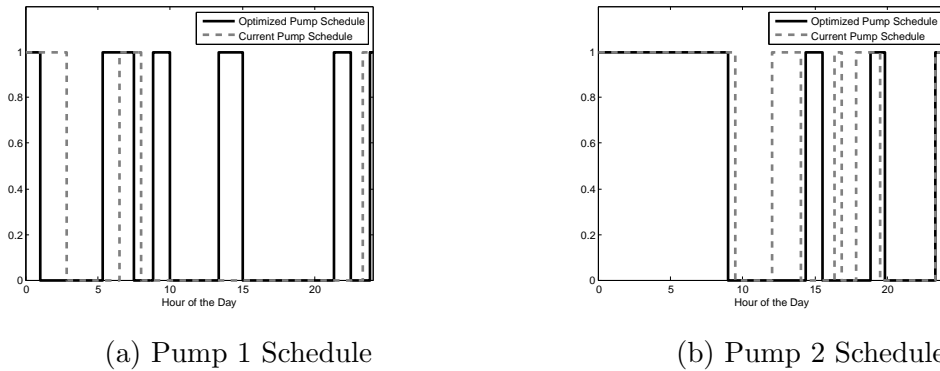


Figure 5: IND1 May 21 Pump Schedules

In the schedules actually implemented in the IND1 network, the pumps are active for a total of 96.6 hours, while in the optimized schedules, the pumps are active for a total of 95.8 hours. However, the optimized schedules use Pump 1 more than in the existing schedules, which results in a significant decrease in cost. Moreover, as shown in Table 4 and in Figure 6, the optimized schedules allow the tanks to operate at lower levels than they would with the current schedules: the average tank level of Tank 1 over the five days (May 21 - May 25) is 1.7m with the optimized schedules, compared to an average level of 2.3m with the current schedules; this is a 26% decrease in average tank level. Similarly, for Tank 2, the average level is 1.5m compared to 2.2m with the existing schedules, which represents a 32% decrease in average tank level (note that the existing schedules that are currently being used by our industry partner to operate IND1 lead to small violations of the desired maximum tank limits). Operating water tanks at lower levels reduces the energy required to fill the tanks, thus lowering the pumping cost. However, there is a tradeoff in system reliability because higher tank levels provide a buffer against changes in system demand. For this reason, an appropriate minimum tank level should be selected for the system under

		IND1											
		Current Schedule					Optimized Schedule						
		May 21	May 22	May 23	May 24	May 25	Average	May 21	May 22	May 23	May 24	May 25	Average
Tank 1		2.12	2.30	2.39	2.14	2.38	2.275	1.58	1.94	1.65	1.79	1.54	1.70
Tank 2		2.04	2.23	2.31	2.06	2.30	2.19	1.43	1.80	1.51	1.62	1.39	1.55
		Poormond											
Tank 1		1.43	1.43	1.43	1.43	1.43	1.43	1.26	1.06	0.82	0.93	1.38	1.09
Tank 2		2.54	2.54	2.54	2.54	2.54	2.54	1.63	1.67	1.53	1.47	1.50	1.56
Tank 3		1.56	1.56	1.56	1.56	1.56	1.56	1.74	1.74	1.71	1.71	1.62	1.70
Tank 4		2.75	2.75	2.75	2.75	2.75	2.75	2.71	2.91	2.88	2.80	2.60	2.78
Tank 5		1.46	1.46	1.46	1.46	1.46	1.46	1.53	1.26	1.18	1.62	1.58	1.43

Table 4: Average Tank Levels




Figure 6: IND1 May 21 Tank Levels

consideration.

Similar results are observed for the Poormond network. However, we note that we do not know the actual pump schedules implemented in the Poormond network. Hence, we created an initial feasible pump schedule before applying the proposed optimization approaches. Therefore, the results for the existing schedule of the Poormond network refer to an initial feasible schedule that we created. As shown in Table 3, the initial schedule has the pumps ON for a total of 442.5 hours, while for the optimized schedules the pumps are ON for a total of 404 hours. The exact pump schedule for May 21 is shown in Figure 7. The optimized schedule results in a 16% reduction in cost over the five days (May 21 - May 25), from €828.54 with the initial schedule down to €693.18 with the optimized schedule. Furthermore as shown in Table 4, the optimized schedule results in 11% reduction in average tank levels. The exact tank levels for May 21 are shown in Figure 8.

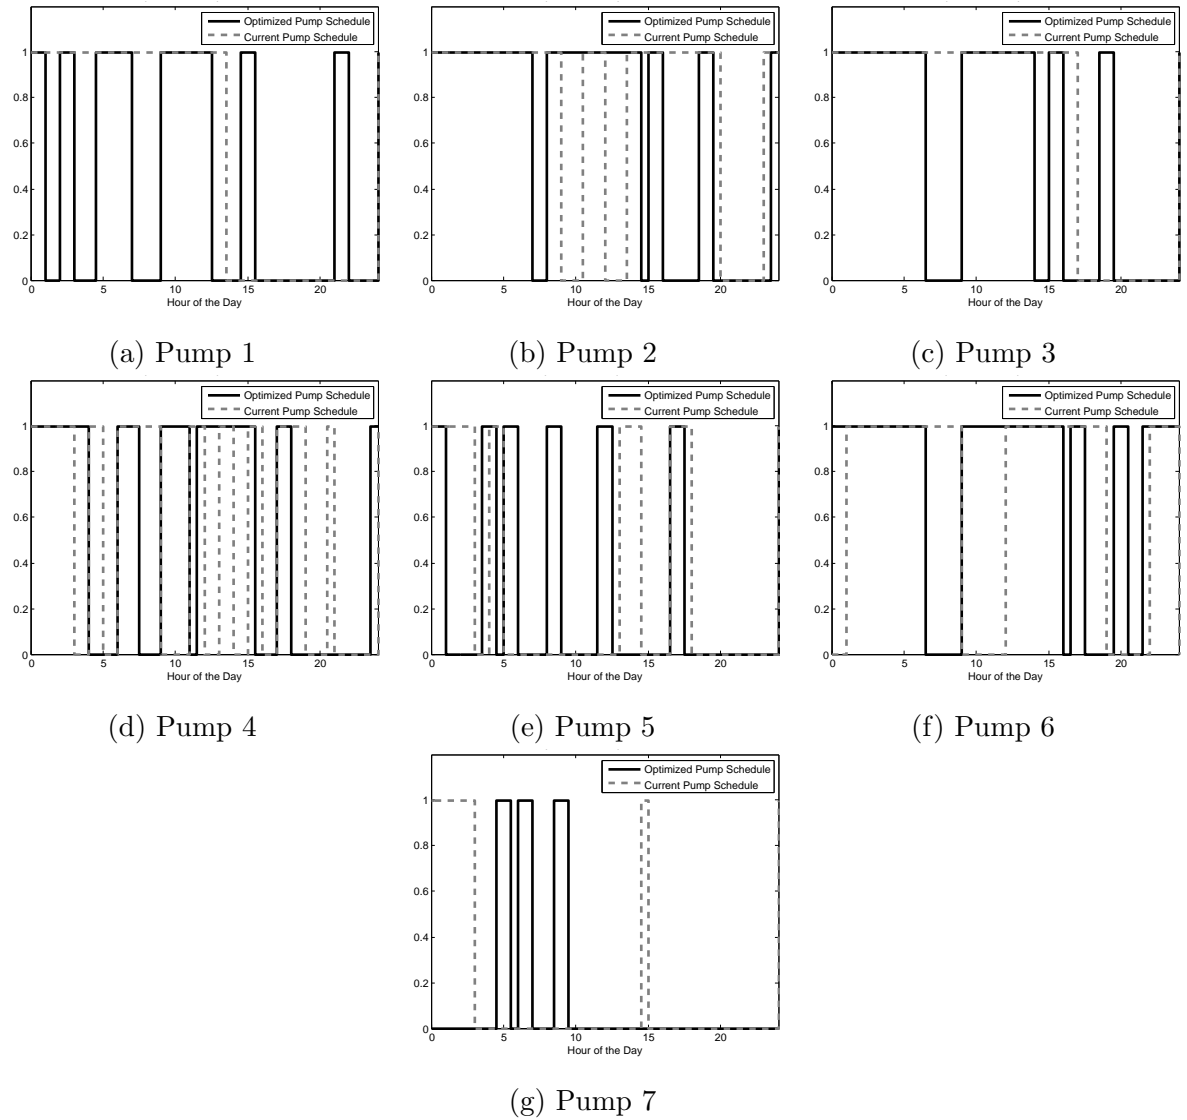


Figure 7: Poormond May 21 Pump Schedules

4 Conclusion

This paper considers the problem of optimizing pump schedules in water networks, which can be formulated as a mixed-integer non-linear program. Due to the non-linearities, the size of water networks, and the discretization of the planning horizon, the problem is not solvable within reasonable computational time using standard optimization solvers. Hence, this paper presents a Lagrangian decomposition algorithm that exploits the structure of the problem. In order to find high-quality feasible solutions, the Lagrangian decomposition is coupled with a simulation-based limited discrepancy search algorithm to provide solutions with optimality guarantees. Computational testing showed that the proposed approach finds high quality solutions in small com-

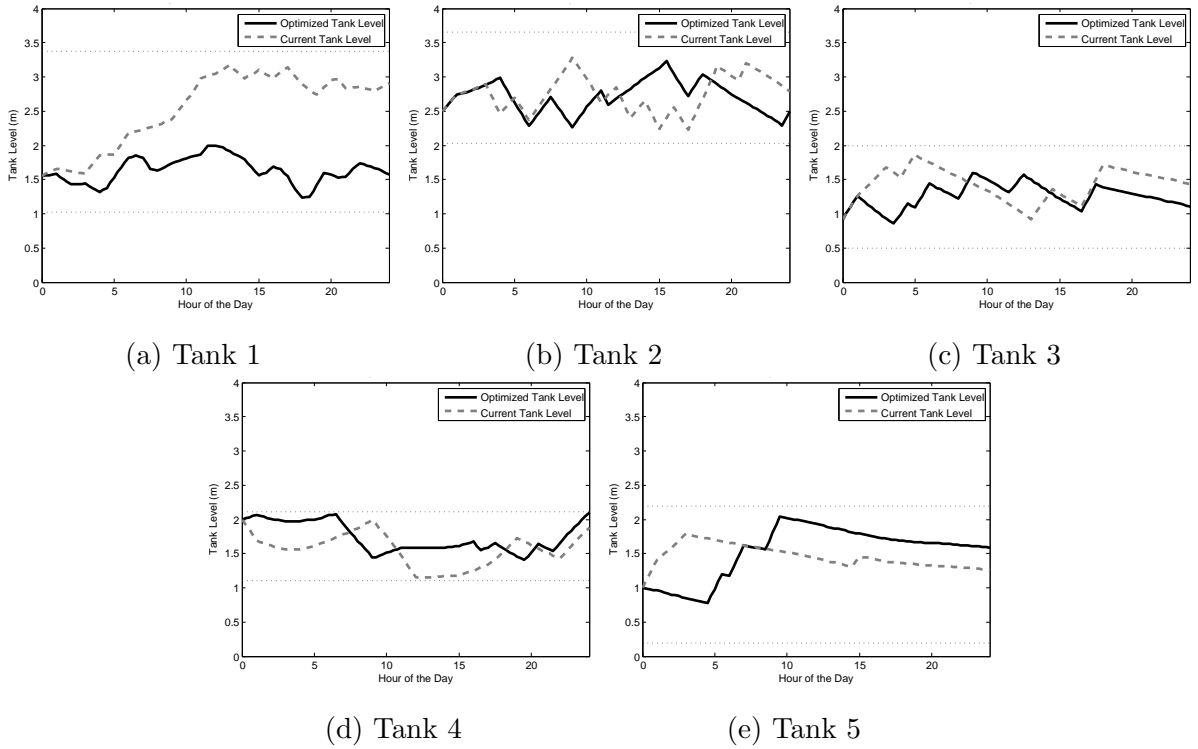


Figure 8: Poormond May 21 Tank Levels

putational time. Furthermore, our testing showed that optimized schedules typically lead to a decrease in average water levels in tanks, and the use of dynamic pricing for electricity tariff leads to significant savings in terms of energy costs.

In future work, we plan to investigate the application of the Lagrangian decomposition and proposed simulation-based limited discrepancy search to other problems that arise in water network planning. Specifically, we plan to apply this method to water pumps that are controlled by tank levels, which are common in water networks. In such a setting, the use of a similar methodology could be applied to optimize the control levels used to activate and deactivate water pumps.

Acknowledgements

We are grateful to Albert Akhriev for implementing parts of the EPANET simulator and to Paulito Palmes for his insights on dynamic pricing. We are also very thankful for the valuable feedback that we received from Sean Mckenna, Martin Mevissen, and Robert Shorten who contributed to the success of this research.

References

- [1] Ampl, 2011. URL <http://www.ampl.com/>.
- [2] Turgut Aykin. Lagrangian relaxation based approaches to capacitated hub-and-spoke network design problem. *European Journal of Operational Research*, 79(3):501–523, 1994.
- [3] Pierre Bonami, Lorenz T. Biegler, Andrew R. Conn, Gérard Cornuéjols, Ignacio E. Grossmann, Carl D. Laird, Jon Lee, Andrea Lodi, Francois Margot, Nicolas Sawaya, and Andreas Wechter. An algorithmic framework for convex mixed integer nonlinear programs. *Discrete Optimization*, 5(2):186 – 204, 2008.
- [4] Bonmin (v. 1.5), 2011. URL <http://projects.coin-or.org/Bonmin>.
- [5] Cristiana Bragalli, Claudia D’Ambrosio, Jon Lee, Andrea Lodi, and Paolo Toth. On the optimal design of water distribution networks: a practical MINLP approach. *Optimization and Engineering*, 13(2):219–246, 2012.
- [6] Lehar M Brion and Larry W Mays. Methodology for optimal operation of pumping stations in water distribution systems. *Journal of Hydraulic Engineering*, 117(11):1551–1569, 1991.
- [7] David Denig-Chakroff. Reducing electricity used for water production: Questions state commissions should ask regulated utilities. Technical report, National Regulatory Research Institute, 2008.
- [8] Ali Diabat, Tarek Abdallah, and Andreas Henschel. A closed-loop location-inventory problem with spare parts consideration. *Computers & Operations Research (to appear)*, 2013.
- [9] Bradley Eck and Martin Mevissen. Fast non-linear optimization for design problems on water networks. In *World Environmental and Water Resources Congress 2013*, pages 696–705, 2013.
- [10] Gideon Eiger, Uri Shamir, and Aharon Ben-Tal. Optimal design of water distribution networks. *Water resources research*, 30(9):2637–2646, 1994.
- [11] Samir Elhedhli, Lingzi Li, Mariem Gzara, and Joe Naoum-Sawaya. A branch-and-price algorithm for the bin packing problem with conflicts. *INFORMS Journal on Computing*, 23(3):404–415, 2011.
- [12] EPANET (v. 2), 2008. URL <http://www.epa.gov/nrmrl/wswrd/dw/epanet.html>.

- [13] Marshall L Fisher. The Lagrangian relaxation method for solving integer programming problems. *Management science*, 50(12):1861–1871, 2004.
- [14] Arthur M Geoffrion. Lagrangian relaxation for integer programming. In *50 Years of Integer Programming 1958-2008*, pages 243–281. Springer, 2010.
- [15] Carlo Giacomello, Zoran Kapelan, and Matteo Nicolini. Fast hybrid optimization method for effective pump scheduling. *Journal of Water Resources Planning and Management*, 139(2):175–183, 2012.
- [16] A Gleixner, Harald Held, Wei Huang, and Stefan Vigerske. Towards globally optimal operation of water supply networks. *Numerical Algebra, Control and Optimization (to appear)*, 2012.
- [17] Fatma Gzara and Erhan Erkut. A Lagrangian relaxation approach to large-scale flow interception problems. *European Journal of Operational Research*, 198(2):405–411, 2009.
- [18] William D Harvey and Matthew L Ginsberg. Limited discrepancy search. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 607–615, 1995.
- [19] Ipop (v. 3.10), 2011. URL <http://projects.coin-or.org/Ipop>.
- [20] Paul W Jowitt and George Germanopoulos. Optimal pump scheduling in watersupply networks. *Journal of Water Resources Planning and Management*, 118(4):406–422, 1992.
- [21] Richard E Korf. Improved limited discrepancy search. In *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI'96)*, pages 1202–1207, 1996.
- [22] Kevin E Lansey and Kofi Awumah. Optimal pump operations considering pump switches. *Journal of Water Resources Planning and Management*, 120(1):17–35, 1994.
- [23] Manuel López-Ibáñez, T Devi Prasad, and Ben Paechter. Ant colony optimization for optimal control of pumps in water distribution networks. *Journal of water resources planning and management*, 134(4):337–346, 2008.
- [24] Gregory McCormick and Roger Powell. Optimal pump scheduling in water supply systems with maximum demand charges. *Journal of water resources planning and management*, 129(5):372–379, 2003.

- [25] Gregory McCormick and Roger Powell. Derivation of near-optimal pump schedules for water distribution by simulated annealing. *Journal of the Operational Research Society*, 55(7): 728–736, 2004.
- [26] National Geographic, 2013. URL <http://news.nationalgeographic.com/news/energy/2013/01/130130-water-demand-forenergy-to-double-by-2035/>.
- [27] Richmond Water Distribution Network, 2013. URL <http://emps.exeter.ac.uk/engineering/research/cws/downloads/benchmarks/>.
- [28] Vilas Nitivattananon, Elaine C Sadowski, and Rafael G Quimpo. Optimization of water supply system operation. *Journal of Water Resources Planning and Management*, 122(5): 374–384, 1996.
- [29] Single Electricity Market Operator, 2013. URL <http://www.sem-o.com>.
- [30] A Burcu Altan Sakarya and Larry W Mays. Optimal operation of water distribution pumps considering water quality. *Journal of Water Resources Planning and Management*, 126(4): 210–220, 2000.
- [31] Hanif D Sherali and Ernest P Smith. A global optimization approach to a water distribution network design problem. *Journal of Global Optimization*, 11(2):107–132, 1997.
- [32] Hanif D Sherali, Ernest P Smith, and Seong-in Kim. A pipe reliability and cost model for an integrated approach toward designing water distribution systems. In *Global Optimization in Engineering Design*, pages 333–354. Springer, 1996.
- [33] Hanif D Sherali, Shivaram Subramanian, and GV Loganathan. Effective relaxations and partitioning schemes for solving water distribution network design problems to global optimality. *Journal of Global Optimization*, 19(1):1–26, 2001.
- [34] IBM ILOG CPLEX Optimization Studio, 2013. URL <http://www-03.ibm.com/software/products/us/en/ibmilogcpleoptistud/>.
- [35] Bogumil Ulanicki, Jens Kahler, and H See. Dynamic optimization approach for solving an optimal scheduling problem in water distribution systems. *Journal of Water Resources Planning and Management*, 133(1):23–32, 2007.

- [36] G. Yu, R.S. Powell, and M.J.H. Sterling. Optimized pump scheduling in water distribution systems. *Journal of Optimization Theory and Applications*, 83(3):463–488, 1994.
- [37] Uri Zessler and Uri Shamir. Optimal operation of water distribution systems. *Journal of Water Resources Planning and Management*, 115(6):735–752, 1989.