Greedy Budget Allocation for Optimizing the Performance of Water Distribution Systems Under Intermittent Supply Conditions

Gal Perelman[1], Lu Xing[2], Mashor Housh[3], Venu Kandiah[4], Barak Fishbain[5], and Ehsan Shafiee[6]

[1,5] Faculty of Civil and Environmental Engineering, Technion – Israel Institute of Technology, Haifa, Israel
[2,6] Xylem Inc., 1 International Drive, Rye Brook, NY 10573, United States
[3] Department of Natural Resources and Environmental Management, University of Haifa, Haifa, Israel
[4] DHI Water & Environment, Inc., 141 Union Blvd #320, Lakewood, CO 80228, United States

[1] *0000-0003-0345-4051, gal-p@campus.technion.ac.il,*
[2] *0000-0002-8881-5487, lu.xing@xylem.com ,*
[3] *0000-0003-4924-2636, mhoush@univ.haifa.ac.il ,*
[4] *veka@dhigroup.com,*
[5] *0000-0003-4211-7445, fishbain@tx.technion.ac.il,*
[6] *0000-0002-5568-146X, ehsan.shafiee@xylem.com,*

Abstract

Intermittent water supply is a prevalent strategy employed in water distribution systems (WDS) facing deteriorating conditions. However, this approach can result in several drawbacks, including insufficient supply, pressure drops, water losses due to leakages, and unequal service levels. Furthermore, these issues often exacerbate when consumers establish private storage facilities and increase the peak demand, which leads to a feedback loop of worsening conditions. Thus, when a budget is available, restoring the system functionality as fast as possible is crucial. The current study presents a method to optimize the rehabilitation of intermittent water supply by improving system design through investments and by operational control settings. The method was developed for the challenge presented in the Battle of the Intermittent Water Supply (BIWS), where the network performance is evaluated through nine different objectives over five years of planning and rehabilitation horizon. The proposed method is based on a greedy optimization approach that was specifically tailored to the challenge of optimizing WDS under extreme hydraulic conditions. To overcome the formidable computational burden in the BIWS challenge, several heuristics are presented for reducing the search space. The results obtained reflect a dramatic improvement in the network performance, with 97.8% of the consumers having continuous supply and water loss reduced from 47% to 23.7% of the total inflow. We also present a generic greedy approach that allows it to be applied to any water network for various decision-making problems.

Keywords: water distribution systems, Greedy algorithm, optimization, Intermittent water supply, Network rehabilitation, Battle of the Intermittent Water Supply

1. Introduction

The provision of a safe and reliable water supply is essential for the well-being and development of communities worldwide. However, in many parts of the world, especially in developing countries, water distribution systems (WDS) face numerous challenges, such as depletion of water sources and degeneration of infrastructure assets. With limited financial resources, water utilities struggle to address these issues directly, and instead, they adopt an alternative strategy of intermittent supply. Intermittent water supply refers to the non-constant availability of water service, where water is provided for limited hours for all or some of the users. Many causes lead WDS to the state of intermittent supply, ranging from sudden, dramatic events such as earthquakes to a sequence of inadequate planning and management decisions (Simukonda et al. 2018). Other factors that impact the ability to maintain continuous supply include rapid urbanization, which is common in developing countries (Bakker et al. 2008), climate change that affects water availability (Miyan 2015), and other instabilities that characterize developing countries. Consumers are the first to suffer from irregular water supply as it disturbs the daily routine and affects the most basic day-to-day actions that rely on water usage. Additionally, irregular supply is associated with health issues and increases the probability of contamination (Ingeduld et al. 2007; Kumpel and Nelson 2016). Other indirect impacts include illegal connections and self-storing, leading to supply inequality (Gottipati and Nanduri 2014) and water waste (Mokssit et al. 2018). While managed intermittent supply attempts to address some of these challenges by maintaining at least partial supply, it cannot serve as a long-term solution. Moreover, this strategy can worsen the system's state by creating a deteriorating feedback loop (Galaitsi et al. 2016; Vairavamoorthy et al. 2008). For example, non-continuous supply prompts consumers to store water independently, which can worsen the pressure drop during the supply hours (de Marchis et al. 2010). Other examples of deteriorating feedback loops include water sources at risk of exhaustion due to leaks and extreme hydraulic states leading to increased energy consumption, which in turn reduces the budget available for system rehabilitation. Therefore, it is essential to develop strategies to optimize planning and rehabilitation of the system toward continuous pressurized operation as well as optimize the intermittent supply operation of the system during the transition period. While several studies reviewed the causes and effects of intermittent supply (Galaitsi et al. 2016; Mokssit et al. 2018; Simukonda et al. 2018), the question of how to recover systems back to continuous pressurized supply has received less attention. The actions that utilities can take to improve the functionality of WDS can be grossly divided into two types: investment actions and operations actions. Investments usually refer to the installation of new facilities like pumps and valves while operation actions refer to the control of valves and pumps. Past studies addressed both aspects. Valves control optimization that maximizes the duration of sufficient pressure and minimizes pressure variations was suggested by (Solgi et al. 2020). A similar study presented by Gullotta et al. (2021), involves both the optimization of valve location and their control setting. Multi-objective approaches were also considered. Ayyash et al. (2024) suggested a two-stage approach where first a design stage is taken to sectorize the network to District Metered Areas (DMAs), and then a control stage for optimizing the scheduling of valves and pumps. Another combined design-operation approach was suggested by Nyahora et al. (2020), which considers a more comprehensive design perspective. The design decisions include pipe replacements and installing new pumps and tanks. All the above studies have developed simulation-optimization based methodologies that greatly depend on hydraulic modeling. Nevertheless, the modeling of intermittent water supply is not trivial and substantially different from pressurized systems.

Intermittent Hydraulic Modeling

A most noticeable aspect of intermittent hydraulic modeling is low pressure, which restricts water consumption and, in some cases, causes pipes to drain and refill according to the intermittent operation cycle (Ingeduld et al. 2007). As a result of the low pressure, the demand-driven analysis assumption that consumers' demands are continuously satisfied is no longer valid. The most common approach to cope with this situation is a pressure-dependent analysis (PDA) (Siew and Tanyimboh 2012). While PDA can simulate the effect of low pressure on demand, it still assumes full pipe flow conditions. Hence, other extreme hydraulic conditions could be developed in the system, causing the hydraulic simulations to be unstable. When hydraulic simulations are needed to evaluate the rehabilitation strategies, these numerical instabilities add more complexity to the already challenging problem of optimizing the WDS performance. For instance, in the popular hydraulic simulator EPANET (Rossman 2000), under insufficient water availability conditions, tanks might be drained beyond their minimum level , leading to simulation errors. Another example is pressure drops to negative values, which EPANET is not designed to handle, thus, the result might not reflect real hydraulic behavior. For example, when leaks in the system are modeled with EPANET emitter coefficients, negative pressure conditions will (artificially) cause the leak to function as a source, delivering water into the system.

Several studies addressed these challenges of intermittent supply hydraulic modeling. Mohan and Abhijith (2020) developed a pressure-driven partial flow model for the hydraulic simulation of full and partial flows. Gullotta and Campisano (2024) presented an approach based on the open-source storm water management model (SWMM) software (Rossman 2009) to model the partial flows in intermittent systems. A review and comparison of more intermittent modeling techniques can be found in Abdelazeem and Meyer (2024). The described hydraulic complexity poses a significant challenge to optimize WDS under intermittent conditions. The implementation of classic mathematic programming optimization methods such as linear and nonlinear programming becomes infeasible. Moreover, due to the extreme states that the system is subjected to, the simulation runtime tends to be relatively long. Thus, utilizing hybrid optimization methods (e.g., evolutionary algorithms) becomes more challenging.

<u>Objectives</u>

Like any other optimization problem, optimal planning and management of intermittent systems have unique objectives (Ilaya-Ayza et al. 2017a; Solgi et al. 2020). On the investment side, a primary objective is to restore the system functionality as quickly as possible, which often involves actions such as leak repairing and the replacement of pipes and pumps. Typically, these objectives may be conflicting. For example, replacing pumps to increase the supply capacity will also increase the system pressure, which in turn will increase the leakage water loss. On the operation side, various objectives must be considered, including supply continuity, supply pressure, and supply equity. Supply equity is particularly important in intermittent systems since significant pressure drops prevent the supply to the high points and locations furthest away from the sources, which leads to inequality in service levels between consumers. Additionally, standard operational objectives such as energy cost and leakage volume reduction need to be considered.

The main conclusion from the above is that optimizing the performance of WDS under intermittent conditions is a highly complex task, requiring coping with nonstandard hydraulic conditions while optimizing multiple conflicting objectives in a high-dimension decision space. Some past studies suggested customized objectives for the case of intermittent supply, mainly supply equity (Gullotta et al. 2021; Hendrickson and Sela 2024), supply pressure and supply hours (Ilaya-Ayza et al. 2017b).

To address these challenges, the current paper presents an optimization methodology developed as part of the Battle of Intermittent Water Supply (BIWS). The BIWS was held as an optimization competition aiming to find the best strategy to improve intermittent WDS performance. In the following sections, a detailed description of the BIWS challenge is provided, followed by a presentation of the developed optimization methodology. Later, the results obtained from the proposed method are presented and discussed, and finally, conclusions and insights are provided.

2. <u>The BIWS Case Study</u>

The BIWS network, depicted in Figure 1, is mainly a gravity-fed system, sourcing water from a natural spring in the south-west edge of the system (R1) with a 200 L/s flow capacity. In addition to the main gravity source, the network incorporates five pumping wells (W1_RI, W2_SA, W3_AB, W4_SM, and W5_PL) with a total flow capacity of 100 L/s. The network consists of 2,859 nodes and 3,231 pipes with 3,591 leaks distributed along them. Other facilities in the network are four storage tanks (T1_CO, T2_PL, T3_MO, and T4_CU) with different volumes located in different parts of the network; a pumping station with two parallel pumps (B_PT1, B_PT2); and control valves that allow the isolation of specific sections of the network. One of these valves is a flow control valve (FCV) that regulates the flow from the south to the north part of the network. Another valve is a pressure sustaining valve (PSV) responsible for maintaining the pressure at the main tank (T1_CO). The third valve is an FCV that restricts the flow from the natural spring to 200 L/s, and the fourth valve is a pressure reducing valve (PRV) that can regulate pressures in the south-west zone of the network (the detailed files that describe the WDS, in inp format, are provided in the supplementary materials S.1). According to the BIWS rules, leakages can be modeled in various ways. In this study, leakages were modeled by splitting pipes (i.e., adding a node) in the location specified in the data files. Another node with an emitter coefficient was added near the location of the pipe split to represent the leak. This leak node is then connected to the node that splits the pipe by adding a check valve. The check valve prevents water flow into the network during negative pressure conditions.

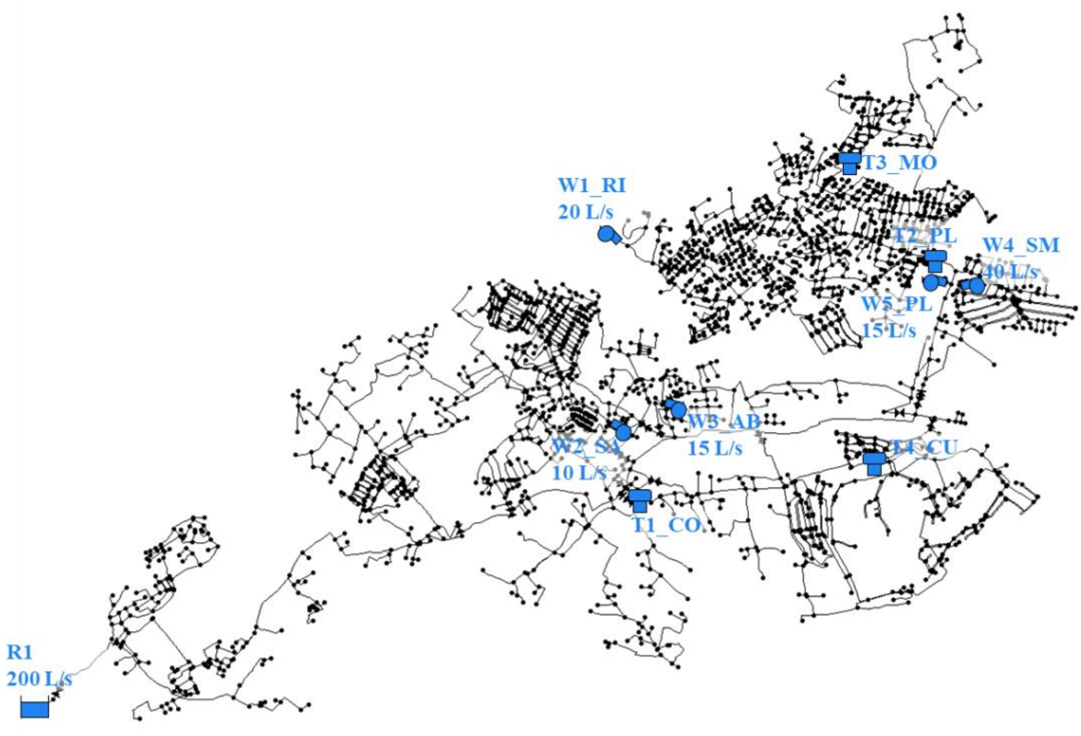

Figure 1 - BIWS network layout with water sources and tanks

The network is encumbered by multiple constraining factors. These include many leakages causing significant water loss and pressure drop. Additionally, the depletion of the underground water level prevents existing pumps from supplying water from wells. This is because the head required to pump from depleted underground water level to the ground elevation exceeds the pump’s shut-off head. Furthermore, some pipes have insufficient diameters, exacerbating pressure loss. The consequences of these conditions include the development of negative pressures, the inability to fill the storage tanks, and inequity in water supply between different zones. As a result, the water supply falls far short of satisfying the common-level service standards. The challenge presented in the BIWS asks to improve the network performance according to nine different indicators (Table 1) during a representative operational period of one week (168 hours).

Table 1 – Optimization objectives

#	Description	Type	Formula
1	Proportion of the number of effective hours a subscriber is served	Max	$I_1 = \frac{\sum_{j=1}^{6}\sum_{i=1}^{N} n_{i,j}}{N \cdot 24 \cdot 364 \cdot 6}$
2	Proportion of subscribers with continuous service	Max	$I_2 = \frac{\sum_{j=1}^{6}\sum_{i=1}^{N} w_{i,j}}{N \cdot 6}$
3	Volume of water leakage	Min	$I_3 = \frac{\sum_{j=1}^{6}\sum_{l=1}^{L_j} V_{l,j}}{\sum_{j=1}^{6}\sum_{s=1}^{S_j} V_{s,j}}$
4	Proportion of volume of water supplied to users	Max	$I_4 = \frac{\sum_{j=1}^{6}\sum_{p=1}^{N} V_{i,j}^{s}}{\sum_{j=1}^{6}\sum_{d=1}^{N} V_{i,j}^{d}}$
5	Level of pressures at consumption nodes	Max	$I_5 = \frac{\sum_{j=1}^{6}\sum_{h=1}^{168}\sum_{i=1}^{N} \max(0, \min(p_{i,h,j}, p_{ref})}{168 \cdot N \cdot 6 \cdot p_{ref}}$
6	Percentage of users supplied continuously	Max	$I_6 = \frac{\sum_{j=1}^{6}\sum_{i=1}^{N} \delta_{i,j}}{N \cdot 6}$
7	Pipe length with negative pressures	Min	$I_7 = \frac{\sum_{j=1}^{6}\sum_{m=1}^{M} L_{m,j}}{6}$
8	Energy consumption of pumps in operation over the whole period	Min	$I_8 = \sum_{j=1}^{6}\sum_{p=1}^{P} E_{p,j}$
9	Level of equity in supply	Max	$SR_{i,j} = \frac{V_{i,j}^{s}}{V_{i,j}^{d}}$ $ASR = \frac{\sum_{j=1}^{6}\sum_{i=1}^{N} SR_{i,j}}{N \cdot 6}$ $ADEV = \frac{\sum_{j=1}^{6}\sum_{i=1}^{N}\lvert SR_{i,j} - ASR\rvert}{N \cdot 6}$ $I_9 = 1 - \frac{ADEV}{ASR}$

- j – Years index
- i – Nodes index
- N – Total number of demand nodes
- $w_{i,j}$ – 1 if consumer i as continuous service pressure in year j
- $V_{l,j}$ – Volume lost by leakage l in year j
- $V_{s,j}$ – Volume supplied by source s in year j
- $V_{i,j}^{s}$ – Volume supplied to consumer i in year j
- $V_{i,j}^{d}$ – Volume demanded by user i in year j
- $P_{i,h,j}$ – Pressure at node i and hour h of year j
- P_{ref} – Min acceptable pressure for quality supply (P_{ref}=20m)
- $\delta_{i,j}$ – 1 if the pressure at node i is larger than P_f for all hours of year j
- P_f – min pressure for supplying all demands (10 m)
- $L_{m,j}$ – Longest negative pressure length of pipe m in year j (max between time steps)
- $E_{p,j}$ – Energy consumption of pump p over year j

3. Methodology

The methodology used to optimize the BIWS case study includes several stages. First, a preprocessing analysis was conducted to gain a comprehensive understanding of the network behavior. Next, a greedy algorithm was developed to find the optimal investment strategy. Last, the system's controls were optimized by using a brute force search.

Preprocess Analysis

This stage included explorative hydraulic simulations to develop a systemized understanding on the impact of different interventions on the objectives. The purpose of this process is to reduce the space of candidates used in the next stage of the Greedy algorithm. One prominent insight is that the existing pumps' configurations tightly constrained the use of groundwater. The difference between groundwater level and ground elevation is larger than the pump's maximum head. As such, physically, pumps cannot pump water into the system. However, since the EPANET mathematical model allows for negative pressure in the system, the pump head in the simulation can appear below the pump's maximum head, due to an unrealistic suction effect that helps lift water from the underground. Activating the pumps in these conditions results in flow rates that exceed the Maximum allowed Flow Rate (MFR) given in the BIWS instructions. The conclusion is that for year 0 (before any investments were made), the use of underground water is strictly limited, except for well B_AB and B_SA that, by setting valve controls, can pump water without exceeding the MFR. Accordingly, pumps were replaced in the first year and were not part of the Greedy algorithm candidates. Another critical insight gained from the hydraulic analysis is that the network's tanks were not functional in their ability to fill and store water effectively. To investigate whether increased storage volume might improve the system performance, several investment scenarios were tested to explore the network hydraulics with larger tanks. This analysis considered future states when underground water is available and leakages water loss is decreased to test if the increased volume will contribute after some of the investments are made. It was observed that even after investing the entire five-year budget, the tanks' volume would not limit the system performance, thus, increasing the tanks' volume was not considered in the investment optimization.

Greedy Algorithm

Given the complexity of the challenge described above, an optimization strategy is essential to address the non-trivial hydraulics conditions and objective functions. Additionally, the optimization approach should take into consideration the computational burden that prevents the use of classical simulation-optimization methods. For the BIWS challenge, we propose a greedy optimization algorithm. This algorithm takes inspiration from Dantzig's greedy approximation algorithm (Dantzig 1957) for the discrete knapsack problems. In the context of the BIWS, the knapsack problem can be defined as follows: Given a set of rehabilitation actions, each with a cost and a benefit, one needs to select a subset of rehabilitation actions so that the total cost does not exceed a given budget while the total benefit is maximized. Dantzig proposed sorting actions in decreasing order of benefit per cost (e.g., marginal benefit) values. This sorting enables the algorithm to prioritize actions that offer the most benefit relative to their cost. By iteratively selecting actions in this manner until the budget constraint is met, the greedy algorithm constructs a solution that is expected to be relatively close to the optimal solution. While this algorithm is straightforward for the original knapsack problem, in which the benefit and the cost of each action are known constants, its implementation for the nonlinear case is challenging (Salhi et al. 1989). In the nonlinear case, such as the one considered herein, the benefit from each action is not constant. Moreover, the cumulative benefit of combinations of rehabilitation actions can have a greater effect than the sum of the individual actions. Therefore, new evaluations of the actions' benefits need to be made after implementing each action. The nonlinearity of the problem raises a major computational challenge, unlike the classic knapsack problem, the determination of actions' benefit requires solving a hydraulic simulation. Since such simulations are computationally expansive, the number of benefit evaluations is limited. The following describes the mechanism of the proposed greedy algorithm and the techniques that have been used to cope with this computational burden.

Initial Candidates Selection

The purpose of this stage is to reduce the search space by eliminating some of the optional investment candidates, to include only investments that have the potential for significant network performance improvement. To accomplish this, a preliminary candidate selection stage was implemented. As stated above, tanks and pumps were not part of the greedy algorithm that focused only on pipe replacement and leak repairs. Each pipe and leak were evaluated based on hydraulic parameters to estimate their potential contribution to the overall network performance. Pipes were evaluated according to their per unit head loss, and leaks were evaluated according to the proportion between their water loss volume and repair cost.

Greedy Budget Allocation

The Greedy algorithm is illustrated in Figure 2 (which also details the greedy improvements described below). The algorithm is initiated with the candidate selection process and setting the budget, which in the case of the BIWS is 650,000 per year. The algorithm iterates through the following steps while the remaining budget is larger than zero: **(1) Evaluate:** each candidate investment is individually implemented in the network, then a hydraulic simulation is executed, and the nine performance indicators are evaluated. The nine indicators are then normalized and converted into a single indicator to calculate the benefit of the investment. The normalization is done based on the prior iteration (current best state) of the network, where the single objective is the sum of all normalized indicators, see Equation (1)

$$Benefit_i = \sum_{j=i}^{9} \frac{I_{i,j}}{I_j^{i-1}} \tag{1}$$

$$MB_i = \frac{Benefit_i - Benefit_{i-1}}{Cost_i - Cost_{i-1}} \tag{2}$$

where $I_{i,j}$ is the value of the indicator j as a result of implementing investment i. I_j^{i-1} is the value of the same indicator in the previous iteration. The change in benefit is the difference between the single normalized objective of the previous iteration and the current iteration of investment i. The change in cost is the additional cost required for the examined investment. Each candidate is scored according to its Marginal Benefit (MB); (2) **Sort**: all examined investments are sorted according to their MB; (3) **Implement:** the investment with the maximum MB is implemented in the network. and the benefit i-1 and cost i-1 are updated; (4) **Update Re-evaluation List**: the list of candidates that will be re-evaluated in the next iteration is updated. That is, not all candidate investments will be re-evaluated to obtain their MB, some candidates will continue to the next iteration with the same MB as evaluated in the current iteration (see next section and Figure 2); (5) **Update Budget:** The last step in every iteration is to update the remaining budget to be the budget from the previous iteration minus the cost of the investments selected in step 3.

Improved Greedy - Evaluations Reduction

While the greedy algorithm described above suggests a systematic method to move toward an improved solution, it is still subject to computational burden. Even after reducing the candidates list, the requirement to evaluate each of the candidates while improving the network with a single investment in every iteration is not tractable with conventional computing resources. Several features were incorporated within the algorithm, aimed at increasing the level of greediness, and therefore at shortening the runtime. Certain investments present multiple options, for example, pipe replacement can be done with various diameters. In theory, each potential diameter could be considered as an individual investment candidate. To reduce the number of evaluated candidates and adhere to the greedy guidelines, pipes are gradually increased. In each iteration, the examined candidate's diameter is increased by one size. Subsequent iterations consider further diameter increases to the next available size, with the additional cost calculated as the difference between the current and the new diameter costs. Another feature is in Step 3, where instead of selecting only the best investment to implement, the algorithm selects a set of best investments. The size of the set is determined dynamically in every iteration, and it depends on the distribution of the MB score. First, the best MB is calculated, and then every investment that has a score above $MAX(MB) \cdot (1 - \tau)$ is added to the set of investments, where τ is one of the algorithm parameters that control the level of greediness. The last feature is Step 4 so that instead of re-evaluating the MB of all the investment candidates in every iteration, only a subset is selected to be re-evaluated. The selection of candidates to be re-evaluated is done by detecting the elements in the system that were affected the most by the last changes (previous iteration). That is, it is expected that there is negligible change in MB score for elements that were not hydraulicly affected (i.e., did not experience a change in flow or pressure because of the last investment). The logic behind this approach is that in a large network, the hydraulic impact of actions like repairing a leakage or increasing pipe diameter is localized and limited to the region around the action. The hydraulic impacts in other parts of the system will be limited. To implement this logic in the greedy algorithm, Step 4 is added. Based on assessing the change in the pipes flow before and after the current investments, the n pipes with the largest change are flagged for MB re-evaluation in the next iteration. The parameter n is selected as a percentage of the total number of candidates when the algorithm starts. Additionally, all the leaks in these pipes are also flagged for MB re-evaluation. The number of selected elements to be MB re-evaluated is a second parameter of the algorithm the controls the level of greediness. Obviously, due to problem nonlinearity, MB re-evaluation of all investment candidates is more accurate, but this will come with the price of longer runtime due to more hydraulic simulations in each iteration. To conclude, the two parameters n, τ can be tuned to set the level of greediness to match a desired runtime.

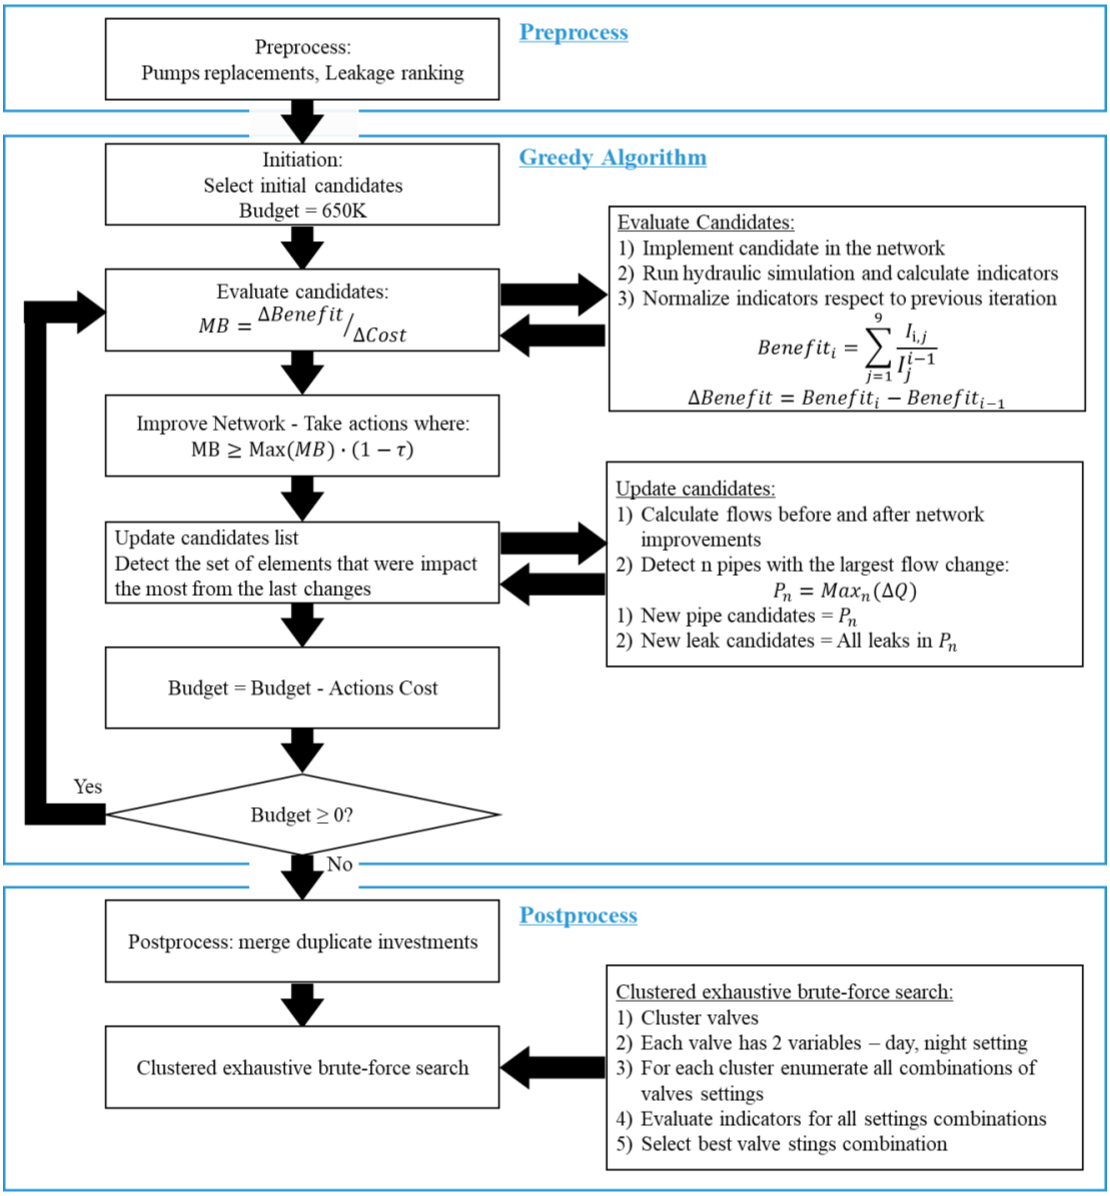

Figure 2 – Optimization strategy flowchart

<u>Postprocess</u>

A drawback of the proposed method is the evaluation of individual investments in a way that does not account for the implementation of other investments. This could result in incompatible investments in the final solution. For example, the greedy algorithm evaluates MB after repairing a leakage and finds that it is best to fix it in the current iteration. In a later iteration (but in the same investment year), the replacement of the pipe of the same leakage is selected. Obviously, since pipe replacement solves all the leakages in the same pipe, it will be a budget waste to repair the leakage. The postprocessing stage finds such duplicates and excludes them from the final solution. The spared budget is used to implement more investments from the last greedy iteration.

Clustered Brute Force Search

The final stage of the optimization process is to set the control policy for pumps and valves. The system consists of 7 pumps, 11 throttle control valves (TCV), 2 FCVs, one PRV and one PSV. Also, the system has 12 isolation valves that can be opened or closed with no partial opening. All valves can be controlled on an hourly basis during the 168-hour optimization which leads to a very high dimensional discrete decision space. For the operation of the pump, it was assumed that pumps should be operated as much as possible since they increase the inflow rate and the system pressures, which improves 7 out of the 9 indicators (exceptions are leakage volume and energy usage). Accordingly, pumps are operated constantly. For valves, the strategy taken is to reduce the decision space based on similar principles as in the second greedy improvement. Observing that the influence of a control change is limited to its local physical vicinity, the control elements were clustered into seven groups (clusters), where each cluster is optimized individually. The optimization strategy used here is an exhaustive brute-force search where an element can be either closed or open. In theory, each hour of the total 168 hours can get a different value, however, to achieve tractability, each element got only two values, one for daytime and one for nighttime. Based on the demand pattern of the system, daytime is defined from 07:00 until 00:00 and nighttime is from 00:00 to 07:00. The clustering of the elements was done based on their topological location and based on hydraulic impact. For instance, valves V_G1, V_G2, and V_G3 regulate the South-East part of the network as an individual DMA. Therefore, these valves affect mainly on this DMA and minimal impact is assumed regarding further parts of the network. Clusters were defined using engineering judgment to divide the network into DMAs, each regulated by specific valves. Each cluster contains up to 5 valves. For a cluster with 5 elements, we get a search space size of $2^{2\cdot 5} = 1024$, thus it can be easily enumerated. Figure 3 shows the clusters that were used for the brute force search.

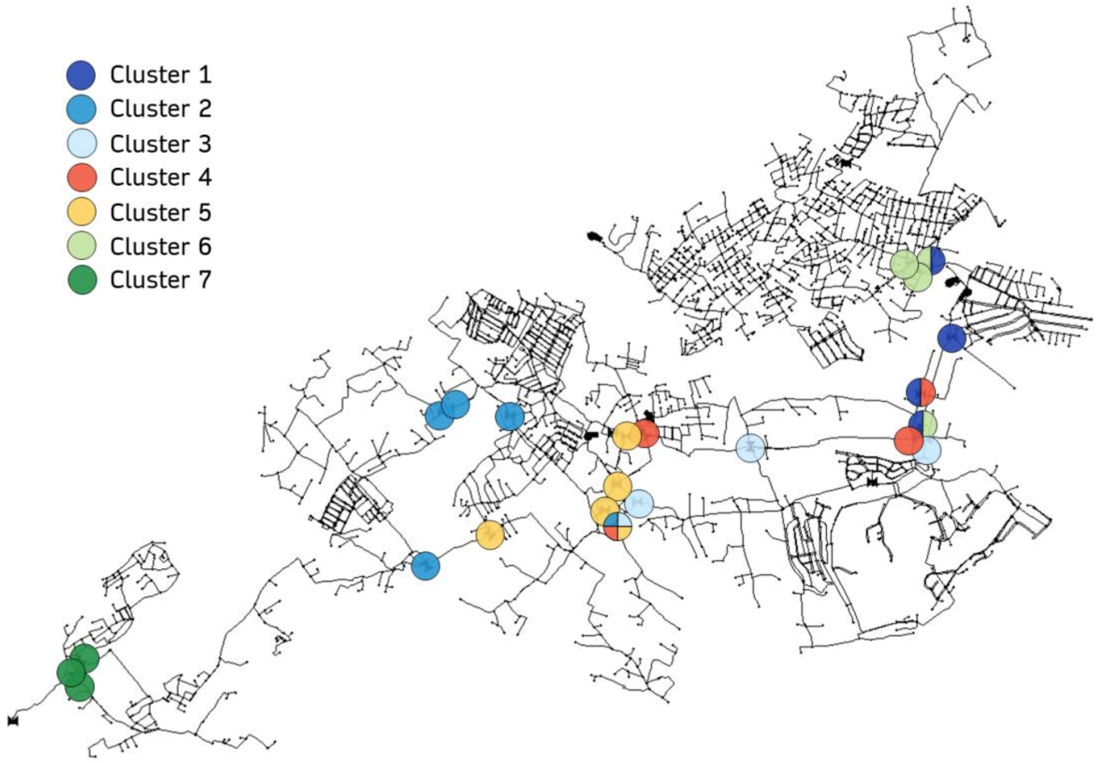

Figure 3 – Valves clusters for brute force exhaustive search. Note: a valve can be a member of multiple clusters as indicated by the multi-coloured circles

4. Results

First, insights from the preprocess analysis were implemented. All pumps installed in water-wells were replaced, and FCVs were installed in the discharge of the pump to avoid MFR deviations. The selection of pumps to be installed was done based on engineering judgment. The guiding principle was to select pumps with flow rates that are slightly smaller than the MFR so they can be continuously operated without exceeding the MFR. Since FCVs are installed in the discharge of pumps, it is guaranteed that flows will not exceed MFR anyway. However, by selecting the pump flow rate to the MFR values, the pressure dissipated in the FCV is smaller, and pumps are more efficient. To select the pump head, the discharge pressure ranges were analyzed assuming that toward the end of the planning horizon pressures in the network would be between P_f (min pressure to satisfy all the supply) and P_{ref} (pressure for quality supply). Accordingly, the pumps were selected by the higher values of the pressures analyzed and with the assumption that negative pressures will be rare. Leakages analysis found that a small number of large leakages is responsible for most of the water loss. Figure 4 shows that while the network contains almost 3600 leakages, 180 of them (5% of all leakages) cause 59.6% of the total leak volume. The worst 20% leakages cause 81.5% of the loss. This nicely adheres to the well-known Pareto principle which states that 80% of outcomes result from 20% of causes (the "vital few"). In terms of repair cost, the worst 600 leakages (16.6% of all leakages) cost more than the total budget for the five years. As such, not all the leakages can be repaired, and the greedy candidates were narrowed to the worst 1000 in terms of water loss per repair cost.

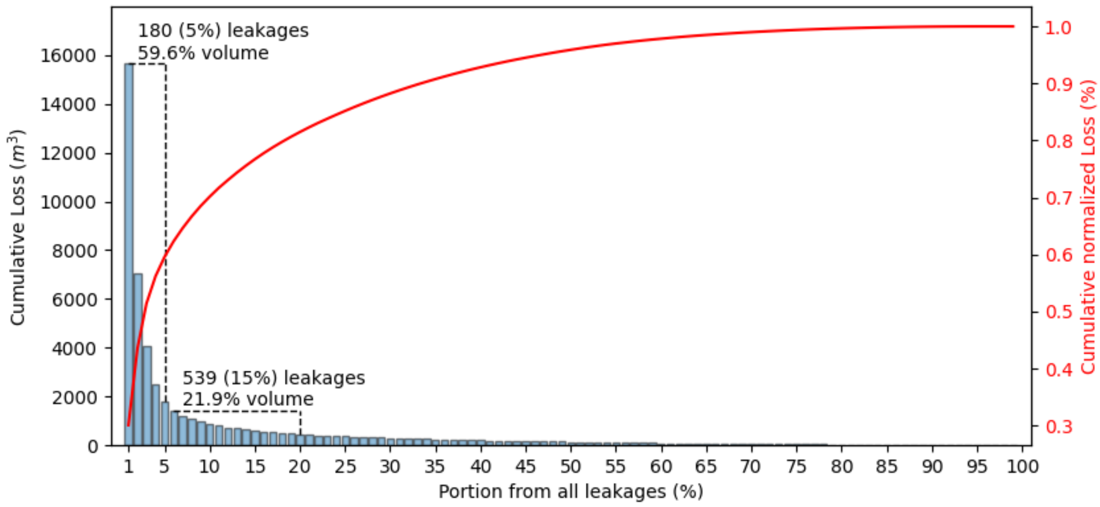

Figure 4 – Leakages water loss volume

Next, a greedy algorithm was run separately for each of the five years. The budget for the first year was reduced due to the investment in pumps and valves. The initial candidate's selection took into consideration: (a) all pipes with an average per unit head loss larger than 3 m/km and (b) the highest 1,000 leakages according to water loss volume per repair cost. The parameter τ that governs the number of investments selected in each iteration was set to 3%, meaning all investments that were ranked with MB larger than 97% of the best investment will be implemented. The number of MB re-evaluations in each iteration (n) was set to be 3% of the total number of candidates. The obtained results are presented in Table 2 and Figure 5. The results show that for all indicators, the most significant improvement is after the first year. The first year includes the pump replacements, which allow the use of other water sources, increasing the inflow rate from 150 L/s to 255 L/s. The increase in the supplied water, as well as the increase in the system's pressure, contribute to indicators such as Effective hours proportion (I1), Continuous service and supply pressure (I2, I6), Supply-demand proportion (I4), and others. Despite the pressure rise, the volume of water loss due to leakages is dramatically decreased due to the repairing of leakages.

Table 2 – Solution Indicators

Year	I1 Max	I2 Max	I3 Min	I4 Max	I5 Max	I6 Max	I7 Min	I8 Min	I9 Max
0	0.649	0.298	0.474	0.412	0.354	0.08	248,594	392	0.355
1	0.926	0.772	0.368	0.739	0.782	0.534	97,588	14,941	0.763
2	0.937	0.767	0.346	0.768	0.799	0.518	103,920	15,175	0.787
3	0.965	0.804	0.283	0.834	0.856	0.539	89,602	15,695	0.869
4	0.971	0.855	0.278	0.856	0.87	0.629	65,960	15,208	0.888
5	0.978	0.871	0.237	0.869	0.882	0.658	60,327	16,371	0.903
Total	0.904	0.728	0.323	0.746	0.757	0.493	110,998	77,784	0.722

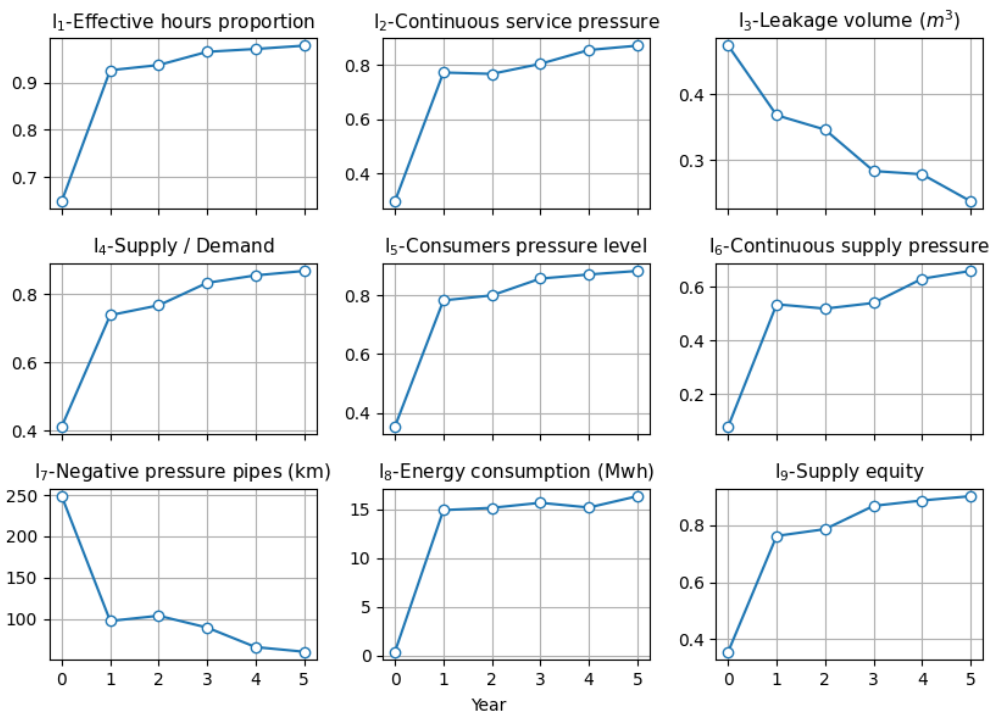

Figure 5 - Indicators values over the 6-year horizon

The best-known results of the BIWS challenge were presented by Marsili et al. (2023), another study by Mottahedin et al. (2023) ranked 3rd while the greedy approach ranked 4th. Figure 6 shows a comparison of the two studies with Marsili et al. (2023) the presented study. Each team is represented by bars of different colors, rows represent the yearly results and columns the performance of the nine indicators. The greedy approach and Mottahedin et al. (2023) present a slight advantage in year 0 due to a control strategy that allowed the operation of well B_AB and B_SA without exceeding the MFR. Marsili et al. (2023) used most of the budget for repairing leakages and replaced only a single pump. This approach results in a dramatically larger number of repaired leakages compared to the two other approaches. A dramatic advantage of both Marsili et al. (2023) and Mottahedin et al. (2023) over the current study is in their control strategy. While the current study considered only day \ night control changes to reduce the search space, the two other teams presented more dynamic strategies that improved dramatically pump energy consumption. Furthermore, the two teams used valve controls to isolate parts of the networks which significantly improved the Negative pressure pipe length (I7). This indicator is calculated as the sum of pipe length that is subjected to negative pressure, where length is calculated at the worst time step (e.g., longest segment with negative during the simulation). However, isolated areas are considered as non-active and excluded from the calculation. Although such isolations show superior performance in the BIWS score it is in doubt how much this indicator really indicates a benefit to the users. For example, in the case of a long pipe, with only one consumer in its edge, and such that in only one-time step, the pressure of the pipe is negative, the indicator will get a poor score, although its functionality is not necessarily poor. Indicator I8 presents a substantial difference between the three approaches. In the current study, all pumps are replaced with better efficiency pumps, which results in lower energy consumption compared to Marsili et al. (2023) for the first three years. In the later years, the control strategy by Marsili et al. (2023) was managed to better even the new pumps suggested in this study. Mottahedin et al. (2023) approach is superior to these two by combining new pumps, and frequent inverters allowing for more dynamic control by adjusting pump speed control.

Overall Mottahedin et al. (2023) presented a similar solution to the current study where in the first year a large portion of the budget is invested in new devices (pumps and valves) and the rest of the budget is spent on prioritized pipe replacements and leak repairs where control is optimized separately from the investment decisions. The winning approach by Marsili et al. (2023) differs in the sense that only a minimal budget is invested in new devices with only one pump replaced and valves installed only at pump discharges to maintain MFR. Despite the advantages of new device installations, the results show that prioritizing pipe replacement and leak repairs is the most dominant strategy leading to the best performance.

It is noted that Marsili et al. (2023) did not detail the I9 score in their paper and stated that it could not be evaluated for individual years. Also, it is noted that Marsili et al. (2023) calculated the energy consumption for the total number of hours in each year and not just for the one-week simulation. The values in Figure 6 were calculated by the authors using the final networks provided by Marsili et al. (2023).

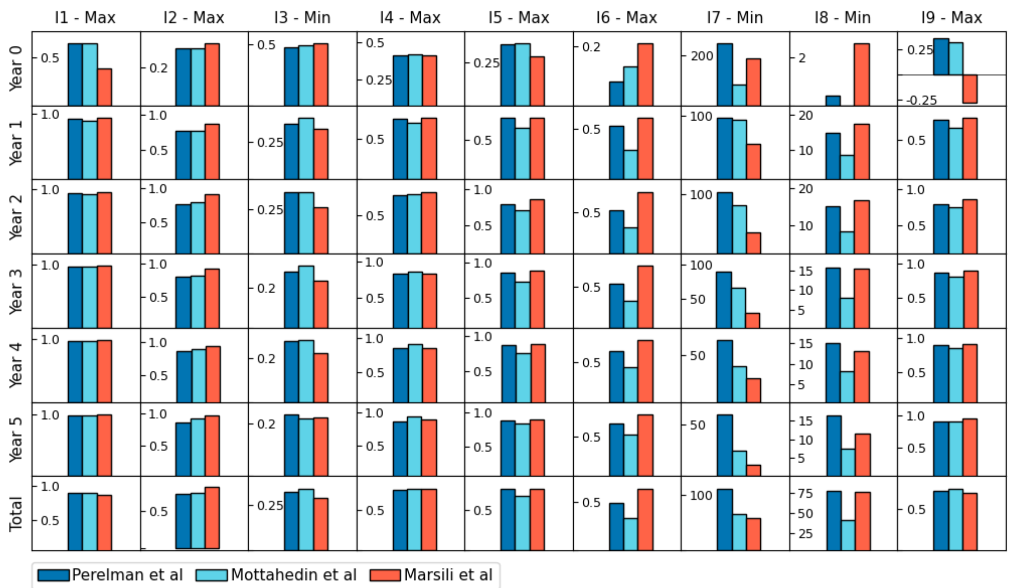

Figure 6 - Indicators values over the 6-year horizon

The greedy algorithm effectively produced a rehabilitation strategy that is reasonable from an engineering perspective. The algorithm prioritized the most critical leakages to be repaired in the early years and identified bottlenecks in pipe diameters. In this context, a 'bottleneck' refers to a segment connecting two pipes with a diameter significantly smaller than that of the connected pipes. Figure 7 presents the leakages in the system by their total water loss volume before any investment was made and by their repair cost. The colors in Figure 7 represent the year in which a leak is repaired where the lightest blue is for repairs in year 1, and the darkest blue is for repairs in year 5. Leaks that were not repaired are marked as red circles. As much as a leak is larger and loses more water, it is easier to locate, and thus, its repair cost is smaller (the functions provided in BIWS instructions, 2022). This relationship between the water loss and cost is very convenient for the greedy algorithm because the two factors are negatively correlated. While, usually, a large benefit requires a large investment, here it is the opposite case. As such, the most urgent leaks have the smallest repair cost, and they are prioritized to be repaired in the first year. Since a small leak loses less water and costs more to repair, it is delayed to later years or not repaired at all. Some exceptions are observed due to strategic locations of leaks on main pipes and due to the suboptimality of the greedy algorithm.

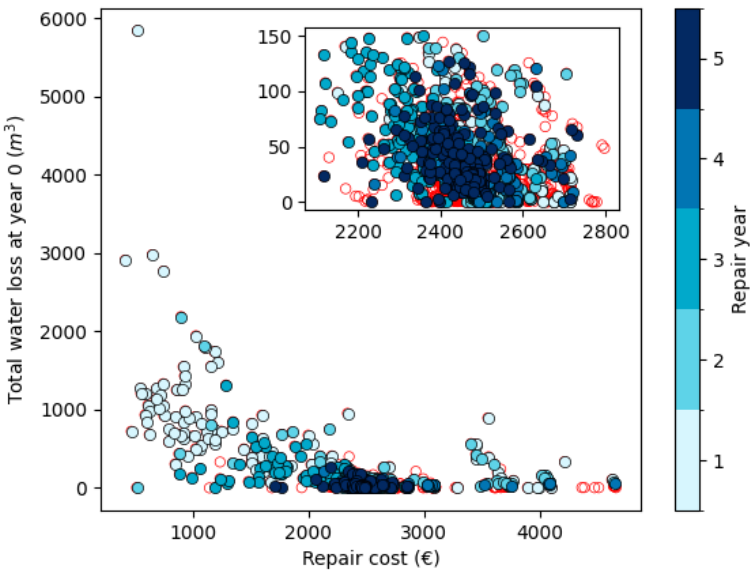

Figure 7 - Network leakages by their water loss, repair cost and year repaired

Among the pipes that were replaced, the greedy algorithm identified those that constituted bottlenecks and prioritized them. Such pipes are mainly pipes with short lengths with small diameters (relative to the downstream demand and neighbor pipes) and pipes that connect separated sections. Figure 8 presents the pipes that were replaced (to a larger diameter) with the same color code as in Figure 7 so that light colors represent earlier replacement. The pipes that were replaced across all years are mostly short pipes that improve the network connectivity. One drawback of the greedy algorithm is its inability to evaluate the contribution of the combination of investments, but only each investment on its own. This leads to suboptimality since it is possible that two individual investments will be ranked with low MB but to fulfill the potential benefit of these investments, two (or more) small investments must be taken. For example, a bottleneck is constructed as a series of two short pipes. Each pipe is evaluated individually, and while the cost is low since these are short pipes, the benefit of each of the investments is also low because the obstacle was not removed. If sets of investments could be evaluated, then the solution would be much closer to the global optimum. As stated above, even the evaluation of single investments was subject to a strict computational burden thus, evaluating sets of investments is intractable with standard computational resources.

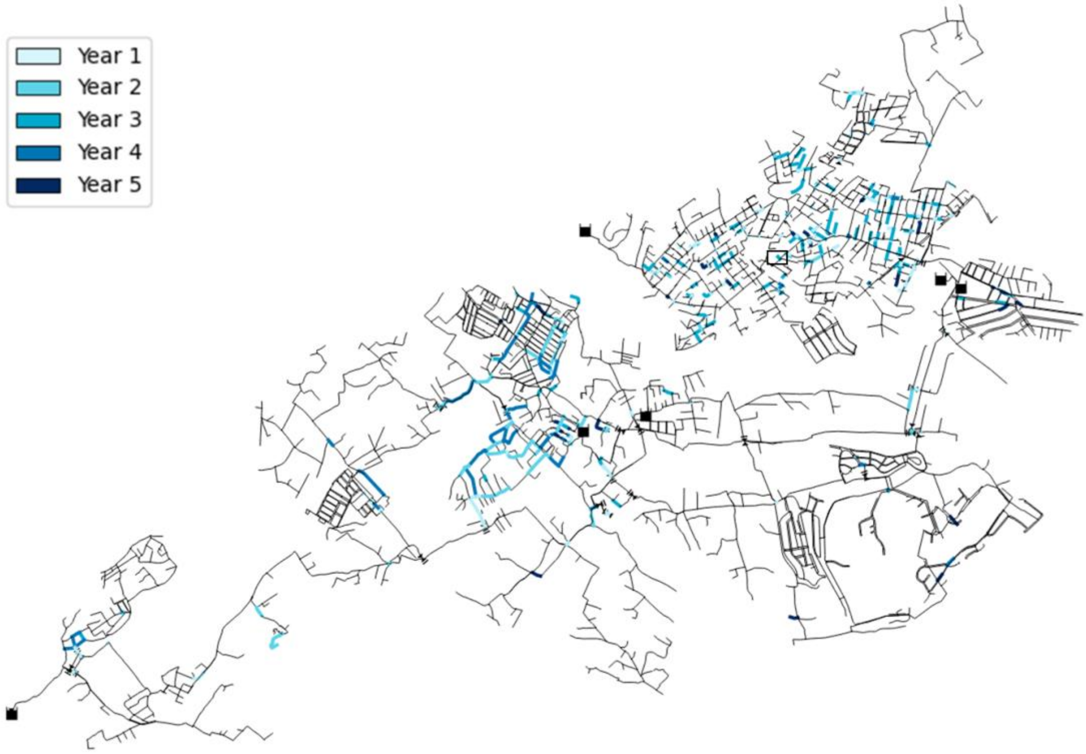

Figure 8 - Year when pipes were replaced

The valves opening policy resulted from the exhaustive brute force search is presented in the Supplementary Material as Figures S.2 -S.7. The network two main valves are V_CO that regulates the flows into the central tank (T1_CO) and V_TR that regulates the flow between the main network and the northern part. While V_CO is configured to be continuously opened throughout all examined years, V_TR is opened in years 0 and 1, closed during night-time at year 2, and continuously closed at years 3-5, generating segregation between the two parts of the network. It is noted that V_TR also includes a check valve that limits the flow direction from the main to the north part only. The conclusion is that the northern part can be satisfied by its own sources, and it is better to address water from the main source to other parts of the network.

<u>Greedy performance</u>

As explained above the greedy algorithm can be tuned using two parameters to trade off greediness and runtime. The parameters that control this tradeoff are the threshold ($\tau = 3\%$) that defines the distance from the best MB for an investment to be implemented and the number of MB re-evaluations per iteration (n=3% of the same year candidates). As explained above, the investments that will be re-evaluated for their MB in the next iteration are those that were affected the most by the latest investments. As such, the greedy algorithm continuously monitors the changes in pipes' flows and flags the most *n* affected pipes. Figure 9 presents this analysis for the most MB re-evaluated pipes, Figure S.8 in the supplementary materials presents the same pipes on the network layout. One can observe that pipes around the tank in the center of the network (T1_CO) were evaluated the most (L1771, L1773, L3247, L1838). This makes sense because this tank is in the "heart" of the system, as it regulates the main gravity source, two wells and most of the consumption in the network. Therefore, the hydraulics around this tank are expected to change dramatically with a significant improvement in the network. Other areas that are intensively re-evaluated are the outlet from the main source (L2452, L2453), the area that connects the north part to the main network (L3171, L1068, L1061, L1063), and other segments along the main transmission line and links between different consumption zones. From a practical perspective, although not all these pipes (and leakages of these pipes) were selected within the greedy algorithm, they are identified as critical elements in the system that can be explored for further improvements.

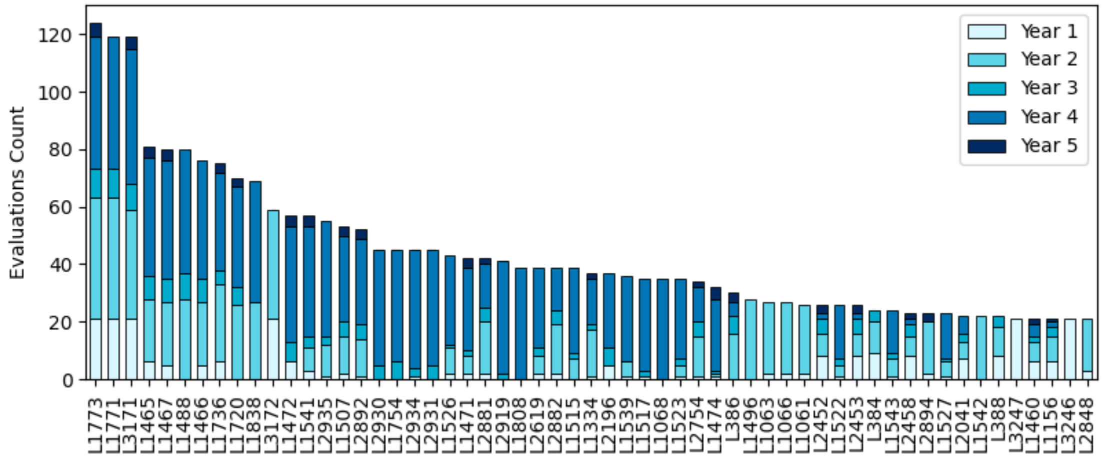

Figure 9 - Most MB re-evaluated pipes in the greedy algorithm by year

Figure 10 presents the number of investments made and cost per iteration of the greedy algorithm. The optional investments were pipe replacements (to a larger diameter) and leak repairs. The best result could be achieved if the greedy algorithm would implement only a single investment in each iteration and then evaluate again all other optional investments to decide the next implementation. Due to the computational burden, this process is accelerated by allowing the algorithm to implement more investments in case their MB is close to the MB of the best investment. While most iterations included less than 15 investments, there were few iterations where many investments were implemented. This testifies to a group of actions with a similar MB rank. The most prominent example is the last iteration of the third year where 198 investments were made. Most of those are repairing leaks with the same repair cost. This illustrates how the acceleration features in the greedy algorithm can improve its performance by reducing the number of iterations. In case no significant priority (in terms of MB) is observed, the algorithm will choose to implement all the top-ranked actions within the remaining budget and terminate the iteration process. However, while this procedure reduces the number of iterations, it can cause a suboptimal decision since it is possible that a new re-evaluation (for example, after implementing the best investment) could change the MB ranking and prioritize the investments differently.

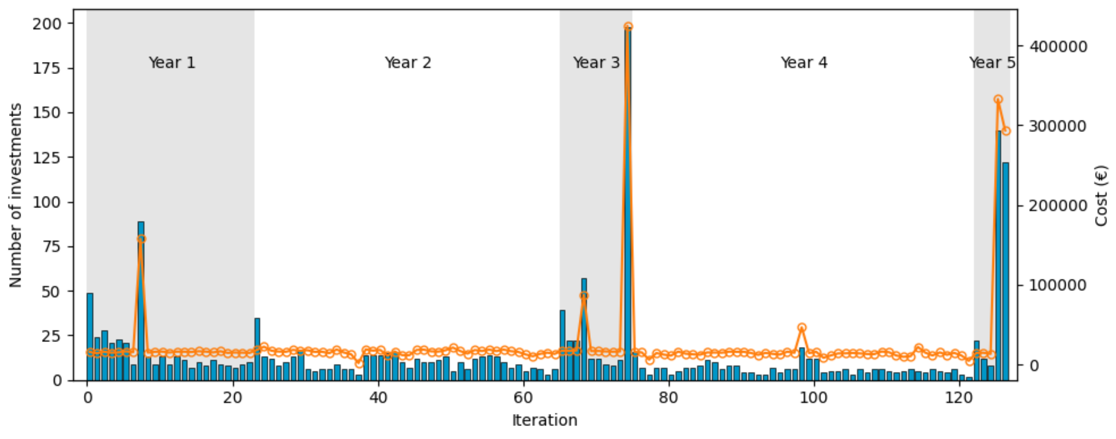

Figure 10 – Investments number and cost per iteration

5. Conclusions

The Battle of the Intermittent Water Supply presented a challenge to optimize the performance of a network under poor conditions through nine objectives. The problem's structure limits the use of standard optimization techniques due to unconventional hydraulic conditions and the inseparability of the problem in time and space. Another challenge is the huge dimension of the search space even when compared to other WDS optimization problems. These properties of the problem make the formulation of mathematical optimization model a very complex task, while the use of classical simulation-optimization methods is limited due to long simulations times. Hence, a robust systematic method is required to cope with the challenge of optimizing the system's rehabilitation. The method presented in this study was devised to efficiently navigate through the problem's complexities. The method starts with a preprocess stage including both hydraulic analysis and engineering judgment to identify highly beneficial early investments and reduce the search space. Next, a greedy algorithm is developed to optimize the rehabilitation budget allocation. Lastly, the operational controls policy is determined by using a clustered brute force search. The method overcomes the computational burden of the problem by reducing the search space in all stages. The preprocessing stage eliminates some optional investments and selects the highest potential candidates for the next stage of the greedy algorithm. Within the greedy algorithm, two mechanisms were implemented to increase the greediness level and accelerate the process. The operational control search is reduced to several clusters that are assumed to be independent. The obtained results are consistent with engineering logic as repairing the largest leaks first, increasing diameter at bottleneck points, and improving the connectivity of different sections in the network. All the indicators were improved dramatically because of the selected strategy. The proportion of continuously served consumers is increasing from 66% to 95% at the end of the planning horizon, the portion of water loss due to leaks decreases from 55% to 32%. Furthermore, 80% of the demand is supplied and the level of supply equity increased to 83% compared to 39% and 44% respectively before the rehabilitation of the system. The proposed greedy optimization approach can be applied to other networks and to any problem of budget allocation under constraints. It suggests robust systematic decision making regardless of the complexity level required to model the problem. As such, the method guarantees continuous improvement toward better solutions, avoiding non-convergence issues.

6. Data Availability

The code and results presented in this study are available on: https://github.com/GalPerelman/BIWS-Paper

7. References

8. Appendixes

Figures S.2-S.7 depict the solution of valve control. In each figure, the upper subplots show the demand pattern and seven subplots show the status of the valves along the 168 hours of optimization. Each subplot presents a temporal visualization of valve operations, structured similarly to a Gantt chart. Each horizontal bar represents an individual valve as identified by their respective labels. The timeline, divided into discrete intervals along the horizontal axis, delineates the operational status of each valve: black sections indicate periods when the valves are open, and white sections represent closures. For example, it can be seen in year 0, valve V_TR is constantly open while valve CV8 is open only during the night time (00:00-07:00).

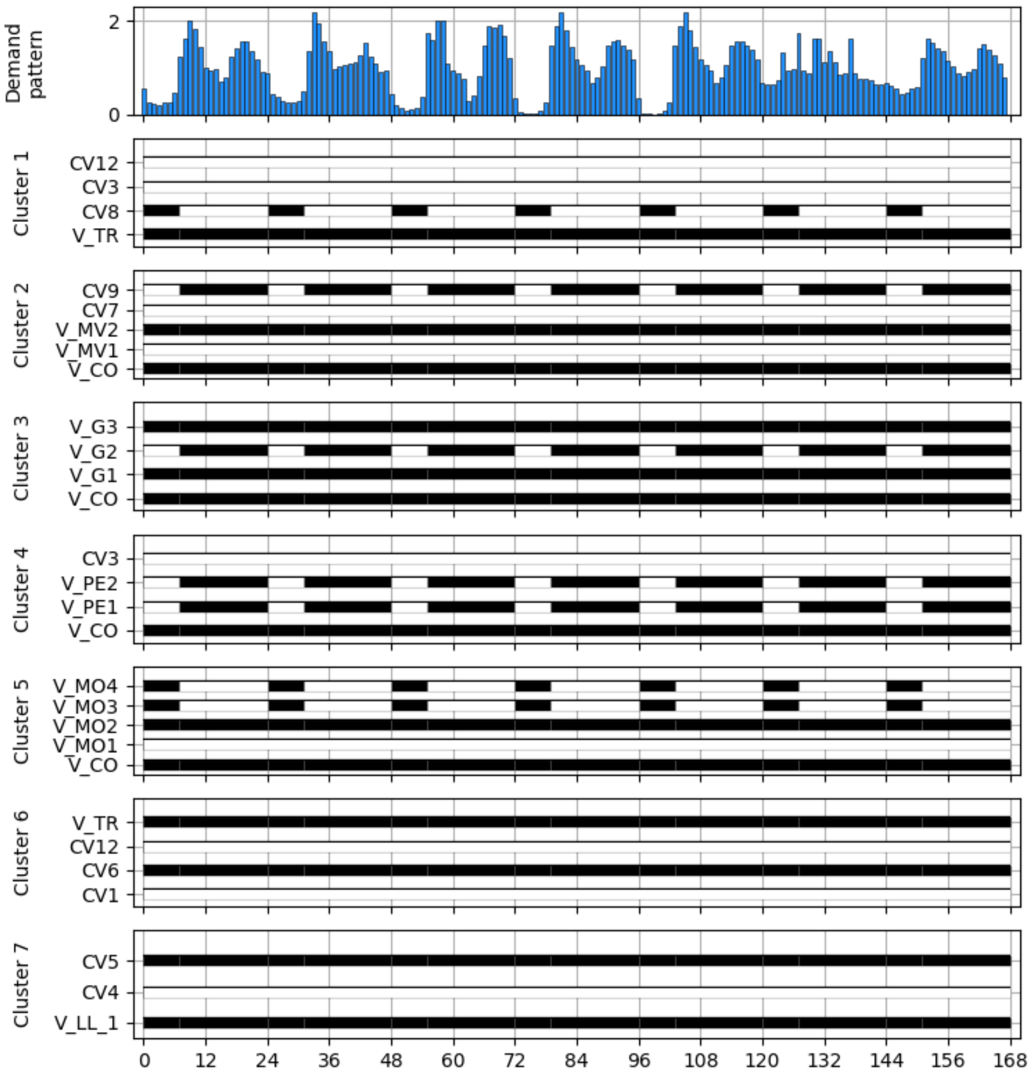

Figure S.2 – Exhaustive search results for valves status at year 0

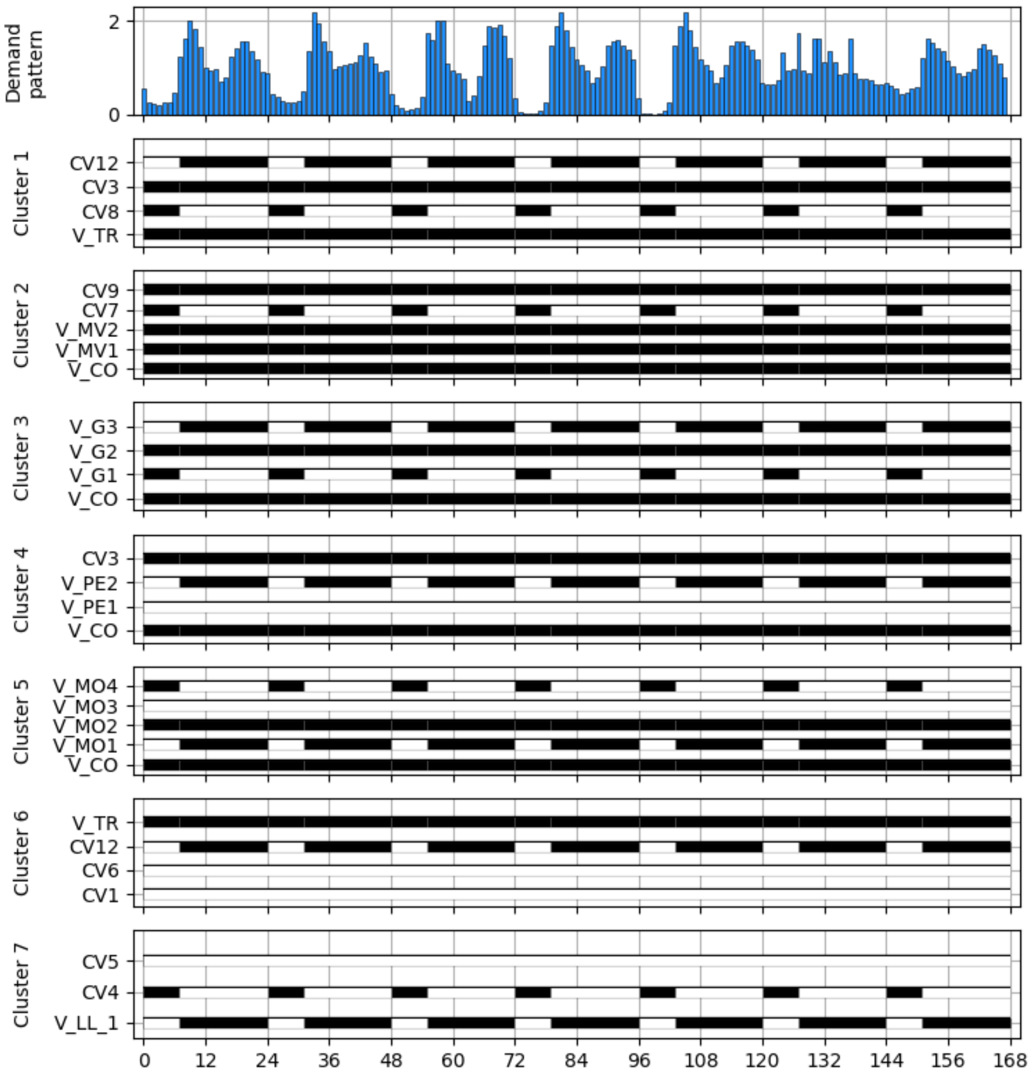

Figure S.3 – Exhaustive search results for valves status at year 1

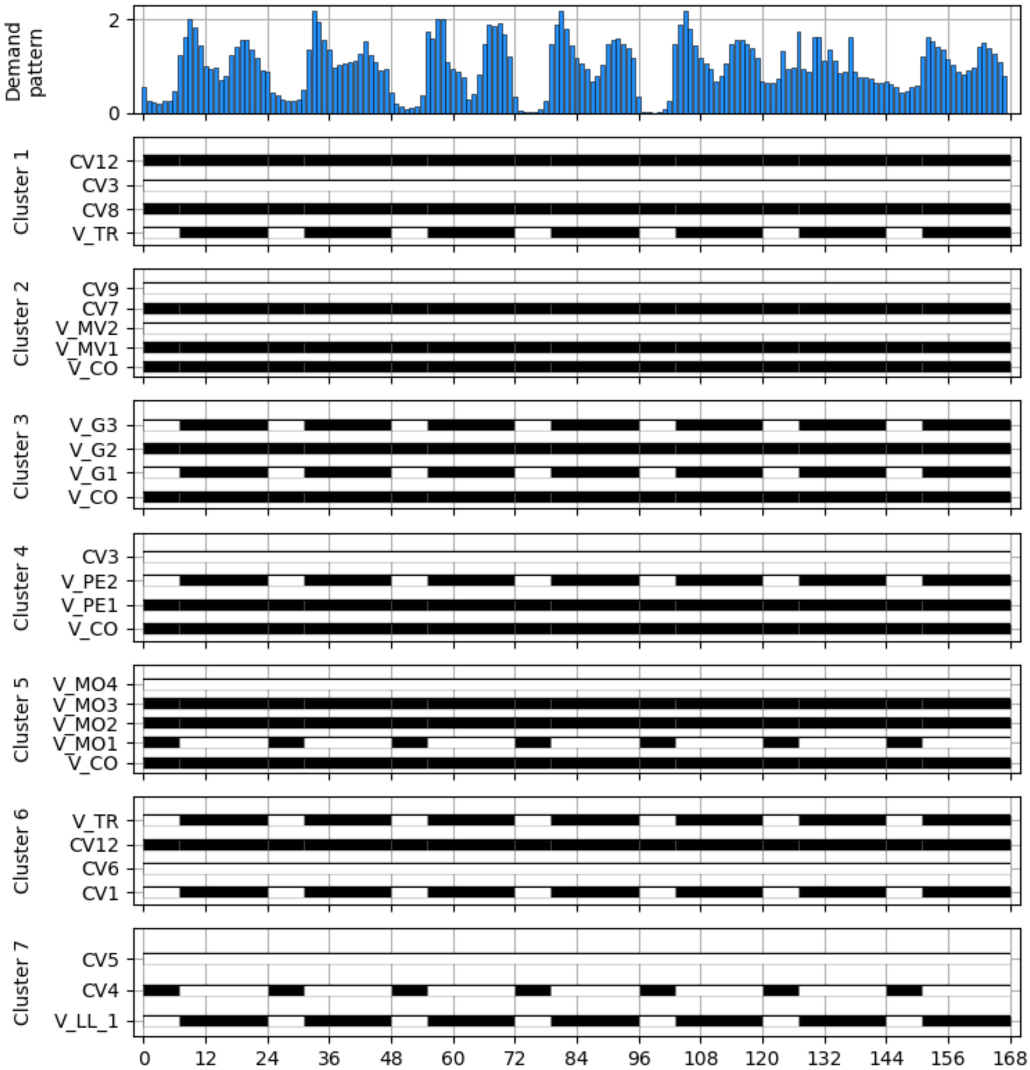

Figure S.4 – Exhaustive search results for valves status at year 2

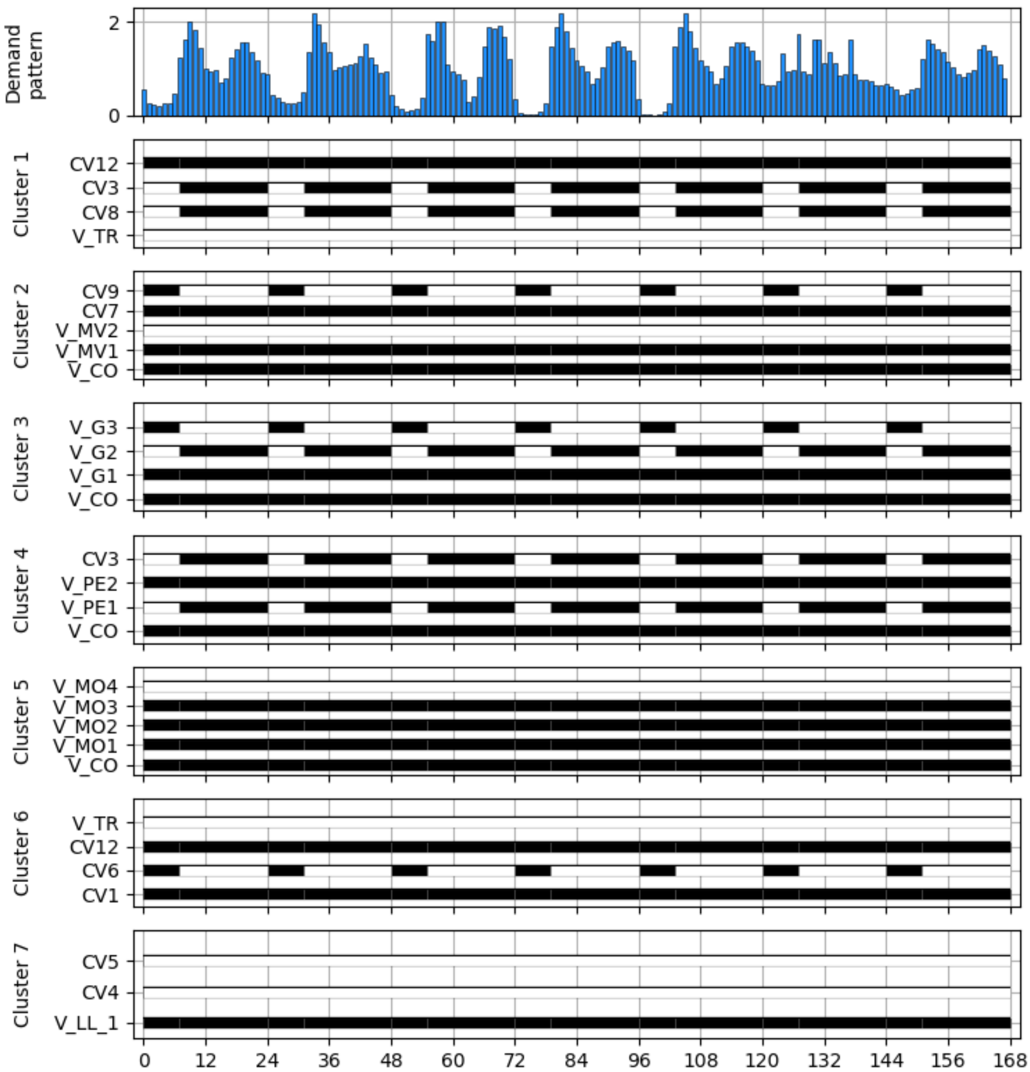

Figure S.5 – Exhaustive search results for valves status at year 3

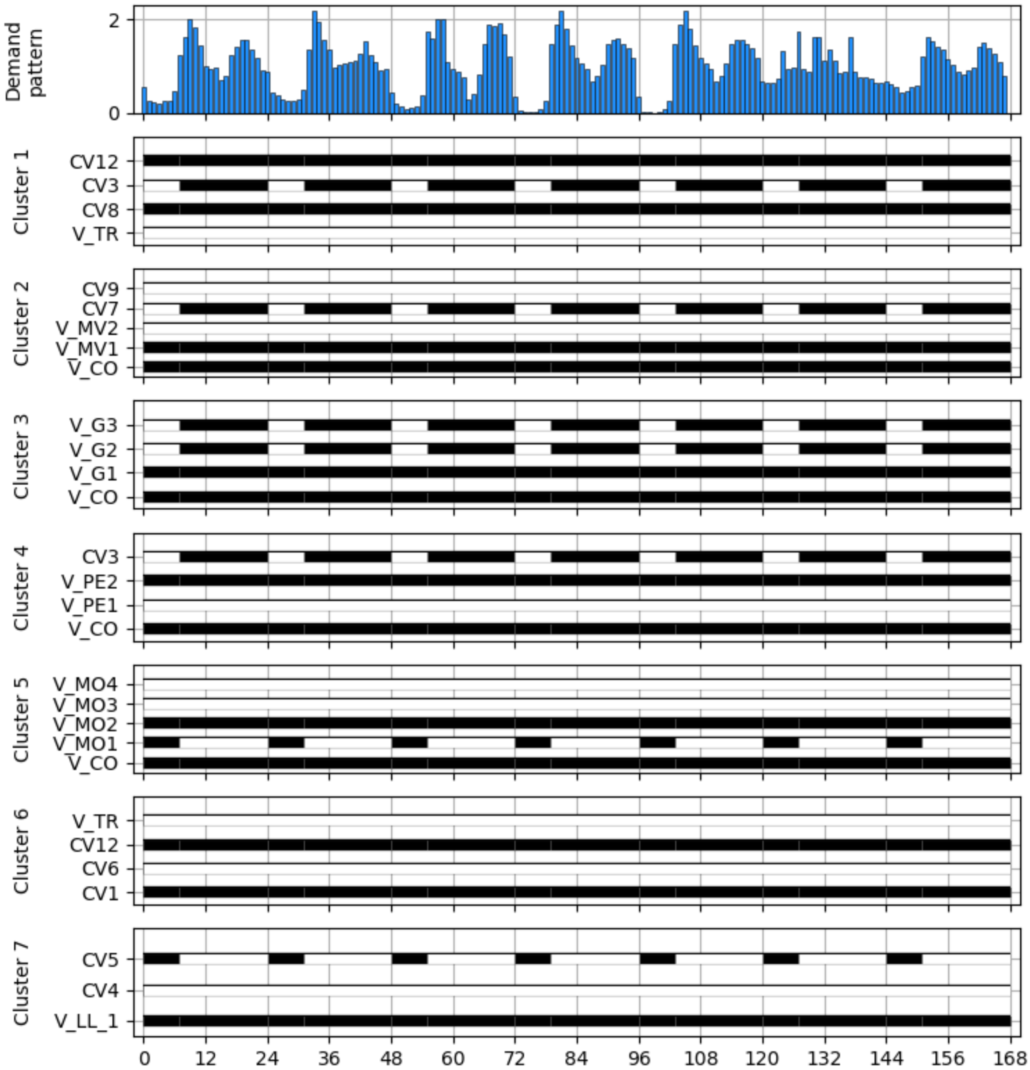

Figure S.6 – Exhaustive search results for valves status at year 4

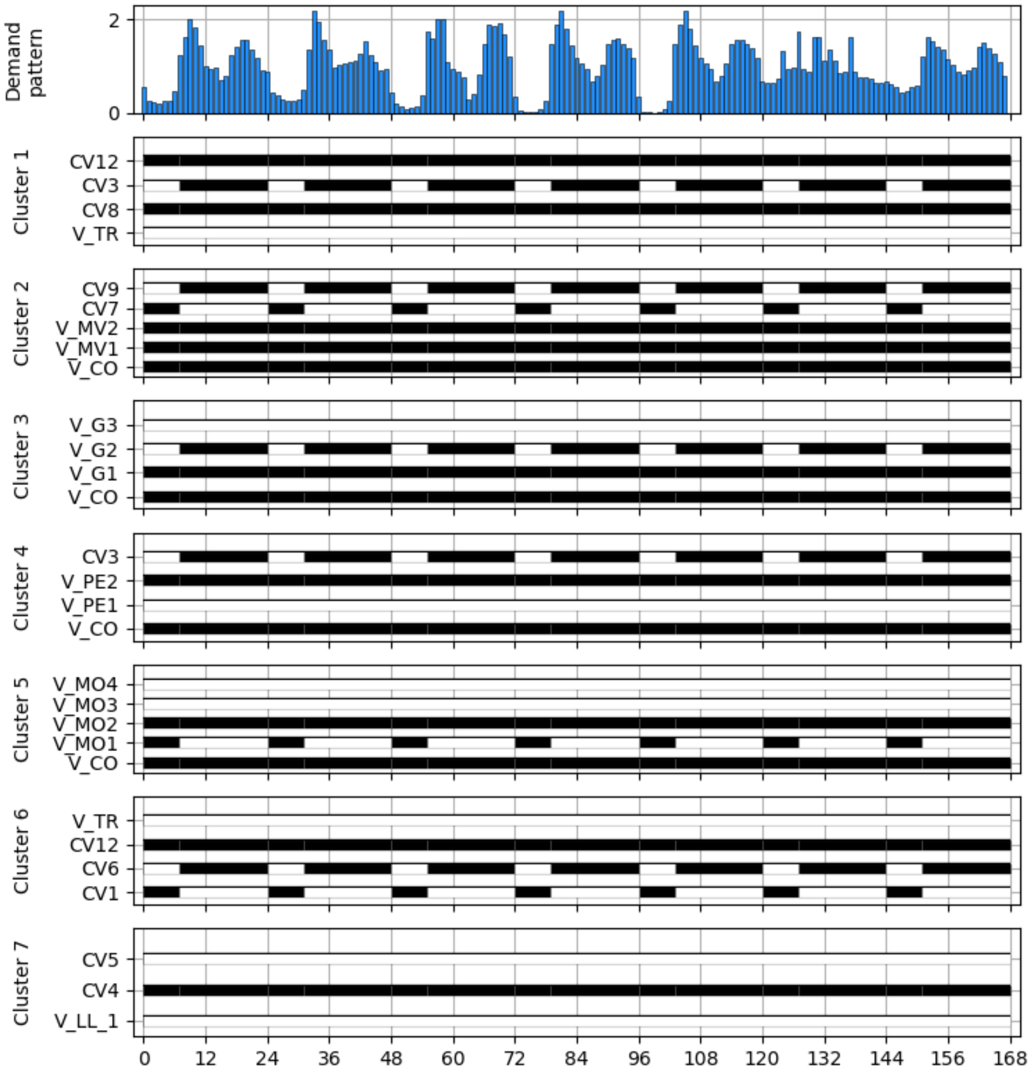

Figure S.7 – Exhaustive search results for valves status at year 5

Figure S.8 visualizes the elements that were most evaluated during the Greedy algorithm. As explained in the methodology section, evaluating each investment action in each iteration is infeasible. The algorithm reevaluate only the elements that were most likely to be affected from the previous iteration investments. Figure S.8 presents all the elements that were evaluated more than 20 times which indicates where most investments were made.

Figure S.8 - Pipes with over 20 evaluations

Figure S.9 shows the level of the main tank (T1_CO) along the 168 hours of optimization for each optimized year. This figure is another metric of the improvement in network performance, where it can be seen that before investments the tank is constantly empty and does not function as it should be. With the progress of rehabilitation investments, the tank behaves more similar to a typical tank in a pressurized system, such that filled during off-peak periods and drained in the peak demand period.

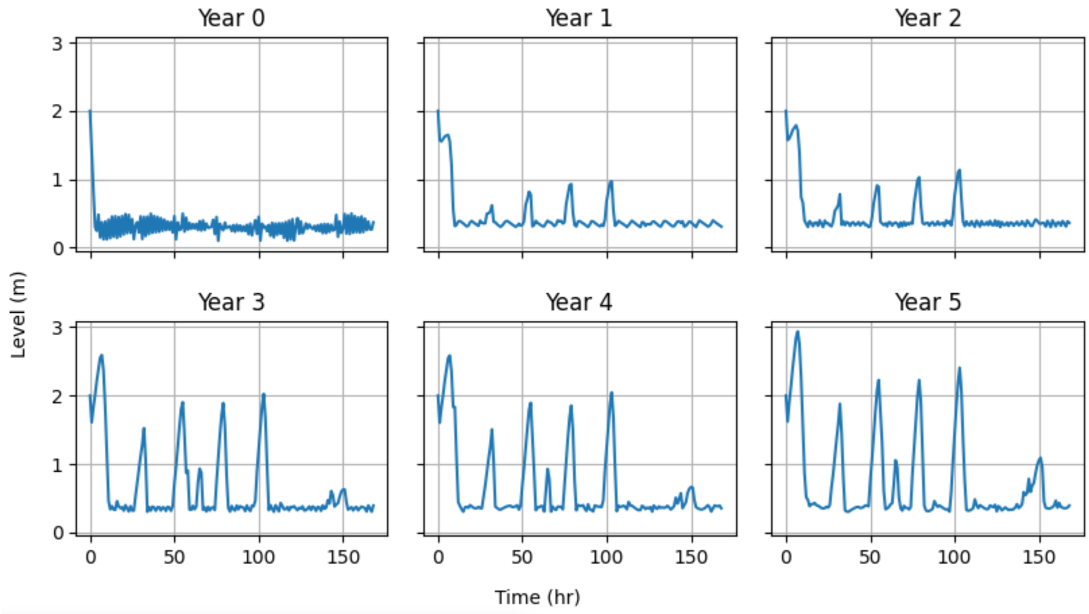

Figure S.9 – Tank T1_CO level over the planning horizon