

Chapter 6

DHV water pumping optimization

Simon van Mourik¹ Joris Bierkens² Hans Stigter¹ Martijn Dirkse³ Karel Keesman^{4†} Vivi Rottschäfer⁵

Abstract:

This contribution investigates the possibilities for optimizing a drinking water network over a horizon of 48 hours, given variable water demands, energy prices and constraints on the pumping strategy and water levels in the reservoirs. Both the dynamic model and goal function are non-linear in the control inputs, the pump flow rates. Since each pump can be switched on or off every 15 minutes and since there are 15 pumps in the system, for a horizon of 48 hours there are $2^{(4 \times 48 \times 15)}$ switching possibilities. Obviously, this problem is too big to solve it in real-time by enumeration. Hence, a decomposition of the problem is needed. Relaxing the constraints and assuming a continuous-time flow rate, allows a (semi)-analytical solution using Lagrangian theory. Furthermore, a numerical solution of the constrained optimization problem is found by using the TomLab PROMPT toolbox. The conversion from a continuous-time pump flow rate to a strategy with on/off switching is also investigated, as well as the possibility of linear feedback control. The resulting trajectories of the pump flow rates and water levels in the reservoirs are realistic and can be physically interpreted.

KEYWORDS: *Dynamic optimization, modeling, feedback control, drinking water network.*

¹Biometris, Wageningen University, The Netherlands

²Leiden University, The Netherlands

³Ecofys Netherlands BV, Utrecht, The Netherlands

⁴Systems & Control group, Wageningen University, The Netherlands

[†]corresponding author: karel.keesman@wur.nl

6.1 Introduction

For the 67th Studygroup Mathematics with Industry held at the University of Wageningen, we worked on a question posed by DHV which is an international group of consulting engineers located in Amersfoort. The question concerned water pump optimization. We were asked to optimize the distribution of drinking water in a region with towns (which require drinking water), reservoirs and pumps (which pump drinking water from one part of the region to another part).

The specific setting we studied, the Grimsby drinking water supply region in Canada, is shown in Figure 6.1. It consists of three towns Smithville, Beamsville and Grimsby. The drinking water demand of each town has a typical pattern that is more or less known in advance. Typical demand curves are given in Figure 6.2 where the demand is known per 15 minute sections per day.




Figure 6.1: A sketch of the Grimsby drinking water supply region with the three towns Smithville, Beamsville and Grimsby, and the pump stations and the reservoirs.

Drinking water is pumped into this supply region through the Grimsby High lift pumping station, located at a certain height H_{Gp} above sea level. With the use of water reservoirs and pumping stations, the drinking water is stored in the region and distributed over the three towns.

The pumping stations at Smithville, Park road and Beamsville are located at different




Figure 6.2: The drinking water demand for each of the three towns.

heights, H_{Sp} , H_{Pp} and H_{Bp} , respectively. The pumping stations each contain a certain number of pumps, see Figure 6.1, which have different capacities and which can either be turned on or off every 15 minutes. Also, the pressure difference *before* and *behind* the pumps, the so-called head loss, determines the operation of the pump. When switching on a pump, this head loss first has to be overcome before the water starts flowing through the pump.

The reservoirs have different capacities as denoted in Figure 6.1. Moreover, there are restrictions on the minimum level, 75 %, and maximum level, 95 %, that the reservoirs are allowed to contain.

Finally, operating the pumps costs energy which in turn costs money. The cost of energy is known; it varies through the day and is different on weekends, see Figure 6.3. Under the given restrictions and water demand, we were asked for an optimal solution such that the cost of energy is minimal. In other words, DHV would like to know, for a period of 48 hours in advance, at which moment the pumps should be turned on or off. Each pump can be switched on or off every 15 minutes. Since there are 15 pumps in the system, there are within the time-frame of 48 hours, $2^{(4 \times 48 \times 15)}$ switching possibilities. Obviously, studying this complete system with all of these possibilities is not possible, so other approaches need to be taken.

One of the pumping strategies, for example, is to fill the reservoirs during the night when the energy is cheapest. In this way, there is sufficient water supply to satisfy the peak in the drinking water demand of the towns in the morning. In this way the pumps

Table 1: Energy prices

Day of the Week	Time	Time-of-Use Period	Time-of-Use Price* (cents/kWh)
Weekends & holidays	All day	Off-peak	3.0
Summer Weekdays (May 1st - Oct 31st)	7:00 a.m. to 11:00 a.m.	Mid-peak	7.0
	11:00 a.m. to 5:00 p.m.	On-peak	8.7
	5:00 p.m. to 10:00 p.m.	Mid-peak	7.0
	10:00 p.m. to 7:00 a.m.	Off-peak	3.0
	Winter Weekdays (Nov 1st - Apr 30th)	7:00 a.m. to 11:00 a.m.	On-peak
	11:00 a.m. to 5:00 p.m.	Mid-peak	7.0
	5:00 p.m. to 8:00 p.m.	On-peak	8.7
	8:00 p.m. to 10:00 p.m.	Mid-peak	7.0
	10:00 p.m. to 7:00 a.m.	Off-peak	3.0

Figure 6.3: The energy prices.

are operated less during the times when the energy is most expensive. However, it is not at all clear that this is the optimal solution since the energy price varies through the day and several other constraints need to be satisfied. Also, filling the reservoirs up to the maximum level could result in a surplus of water stored in the reservoirs. Moreover, how to operate the pumps such that this strategy is fulfilled, is also not known.

We use several analytic and computer aided approaches to tackle the problem. First, we use the fact that the Grimsby drinking water supply region can be split up into several independent modules (OPIRS in Figure 6.1). In section 2, we study one such a module analytically. For this module, a set of algebraic-differential equations is derived which is then optimized by using a Lagrange multiplier. In section 3, we analyze optimal pump rates for the pump stations. Then, in section 4, we develop a method that, given a certain flow rate going to a pump station, determines the combination of which pumps should be switched on and which ones off to give this flow rate. Finally, using this result, a feedback controller is proposed in section 5.

6.2 Analytic approach to flow control

6.2.1 Modular approach to water pump optimization

In a network of water pumps, reservoirs, and supply regions, it is possible to identify a general module. The entire network can then be interpreted as a network of such modules. The general module is displayed in Figure 6.4.

The module consists of n water pumps that can be switched either on or off, with




Figure 6.4: A module consisting of a set of n water pumps, a pipe with resistance R and a reservoir.

no intermediate states. The flow through pump i is given by $F_i \geq 0$ in m^3/h . The sum of these flows is necessarily equal to F_{in} . Behind each pump, the water flows through a pipe which has a characteristic resistance R (h^2/m^5). The flow is then split into the flow demand by the supply region (or to another module), $F_{out} (\geq 0)$, and a flow F_{res} into the reservoir. The reservoir flow F_{res} can be negative, representing a flow from the reservoir to F_{out} . We have the following continuity equation

$$F_{in} = \sum_{i=1}^n F_i = F_{out} + F_{res}. \quad (6.1)$$

The water height above sea level in the reservoir is denoted by L_{res} in m . It can thus be compared to the water level just before the pumps, L_{in} , also absolute above sea level. The water level satisfies the differential equation

$$A \frac{dL_{res}}{dt} = F_{res}, \quad (6.2)$$

where A is the surface area of the reservoir. The head H in m (a measure for pressure)

required to transfer the water from L_{in} to L_{res} is given by

$$H = L_{\text{res}} - L_{\text{in}} + R \left(\sum_{i=1}^n F_i \right)^2 = L_{\text{res}} - L_{\text{in}} + R F_{\text{in}}^2. \quad (6.3)$$

For all the pumps that are switched on, a nonlinear relation holds between the head over the pump and flow F_i through the pump. This relation is given by

$$H = -\alpha_i F_i^2 + \gamma_i, \quad i = 1, \dots, n. \quad (6.4)$$

Here (α_i, γ_i) are positive constants that characterize pump i for $i = 1, \dots, n$. The pressure over all pumps is equal, which explains why H is independent of i . If a pump i is switched off it gives rise to the flow $F_i = 0$.

The requested flow F_{out} and the water level L_{in} are assumed to be given functions of time. If we decide which pumps are turned on, equations (6.1), (6.3) and (6.4) give $m + 2$ relations in the $m + 2$ unknowns $F_{\text{res}}, (F_i), H$, where m is the number of pumps that are switched on. Roughly speaking, by the implicit function theorem this determines F_{res} locally as a function of $L_{\text{in}}, L_{\text{res}}$ and F_{out} . The dynamics of L_{res} are then described by (6.2), and from L_{res} all the other dynamics follow.

6.2.2 Analytic approach to flow control

The energy price is given by $c(t)$ as a function of time (in ct/Kwh). We have $P = P(F_{\text{in}}, H)$ which expresses the power consumed by the pumps in kW . An approximation for P is

$$P(F_{\text{in}}, H) = k F_{\text{in}} H, \quad (6.5)$$

where $k > 0$ is a constant that relates to the efficiency of the pumps. In practice k depends on F_{in} . Altogether, this enables us to express the pumping power in monetary units as a function of time. Our goal is to minimize the total monetary costs over a time span T . To make analysis possible, we make the following simplifying assumptions.

Assumption 6.1. We can obtain any flow F_{in} by switching pumps on or off.

Assumption 6.2. There are no constraints on allowed reservoir levels.

There are some objections. Since pumps are only switched on/off at discrete times (e.g. at most every 15 minutes), and since pumps have their limits, Assumption 6.1 does not hold in practice. Furthermore, water levels should be kept in a [75 %, 95%] range of the reservoir capacity, so Assumption 6.2 can not hold in practice. We discuss these objections later.

For any given F_{in} , we can in principle compute the resulting head through either equation (6.3) or, in case we switch only one pump on, (6.4). It makes more sense to use (6.3) since the required head is expressed through this equation, which is of crucial importance. Equation (6.4) actually loses its meaning when we use assumption 6.1: we assume that the given head/flow combination can be delivered by a certain combination of pumps and are thus indifferent about the specific pump characteristics.

Next we derive an optimization problem to minimize the required energy. We treat F_{in} as the control variable and use the notation $u = F_{\text{in}}$, $x = L_{\text{res}}$.

The total energy used over time span T is given by

$$E[x, u] = \int_{t_0}^{t_0+T} c(t)P(F_{\text{in}}, H) dt = \int_{t_0}^{t_0+T} c(t)ku(t) \left(x(t) - L_{\text{in}}(t) + Ru(t)^2 \right) dt, \quad (6.6)$$

where we used (6.3) to express H as a function of (u, x) , and (6.5) to calculate the power used by the pumps. Equation (6.2) translates into the constraint

$$\dot{x}(t) = \left(u(t) - F_{\text{out}}(t) \right) / A, \quad t \in [t_0, t_0 + T]. \quad (6.7)$$

We require that, after time T , the reservoir level x is equal to its starting value at t_0 , that is

$$x(t_0) = x(t_0 + T). \quad (6.8)$$

Using Lagrangian multiplier $\lambda(\cdot)$ to include the constraint (6.7) we obtain the following Lagrangian,

$$\begin{aligned} L(x, u, \lambda, \mu) &= E(x, u) + \int_{t_0}^{t_0+T} \lambda(t) \left((u(t) - F_{\text{out}}(t)) / A - \dot{x}(t) \right) dt + \mu(x(t_0 + T) - x(t_0)) \\ &= \int_{t_0}^{t_0+T} c(t)ku(t) \left(x(t) - L_{\text{in}}(t) + Ru(t)^2 \right) + \lambda(t) \left((u(t) - F_{\text{out}}(t)) / A - \dot{x}(t) \right) dt \\ &\quad + \mu(x(t_0 + T) - x(t_0)) \\ &= \int_{t_0}^{t_0+T} c(t)ku(t) \left(x(t) - L_{\text{in}}(t) + Ru(t)^2 \right) + \lambda(t) \left(u(t) - F_{\text{out}}(t) \right) / A \\ &\quad + \dot{\lambda}(t)x(t) dt \quad + (\mu - \lambda(t_0 + T))x(t_0 + T) - (\mu - \lambda(t_0))x(t_0), \end{aligned}$$

where we used partial integration in the last step. At an extremum, small variations of x (with fixed boundary values $x(t_0) = x(t_0 + T)$), u , λ and μ should have no influence on the value of the Lagrangian. By formally differentiating with respect to x we obtain the condition

$$c(t)ku(t) + \dot{\lambda}(t) = 0, \quad t \in [t_0, t_0 + T]. \quad (6.9)$$

Differentiating with respect to u gives

$$c(t)k\left(x(t) - L_{\text{in}}(t) + 3Ru(t)^2\right) + \lambda(t) = 0, \quad t \in [t_0, t_0 + T]. \quad (6.10)$$

From (6.10) we derive that, for $t \in [t_0, t_0 + T]$,

$$u(t) = \begin{cases} \sqrt{\frac{1}{3R}\left(L_{\text{in}}(t) - x(t) - \frac{\lambda(t)}{c(t)k}\right)}, & \text{if } c(t)k(L_{\text{in}}(t) - x(t)) - \lambda(t) \geq 0, \\ 0, & \text{otherwise.} \end{cases} \quad (6.11)$$

Recall the differential equation (6.7) for x . We have now obtained the coupled set of differential equations

$$\begin{aligned} \dot{\lambda}(t) &= -c(t)ku(t), \\ \dot{x}(t) &= (u(t) - F_{\text{out}})/A, \end{aligned} \quad t \in [t_0, t_0 + T],$$

subject to boundary condition (6.8).

As an example we solved these equations numerically for the Smithville reservoir, starting from $t_0 = 0$ over a time period $T = 24$ h and assuming a fixed water level L_{in} . The solutions are depicted in Figure 6.5.

We see that there is no flow when energy is most expensive. Unfortunately the reservoir limits [230, 240] are significantly exceeded. So Assumption 6.2 is strong. Otherwise the results seem realistic, which gives confidence in our methods. It is interesting that the flow rate is not periodic. Furthermore it is of interest why the function of flow rate with respect to time has (at certain time intervals) the form of the square root function.

The violation of assumption 2, required for a global optimum, implies that the water level constraint reduces optimality of the controlled system. This conclusion can help for future design considerations.

In future research, reservoir limits can be included in the optimization either by using slack variables or by using a penalty function. The same can be done to include maximum flow rates, in order to relax assumption 6.1. Another interesting topic of further research is to optimize a coupled set of modules, so that an entire network of pumps, pipes and reservoirs can be controlled optimally.

6.3 Optimal Pump Rates for Four Stations

6.3.1 A simple model

In this section we calculate optimal pump rates for the four pump stations at Park Road, Grimsby High, Smithville, and Beamsville. First, a model is defined for the dynamic




Figure 6.5: Upper left: Water demand F_{out} in Smithville. Upper right: Energy costs $c(t)$ per kWh. Bottom left: Optimal water flow $u = F_{\text{in}}$ through pumps. Bottom right: Reservoir level $x = L_{\text{res}}$ in case of optimal flow.

behavior of the water levels in the four corresponding reservoirs. This model can easily be obtained by performing a mass-balance equation for each reservoir. Denote with $x_i(t)$ the water level in each of the reservoirs so that the following 4-state model can be defined:

$$\begin{aligned}
 \dot{x}_1(t) &= \frac{u_1(t) - d_{Gr}(t) - u_2(t) - u_4(t)}{A_{PRdRes}} \\
 \dot{x}_2(t) &= \frac{u_2(t) - u_3(t)}{A_{SmRes}} \\
 \dot{x}_3(t) &= \frac{u_3(t) - d_{SmV}(t)}{A_{SmTow}} \\
 \dot{x}_4(t) &= \frac{u_4(t) - d_{BV}(t)}{A_{HxRes}}
 \end{aligned} \tag{6.12}$$

where $d_{Gr}(t)$, $d_{SmV}(t)$, $d_{BV}(t)$ are the demand curves for Grimsby, Smithville, and Beamsville, respectively (these are assumed known and given) [m^3/hr]; A_{PRdRes} , A_{SmRes} , A_{HxRes} , A_{SmTow} are the local surface areas of the three reservoirs and Smithville Tower [m^2]; $u_1(t)$, $u_2(t)$, $u_3(t)$, $u_4(t)$ are the pump rates at Grimsby High Pumping station, Park Road, Smitville Pump-

ing Station, and Beamsville pumping station, respectively [m^3/hr]. Finally, the states $x_1(t)$, $x_2(t)$, $x_3(t)$, $x_4(t)$ are the water levels in the three reservoirs and at Smithville tower.

An important note on the required pump rates (to be solved for) is the assumption of a continuous variable $u_i(t)$ for all four pumping stations. This is a simplification which could be relaxed at a later stage of the project, but for now it is very convenient to allow a continuous variable since the optimal control algorithm at our disposal can directly be applied. A realization of the optimal pump rates is deferred to section 6.4. In addition to the above dynamic constraints we also introduce four state constraints on the water levels. After some discussion with the problem owner we decided to maintain the water levels in a bandwidth of 25% to 95% of the maximum levels allowed. With regard to the pump rates it should be noted that we assume only positive values of the inputs $u_i(t)$ for a simple reason: the pump rates are not allowed to pump in reverse direction.

6.3.2 The goal function

The provided problem description included a clear goal, namely to minimize on the electricity price for operation of the four pumping stations. The following table of electricity prices was included:

Electricity Price [ct/kWh]	Time Slot
3.0	22:00 – 7:00 hr (next day)
7.0	7:00 – 11:00 and 17:00 – 22:00 hr
8.7	11:00 – 17:00 hr

The trade-off that needs to be optimized in this case is storage of water in the reservoirs that can be stocked at cheap electricity time-slots, whilst not increasing the water levels too much since additional head builds up when the reservoirs are filled with water and this hampers the pumps in their task (thereby reducing the flow rates). For each pump station the consumed power is

$$P(t) = C H(t) u(t)$$

where $P(t)$ is the power [kW], C is a constant characterizing the pump efficiency, and $u(t)$ is the flow rate [m^3/hr]. The hydrologic head as experienced by a pump is given by

$$H(t) = \Delta L + x(t) + R u^2(t)$$

where ΔL is the elevation difference between two pump stations, $x(t)$ is the water level in the reservoir, and R is the hydrologic resistance of the piping network. Let $p_E(t)$ denote

the pricing of electricity [ct/kW]. Then our problem is to minimize the total monetary costs over 24 hours:

$$\int_0^{24} p_E(\tau) P(\tau) d\tau$$

6.3.3 Results

To obtain some first results the above problem was programmed in Matlab, making use of the so-called TomLab PROPT toolbox for optimal control. The software allowed all constraints (both input and state constraints) to be included. To force a cyclic solution, and not to obtain so-called ‘greedy control’, terminal constraints were included so that the final water levels in the reservoirs are exactly the same as the initial water levels. In Figure 6.6 the optimization results are presented in three graphs. In the first graph we see the water-levels in the reservoirs. It is immediately clear that Smith Tower with a relatively small capacity is used as a storage during off-peak hours and this clearly pays off in terms of electricity use.

Grimsby High lifting station has the highest pump rates which can be expected since it has such a central position as a gateway to the three communities. It is clear from the results that our calculated strategy anticipates on low electricity prices by pumping intensively during the off-peak hours. Also, the pumps do not switch off completely in the most expensive hours, indicating that hydrologic head buildup is circumvented.

Of course, the above results are just a starting point that should be elaborated upon at a later stage. More refinement in, for example, the hydrologic resistance R for the piping networks could be taken into account and, also, the on-off switching nature of the controls.

6.4 Conversion of continuous flow rates into pumping combinations

6.4.1 Introduction

In the previous sections, continuous-time methods are used to control drinking-water supply systems. These methods assume that the flow rates for each pumping station is a continuous control input that can be controlled directly. However, the on- and off switching of the pumps make it a discrete quantity. In this section, a method is developed that computes for any given continuous-time flow rate a combination of switched on pumps,




Figure 6.6: Optimization results for pump-stations

that results in a flow rate that is most similar to the one that was given. Throughout this section, we assume that there are no transient effects, i.e. when a pump is switched on or off, the resulting flow is immediately in steady state.

6.4.2 Modular flow model for given pump states

First, a single pump is considered. The pumping pressure can be represented by a quantity in meter. The head H is

$$H = \frac{p}{\rho g}, \quad (6.13)$$

with p the pressure, ρ the density and g the acceleration of gravity. The sum of the head and the physical height difference (generalized head) determines the flow rate. This relation can be inverted: if the flow rate is given, then the generalized head can be calculated via a Bernoulli equation.

A pump P is considered as an object with two member functions: $P.\text{head2flow}(H)$ calculates the flow rate for a given generalized head H , and $P.\text{flow2head}(F)$ calculates the generalized head for a given flow rate. The same analysis applies to pipes. Pumps and pipes are examples of a network. Each network object N has member functions $N.\text{head2flow}(H)$ and $N.\text{flow2head}(F)$.

Networks can be build recursively from parallel connections and serial connections, and we treat them separately.

- For a parallel connection N with subnetworks $S[1], \dots, S[n]$, the flow resulting from a given head is calculated by adding the flows through the subnetworks:

$$N.\text{head2flow}(H) = \sum_{i=1}^n S[i].\text{head2flow}(H) \quad (6.14)$$

The function $H = N.\text{flow2head}(F)$ is now evaluated by iteratively searching H such that $N.\text{head2flow}(H) = F$. For this, we used an algorithm that solves one nonlinear equation with one unknown.

- For a serial connection N with subnetworks $S[1], \dots, S[n]$, the head over network N is calculated by adding the heads over the subnetworks:

$$N.\text{flow2head}(F) = \sum_{i=1}^n S[i].\text{flow2head}(F) \quad (6.15)$$

The function $F = N.\text{head2flow}(H)$ is evaluated by iteratively searching a flow F such that $N.\text{flow2head}(F) = H$. Note that the generalized head over a network N between two reservoirs is known, because it equals the height difference ΔL (m) between the reservoir levels.

Now, for any given pump state, the resulting flow rates and the pumping pressures can be found by evaluating $N.\text{head2flow}(\Delta L)$. This gives a table with all possible flow realizations and their corresponding pump states.

	α	γ	max. head	max. flow
pump 1	-1/2	2	2	2
pump 2	-1/3	3	3	3

Table 6.1: Characteristics of pumps in calculation example

flow	Pump states		head of pumps	head of pipe
	pump 1	pump 2		
0	off	off	0	0
1	on	off	1.5	-0.5
1.5	off	on	2.2	-1.2
-	on	on	-	-

Table 6.2: Possible flows in calculation example; if both pumps are on, there is no solution for the head and the flow rate

6.4.3 Model application

As an example, we investigate a pumping station with two unequal pumps. Assume that the pumps both satisfy $\text{flow}^2\text{head}(F) = -\alpha F^2 + \gamma$, but with different characteristics α and γ , listed in Table 6.1. These pumps are connected in parallel, and the pumping station is connected in series with a pipe satisfying $\text{flow}^2\text{head}(F) = -0.5F^2$. Assume that this network is connected with two reservoirs having a water level difference of 1. The flow rates that can be realized are shown in Table 6.2. There are two pumps that can be either on or off, resulting in four possibilities. It is interesting that there is no solution if both pumps are turned on. The reason is that the largest pump would produce a larger head than the maximum allowed for the smallest pump (Table 6.1).

The drinking water supply system for Smithville, Beamsville and Grimsby is shown in Figure 6.7. To apply our algorithm, the system is divided in five subsystems. If the dynamics of the reservoir levels is considered small compared to the static height differences between the reservoirs, then subsystem 4 and subsystem 5 in Figure 6.7 are driven by a fixed height difference. These systems are independent on the subsystems 1 – 3 (provided that shared reservoirs are not empty). Subsystems 1 – 3 are not independent, because they depend on the generalized head in connection point C. To solve this, we make use of an extra equation that expresses mass conservation

$$F_1 - F_2 - F_3 = 0. \tag{6.16}$$




Figure 6.7: Network representation of supply to Smithville, Beamsville and Grimsby, with subsystems 1 – 5

6.4.4 Discussion

By applying recursion, the multivariate problem can be solved using a numerical method that solves scalar equations. The recursive algorithm applies to a general class of drinking water supply systems, including the Smithville, Beamsville and Grimsby situation. The algorithm was applied to a calculation example with unequal pumps. In this example, one pump state combination was impossible, because it resulted in a larger head than possible for the one pump. This boundary indicates that the algorithm should be used with caution.

The result is a table that contains all possible steady states for the flow rates in the system. This table can be used to approximate a continuous-time control input (flow rate) by a discrete one (pump combination).

6.5 Local linear feedback control

6.5.1 Problem definition

The starting point of this section is the optimal state and input trajectories that were derived in 6.3. An optimal control method is called open loop, which means that it

computes the trajectories beforehand. This provides an ideal choice of input for the undisturbed system, but in general the robustness with respect to model errors and disturbances, such as deviations from the predicted water demand, is hard to guarantee. For example, although the optimal input trajectories are refreshed each 15 minutes, which can be seen as state feedback control, there is no 'integrating action'. This means that the for a model error or disturbance that generates an output error that is constant in time, the controller keeps responding in the same way and thus keeps making the same error. More concrete, if the optimal controller keeps predicting an input that results in a too low water level in the basins, there is no mechanism that adjusts the input for that constant output error.

6.5.2 Approach

A feedback controller is proposed to make the optimally controlled system more robust against the above mentioned perturbations. The purpose of this controller is to drive the system output (water height in the basins) to the output trajectory that was predicted by the optimal controller by using a feedback mechanism that adjusts the input (pumping power). There are many ways to design a feedback controller, and they come from different fields. We choose model based linear feedback control, since this is a widely successfully applied approach that has a strong mathematical foundation. Further, it is a textbook subject and design tools are widely available.

The approach is the following. First, since linear control design is mathematically only possible for linear systems, the system is linearized around the optimal state and input trajectories. The new variables are the deviations from the optimal variables, and they are defined as

$$\begin{aligned}\tilde{x}(t) &= x(t) - x_{opt}(t) \\ \tilde{u}(t) &= u(t) - u_{opt}(t) \\ \tilde{y}(t) &= y(t) - y_{opt}(t),\end{aligned}\tag{6.17}$$

where x is the water height in the basins, y the measured output (for example the water height in some of the basins), and u the pump rates. The subscript "opt" refers to the optimal trajectories derived in section 6.3. Inserting (6.17) in the model equations (6.12) gives a system of the form

$$\begin{aligned}\frac{dx(t)}{dt} &= Ax(t) + Bu(t) + Ee(t) \\ y(t) &= Cx(t),\end{aligned}\tag{6.18}$$

with e the disturbances in the water demand curves $d(t)$, and A , B , C and E system matrices. The system (6.18) is in a form that allows the design of a feedback controller, e.g. via H_∞ theory.

6.5.3 Discussion

It is shown that given an optimal input and state trajectory, a robust linear feedback control design is possible. The design itself is omitted, but as mentioned before this is a textbook subject. The controller acts locally in time, as opposed to the optimal controller, so it does not look ahead to save energy costs. For example, it does not shift pumping duties to the night time because the power is then cheaper. In theory, each time that the optimal trajectory is refreshed, the trajectories that the state and input are linearized around changes, resulting in a different controller each 15 minutes. So A , B , C and E change. However, since the controller will be robust against model errors, this is not necessary as long as the differences stay reasonably small.

6.6 Conclusions

Given the fluctuating energy prices and the drinking water consumption in the Grimsby area, where each of the 15 pumps are either switch on or off, in total $2^{4 \times 48 \times 15}$ possible trajectories result. Hence, a solution via enumeration is infeasible and thus there is a need for approximations.

At first, we consider smooth (continuous-time) pump functions and one head-flow relationship per pumping station. Hence, a set of algebraic-differential equations with constraints result. This allows us to use Lagrangian theory for dynamic systems, also known as the minimum principle of Pontryagin. After all, a two-point boundary value problem (TPBVP) in terms of the states and co-states results. Recall that in this problem, the flow generated by a pumping station is the control input and the height in a reservoir is the state of the system. Given the energy-related goal function, together with the input and state constraints, numerical solutions to the TPBVP have been found. If, however, we consider the unconstrained problem for a single pumping station configuration, (semi-)analytical solutions result.

In a second step and given a required head-flow combination as a function of time, as found after solving the TPBVP, an optimal pump configuration can be selected. Careful analysis of the Grimsby region, under quasi-steady state assumptions, also shows four archetypical modeling problems, which can be solved in a modular approach.

The direct link between energy costs and the flow signals allow a direct physical

interpretation. The unconstrained problem for Beamsville, while considering only the running costs, clearly shows that enlarging the Hixon reservoir and pumping capacity is profitable.

Further research is needed to analyze the problem with respect to the sub-optimal solutions found in this work, to set-up a generic framework for the dynamic optimization of any drinking water network and to come up with real-time solutions.