

Flow balancing in a wastewater network by optimising storage in pumping stations

C.N. Burt^{1,2}, B. Candy¹, R. D'Arpa¹, E. Lim Koh¹,
P. Meng¹, D. Salamone¹, T. Sondalini¹ and P. Taylor¹

1. *Water Corporation, Perth, Australia*

2. *Department of Mathematics and Statistics, The University of Melbourne*

Email: phillip.meng@watercorporation.com.au

Abstract: We have an urban wastewater conveyance and treatment network that requires the assets to cope with the peak system loading, even if for just a brief period of time per day. This results in an underutilised asset base, and unbalanced inflow into the treatment plant; the latter results in suboptimal conditions for the systems that treat the water. We propose to use the vast, existing wastewater network and reticulation as an 'equalizer tank' in order to smooth the demand at the treatment plant. The purpose is to achieve improved performance of our existing assets without further investment. In this paper, we describe a proof-of-concept algorithm to demonstrate the ability to control the diurnal load pattern of wastewater conveyance in a live network, along with results from a live pilot.

The benefits of such capability are two-fold: potential deferral of capital investment for asset upgrades, and improved operating conditions and costs. Equaliser tanks are commonly used to both improve operating conditions at the treatment plant and to provide a mechanism to control high volume events. However, these benefits could be achieved using the network itself. By controlling the flow and effectively flattening it across the day, the requirement to handle the peaks in flow diminishes and reduces the need to increase capacity of pipes, pumping stations and the treatment plant itself to meet the peak demand.

The problem of deciding how much flow to permit from each pumping station is fairly trivial. However, in our network setting we only control the pumping station flow indirectly via level set-points in the well. That is, the pump will only activate when the water reaches a cut-in set-point, and will stop pumping when the water reaches the cut-out set-point. The pump rate is usually fixed-speed, and is not expected to change during or between pump cycles. Additionally, we have only indirect measurement of the flow itself, either by the hours of pumping and pumping rates or by the levels and pumping times. Furthermore, the control system is limited to 12 settings per day and 4 settings per year, restricting us to a static solution for a stochastic problem.

In this paper we present an algorithm for obtaining a complete network set-point solution in efficient runtime. The algorithm guarantees finding feasible solutions when they exist, and is guided by the optimal solution for the deterministic version of the problem. We have modelled the wastewater network in a hydrology modelling tool which provides some basic simulation functionality. We present the results of a pilot we conducted on a subset of our wastewater network. We monitor asset performance and H₂S levels and report the differences. We demonstrate that flattening of the diurnal curve is possible to achieve in a controlled manner.

Keywords: *Wastewater flow balancing, equalising, optimisation, dynamic algorithm, stochastic*

1. INTRODUCTION

The daily inflow of wastewater into an urban wastewater treatment plant typically follows a bi-modal curve, sometimes referred to as diurnal peaks (see Figure 1) (see, e.g. Czachorski 2017). The problem is to achieve a flattened inflow through flow balancing. Flow balancing is an industry standard strategy for smoothing the inflow to treatment plants and is typically achieved through the use of storage tanks (e.g. see Demey 2001; Aeration Industries International 2019). In this paper, we propose to smooth the inflow at treatment plants by performing flow balancing within the network itself, through the design and implementation of a data-driven optimisation algorithm to control the safe storage and discharge of waste water. We present the results of a pilot study on an isolated branch of the network (see Figure 2).




Figure 1. A typical daily inflow curve depicting bi-modal behaviour.

In one case study we can defer capital investment in the order of \$60 million dollars for at least 10 years. A secondary driver is the benefit obtained through the smoothed diurnal curve at the treatment plant itself, which leads to a more efficient and cost-effective biological process (e.g. see Aeration Industries International 2019, Goel 2005).




Figure 2. A schematic of the network considered in our pilot study. Four pumping stations pump via pressure main (red) to a gravity main (black) to a treatment plant. Users of wastewater services in the catchment feed into the pumping stations or gravity main directly via reticulation network.

are controlled by levels in the well. In this setting, each pump operates only when the wastewater level reaches the so called ‘cut-in’ level, and pumps at a fixed speed until the wastewater level is reduced to the ‘cut-out’ level. As a result, the outflow rate, quantity and timing of the pumping stations can only be controlled indirectly by altering the cut-in and cut-out levels. Adding to this challenge is the fact that the controller is limited to only 12 cut-in and cut-out levels per day (fixed per season).

In this paper, we propose an algorithm that:

- automatically scales to larger and other networks,
- is guaranteed to find a feasible solution if one exists,
- is guaranteed to determine if no feasible solution exists,
- is conservative against stochastic fluctuations, and
- safely discharges the stored wastewater such that no new peaks are formed.

We present the results from a proof-of-concept where we implemented the algorithm and tested the result in both a simulation and real-world, practical pilot. But first, we describe the necessary background in the following section. Then, we formally define the problem and describe the results of the literature survey, followed by a description of the algorithm.

The key driver for this work is the total cost of the assets for the conveyance and treatment of wastewater. The capacity of the entire network, including the treatment plant, must be scaled to meet the peak demand. Thus the network is underutilised for over 85% of the day, representing an opportunity to attenuate the impact of peak wastewater inflow by storing it in residual capacity within the network itself. By smoothing the inflow we can defer installation of larger assets.

However, the problem is not trivial. While the inflow pattern follows a predictable shape, it is stochastic (Penn, 2017) and can vary significantly from day to day, even after controlling for wet weather. For our wastewater network system in current state, the control system used in transfer pumping stations is limited to a static (i.e. fixed) profile. Furthermore, the profiles are also fixed to a small number of implementable seasons. This leaves us with a static day-to-day solution for a stochastic problem. A second challenge arises from the nature of the pumping stations themselves, which predominantly have fixed speed pumps that

2. BACKGROUND

Metropolitan Perth wastewater is transferred from customers to treatment facilities via a network of reticulation pipes, gravity mains and pressure mains (e.g. see Figure 2). Wastewater pump stations receive inflow via gravity-fed reticulation pipes into local well-storage with defined pump start (cut-in) and stop (cut-out) level set points. The pumped wastewater is transported via pressure mains to the next part of the network, cascading through the network of mains and pump stations towards the treatment plant.

Assets for the wastewater treatment network must be designed to be capable of handling the peak flow and total daily volume. The characteristic diurnal flow profile across the network has an overnight minimum and two peaks: one in the morning and one in the late-afternoon. The requirement to handle these peaks means that for the majority of a day, the capacity of the system is underutilised; while for only two short time periods the system is under strain. It has been demonstrated that if the load across a day could be smoothed out, this would benefit the wastewater treatment operation and allows asset capacity upgrades to be deferred (e.g. see Aeration Industries International 2019, Goel 2005). This is most commonly achieved through the use of flow balancing or equalising tanks, which can be very large (~50-200ML) and store the bulk of the daily peaks. However, we could not find instances of this being achieved using the existing network as storage.

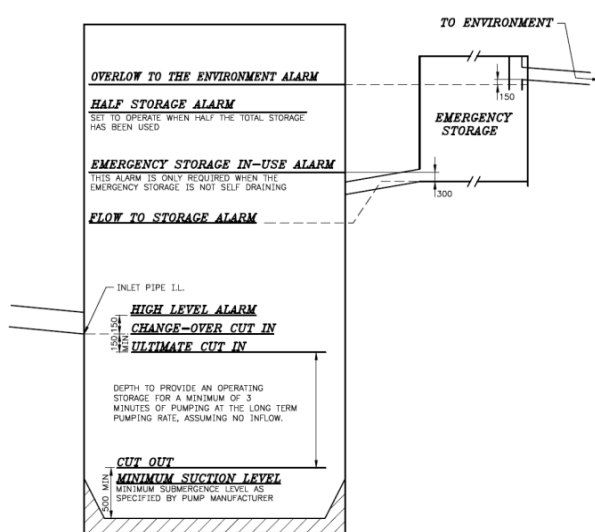


Figure 3. Schematic of pumping station, illustrating the inflow inlet, outflow to emergency storage, and standard cut-in and cut-out levels (Water Corporation 2019).

Pump stations are equipped with instrumentation for wastewater level sensing, programmable logic controllers (PLC) for local control of the pumps based on the level sensing, as well as external control and configuration through Supervisory Control and Data Acquisition (SCADA). Most pumps in the network are not variable speed, and are therefore operated as either on or off, alternating between pump units (where installed in a dual configuration). A pump turns on if water in its wet well reaches the ‘ultimate cut-in’ level, and turns off if water reaches the ‘cut-out’ level (see Figure 3). Control and configuration of these cut-in and cut-out levels is limited (in our system) to 12 diurnal level control set points. Our SCADA controller is limited to four seasons, which effectively limits the solution to four settings/modes per year. Our analysis of the seasonality confirms that the best use of these seasons is to cover wet and dry, weekday and weekend.

As part of the network design planning, a key metric used is time to overflow (TTO), which is a calculation based on the remaining storage in the system in combination with the characteristic inflow curve, and is used to ensure sufficient capacity is present in the system (translated to time) for operations crews to react to and resolve faults.

3. METHOD

Our approach to the flow-balancing problem is to perform the balancing within the existing network where possible. This leads to our formal problem definition:

Given

- a set of pumping stations (and associated wet-wells),
- a set of gravity mains and pressure mains that make up the wastewater reticulation network,
- elevation level to volume mapping for the reticulation network upstream of each pump station,
- the maximum level associated with each wet-well such that the maximum does not violate time to overflow requirements, and
- the expected, worst case ‘demand’ (i.e. inflow curve) into each pumping station in the network for the given season;

determine the cut-in and cut-out levels for each pumping station such that the maximum flow capacity is minimised at the wastewater treatment plant such that:

- the frequency of pumping is not increased,

- the overflow time limit associated with the catchment area is not violated,
- minimum flow constraints are met,
- flow equilibrium throughout the network is consistent, and
- no use is made of emergency storage.

To the best of our knowledge, there is no other published work that addresses this problem or application. After an extensive literature search in the industrial space, we determined that there are many instances where equalisation has been achieved through the use of tanks or storage pipes, but not in an integrated network approach. The former problem provides a trivial optimisation problem, and therefore does not warrant publication of the solution. We were not able to find a class of problems in the academic literature that provided a match for the whole or subset of our problem. This is likely to be due to the complexities arising from limitations and uniqueness in our SCADA set-up, rather than from the network problem itself (which is trivial).

Our approach for developing an algorithm to solve this problem was based on the need for a computationally efficient and scalable algorithm, and was guided by the tenets of optimisation: that we can prove if no feasible solution exists; that if a feasible solution exists, we can find it; and, that we always know the quality of our solutions with respect to optimality. Since the control system is static, the settings are fixed for a given mode; we are forced to provide a static solution that accounts for the stochastic and dynamic nature of the problem while guaranteeing that the new inflow curve does not exceed the height of the characteristic inflow curve.




Figure 4. The target volume is highlighted in dark blue. The light blue line intersects with the y-axis to give the target inflow rate. The intersection of the light blue line and characteristic curve gives the corresponding time.

The key to solving this problem efficiently lies in the decomposition of the network problem. We decompose the problem by observing that a target volume in a peak corresponds to a target inflow rate, and this subsequently corresponds to a target inflow rate at the upstream pumping stations. The intersection of the target inflow rate and the characteristic curve gives the start time from which we should store the water at that location (see Figure 4).

As a part of our input data set-up, the characteristic daily inflow curve (e.g. see Figure 1) for each pump station was computed using pump station sensor time series data for the past two years—data older than this was considered to be unrepresentative of the current-state system behaviour and would bias the result. We performed a seasonality analysis of these curves to confirm the seasonality in both the day of the week (weekday or weekend) and month of the year for the four SCADA settings.

Our Operations Centre provided a time-based constraint, time to overflow, which should not be violated when modifying the pump station operation parameters in this optimization. Remaining storage, relative to time to overflow, was calculated by using the storage volume associated with the time to overflow and solving for when the accumulated inflow of water matches this storage (assuming no pump operation). Thus the result is a nonlinear curve that indicates the critical times in the operation, corresponding to the periods of peak inflow (see Figure 5). We also adopted an additional constraint by ensuring that our solution did not use the designated emergency storage, as this would incur additional cleaning costs. Together, this curve and limit formed the discharge limit for each discharge time.




Figure 5. An illustration of how remaining storage changes over the course of the day, depending on the inflow.

To begin the algorithm, we first assess the feasibility of discharging the water in the

lesser inflow period between the two peaks by first calculating the volume difference to the target inflow line. Then, starting with the most downstream station in the catchment, we dynamically calculate the starting time for the storage period for each station in the network by factoring in the travel time for each link. We then calculate new inflow curves for each station by starting with the leaf nodes and deducting the stored volume; this change in inflow cascades down the network, with further deductions made at subsequent stations.

We obtain the set-point levels for the storage period by ensuring the cut-out level is above the required storage volume for that station. We obtain the remaining discharge set-points by solving a linear program for each station to determine the quantity of volume discharged at nominated time intervals. For the case where pumping stations connect in parallel to the final pumping station (i.e. the scenario we have for our pilot), the linear program is defined as follows. Let $y_{s,t}$ be the quantity of water discharged at pumping station s at time t ; D_t^{max}/D_t^{min} is the maximum/minimum allowable discharge volume at time t ; V_s is the stored volume at pumping station s ; and $\delta, \rho, \delta_t, \rho_t$ are deviation variables. We minimise the deviation from meeting the target bounds, while ensuring the total volume of water is discharged.

$$\begin{aligned}
& \min \rho + \delta \\
& st \sum_s y_{s,t} - \delta_t \leq D_t^{max} \quad \forall t, \\
& \sum_s y_{s,t} + \rho_t \geq D_t^{min} \quad \forall t, \\
& \sum_t y_{s,t} \geq V_s \quad \forall s, \\
& \delta_t - \delta \leq 0 \quad \forall t, \\
& \rho_t - \rho \leq 0 \quad \forall t, \\
& y_{s,t}, \delta_t, \rho_t, \delta, \rho \in \mathcal{R}
\end{aligned}$$

Since a linear program is used to find the remaining part of the solution, we are guaranteed to find a feasible solution (for the given time intervals and static characteristic inflow curve) if it exists. If no solution is found for the linear program, then no solution exists for the problem as defined. For the dependent case, we use a different model, which is beyond the scope of this paper to describe here as it was not used for the pilot.

4. RESULTS

We extracted asset and sensor data from in-house databases and our wastewater network model (see Section 4.1). This data was fed into our algorithm to obtain the set-points. The set-points were tested in a simulation in the network model. We implemented the algorithm in Python 3.6 using SoPlex v.2.1.0 (Gleixner *et. al.* 2018) to solve the underlying linear program. The experiments were generated on a standard desktop machine, using calls to the NEOS Server (Czyzyk *et. al.* 1998; Dolan 2001; Gropp *et. al.* 1997) to access SoPlex. The entire optimisation algorithm runs in under 10 seconds if run locally.

4.1. Model Validation

We created a hydraulic network model of the catchment in Innovyze Infoworks ICM v.9.5 (Innovyze 2019). This model is based on physical characteristics, such as gradient and size of pipes, and is therefore capable of performing an accurate, optimistic simulation of wastewater flowing through the network. We added our proposed set-points into the network model and performed this simulation. The primary objectives of this simulation step were to check the flow rates through the system and if any depth limits or other constraints were violated. The simulation result illustrated the performance we could expect from a real-world pilot, and provided confidence in moving forward with a pilot. The results indicated that the set-points were sufficient to reduce the first peak to the height of the second and that the discharge would not create a new peak (see Figure 6).




Figure 6. Inflow curve generated from simulation in Infoworks ICM (orange) compared with the characteristic inflow curve used in the simulation (blue).

4.2. Pilot

We implemented our solution in a live sewer district in the Perth Metropolitan region for six operational days. The district contained six pumping stations, four of which we could control using SCADA and two which were private pumping stations (where no level is measurable, and as such inflow was assumed to be a static volume). We obtained the inflow curves for the median of the respective season for optimisation. The algorithm returned multiple optimal solutions, but we chose to implement the solution that involved only one pumping station in order to push the limits of that site and get a stronger sense of whether odour would be an issue. This, however, increases the risk for that particular site by decreasing the time to overflow, while having no risk change for the remaining stations. In an ongoing implementation we would opt for a solution that shares the overall risk amongst all pumping stations, where appropriate. The distributed implementation solution would allow the same amount of flow-balancing but with lessened effects of odour and material build-up, as suggested by wastewater conveyance engineers.

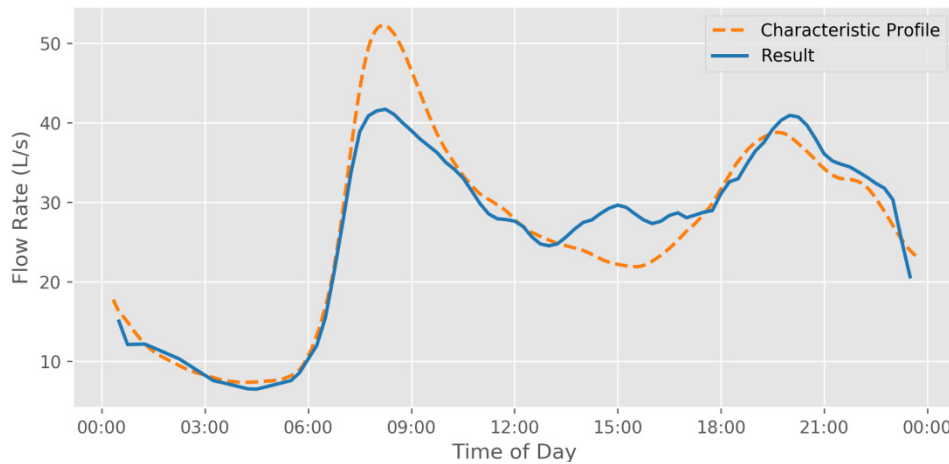


Figure 7. The result from one day of the pilot. The result is in blue while the characteristic inflow curve is in orange (dashed).

In Figure 7, we plot the resulting inflow curve from a single day in the pilot against the characteristic historical inflow curve. The result shows that we successfully reduced the first peak to the height of the second peak and safely discharged the stored water in the early afternoon without causing a new peak. The change in height of the discharge indicates that there is a lot of potential to discharge more water during that period, and that fewer set-points could be used to achieve a similar result.




Figure 8. Histogram of odour readings. Blue results are from the 2 days preceding the pilot, while the orange are from the 6 day pilot.

One potential risk is the increase in odour. Increased odour emanating from the network due to the greater storage of wastewater than usual was monitored. No increase was detected upstream from the surcharged site. The only site showing a potential increase in hydrogen sulphide gas was the discharge point of the surcharged site itself. This result is expected since the septic conditions in the systems are increased. However, the control data set was too small to draw conclusions, and we will collect more data in future studies. The number of pump starts increased on average but were less than the maximum number of starts under normal operations.

4.3. Discussion

The peak reduction in the 6 day pilot ranged from 5 to 19.5%, with an optimal solution of 18% for the deterministic characteristic inflow curve (weekday) and 12% for the weekend. The fact that we could outperform the optimal solution (i.e. obtain a better solution than expected from the linear programming output) is a reflection of the stochastic nature of the inflow, while the optimal solution was calculated using the characteristic curve. We believe we can obtain a better worst-case performance by optimising over the

90th percentile characteristic curve instead of the median, and allowing for some ‘cushioning’ on either side of the peak.

5. CONCLUSION

In this work, we have shown that we can safely displace water from the daily inflow peaks by discharging it in periods of lesser inflow. From this trial, it is conceivable that we could flatten the peaks further in future trials. The ability to achieve a better flow-balance in the wastewater network provides alternate opportunities to asset investment and replacement that are important for a cost conscience organisation. Our solution considers the existing capability of system and carefully computes the gains that can be made in a controlled manner that limits any side effects from operating under conditions other than the original design.

This solution provided enhancements in optimising the flow-balancing but analysis of the results indicates improvements can still be made. A greater number of set points to provide greater control with the storage and release of wastewater has been proposed, and is set to be enabled in the near future. This will provide more flexibility and less discretisation error in the optimisation. Importantly, this algorithm is designed to roll out to a complete network and since the calculations at each step are computationally efficient, the complete algorithm will roll out in a computationally efficient way. The calculations are fast enough that we could, in a future version, predict the inflow curve from a limited sample of the day’s data and compute the set-points in time to reduce the first peak. We are currently developing an algorithm to perform such a prediction.

ACKNOWLEDGMENTS

This project has been made possible through the continued support and collaboration from many stakeholders throughout the Water Corporation, including Matt Walsh, Wayne Bagg, and Shannon Arnott.

REFERENCES

- Aeration Industries International (2019). <https://www.aireo2.com/en/applications/equalization-basin/> [accessed 15/07/2019]
- Czachorski, R. (2017). What is diurnal flow? <https://www.h2ometrics.com/what-is-diurnal-flow/> [accessed 15/07/2019]
- Czyzyk, J., Mesnier, M. P., and Moré, J. J. 1998. The NEOS Server. *IEEE Journal on Computational Science and Engineering* 5(3), 68–75.
- Demey, D., B. Vanderhaegen, Henk Vanhooren, J. Liessens, L. Van Eyck, Lisa Hopkins, and P. A. Vanrolleghem. Validation and implementation of model based control strategies at an industrial wastewater treatment plant. *Water science and technology* 44, no. 2-3 (2001): 145–154.
- Dolan, E. 2001. The NEOS Server 4.0 Administrative Guide. Technical Memorandum ANL/MCS-TM-250, Mathematics and Computer Science Division, Argonne National Laboratory.
- Gleixner, A., Bastubbe, M., Eifler, L., Gally, T., Gamrath, G., Gottwald, R., Hendel, G., Hojny, C., Koch, T., Lübbecke, M., Maher, S., Miltenberger, M., Müller, B., Pfetsch, M., Puchert, C., Rehfeldt, D., Schlösser, F., Schubert, C., Serrano, F., Shinano, Y, Viernickel, J. M., Walter, M., Wegscheider, F., Witt, J. T. and Witzig, J., (2019). *SoPlex* <https://soplex.zib.de> http://www.optimization-online.org/DB_HTML/2018/07/6692.html
- Goel, R. K., Flora, J., and Chen, J.P. (2005). Flow equalization and neutralization. *Physicochemical Treatment Processes*, 21–45. Humana Press.
- Gropp, W. and Moré, J. J. 1997. Optimization Environments and the NEOS Server. *Approximation Theory and Optimization*, M. D. Buhmann and A. Iserles, eds., Cambridge University Press, 167–182.
- Innovyze Infoworks (2019). <https://www.innovyze.com/en-us/products/infoworks-icm> [accessed 15/07/2019]
- Penn, R., et al. Simulation method for stochastic generation of domestic wastewater discharges and the effect of greywater reuse on gross solid transport. *Urban water journal* 14.8 (2017): 846–852.
- Water Corporation (2019) Design Standards DS 51: The design and construction of wastewater pumping stations and pressure mains 4.5 to 180 Litres per second capacity. <https://www.watercorporation.com.au/-/media/files/suppliers/resources/design-standards/ds51-designconstructionofwastewaterpumpingstationspressuremains4to90litrespersecondcapacity.pdf> [accessed 15/07/2019]