Including Hydraulics in Optimising Groundwater Abstraction Scheduling

Sonam Tshering

Sarang Kulkarni and Guanglu Zhou
Industrial Engineering
Curtin University

Phillip Meng and Mike Canci
CEED Client: Water Corporation

Abstract

The Advanced Analytics Team at the Water Corporation has developed a mathematical model for optimising the abstraction schedule of groundwater from various bores to minimise the chemical treatment costs and maximise the production volume while adhering to all relevant business rules. The model generates a weekly production schedule in the form of sets of bores to be operated across an annual planning horizon. Despite its success, operational issues have surfaced over time. The model is affected by dynamics in the groundwater bore collection pipe network. Hydraulic effects, varying with bore combinations, result in lower actual flow rates than expected, affecting production targets. The project aims to reconcile the disparities between specified and actual flow rates by incorporating hydraulic effects in the optimisation model and is focussed on a single treatment plant. To achieve the objective, a comprehensive methodology that considers hydraulic effects in the bores has been developed using Mixed Integer Linear Programming. The model, implemented in Python and solved using Gurobi, is validated using historical data from the past two years and shows improved yield. This approach enhances the scheduling model's precision by leveraging hydraulic insights, reinforcing the Water Corporation's commitment to optimal groundwater resource management.

1. Introduction

Water Corporation uses mixed integer linear programming (MILP) to schedule the operations of bores in the Perth Metropolitan region (Marques et al., 2022). The model treats the flow rate as a parameter and generates a week-by-week abstraction schedule for the entire water year to ensure that the abstracted volumes meet the demand profile, water quality limits are not exceeded, and licensing constraints are adhered to while also considering a host of other requirements.

Based on the schedule, water is drawn from active bores through a shared main pipeline to the water treatment plant. While the model greatly simplified the scheduling process and brought a strategic perspective to resource allocation and managing operational constraints, a consistent observation emerged - actual flow rates were consistently lower than expected. Specifically, when multiple bores connected to the main were active, the hydraulic effects in the network led to actual flow rates being lower than if each bore operated independently (WC, 2023c). Consequently, the bore flow rate, used as a parameter in the optimisation model, is not representative of the actual flow rate, given that bore selection occurs without accounting for this hydraulic effect.

The project aims to find a methodology to incorporate the hydraulic effects in the optimisation model. The challenge from an optimisation perspective is that the hydraulic effects depend on the bore selection, and the optimal bore selection depends on the hydraulic effects. Additionally, it is important to note that most of the bore pumps in the network are constant-speed drives. These drives do not allow for varying flow rates, restricting operational flexibility. Operating pumps at higher hydraulic heads beyond their rated capacity is also not advisable, as it would result in higher energy costs and increased losses.

1.1 Optimal Scheduling of Ground Water Production

Water distribution and scheduling optimisation is a widely studied topic; however, the optimisation of groundwater production in bore networks remains relatively underexplored. The Water Corporation, in collaboration with Curtin Research Student Amanda De Azevedo Marques, has addressed this gap through their work on the optimal scheduling of groundwater production.

Scheduling in process systems generally addresses the issues of assignment, sequencing and timing and are generally formulated as MILP problems, because such problems often involve making decisions (represented as discrete variables) about how to allocate resources or perform tasks subject to certain rules (Pinto & Grossmann, 1998). The team utilised such a MILP framework to tackle groundwater scheduling issues at Water Corporation. They approached the problem by considering the flow rate in a bore as a parameter and developed a time-indexed model that divided the water year into 53 discrete events. By consistently applying mass balance principles, they accurately represented inflows and outflows within each discrete event. This approach also allowed them to control water quality parameters and decision variables across all events. The resulting model was able to generate a bore production schedule for a treatment plant for an entire year, managing over 48 thousand constraints and 78 thousand decision variables, with computational times ranging from 1 to 30 minutes. The model also enabled the team to understand the operational flexibilities of the groundwater production process. However, hydraulic effects in the bores were not considered in the model. This results in the parameter being decision dependent.

2. Including Hydraulics in the Optimisation Model

The Water Corporation has modelled its groundwater distribution systems using InfoWorks WS Pro. This software enables us to simulate and determine accurately the flow rate in a bore for any bore configuration. We capitalised on this capability to reformulate the problem as a combinatorial optimisation challenge, which can be defined as the search for the best possible permutation of a given set of objects according to a given objective function (Salvagnin, 2014). It involves enumerating all possible bore combinations and using the hydraulic model to create a lookup table. We then use Mixed-Integer Linear Programming (MILP) to select the optimal bore combination set for each period. This method is well-documented in the literature, albeit in different contexts. For instance, Pinto and Grossmann (1998) explored a single-unit assignment in a multistage scheduling problem, where a binary variable x_{ik} indicates whether product *i* is processed in time slot *k*.

$$\sum_{i} x_{ik} = 1, \qquad \forall k$$

Our approach mirrors this by aiming to select and utilise an optimal bore combination set for groundwater production on a weekly basis. This methodology is preferred mainly because it ensures that the solution is both optimal and global within the defined constraints, by exhaustively exploring all feasible combinations.

2.1 Hydraulic Model Automation and Network Segmentation

With 33 bores in the network, there are more than 8.5×10^9 different bore configurations possible. Manually simulating the InfoWorks WS Pro model for all these configurations is impossible. An InfoWorks expert at Water Corporation helped us automate this process. The tool is based on a Ruby script to read one or more pump configurations from a CSV file, execute simulations for each configuration and export summary results to a CSV file to be used as input in the optimisation model. The script is integrated with WS Pro Exchange to automate and manage simulations as an external process.

However, simulating 8.5×10^9 configurations, even with automation, presents a formidable challenge. We addressed this issue by employing a divide-and-conquer strategy, thereby decomposing the problem into two smaller, more manageable sub-problems. This approach significantly reduced the input dataset. Specifically, we identified that the network consists of two independent sections, comprising 18 and 15 bores each, which converge at the Treatment Plant. By treating the problem in this manner, we were able to reduce the input dataset to 294,912 combinations. Initial experiments indicate that simulating all bore configurations using this approach to create the lookup table requires approximately 1 hour and 40 minutes of CPU time.

2.2 The Model

Based on the selected methodology, a mathematical model was developed to represent the project's objectives, constraints, and bounds. This model ensures that the project's goals are met, and all relevant and business rules and requirements are addressed.

The model considers several sets: Regions include two regions - Region A with 18 bores and Region B with 15 bores, each bore region denoted as $r \in R$; Weeks span 53 weeks, with each period (week) represented as $t \in T$; Licenses include four licenses, denoted as $l \in L$; and there are 33 bores, denoted as $b \in B$. The model also encompasses quality parameters, represented as $q \in Q$, and includes streams for treatment and for bypass, denoted as $s \in S$; and the set of bore combinations from both regions, denoted as $i \in C$.

The primary decision variables are: $combination_{i,t}$, a binary decision variable that represents if combination *i* is active in week *t*; δ_t^+, δ_t^-, which are the volumes abstracted above and below the demand target in week t, respectively; and variables representing how much water is sent to the treatment stream and how much is bypassed.

The objective function in its general form is designed to maximise production while minimising the deviations from the weekly demand target.

$$Max \quad \sum_{t \in T} \sum_{i \in C} combination_{i,t}\, Weekly\_Flow_i - M \sum_{t \in T} (\delta_t^+ + \delta_t^-)$$

where: $Weekly\_Flow_i$ is the total flow for the combination i.
M is the penalty coefficient.

δ_t^+ and δ_t^- are then used in the demand constraint to ensure we get a feasible solution.

$$\sum_{i \in C} combination_{i,t}\ Weekly\_Flow_i + \delta_t^+ - \delta_t^- = Demand_t \quad \forall\, t \in T$$

By avoiding a strict demand requirement, it enables the model to consider both slack and excess production. This is necessary because the total flow is a function of the bore configuration and cannot be adjusted to meet demand targets. However, we do set a bound on how much we can go over and under the demand target. This increased flexibility enhances the likelihood of identifying a solution that satisfies the overall objectives while still meeting the demand requirements.

We then use the following constraint to ensure that only one bore combination set from each segmented region is used for production in each week.

$$\sum_{i \in C_r} combination_{i,t} \leq 1 \quad \forall\, t \in T, r \in R$$

A critical constraint in our model is the treatment stream constraint, which ensures that quality parameters such as alkalinity, TDS, hardness, UV, turbidity, and iron (Fe) remain within predetermined limits. Additionally, constraints are established for managing the allocated and consumed volumes of water for each bore as well as for a group of bores managed under a license. These constraints collectively facilitate the development of an operational plan that optimally blends bore water according to the capacity of each treatment plant while ensuring compliance with water quality limits and production targets throughout the year. Furthermore, the operational plan accommodates plant shutdowns and bore maintenance requirements. When a bore is under maintenance for a period, all combinations involving that bore become unusable for that period. This formulation simplifies the challenge to assigning a bound of 0 on the binary variable $combination_{it}$ for all combination i linked to the bore under maintenance in that period t.

3. Results and Discussion

The mathematical model was implemented in Python using Gurobi as the solver. Water Corporation uses Gurobi for all its complex optimisation needs due to its exceptional performance and advanced capabilities (Marques et al., 2022).

The Gurobi Python API was utilised to define model variables, constraints, and the objective function. For testing, we used historical data pulled from the PI System. Bore flow data spanning the past two years were processed to create the lookup table. During this period, 1556 unique bore combinations were used for groundwater abstraction.

Model runs using the above data show that it can schedule bores with an average production efficiency of around 99%, as illustrated in Figure 1. The final optimality gap is achieved within approximately 400 to 600 seconds, after which it stabilises, and no further improvements are observed. Given the limited lookup dataset, extending the runtime further does not significantly improve the solution. We were also able to ensure that the quality of the water produced meets the limits while ensuring that the treatment cost is minimised. This requires us to modify the objective function to include a penalty on the stream flow so that only the required amount of water is sent to the treatment stream, and the rest is diverted to the bypass.

The model built using the test data has 9122 constraints and 82680 variables, of which 212 are continuous and 82468 are binary. The model was executed on a laptop with an Intel(R) Core i5-7360U CPU @ 2.30GHz, featuring 2 physical cores and 4 logical processors and Gurobi version 11.0.2.

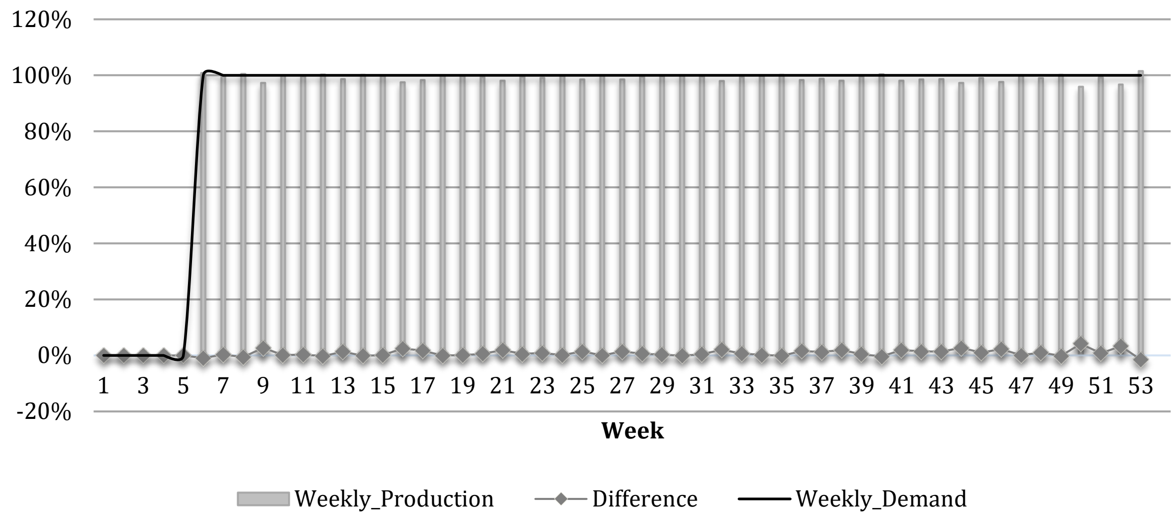

Figure 1 Normalised Weekly Production vs. Demand and Difference, showing the production efficiency from a sample run.

Although the results are promising, scaling issues are present in the model, as shown in the following model statistic:

Matrix coefficient range: [0.12001, 3.33331e+11]

The matrix coefficient range indicates a large difference between the smallest and largest values in the constraint matrix. This can potentially lead to numerical issues; however, the solution quality statistics reveal no violations in constraints, bounds, or integrality.

Nevertheless, we have identified the specific constraint causing the scaling issue and will reformulate it to enhance the model's numerical stability and ensure the reliability of the optimisation results.

4. Conclusions and Future Work

Results from the model run using historical data show that hydraulic effects can be effectively integrated into the optimisation model by framing it as a combinatorial optimisation problem. The runtime required to achieve the desired production efficiencies is reasonable, and the computational requirements are manageable.

We plan to test the model with a comprehensive dataset obtained from the hydraulic model. While we anticipate similar results, the computational requirements and runtimes for this expanded dataset remain uncertain.

Additionally, we plan to incorporate energy data for operating the pumps, available from the hydraulic model simulations, into the model and see how we can use it for energy savings and enhance the optimisation process further.

5. Acknowledgements

I am deeply thankful to my mentors, Phillip Meng and Mike Canci, and my supervisors, Dr. Sarang Kulkarni and Dr. Guanglu Zhou, for their unwavering support, guidance, and encouragement throughout the project. Their expertise and insights have been instrumental in navigating the challenges and achieving the project objectives. I would also like to thank Amanda McGillivray and Penny Van Reenen for their valuable input and feedback, and Bernard Candy for his work on the automation of the hydraulic model. My sincere thanks also go to the Water Corporation Research and Development program for funding and supporting this project, and to CEEDWA and Curtin University for providing this opportunity.

6. References

Marques, A. D. A., McGillivray, A., & Meng, P. (2022). Advanced Water Production Scheduling Using Mathematical Optimisation.

Pinto, J. M., & Grossmann, I. E. (1998). Assignment and sequencing models for the scheduling of process systems. Annals of Operations Research, 81(0), 433-466.

Salvagnin, D. (2014). Detecting and exploiting permutation structures in MIPs. Integration of AI and OR Techniques in Constraint Programming: 11th International Conference, CPAIOR 2014, Cork, Ireland, May 19-23, 2014. Proceedings 11,

WC. (2023c). Including Hydraulics in Optimising Groundwater Abstraction Scheduling. https://nexus.watercorporation.com.au/otcs/cs.exe/app/nodes/163948716