Integer Programming Optimization of Production Well Placement

Rishi Adiga, John O'Sullivan, Andy Philpott
Department of Engineering Science
University of Auckland
New Zealand
radi130@aucklanduni.ac.nz

Abstract

Geothermal power generation is not keeping pace with other renewable energy technologies. This is due to a number of factors, including the industry's high capital cost, of which wells account for a significant portion. Hence, it is imperative to maximize value from wells drilled by selecting them optimally. An important technology used when making well placement decisions is computer simulation of production. This is usually done manually, with experts creating reservoir models, simulating wells at candidate feedzones and comparing the predicted production scenarios. Manual selection in this manner is slow and labor intensive.

Various heuristics have been investigated to try and automate this process, mainly based on gradient descent and stochastic search methods. However, no strict form optimization that guarantees the best solution has been attempted for the complex problem of selecting multiple production wells to maximize value. This paper uses Mixed Integer Programming (MIP) to address this problem. An economic model was created to calculate Net Present Values (NPVs) for a set of candidate wells and the interactions between them using AUTOUGH2 simulation results of an example geothermal system. Binary decision variables were used in the optimization to select the combination of wells that would maximize total NPV.

1 Introduction

1.1 Motivation

The use of renewable forms of energy is growing globally, but geothermal is lagging behind other forms, with a 2015 average growth rate of 2.4%, compared to an average across all renewable sources of 12% (REN21 2016). One of the main reasons for this is that geothermal requires a much higher capital investment than the rest, a significant portion of which can be attributed to the cost of drilling wells. In Iceland, for example, the costs associated drilling and constructing wells comprise 34% of total capital expenditure (Gehringer and Loksha 2012). Also, Blankenship et al. estimate that drilling related expenses can exceed 50% of total plant costs (Blankenship et al. 2005).

Along with high upfront costs, geothermal ventures also involve high degrees of risk. Well drilling can be a hit-and-miss activity; a global study on the success of geothermal wells conducted by the IFC estimates a success rate of about 50% for the first well in a field (IFC 2013). The success rate improves as more wells are drilled in a field, but even over the first 30 wells the studys estimate for cumulative success rate is only about 70%. Well costs can be a make-or-break factor in a geothermal project, and improving success rates for wells will bring large gains in reducing capital sunk into unproductive wells.

The IFC report also found that while the success rates of exploration phase wells have been increasing notably over the years, those of development wells and operational wells have not. This suggests that while methods for collecting information have been progressing, those for decision making with that information have not. This project focuses on development phase and operational wells, aiming to improve well placement decisions using numerical simulation on AUTOUGH2 (Yeh, Croucher, and O'Sullivan 2012), with Mixed Integer Optimization (MIP) models solved in Gurobi (Gurobi optimizer reference manual 2017).

1.2 Background

The use of numerical simulation as a tool for resource estimation and to inform drilling and production decisions has become increasingly common. Reservoir models are created and calibrated based on observations and field data such as topological measurements, magnetotelluric (MT) surveys, and exploratory well data, in a process known as natural state modeling. A natural state simulation will be run for millions of years from some initial state, till it converges to a steady state representing the current reservoir and matching the available data. A calibrated natural state model is then used as the initial state for future simulation of production. Both natural state and future simulation modeling are done manually, in the sense that expert modelers calibrate the models and select the conditions and parameters to run them based on technical knowledge and experience.

Manual selection of well locations is very time and labor intensive, especially for large, high fidelity models that take hours or even up to weeks to run. This research attempts to create a framework for automating the future simulation process and arrive at optimal drilling recommendations, given a calibrated natural state model. This has the benefits of formalizing the definition of possible options and the selection of the best one, insofar as the numerical model is representative of the physical system. Such an approach would reduce the human effort involved, as well as dependence on human expertise and the effect of human error. The simulation software used was AUTOUGH2.

1.3 Previous Work

Over the past few years, there have been many attempts to formulate theoretical frameworks, or use mathematical techniques to inform well placement decisions. They have generally focused on using metaheuristics to find good solutions, and fall broadly into two categories: gradient-based methods, and stochastic search algorithms. Stochastic here refers to the mechanism for searching the solution space.

A common stochastic method used is Particle Swam Optimization (PSO). Ansari et al. (Ansari, Hughes, and White 2014) used PSO to select locations for 4 production and 4 re-injection geothermal wells out of a set of 11 existing but abandoned wells in the US Gulf Coast. Onwunalu and Durlofsky also used PSO, but with Well Pattern Optimization (WPO) on an oil field (Onwunalu and Durlofsky 2011), essentially selecting parameters that specify the well patterns that encode the potential solutions.

WPO has also been used with Genetic algorithms (GA); Ozdogan et al. used a hybrid genetic algorithm in a WPO, with a fixed well pattern to reduce the solution space (Ozdogan et al. 2005). GAs themselves have been quite commonly used for well placement optimization and not just with WPO, for example by Montes et al. (Montes and Bartolome 2001), who developed and tested a GA on two example reservoirs. Another stochastic method that has been used in this area is Simulated Annealing (SA). Beckner and Song used SA with a Travelling Salesperson formulation to optimize well placement and scheduling on an example petroleum field (Beckner and Song 1995).

Many gradient based methods have also been used for the well placement problem. Sarma and Chen use an adjoint based gradient method on a continuous approximation of some example oil reservoirs (Sarma and Chen 2008). There have also been combinations of these methods; Bangerth et al. used a Simultaneous Perturbation Stochastic Approximation, which is a stochastic version of a steepest descent algorithm, and compared it to a Finite Difference gradient method and a SA method (Bangerth et al. 2006).

Though these approaches all have their advantages and disadvantages, none of them guarantee optimality (with respect to the numerical model). They all aim to find good solutions with as few simulation runs as possible. Helgason et al. (Helgason, gst Valfells, and Jlusson 2017) ranked all blocks in an example reservoir by NPV to find an optimal location. This is essentially a grid search enumerating over the entire solution space and choosing the best one, but it is guaranteed to be optimal if only one well is being selected. However, no one has used a method that guarantees optimality for the complex problem of selecting multiple production wells. This paper attempts to do so while keeping the number of simulations low.

2 Problem Definition

2.1 Numerical Model

A relatively small simulation model was used, based on a geothermal system in Indonesia. It is 16km by 14km in area, and extends 4km below the surface. The reservoir is intersected by four high permeability faults and covered by a low permeability clay cap. It was discretized into 8195 blocks and 528 nodes, in 483 columns and 19 rock layers. Its natural state was calibrated with 3 deep up-flows (a type of boundary flux condition) and 47 defined rock types, to match synthetic downhole temperature data generated for some exploration wells. Future simulation runs considered a 25 year production timeframe, and took approximately a minute on to run to completion on a standard Windows desktop machine. The wells in these simulations used a deliverability model, based on a fixed productivity index (PI). AUTOUGH2 produces listing files to store the results of these simulations, from which production time histories were extracted and processed using PyTOUGH (Croucher 2015) modules in Python.

2.2 Conceptual Framework

MIPs require a linear objective function and constraints. The numerical simulations are highly non-linear, complex and cannot be used as black-box models within the MIP optimization, so a surrogate model was required for translating simulation output to parameters in the optimization. It had to be capable of representing the effect of a well extracting from a feedzone on all other possible feedzones, as every well can change the temperature and pressure distributions and flow pathways in the reservoir. With a small simulation model like the one used, simulating all possible solutions (feedzone combinations) can be done for a small enough solution space, but this is impractical for larger models with longer solve times. As such, the surrogate model had to be able to represent all possible combinations without simulating each one.

The surrogate model was created from the outputs of a set of simulation scenarios. Each simulation had wells placed at candidate feedzones, and time histories of well mass flows and enthalpies were recorded and multiplied to get heat flow predictions. The fluid harvested from the wells should actually depend on the type of power plant installed. Dry steam plants require steam to directly turn the generator turbines, flash steam plants depressurize hot liquid to convert it to steam before driving the turbines, and binary cycle plants can use liquid at lower temperatures to heat a secondary working fluid with a lower boiling point, and use its steam to drive the turbines.

There are also other issues, such as heat loss during extraction, and possible re-injection of used fluid back into the reservoir. These were all ignored to simplify the problem, and heat flow was used as the production quantity rather than steam flow or temperature regulated mass flow, assuming a direct conversion from heat to electrical energy with a fixed generation efficiency. Production start times and limits on extraction were also excluded, as the main aim of this model was to make well placement decisions, not production management decisions. A simple NPV calculation was used, multiplying the heat flows by the generator efficiency and an electricity price to get cash flows, which were then discounted annually and summed to give a single monetary value to each candidate solution.

The generator efficiency was set to 12%, the global average conversion efficiency for geothermal plants as of 2012, according to Moon and Zarrouk (Moon and Zarrouk 2012). The electricity price used was the marginal cost of new generation in 2012 as per the MBIE (MBIE 2013). It doesnt matter that a New Zealand electricity price was used even though the example field is based on one in Indonesia, as the objective of this work was to test the approach rather than find a specific solution. The discount rate was set arbitrarily at 10%. Plant and well costs were neglected at first, though a cost model was included later on.

Candidate feedzones were selected based on simple physical cutoffs for temperature, depth and permeability. These cutoffs were somewhat arbitrary, and were meant to demonstrate that simple, programmable criteria can be used to define a set of candidate feedzones with minimal manual inspection. They filtered out 41 feedzones. Since the reservoir model is 3D, these refer to blocks in the model, and not to geographical surface locations. The number of wells desired was also limited to four, so the solution space was every combination of four wells out of the 41 candidate locations. This is somewhat reflective of reality, where the number of wells drilled is limited by plant capacity.

3 Method

3.1 Additive Interaction Model

The surrogate model created was called the Additive Interaction model, because it considered the effect of extracting from each candidate feedzone on the potential resource available to all the others individually. This was done by running simulations with wells producing from all the candidtate feedzones, but only one with a normal PI (the main well feedzone) and the rest (observer well feedzones) with reduced PIs, so they would producing insignificant mass flows. 41 simulations were run in total, one with each of the candidates as the main and the rest as observers.

Despite the very small mass flows, the decays in production from the observer feedzones were indicative of the effect of extraction from the main feedzone on them, and were scaled back up and discounted to give NPV penalties representing how much the main feedzone production takes away from the observer feedzones' potential values. Since all the observers have wells with very small PIs which extract negligible amounts of resource, their effects on each other can be ignored, and thus the main feedzone well's effects can be isolated.

Operationally, this was done by dividing the observer wells' PIs by a scale factor for the simulations, and then multiplying the extracted mass flows back up by the same scale factor. The mass flows were then multiplied by the well enthalpies to get apparent heat flow curves, which were shifted by the baseline value (zeroed) to get heat losses, shown for an example well in Figure 1 below. These losses are how much potential heat flow observer feedzones lose due to the main well's production, and were then converted to cash flows and discounted to get the NPV penalties.

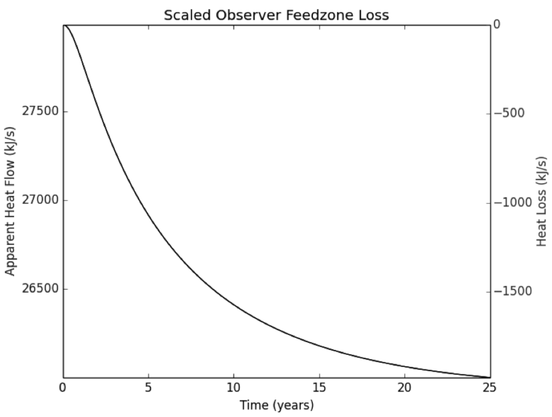

Figure 1: Example observer well apparent heat flow (blue axis) and heat flow loss (red axis)

The MIP formulation is given overleaf. The decision variables z form a 41x41 matrix, of which the diagonals select feedzones, and the off-diagonals select interactions between the selected feedzones. The objective function coefficients f also form a matrix of the same size, containing the calculated NPVs. The diagonals contain positive values (main NPVs), and the off-diagonals contain negative values (observer NPV penalties). The objective function maximizes the total NPV from all selected wells. Constraint C1 ensures that the effects of all selected wells on each other are included (if wells i and j are both on, then the NPV penalty of well j on well i must be included). Constraint C2 limits the number of wells selected to 4.

Maximize:

$$\Sigma_{i=1}^{41}\Sigma_{j=1}^{41} f_{ij} z_{ij}$$

Subject to:

$$z_{ij} \geq z_{ii} + z_{jj} - 1 \tag{C1}$$

$$\Sigma_{i=1}^{41} z_{ii} \leq 4 \tag{C2}$$

Where z and f are defined as:

$$z_{ii} = \begin{cases} 1, & \text{if well } i \text{ is on} \\ 0, & \text{otherwise} \end{cases}$$

$$z_{ij,i\neq j} = \begin{cases} 1, & \text{if well } i \text{ is influenced by well } j \\ 0, & \text{otherwise} \end{cases}$$

$$f_{ij} = \begin{cases} \text{NPV of well } i, & \text{if well } i = j \\ \text{NPV penalty of well } j \text{ on well } i, & \text{otherwise} \end{cases}$$

3.2 Results

The results from optimizing with this surrogate model are given below in Table 1. The Additive Interaction model's predicted NPVs for the optimal feedzones are compared against those calculated from directly simulating wells producing at the four feedzones together in AUTOUGH2. The direct simulation NPVs were calculated the same way - extracting production curves, converting them to cash flows and discounting.

Feedzone	Surrogate Model	Simulation	% Difference
36	\$ 30.1 mn	\$ 31.3 mn	3.7
37	\$ 30.3 mn	\$ 31.2 mn	2.9
40	\$ 35.0 mn	\$ 36.0 mn	2.7
41	\$ 34.4 mn	\$ 35.4 mn	2.8
Total	\$ 130 mn	\$ 134 mn	3.0

Table 1: NPV comparison of Additive Interaction model to direct simulation for feedzones in the optimal solution

The deviation of the surrogate model NPVs from those calculated from directly simulation is very low, showing that this method can be used with a high degree of accuracy, at least for this numerical model. Though the NPV estimation of the surrogate model was shown to be very accurate for the set of feedzones deemed optimal, there is no guarantee that this set is also optimal with respect to the simulation model. If the Additive Interaction model isnt accurate over the whole solution space (all combinations), then it is possible that the true optimal solution might have been overlooked.

Checking this required simulating wells at every combination of feedzones in AUTOUGH2, calculating the resulting NPVs and comparing with the surrogate model predictions. Doing so for every four out of the 41 candidate feedzones would require 101,270 simulations in total. To save runtime, this was done for a reduced solution space instead; testing every four well combination out of a set of 20 candidate locations. This required 4845 simulation runs. These 20 candidate locations were defined in the same way as the original 41, but with higher cutoffs for the temperature and maximum permeability. The total NPVs (sum of feedzone NPVs) for all combinations were calculated both from the surrogate model and directly from simulation, then ranked and compared. The optimal solution for the surrogate model was also optimal for the simulation, and the NPV errors were consistently small, being less than 4% across all 4845 combinations.

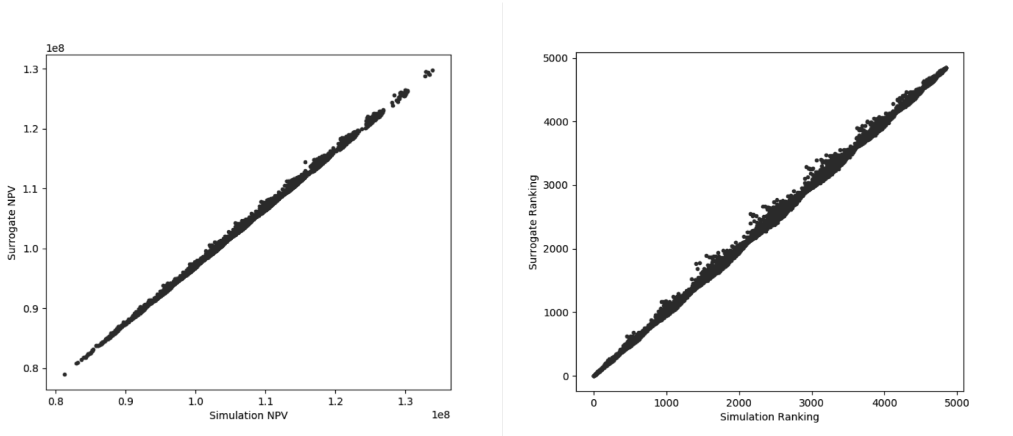

(a) NPV comparison

(b) Ranking comparison

Figure 2: Solution comparison of Additive Interaction model to direct simulation

Plotting the NPVs and rankings from the simulation against those from surrogate model for the whole set, as in Figures 2a and b above, shows an almost linear trend. Correlations between the surrogate model and direct simulation were calculated for both the rankings and NPVs, and both were above 99%. There are bunches of local clustering, with solutions grouping together in bands that can be clearly ordered. Within these groups however, similar solutions can be "swapped", in the sense that one is slightly better in the numerical model, but the surrogate model predicts the order the other way around.

Similar solutions within clusters generally only differ by one or two feedzones being a few blocks away, and generally have at least three of the four quite physically close to each other. As such, the surrogate model's optimal solution is not guaranteed to be optimal for the simulation, but it will be near-optimal, being the same vicinity and having a very similar output to the true optimal solution.

3.3 More Wells

The next step was to extend this to larger numbers of well selected. The same optimization was run multiple times with the NPVs from the surrogate model, but with the limit on the maximum number of selected wells in constraint C2 gradually increased from four up to 15. The MIP only ever chose 11 wells; even when the maximum well limit was 12 or above. For these cases, it didnt select as many wells as it could have. This was because at that point, the penalties from additional wells began to outweigh their own NPV contributions. Gurobi also took longer to solve the MIP when the well limit was increased. The MIP would solve in under a second for the four well case, but took several minutes for the ten well case. This is because there are far more possible combinations of ten feedzones than there are of four, so the solution space covered during the solve is considerably larger.

The optimal feedzones selected in the four well limit scenario remained in the optimal selection as the well limit was increased, with others being added to the selection. For all the limit scenarios, the model outputs were tested by comparing the NPVs to those calculated from running AUTOUGH2 simulations with wells producing from all the selected feedzones. The percentage error of NPV from the surrogate model compared to that from directly simulating the wells was plotted against the number of wells in Figure 3 below, for the four wells that remained optimal in all scenarios, and for the NPV sum over all wells.

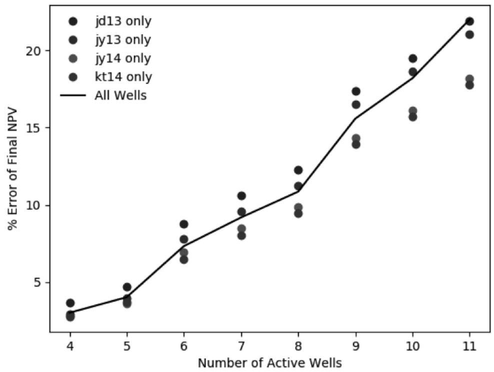

Figure 3: NPV errors for optimal wells vs. number of optimal wells

While the model is very accurate for small number of well, the discrepancy from the direct simulation values grows quite large as the number of wells increases, possibly nonlinearly. It reaches about 10% for eight wells, and about 20% for ten wells. This is likely relative to model size though; applying this procedure to a larger reservoir model will probably still give a small error for the ten well case and the error might not become significant till the number of selected wells reaches 20. The errors are also related to location; feedzones jy14 and kt14 have lower errors than the other two because they are deeper in the reservoir. The wells added as the well limit was increased were in shallower feedzones and therefore were further away and had less effect on these two feedzones than on the other two.

3.4 Additive Interaction Model with Well Costs

Having shown that the model is accurate for small numbers of wells and when used with the optimization gives solutions that are optimal or near optimal, another optimization was done with drilling costs included in the economic calculations. This was done to get some results reflective of a real world scenario. An existing cost model was taken (Lukawski et al. 2014) and used to define well drilling cost as a function of depth for each candidate feedzone. Their model is given below, with the cost being in USD.

$$\text{Cost} = 1.72 * 10^{-7} * \text{Depth}^2 + 2.3 * 10^{-3} * \text{Depth} - 0.62$$

These were then treated as upfront costs and were not discounted. Well maintenance and other infrastructure costs were not included either. The constraint in the previous MIP formulation that limited the number of wells (C2) was removed. Instead, the well costs were included in the objective function, and were used to limit the number of wells. The new formulation is given below.

Maximize:

$$\Sigma_{i=1}^{41}\Sigma_{j=1}^{41} f_{ij} z_{ij} - \Sigma_{i=1}^{41} c_i z_{ii}$$

Subject to:

$$z_{ij} \geq z_{ii} + z_{jj} - 1 \tag{C1}$$

Where:

$$c_i = \text{Drilling cost for well } i$$

The optimal solution selected included nine wells in total; the same nine wells as were selected when the well limit was set to nine without costs imposed. If the costs set were greater, fewer wells would have been chosen, and if they were lower, more wells would have been chosen. The costs didnt affect which wells are chosen because the best well blocks are all quite close to each other (were at similar depths). If they were more spread out, of if surface topology varied more drastically, the well costs might have been a bigger determinant. Another factor was that directional drilling and using multiple feed zones in a single wellbore were not considered. If they were included in the cost model, the clustering effect of the blocks selected would have been further accentuated.

3.5 Alternative Numerical Model

An issue with the results so far is that the optimal selections for all the scenarios tested tend to pick feedzones clustered together in close proximity. This is what would be expected from a greedy heuristic, and is likely due to the fact that the model calibration defined a very "leaky" reservoir, meaning the rock permeabilities are quite high. This is not necessarily a bad thing as the model did fit the data it was initially calibrated to in the natural state modelling, and this is where the non-uniqueness of these numerical simulations should be noted, as a different natural state calibration could also fit that data. An interesting question to ask, though out of the scope of this paper, is which future simulation prediction should be trusted, given multiple, equally valid natural state models?

That aside, another calibration of the same reservoir model (same geometry and discretization) was used for testing if the greedy approach is always right. It was calibrated to the same data, but using different deep upflow boundary conditions, and different, much lower rock permeabilities. The effect of the lower permeabilities is that it is harder for fluid to flow through the reservoir, so there is greater local pressure loss due to extraction from any feedzone. As such, choosing feedzones quite close to each other should be un-economical, as there is less resource available for each to draw from, and so they should be more spaced out.

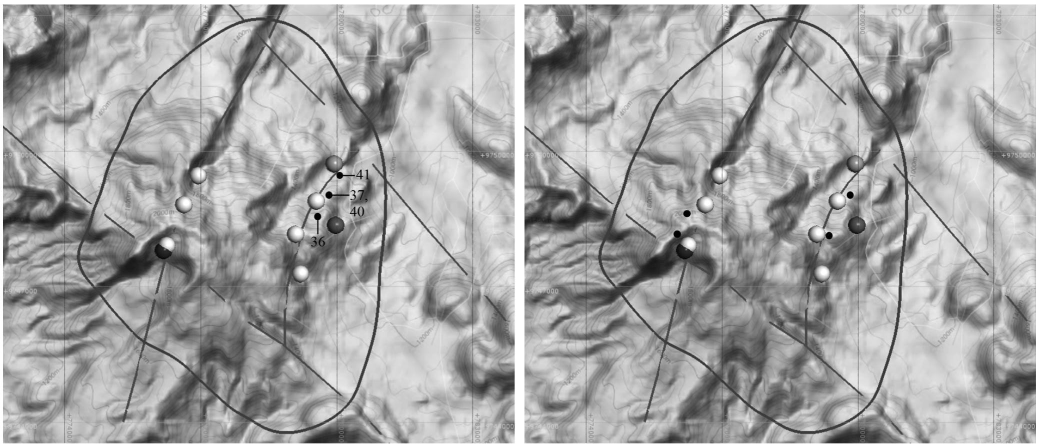

(a) Original Calibration (b) Alternative Calibration

Figure 4: Optimal Solution Comparison

This is exactly what was seen in the optimal selection from using the MIP without well costs on this calibration. The optimal feedzone surface locations for both calibrations are shown above in Figures 4a and b as the small black dots; those from the original calibration are also labelled with the feedzone IDs for reference. The other circles are various surface features. For the original calibration, the optimal feedzones are all clustered together on the most permeable of the faults (light grey lines), but for the alternative one they are spread out on two different faults on opposite sides of the reservoir.

4 Summary

4.1 Conclusion

This paper aimed to streamline and formalize the future simulation process, to optimally select multiple production wells, with as few simulation runs as possible. MIP formulations were used to achieve this using Gurobi as the solver. First, a few simple rules were made to define a set of candidate feedzones as production well targets in AUTOUGH2 simulations. NPVs calculated from simulation outputs were used as coefficients in the objection function for the optimization. A surrogate model was developed to allow the optimization to explore the entire solution space without having to run a large number of AUTOUGH2 simulations. It simulated a well at each feedzone individually (a total of 41 runs) and got NPVs for each well, and NPV penalties for the effect of each well on every other candidate feedzone.

It was very accurate when compared to NPVs calculated from well directly simulated together, for small numbers of well selected. When compared against direct simulation for all possible solutions in a reduced solution space, the solution NPVs and rankings for the model were very strongly correlated with those from the simulation. The model optimal solution is guaranteed to find solutions to be near optimal, if not optimal, for the simulation. As the limit on wells chosen increased from four to 11, the NPV error of the model went up from less than 4% to over 20%. The MIP formulation was extending to include a basic cost model, and finally was the method was tested on another version of the same reservoir with a different calibration, showing that non-greedy solutions can also be found.

4.2 Future Work

The Additive Interaction model and its accompanying MIP formulation will be extended to include scheduling optimization, as not all wells are planned and drilled from the start in a real project. This is being attempted with the current model by shifting the Additive Interaction model's output production curves in time based on estimates of how much potential resource a feedzone has lost due to the previous operation of other feedzones, which has the benefit of keeping the number of simulation runs low. These shifted curves will be converted to NPVs and used to inform the optimization, perhaps with a column generation approach to tackle the huge complexity of the problem.

Another approach that will be tried is to stagger the simulations and the decisions in a multi-stage process. For example, this could go as follows: Solve an initial MIP to select the first four wells drilled, then run a simulation with those wells producing for a phase A period. Process the simulation outputs from the end of the phase A period and use in a MIP to select the next four wells, and continue the simulation into a phase B period with the new wells producing, and so on. This is also more reflective of a real world scenario, where wells are typically drilled in stages. Once developed, the full well placement and scheduling optimization framework will be tested on more developed and realistic reservoir models. It is planned to use it with an AUTOUGH2 model of the Ohaaki geothermal system in the central north island, which has been continually developed over decades and for which there is an abundance of data and usage experience.

After thoroughly testing the fully deterministic model, the next step would then be to incorporate uncertainty in the simulations and carry it through to the optimizations. Uncertainty can be introduced into the system at a number of levels: in the well PIs, in the economic parameters (generator efficiency, electricity price, discount rate), in the model calibration as mentioned previously, or even in the geometry and discretization of the numerical model. Various optimization techniques will be explored for dealing the inherent uncertainty of the problem to give robust optimal recommendations.

Acknowledgments

I wish to thank the University of Auckland and the Geothermal Institute for enabling this research to go forward.

References

Ansari, Esmail, Richard Hughes, and Christopher D. White. 2014. “Well Placement Optimization for Maximum Energy Recovery from Hot Saline Aquifers.” *Thirty-Ninth Workshop on Geothermal Reservoir Engineering*. Stanford.

Bangerth, W., H. Klie, M. F. Wheeler, P. L. Stoffa, and M. K. Sen. 2006. “On optimization algorithms for the reservoir oil well placement problem.” *Computational Geosciences*, pp. 303–319.

Beckner, B. L., and X. Song. 1995. “Field Development Planning Using Simulated Annealing - Optimal Economic Well Scheduling and Placement.” *SPE Annual Technical Conference and Exhibition*. Dallas.

Blankenship, D. A., J. L. Wise, S. J. Bauer, A. J. Mansure, R. A. Norman, D. W. Raymond, and R. J. LaSala. 2005. “Research Efforts to Reduce the Cost of Well Development for Geothermal Power Generation.” *40th U.S. Symposium on Rock Mechanics.* Anchorage.

Croucher, Adrian. 2015. *PyTOUGH user’s guide.* Auckland: The University of Auckland.

Gehringer, Markus, and Victor Loksha. 2012. *Geothermal Handbook: Planning and financing power generation.* Washington, DC: Energy Sector Management Assistance Program, World Bank.

Gurobi Optimization, Inc. 2017. *Gurobi optimizer reference manual.* Gurobi Optimization, Inc.

Helgason, Dagur, gst Valfells, and Egill Jlusson. 2017. “Algorithm for Optimal Well Placement in Geothermal Systems Based on TOUGH2 models.” *42nd Workshop on Geothermal Reservoir Engineering.* Stanford.

IFC. 2013. “Success of Geothermal Wells: A Global Study.” Technical Report, International Finance Corporation, World Bank, Washington, DC.

Lukawski, Maciej Z., Brian J. Anderson, Chad Augustine, Louis E. Capuano, Koenraad F. Beckers, Bill Livesay, and Jefferson W. Tester. 2014. “Cost Analysis of Oil, Gas, and Geothermal Drilling.” *Journal of Petroleum Science and Engineering*, pp. 1–14.

MBIE. 2013. “New Zealand’s Energy Outlook.” Technical Report, Ministry of Business, Innovation and Employment.

Montes, Guillermo, and B. P. Bartolome. 2001. “The Use of Genetic Algorithms in Well Placement Optimization.” *SPE Latin American and Caribbean Petroleum Engineering Conference.* Buenos Aires.

Moon, Hyungsul, and Sadiq J. Zarrouk. 2012. “Efficiency of geothermal power-plants: a worldwide review.” *New Zealand Geothermal Workshop 2012.* Auckland.

Onwunalu, Jerome, and Louis Durlofsky. 2011. “A New Well-Pattern-Optimization Procedure for Large-Scale Field Development.” *SPE Journal*, pp. 594–607.

Ozdogan, Umut, Akshay Sahni, Burak Yeten, Baris Gulyaguler, and Wen H. Chen. 2005. “Efficient Assessment and Optimization of A Deepwater Asset Using Fixed Pattern Approach.” *SPE Annual Technical Conference and Exhibition.* Dallas.

REN21. 2016. “Renewables 2016 Global status report.” Technical Report, Renewable Energy Policy Network for the 21st Century, Paris.

Sarma, Pallav, and Wen H. Chen. 2008. “Efficient Well Placement Optimization with Gradient-Based Algorithms and Adjoint Models.” *SPE Intelligent Energy Conference and Exhibition.* Amsterdam.

Yeh, Angus, Adrian Croucher, and M. J. O’Sullivan. 2012. “Recent developments in the AUTOUGH2 simulator.” *TOUGH Symposium.* Berkeley.