

INTEGER PROGRAMMING OPTIMIZATION OF PRODUCTION WELL PLACEMENT

Rishi Adiga¹, John O'Sullivan¹ and Andy Philpott¹

¹The University of Auckland

radi130@aucklanduni.ac.nz

Keywords: *Well placement, optimization, resource estimation, resource management, TOUGH2, Python, PyTOUGH, Gurobi*

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 locations and comparing the predicted production scenarios.

Manual selection in this manner is slow and labor intensive, requiring weeks or months for expert modelers to make recommendations. Recently, various heuristics have been investigated to try and automate this process. Examples include Particle Swarm Optimization, Genetic Algorithms, Simulated Annealing, and Gradient descents. 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 Integer Programming to address this problem. An economic model was used 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. Unlike with heuristics, the solution is guaranteed to be optimal, at least with respect to the economic model. It was also found to be optimal with respect to the AUTOUGH2 model for the reservoir used.

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% [1]. 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 [2]. Also, Blankenship *et al.* estimate that drilling related expenses can exceed 50% of total plant costs [3].

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 [4]. The success rate improves as more wells are drilled in a field, but even over the first 30

wells the study's 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 [4]. This suggests that while methods for collecting information have been progressing, those for decision making with that information have not. This research focuses on development phase and operational wells, aiming to improve well placement decisions using numerical simulation and optimization.

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 from observations and field data such as topological measurements, MT surveys, and exploratory well data, in a process known as natural state modeling. A calibrated natural state model is then used as the initial state in 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 [5].

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 Swarm Optimization (PSO). Ansari *et al.* [6] 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 [7], 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 [8]. GAs themselves have been quite commonly used for well placement optimization and not just with WPO, for example by Montes *et al.* [9], 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 [10].

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 [11]. There have also been combinations of these methods; Bangert *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 [12].

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.* [13] 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 that while keeping the number of simulation runs low.

1.4 Overview of Integer Programming

A Mixed Integer Programming (MIP) model with binary decision variables is used for the optimization of well location, where binary decision variables model the selection of which wells to drill. Solution methodologies that deliver exact solutions for MIP models have made huge advances over the last twenty years, and now it is routine to solve such models with thousands of binary variables. The most popular solution approaches for MIP models use repeated application of well-known algorithms for Linear Programming (LP) models. In LPs, the objective function and constraints are linear, but variables can take fractional values. MIP solvers sequentially add extra constraints (cutting planes) to LPs that preclude fractional optimal solutions, and enumerate different ways of fixing decision variables to binary values (branch-and-bound) in a search procedure that eventually yields a provably optimal solution. We use the state-of-the-art solver Gurobi [14] for solving our MIP models.

Since the decision variables were whether or not to have a well at a given location, the coefficients of these variables in the objective function had to carry the information used to compare them. Net Present Value (NPV) was used to do this, and was determined from production time histories extracted from the simulation results, using a simple economic calculation. The optimization chooses wells "blindly", in the sense that it does not know the structural information of the

reservoir, and is based solely on the effect of that structure in producing the NPVs entered to it.

The simplest scenario of selecting only one well doesn't need a MIP, since the NPVs can just be compared and the highest selected, as was done by Helgason *et al.* The more complex scenario of selecting multiple wells in this fashion warrants integer programming, but also requires information about the interactions between wells. This study considered 41 candidate locations for a total of four wells. Simply simulating every possible combination of these wells in AUTOUGH2 is computationally intense approach, so a surrogate model was required for capturing all these interactions, from as few simulations as possible. The surrogate model and the MIP formulations created are discussed in Section 2.

2. METHOD

2.1 AUTOUGH2 Model

The simulation model used was a relatively small one, and was originally based on a geothermal system in Indonesia. It includes a recharge area that is 16km by 14km wide, and extends between 3.5km to 4km below the surface; the system is under the slopes of a volcano so the surface topology varies quite a bit in elevation. The reservoir is intersected by four faults that essentially bound it. Two main faults (higher permeability) run in a near northeast-southwest direction, and two lesser (lower permeability) ones run northwest-southeast. The reservoir is covered by a low permeability clay cap. The numerical model was discretized into 8195 blocks and 528 nodes, in 483 columns and 19 rock layers, plus an atmospheric layer.

Its natural state was calibrated with 3 deep up-flows and 47 different defined rock types, to match synthetic down-hole temperature data generated for exploration wells. Future simulation runs were set up to consider a 25 year production lifespan, and took approximately one minute on to run to completion using AUTOUGH2 on a standard Windows desktop machine. The wells in these simulations used a deliverability model, with a fixed productivity index (PI). AUTOUGH2 produces listing files to store the results of these simulations, from which production time histories for were extracted and processed using PyTOUGH [15] modules in Python.

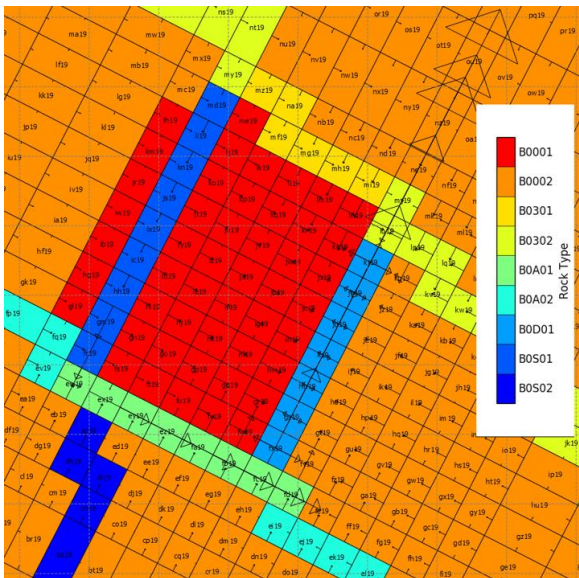


Figure 1: Horizontal slice through the numerical model, showing the fault structure.

2.2 Surrogate Model

As discussed above, the IP approach requires a linear objective function and constraints. The full AUTOUGH2 simulations are 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 using a well on other possible wells, 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 (well 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 well combinations without simulating each one.

The surrogate model was created from simulation results of a subset of possible solution scenarios. Each simulation had wells placed at candidate locations, 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. Start times and limits on extraction were also excluded, as the main aim of these models was to make well placement decisions, not operational 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 on an annual basis 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 [16]. The electricity price used was the marginal cost of new generation in 2012 as per the MBIE [17]. It doesn't 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. The blocks for placing wells in the simulation runs 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 well blocks with minimal manual inspection. They filtered out 41 blocks in the model to serve as candidate well target locations. Since the reservoir model is 3D, the "candidate locations" 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.

2.3 Additive Interaction Formulation

This surrogate model considered the effect of extracting from one well on the potential resource available for all other candidate wells. This was done by running simulations with all the wells on and producing, but only one with a normal PI (the main well) and the rest (observer wells) with reduced PIs, so they would produce insignificant mass flows. 41 simulations were run in total, one with each of the candidate wells as the main well and the rest as observers.

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

Operationally, this was done by dividing the observer wells' PIs by a scale factor for the AUTOUGH2 simulations, and then multiplying the extracted mass flows back up by the same scale factor. The mass flows were then multiplied by the 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 2 below. These losses are how much potential heat flow observer wells lose due to the main well's production, and were discounted to get the NPV penalties.




Figure 2: 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 41×41 matrix, of which the diagonals select wells, and the off-diagonals select interactions between the selected wells. The objective function coefficients f also form a matrix of the same size, containing the calculated NPVs. The diagonals contain positive values (main well NPVs), and the off-diagonals contain negative values (observer well 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:

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

Subject to:

$$z_{ij} \geq z_{ii} + z_{jj} - 1, \quad (C1)$$

$$\sum_{i=1}^{41} z_{ii} \leq 4, \quad (C2)$$

Where z and f are defined as:

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

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

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

The results from optimizing with this surrogate model are given below in Table 1. The model's NPVs for the optimal well blocks are compared against those calculated for each the blocks from directly simulating the scenario with all of them producing using AUTOUGH2.

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

Well	Simulation	Surrogate	% Difference
jd13	3.13E+07	3.01E+07	3.7
jy13	3.12E+07	3.03E+07	2.9
jy14	3.60E+07	3.50E+07	2.7
kt14	3.54E+07	3.44E+07	2.8
Total	1.34E+08	1.30E+08	3.0

The surrogate model slightly under-predicts NPV. This is because in reality, when multiple wells are producing, their negative effects on each other are not strictly additive, as this model assumes. Main wells would not only negatively impact observer wells, but also reduce the negative impacts of other main wells on the observer wells, as illustrated conceptually in Figure 3 below. This model accounts for the effects of all the wells on each other (the solid arrows), but not the effects of each well on the other wells' interactions (dotted arrows).




Figure 3: Conceptual schematic of interactions.

Despite this, the deviation of the optimal well NPVs from those calculated by directly simulating the 4 wells together is very low, showing that this model can be used with a high degree of accuracy. Though the NPV estimation of the surrogate model was shown to be very accurate for the set of wells deemed optimal, there is no guarantee that this set of wells is also optimal with respect to the simulation model. If the Additive Interaction model isn't accurate over the whole solution space (all well combinations), then it is possible that the surrogate model could have overlooked a solution that is optimal for the simulation model.

Checking this required simulating every combination of wells in AUTOUGH2, calculating the resulting NPVs and comparing with the surrogate model predictions. Doing so for every 4 wells out of the 41 candidate locations 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 previous 41 had been, but with higher cutoffs for the temperature and permeability. Since this candidate set was a subset of the original one, the surrogate model's optimal solution was the same. The total NPVs for all combinations were calculated both from the surrogate model and directly from simulation, then ranked and compared. The top 20 combinations from direct simulation are given below in Table 2.

Table 2: Total NPVs and rankings for the Additive Interaction model and direct simulation for the top 20 well combinations.

Simulation		Surrogate Model		NPV % Diff
Rank	Value	Rank	Value	
1	1.34E+08	1	1.30E+08	3.0
2	1.33E+08	5	1.29E+08	3.4
3	1.33E+08	4	1.29E+08	3.2
4	1.33E+08	3	1.29E+08	2.9
5	1.33E+08	2	1.30E+08	2.5
6	1.33E+08	6	1.29E+08	3.0
7	1.30E+08	10	1.26E+08	3.1
8	1.30E+08	11	1.26E+08	3.1
9	1.30E+08	8	1.26E+08	3.0

10	1.30E+08	9	1.26E+08	3.0
11	1.30E+08	13	1.26E+08	3.2
12	1.30E+08	17	1.26E+08	3.3
13	1.30E+08	7	1.26E+08	2.7
14	1.30E+08	18	1.26E+08	3.3
15	1.30E+08	16	1.26E+08	3.2
16	1.30E+08	14	1.26E+08	3.1
17	1.30E+08	19	1.26E+08	3.3
18	1.30E+08	20	1.25E+08	3.3
19	1.30E+08	15	1.26E+08	3.1
20	1.30E+08	12	1.26E+08	2.8

It can be seen from this table that the well combination found to be optimal with the surrogate models inputs to the MIP are indeed the optimal combination when calculated using the AUTOUGH2 simulation. The NPV errors are also consistently small, with all of them being less than 4% across all 4845 combinations. Despite being a very good surrogate model, it doesn't give the exact same rankings for the solutions as the direct simulation. For example, the second best solution as per the surrogate model is actually the fifth best from the simulation, and vice versa.

Plotting the NPVs and rankings from the simulation against those from surrogate model for the whole set, as in Figures 4 and 5, 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%. However, there are bunches of local scattering. This is a clustering effect, with solutions grouping together in bands that can be clearly ordered. Within these groups however, similar solutions get "swapped", in the sense that one is slightly better in the numerical model, but the surrogate model predicts the order the other way around.

These bands are visible in both plots, but are exaggerated in the rankings plot, due to the smaller plotting scale. They are also shown in Table 2 as well, separated by thick red boundaries between them. The top 5 solutions are common to both the surrogate model and direct simulation, despite not being in quite the same order. The next 8 solutions and the 7 after them (in simulation rank) form two more bands respectively, with almost all the solutions in them being common to both the simulation and surrogate model.




Figure 4: Solution NPV comparison of Additive Interaction model to direct simulation.




Figure 5: Solution ranking comparison of Additive Interaction model to direct simulation

Similar solutions within clusters might only differ by one or two wells being a few blocks away, and generally have at least 3 of the 4 wells quite physically close to each other. For example, the solutions with simulation rank 2 and 5 (surrogate rank 5 and 2 respectively) have two well blocks in common, circled in black in Figure 6 below. The other two wells are also very close and are only shifted by a few blocks, circled in red and blue. For this reason, it doesn't matter if the surrogate model's optimal solution isn't exactly the same as that for the simulation, because it will be in the same vicinity and have a very similar output.

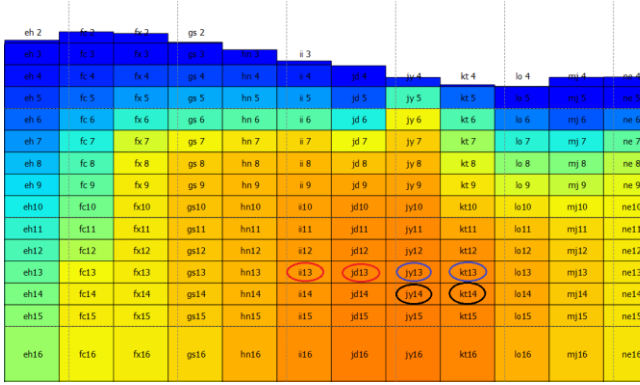


Figure 6: Vertical slice through the eastern main fault showing temperature output from the AUTOUGH2 model with well locations for solutions 2 and 5.

So far, optimizing with the Additive Interaction model has been shown to give optimal or near optimal well locations with a high degree of accuracy for the 4 well case. 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 4 up to 15.

The optimization only ever chose 11 wells; even when the maximum well limit was 12 or above. For these cases, it didn't 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, so the optimization would choose not to add them. It is also worth mentioning here that Gurobi took longer to solve the MIP when the well limit was increased. The optimization would run in under a second for the 4 well case, but took several minutes for the 10 well case. This is because there are far more possible combinations of 10 wells than there are of 4 wells, so the solution space covered during the solve was much larger.

The optimal wells selected in the 4 well limit scenario remained in the optimal selection as the well limit was increased, with other wells being added to the selection. For all the limit scenarios, the model outputs were tested by comparing the NPVs to those calculated by running AUTOUGH2 simulations with all the selected wells producing. The percentage error of NPV from the surrogate model compared to that from directly simulating the wells was plotted against the number of wells, for the 4 wells that remained optimal in all scenarios, and for the NPV sum over all wells:

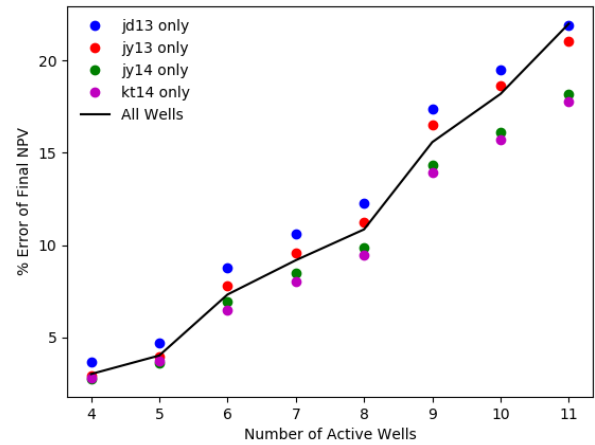


Figure 7: NPV errors for optimal wells vs. number of optimal wells.

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 8 wells, and about 20% for 10 wells. The errors are related to well location. Wells *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 regions and therefore were further away and had less effect on these two wells than on the other two.

2.4 Additive Interaction Formulation 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 final results reflective of a possible real world scenario. The cost model made by Lukawski *et al.* [18] was used to define drilling cost as a function of depth for each candidate well. Their model is given below, with C being the cost in USD.

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

Their paper is from 2014, and since the electricity cost estimate used was from 2012 and was for New Zealand, the calculated well drilling costs were converted to NZD [19] and deflated back to 2012 values [20] for the sake of consistency. These were then treated as upfront costs and were not discounted. Well maintenance and steam gathering 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 used to limit the number of wells. The new formulation is given below:

Maximize:

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

Subject to:

$$z_{ij} \geq z_{ii} + z_{jj} - 1, \quad (C1)$$

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 didn't 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, or 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. SUMMARY

3.1 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. Integer Programming formulations were used to achieve this using Gurobi as the solver. First, a few simple rules were made to define a set of candidate well blocks as potential production wells 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 full reservoir simulations. It simulated a well at each location 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 location. 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. However, as the limit on wells chosen increased from 4 to 11, the NPV error of the model went up from less than 4% to over 20%

3.2 Future Work

The Additive Interaction model can be improved so it is more accurate for larger numbers of wells selected, by trying to include the effects of wells on other wells' effects on each other (see Figure 3). Once this surrogate model is found to be sufficiently accurate for the example reservoir used, it will be tested on a more developed and realistic reservoir model. The next step would then be to incorporate uncertainty in the simulations and carry it through to the optimization. Finally, the framework can also be expanded to include well scheduling decisions as well.

ACKNOWLEDGEMENTS

The University of Auckland and the Geothermal Institute are thanked for enabling this research to go forward.

REFERENCES

- [1] "Renewables 2016 Global status report," Renewable Energy Policy Network for the 21st Century, Paris, 2016.
- [2] M. Gehringer and V. Loksha, *Geothermal Handbook: Planning and financing power generation*, Washington, DC: Energy Sector Management Assistance Program, World Bank, 2012.
- [3] D. Blankenship, J. Wise, S. Bauer, A. Mansure, R. Normann, D. Raymond and R. LaSala, "Research Efforts to Reduce the Cost of Well Development for Geothermal Power Generation," in *40th U.S. Symposium on Rock Mechanics*, Anchorage, 2005.
- [4] "Success of Geothermal Wells: A Global Study," International Finance Corporation, World Bank, Washington, DC, 2013.
- [5] A. Yeh, A. Croucher and M. O'Sullivan, "Recent developments in the AUTOUGH2 simulator," in *TOUGH Symposium*, Berkeley, 2012.
- [6] E. Ansari, R. Hughes and C. D. White, "Well Placement Optimization for Maximum Energy Recovery from Hot Saline Aquifers," in *Thirty-Ninth Workshop on Geothermal Reservoir Engineering*, Stanford, 2014.
- [7] J. E. Onwunali and L. Durlafsky, "A New Well-Pattern-Optimization Procedure for Large-Scale Field Development," *SPE Journal*, vol. 16, no. 03, pp. 594-607, 2011.
- [8] U. Ozdogan, A. Sahni, B. Yeten, B. Guyaguler and W. H. Chen, "Efficient Assessment and Optimization of A Deepwater Asset Using Fixed Pattern Approach," in *SPE Annual Technical Conference and Exhibition*, Dallas, 2005.
- [9] G. Montes and B. P. Bartolome, "The Use of Genetic Algorithms in Well Placement Optimization," in *SPE Latin American and Caribbean Petroleum Engineering Conference*, Buenos Aires, 2001.
- [10] B. L. Beckner and X. Song, "Field Development Planning Using Simulated Annealing - Optimal Economic Well Scheduling and Placement," in *Field Development Planning Using Simulated Annealing - Optimal Economic Well*, Dallas, 1995.
- [11] "Efficient Well Placement Optimization with Gradient-Based Algorithms and Adjoint Models," in *SPE Intelligent Energy Conference and Exhibition*, Amsterdam, 2008.
- [12] W. Bangerth, H. Klie, M. F. Wheeler, P. L. Stoffa and M. K. Sen, "On optimization algorithms for the reservoir oil well placement problem," *Computational Geosciences*, vol. 10, no. 3, pp. 303-319, 2006.
- [13] D. Helgason, Á. Valfells and E. Júlíusson, "Algorithm for Optimal Well Placement in Geothermal Systems Based on TOUGH2 models," in *42nd Workshop on Geothermal Reservoir Engineering*, Stanford, 2017.

- [14] "Gurobi optimizer reference manual," Gurobi Optimization, Inc., 2017.
- [15] A. Croucher, "PyTOUGH user's guide," University of Auckland, Auckland, 2015.
- [16] H. Moon and S. J. Zarrouk, "Efficiency of geothermal powerplants: a worldwide review," in *New Zealand Geothermal Workshop 2012*, Auckland, 2012.
- [17] "New Zealand's Energy Outlook," Ministry of Business, Innovation & Employment, 2013.
- [18] M. Z. Lukawski, B. J. Anderson, C. Augustine, L. E. Capuano, K. F. Beckers, B. Livesay and J. W. Tester, "Cost Analysis of Oil, Gas, and Geothermal Drilling," *Journal of Petroleum Science and Engineering*, vol. 118, pp. 1-14, 2014.
- [19] "Yearly Average Rates," NZFOREX, [Online]. Available: <http://www.nzforex.co.nz/forex-tools/historical-rate-tools/yearly-average-rates>. [Accessed August 2017].
- [20] "Inflation calculator," Reserve Bank of New Zealand, [Online]. Available: <https://www.rbnz.govt.nz/monetary-policy/inflation-calculator>. [Accessed August 2017].