An update on the Waiwera geothermal flow simulator: development and applications

Adrian Croucher[1], Michael J. O'Sullivan[1], John O'Sullivan[1], Angus Yeh[1], John Burnell[2] and Warwick Kissling[2]

[1]Department of Engineering Science, University of Auckland, Private Bag 92019, Auckland 1142, New Zealand

[2]GNS Science, Private Bag 30-368, Lower Hutt 5040, New Zealand

a.croucher@auckland.ac.nz

Keywords: *Reservoir models, numerical modelling, flow simulator*

ABSTRACT

Waiwera is a new open-source, parallelised geothermal flow simulator, developed initially as part of the Geothermal Supermodels research programme. Recent progress on Waiwera development has included implementation of a parallelised MINC method for simulating fractured media, a zone system for easier specification of model parameters over designated parts of the model mesh and user-customizable simulation output including output from sources. Waiwera has adopted a new, modern build system and also uses a completely new unit-testing system, enabling easy installation and testing on a range of platforms.

As well as describing these new features we also demonstrate Waiwera's performance on test problems and geothermal reservoir models.

1. INTRODUCTION

Waiwera is a new open-source, parallelised geothermal flow simulator (Croucher et al., 2015, 2016, 2017). Its initial development formed a major component of the Geothermal Supermodels programme, a four-year New Zealand-based research programme aimed at developing integrated geothermal modelling tools.

Waiwera builds on our experience with developing and using the TOUGH2 (Pruess et al., 1999) and AUTOUGH2 (Yeh et al., 2012) simulators, but starting from a completely new structured and object-oriented code-base, and leveraging the advanced capabilities of the PETSc computational library (Balay et al., 2018). As one of the aims of the Geothermal Supermodels project was to develop the capability to run very large models, Waiwera is parallelised via the Message Passing Interface (MPI), so that it can run not only on multi-processor personal computers but also on large distributed-memory compute clusters.

Since the completion of the Geothermal Supermodels project, development work on Waiwera has continued. Here we give an update on recent development progress and present results from tests and applications.

2. IMPLEMENTATION PROGRESS

2.1 Modelling fractured media via MINC

Waiwera now includes capability for modelling flow in fractured media using the Multiple INteracting Continua (MINC) method (Pruess and Narasimhan, 1985), in which flow between fractures and the rock matrix is simulated by using additional finite volume cells, effectively nested inside the original model mesh cells. If there is only one additional MINC cell inside each original cell, the method corresponds to a 'dual porosity' approach.

TOUGH2 also allows the MINC method to be used to represent fractured media, but requires the user to add the extra MINC cells to the mesh explicitly. By contrast, in Waiwera the MINC cells are added internally at run time. The user need only specify which zones in the mesh are to have MINC applied, and the MINC parameters (fracture spacing, number of MINC cells per original cell etc.) for each zone.

MINC processing of the mesh is carried out in parallel so that the computational time required is kept small, even for large models. However, when MINC is applied only to some parts of the mesh and not others, this can result in reduced parallel scalability, because of load balancing issues (i.e. some processes ending up with more cells than others). Current development work is focused on applying a re-balancing technique to partial MINC meshes to improve scalability.

Typically, MINC model production runs are started from a natural-state solution computed using a compatible single-porosity model (as MINC effectively only adds storage terms which have no effect on the natural state). As a convenience, Waiwera offers the ability to start a MINC simulation directly from the output of a single-porosity run on the same mesh. Initial conditions in both fracture and rock matrix cells are taken from their corresponding single-porosity cells.

Waiwera can also optionally automatically compute the appropriate matrix rock porosities, if they are not specified explicitly, in order to make the total void fractions in the MINC model consistent with the porosities in the original single-porosity model (given the fracture volume fractions and fracture porosities).

2.2 Mesh zones

Some simulation parameters (e.g. rock types, boundary conditions) are applied only to particular parts of the model mesh, so a means is needed for specifying these subsets of the mesh. One way is to provide an explicit list of the appropriate mesh cells, which is the only way it can be done in TOUGH2. This method has the advantage that arbitrarily complex shapes can be defined, but it also has some disadvantages.

One drawback is that the specification of even a very simple mesh subset (e.g. all the cells in a particular elevation range) may require a long list of cells, each of which has to be identified and added to the model input (though this can be automated via scripting). This becomes particularly cumbersome for very large models, substantially increasing model input file sizes and decreasing readability. Another drawback is that the specification of a mesh subset is mesh-dependent, so that if for example the mesh is refined, or even just re-ordered, the specification of the subset also has to be updated.

Waiwera now has the ability to define "mesh zones" in the simulation input. A zone can be defined as above, via a list of cells, but there are also two other zone types currently available. The first is a "box" zone, defined by ranges in the cell centre co-ordinates. Ranges for any of the co-ordinates can be omitted, to make the zone extend over the whole mesh in that direction. There is also a "combination" zone type, which combines other zones to form a new zone. Zones can be combined by addition (union), subtraction (difference) and multiplication (intersection), to form more complex shapes.

Mesh zones can then be used for specifying rock types, MINC zones etc. They can be re-used for different purposes to avoid repetition in the input.

2.3 Deliverability thresholds

Many geothermal production models contain wells with pre-specified (measured) production rates. However, if the local model permeability is too low, it may be insufficient to maintain the specified flow rate without the pressure approaching zero, which can stall the simulation.

This can be a problem particularly during inverse modelling, in which the model permeabilities are modified automatically by the inverse modelling algorithm in an attempt to match field data and calibrate the model. In this case, stalled simulations can also hold up the (usually lengthy) inverse modelling process.

One approach to circumventing this problem is to put the well on deliverability if the pressure drops below a specified threshold. In this case the measured flow rate is treated effectively as a target rather than an absolutely prescribed value. The appropriate productivity index for the deliverability control is calculated automatically so that the flow rate remains consistent when the deliverability control is switched on. If the pressure later rises back above the threshold value, the well is allowed to revert to using the specified flow rates.

2.4 Model output

Waiwera produces two kinds of output, a high-level log file in YAML format and the main simulation results file ("heavy data") in HDF5 format. This file contains a dataset recording the fluid properties in each model mesh cell over time, and now also an additional dataset recording source properties over time (analogous to the TOUGH2 "GENER" table). This is needed for sources whose properties (e.g. flow rates and enthalpies) are not pre-specified but are computed by source controls (e.g. deliverability) as the simulation progresses.

Customisation of the Waiwera output datasets is now also possible, so that only the fields needed for a particular problem need be output. This can give considerably reduced output file sizes, particularly for large models. (The only restriction on this customisation is that the main fluid properties dataset must always contain the primary thermodynamic variables, to ensure that it is always possible to restart another simulation from it.)

2.5 Build system, testing and deployment

Until 2018 we used a simple build system for compiling Waiwera based around the commonly-used 'make' tool, supplemented by a custom-written script for identifying the source code dependency graph. However a system of this kind does not offer any help with handling other libraries that the code depends on, i.e. making sure these libraries are available on the system, building them if they are not, and linking them in.

As a result, we have now switched to using Meson (www.mesonbuild.com), a recently-developed build system which handles not only source file dependencies and checking for and building third-party libraries, but has other useful features including very fast execution and a built-in framework for unit testing (i.e. testing of individual subroutines in the code).

We also developed a new library called Zofu for creating Fortran 2003 unit tests, designed to integrate with Meson and replace the older FRUIT system we previously used for unit testing. We have released Zofu as a separate open-source project (https://github.com/acroucher/zofu).

We are also developing a deployment system for Waiwera using the Ansible software (www.ansible.com), which will handle checking for (and installing if necessary) tools needed such as compilers, MPI etc, simplifying installation on Linux personal computers and compute clusters. We also plan to offer "containerised" builds of Waiwera using Docker (www.docker.com), which we envisage being the easiest way to run Waiwera on MS Windows and Macintosh machines.

2.5 Licensing and release

We plan to release Waiwera as free open-source software (under the GNU LGPL license) in late 2019, available for download at http://waiwera.github.io.

3. TESTS AND APPLICATIONS

3.1 MINC column production test

As well as unit tests, Waiwera includes a suite of small benchmark test problems for verifying correct running of the code as a whole. We searched for standard benchmark tests for MINC models, to verify Waiwera's new MINC capability (see section 2.1), but did not find anything particularly suitable.

Hence we created a MINC benchmark test, setting up a simple MINC production model on a 10-cell 1-D vertical column mesh. Apart from the added MINC cells, this is the same model described by Croucher et al. (2016) to demonstrate Waiwera's improved phase transition behaviour, and the initial conditions for the MINC production run are taken from the steady-state (single porosity) solution computed therein.

The model mesh extends from elevation 0 m down to -1000 m. Hot water at 240°C is injected into the base of the model at 10^{-3} kg/m^2/s and atmospheric boundary conditions are applied at the top (1 bar, 20°C). The single-porosity rock has permeability 100 mD and porosity 0.1. Linear relative permeability curves and a pure water equation of state are used. The steady-state solution contains a steam zone extending over the upper 300 m of the mesh.

For the production run, fluid is extracted at 2.5×10^{-3} kg/m^2/s from a single well at elevation -350 m. MINC processing is applied between elevations -100 m down to -600 m. Two levels of matrix rock cells are added with volume fractions of 0.3 and 0.6, so the fracture volume fraction is 0.1.Three sets of fracture planes are used, with fracture spacing 5 m. The fracture rock has permeability 100 mD and porosity 0.5,

while the matrix rock has permeability 0.001 mD and porosity 0.056. The production simulation is run for 90 days, with the time step set initially to 0.5 days and not allowed to exceed 3 days.

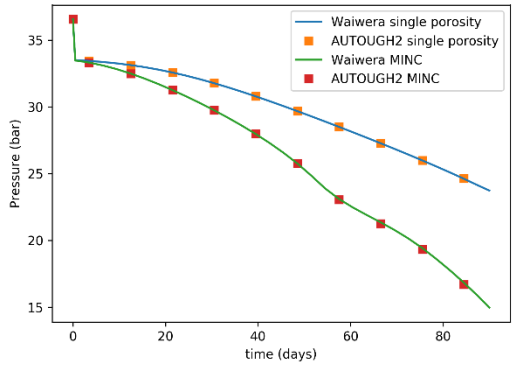

Figure 1: Pressure history at production well for MINC benchmark test

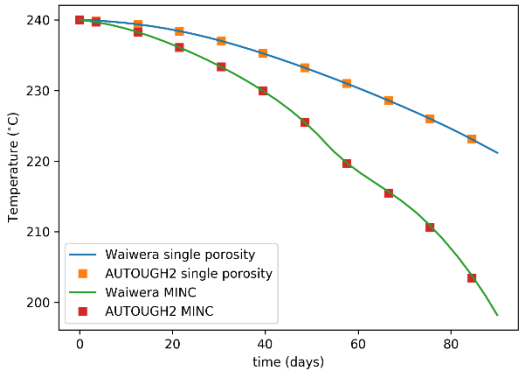

Figure 2: Temperature history at production well for MINC benchmark test

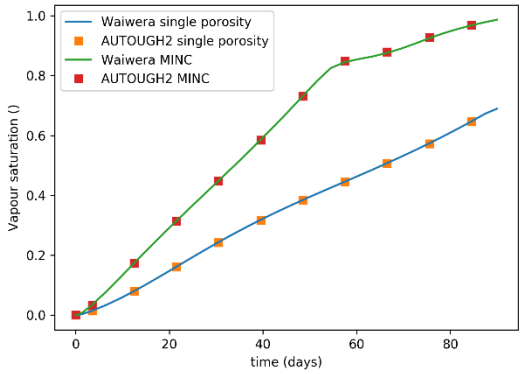

Figure 3: Vapour saturation history at production well for MINC benchmark test

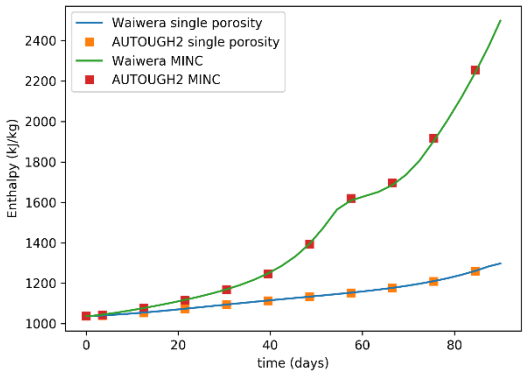

Figure 4: Enthalpy history at production well for MINC benchmark test

Figures 1 – 4 show the modelled time histories of pressure, temperature, vapour saturation and enthalpy at the production well. Results from both Waiwera and AUTOUGH2 are shown for comparison. It can been that the two simulators give results in close agreement.

Corresponding results from a single-porosity production run are also given, to show the impact of the MINC zone on the solution. The inclusion of the MINC zone significantly increases pressure drawdown, decreasing temperatures but increasing vapour saturations at the well. The sudden change in gradients of pressure, temperature, vapour saturation and enthalpy at 54.5 days is caused by the onset of boiling in the cell directly below the production well (which does not occur in the single-porosity version).

3.2 Parallel strong scaling test

In a parallel "strong scaling" test, a model is run on different numbers of parallel processes, measuring the wall clock time taken in each case. The ideal behaviour is linear scaling, in which doubling the number of processes results in a halving of computation time. Various factors may cause scaling behaviour to be less than ideal, including the details of the software's parallel implementation, parallel computing hardware limitations (e.g. memory bandwidth and interconnect speeds between processes) and the details of the problem itself.

For good strong scaling behaviour, the problem must be large enough to benefit from parallelisation. Smaller problems will usually not scale well, because the data communication overhead between parallel processes is too great compared with the computational work carried out on each process. The PETSc documentation recommends problem sizes of at least 10,000 – 20,000 unknowns per process for satisfactory scaling.

We created a strong scaling test for Waiwera based loosely on problem 6 from the 1980 Geothermal Model Intercomparison Study (Molloy, 1981). This model simulates production from a simplified 3-D geothermal reservoir. The original mesh is very coarse, containing only 125 cells. We created a refined version by decreasing the horizontal cell size from 1000 × 800 m to 62.5 × 62.5 m, increasing the depth of the model from 1.8 km to 3 km, and increasing the number of layers from 5 to 160. This gives a mesh with 819,200 cells.

As in the original version, a pure water equation of state is used (with two unknowns per cell), so the total number of unknowns is approximately 1.6 million. To keep the number of unknowns per parallel process above 10,000 – 20,000, the number of processes should therefore be no more than 40 – 80, after which we cannot expect good strong scaling behaviour.

The original model used prescribed initial conditions specific to its 5-layer mesh, which were not a true equilibrium solution and certainly would not be in equilibrium if transferred to a refined mesh. As the study was focused on scaling behaviour rather than matching the original results, there was in any case no need for consistency with the original model parameters. Hence, we decided to follow usual practice and generate initial conditions by running the model without production to a steady-state solution.

We also altered the rock structure and boundary conditions of the model to give a steady-state solution with a single upflow near one corner. Atmospheric boundary conditions are retained over the top surface, and at the base of the model inflow boundary conditions are applied, with water at enthalpy 1250 kJ/kg injected at 0.04 kg/km^2/s in the outfield, 2.4 kg/km^2/s under the upflow, and 0.64 kg/km^2/s in between.

The original model contained only one production well, but with a highly refined mesh it is necessary to spread the production over multiple wells, otherwise pressures will approach zero as production proceeds. We used a total of 72 production wells evenly distributed over a production area 750 m × 750 m in the horizontal and 100 m thick. The total production rate was 24 kg/s at the beginning of the simulation, increasing to 60 kg/s after 2 years, 96 kg/s after 4 years and 144 kg/s after 6 years. 10 years of production were simulated with a constant model time step size of 0.05 years (total 200 time steps).

As well as the single-porosity version we also created a MINC version of the problem, with three additional rock matrix levels added over the entire mesh. This quadruples the problem size to a total of approximately 6.5 million unknowns. To reduce total computational time used by the study, the MINC version of the model was run for only 50 time steps (which is sufficient for investigating scaling behaviour).

The scaling study was carried out on the New Zealand eScience Infrastructure (NeSI) "Māui" supercomputer, a Cray XC50 machine with Skylake Xeon nodes and Cray Aries interconnect in Dragonfly topology. Waiwera was built on Māui using GCC compilers and the model was run using 32, 64, 128 and 256 parallel processes. With a model of this size, reducing the number of processes all the way down to 1 (serial) is not practical because the run times become very long. Hence we followed the practice of e.g. Jung et al. (2017) and evaluated the speedup relative to the smallest number of processes used (in this case 32).

Figure 5 shows the scaling results, plotting the speedup (i.e. run time divided into the 32-process run time) for both the single-porosity and MINC versions of the model. The ideal linear scaling is also shown for comparison. The single-porosity Waiwera results show near-ideal scaling up to around 100 processes for this problem, which meets and exceeds the 40 – 80 processes suggested by the PETSc documentation. For the MINC results we can expect good scaling up to at least around 160 processes, and this is also achieved, with the MINC model scaling slightly better than the single-porosity version and still showing close to ideal scaling at 256 processes.

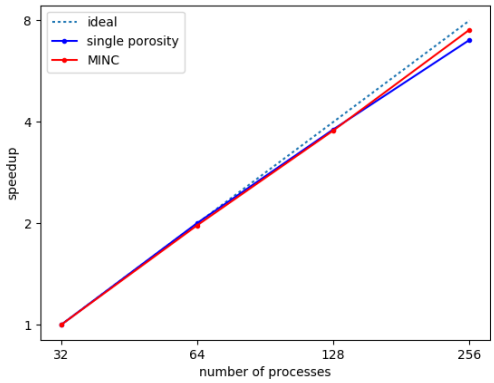

Figure 5: Parallel strong scaling results

3.3 Full-scale production model

Waiwera has been tested on several full-scale production models to date. The most complex of which is particularly challenging because forecasts are required not only for the deep geothermal reservoir but also for geothermal surface conditions. To meet the requirements of the near-surface forecasts the modelling project is now using higher-resolution meshes with cells sizes down to 50m × 50m × 12m. The highest resolution cells are located in the central part of the geothermal system which also coincides with the area of intense surface activity. The resulting mesh has over 850,000 cells and includes a detailed representation of geological and structural model (Figure 6).

Despite the complexity of this modelling project, the general approach is the same as other geothermal modelling projects. A natural state model has been built and calibrated which has then been used to create a production history model which has also been calibrated. Finally forecasts are made using a suite of future scenario models. Both the production history and future scenario simulations are carried out using MINC models which allow important aspects of fracture flow to be represented. Two levels of MINC matrix cells are added over the central part of the models, covering the geothermal system. The outer recharge areas use the single porosity approach. The resulting models have over 2 million cells.

All of the simulations are run using Waiwera's air-water equation of state as it is vitally important for forecasting both the deep geothermal production and the shallow near-surface behaviour that the changes in shallow water levels and the size and distribution of the shallow boiling zone are accurately represented. As this equation of state has three unknowns per cell, the models have a total of approximately 6 million unknowns.

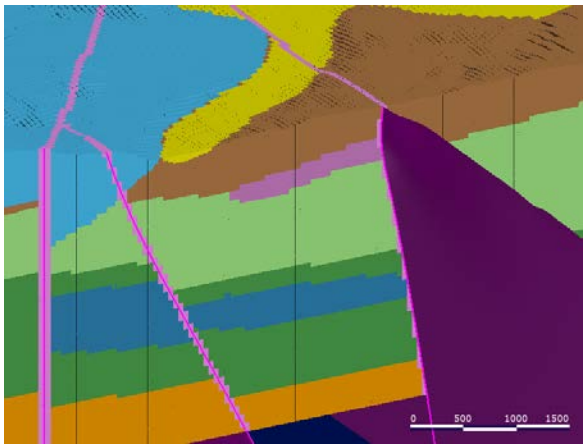

Figure 6: High resolution Waiwera model of a full-scale production model. Structures included shown as magenta surfaces and in the model indicated by pink blocks.

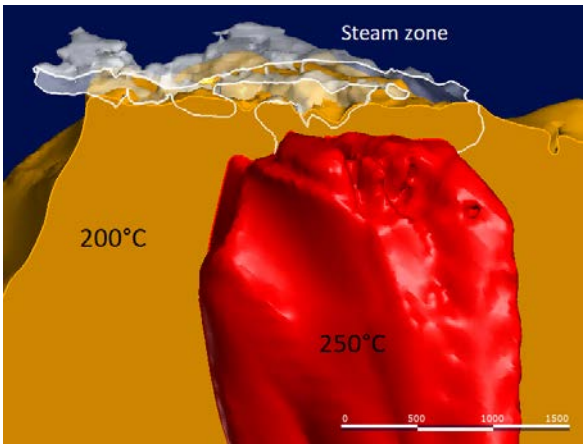

Figure 7: Forecast temperature and steam zone distribution in full-scale production model.

Running models of this size is not currently practical using serial simulators such as AUTOUGH2. Even estimating the speed-up achieved by Waiwera over AUTOUGH2 is difficult, because AUTOUGH2 would take several months to run the natural state model and its convergence behaviour on this model is poor. By contrast with Waiwera the model can be run in under a day using 160 parallel processes on NeSI's "Māui" supercomputer. Production history and future scenario simulations are even more time-consuming due to the extra cells used in the MINC approach but they can be run in 2-3 days using Waiwera on Māui.

As well as running the very large high-resolution models of the system, Waiwera has been used to run hundreds of standard resolution models simultaneously to provide uncertainty quantification of the model forecasts. These models have approximately 70,000 cells, or over 200,000 when MINC is applied to the production history and future scenario models. Using Waiwera on Māui the complete sequence of natural state, production history and future scenario can be carried out in approximately six hours on 40 parallel processes. This means that hundreds of parameter sets can be evaluated over a few weeks. By contrast the complete sequence takes more than two weeks to run for a single parameter set using AUTOUGH2.

In summary, this complex geothermal modelling project has only been possible thanks to the capabilities of the Waiwera simulator. At the same time the project has provided an excellent, full-scale production dataset for rigorously testing Waiwera. Finally, using NeSI's Māui supercomputer has shown that powerful parallel processing capability combined with a scalable geothermal simulator are very effective at modelling the complex multiphysics found in geothermal systems and providing meaningful forecasts for supporting strategic decision-making.

4. CONCLUSIONS

Significant new capability has recently been added to the Waiwera code, including a parallel implementation of MINC processing for modelling flow in fractured media. Benchmark test results confirm that this gives results in agreement with TOUGH2.

Results from a parallel strong scaling study have demonstrated Waiwera's ability to scale to large problems, with near-ideal scaling behaviour up to and exceeding the expected limits.

Waiwera has also been applied to a large, complex geothermal reservoir model, showing its ability to solve realistic problems on a scale that is not practical using AUTOUGH2, but requires parallelisation.

ACKNOWLEDGEMENTS

Waiwera was initially developed as part of the "Geothermal Supermodels" project (grant C05X1306) funded by the New Zealand Ministry of Business, Innovation and Employment (MBIE).

Contact Energy Ltd have also provided funding throughout the project and have contributed support in terms of access to data and collaboration.

The authors wish to acknowledge the use of New Zealand eScience Infrastructure (NeSI) high performance computing facilities as part of this research. New Zealand's national facilities are provided by NeSI and funded jointly by NeSI's collaborator institutions and through MBIE's Research Infrastructure programme.

REFERENCES

Balay, S., Abhyankar, S., Adams, M., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Dener, A., Eijkhout, V., Gropp, W., Karpeyev, D., Kaushik, D., Knepley, M., May, D., Curfman McInnes, L., Mills, R., Munson, T., Rupp, K., Sanan, P., Smith, B., Zampini, S., Zhang, H.: *PETSc Users Manual*, Revision 3.10. Technical Report ANL-95/11 Rev 3.10, Argonne National Laboratory, Illinois (2018).

Croucher, A.E., O'Sullivan, M.J., O'Sullivan, J.P., Pogacnik, J., Yeh, A., Burnell, J. and Kissling, W.: Geothermal Supermodels Project: an update on flow simulator development. *Proc. 37th NZ Geothermal Workshop*, 18 – 20 November 2015, Taupo, New Zealand (2015).

Croucher, A.E., O'Sullivan, M.J., O'Sullivan, J.P., Pogacnik, J., Yeh, A., Burnell, J. and Kissling, W.: Geothermal Supermodels Project: an update on flow simulator development. *Proc. 38th NZ Geothermal Workshop*, 23 – 25 November 2016, Auckland, New Zealand (2016).

Croucher, A.E., O'Sullivan, M.J., O'Sullivan, J.P., Pogacnik, J., Yeh, A., Burnell, J. and Kissling, W.: Geothermal Supermodels Project: an update on flow simulator development. *Proc. 38th NZ Geothermal Workshop*, 22 – 24 November 2017, Rotorua, New Zealand (2017).

Jung, Y., Pau, G.S.H., Finsterle, S. and Pollyea, R.M. TOUGH3: a new efficient version of the TOUGH suite of multiphase flow and transport simulators. *Comp. Geosci*. 108, 2 – 7 (2017).

Molloy, M.W.: Geothermal reservoir engineering code comparison project. *Proc. Special Panel on Geothermal Model Intercomparison Study*, December 16-18, 1980, Stanford University, Stanford, California (1981).

O'Sullivan, J.P, Croucher, A.E., Yeh, A. and O'Sullivan, M.J.: Experiments with Waiwera, a new geothermal simulator. *Proc. 39th NZ Geothermal Workshop*, 22 – 24 November 2017, Rotorua, New Zealand (2017).

Pruess, K. and Narasimhan, T.N.: A practical method for modeling fluid and heat flow in fractured porous media. *Soc. Pet. Eng. J.* 25, 14 – 26 (1985).

Pruess, K., Oldenburg, C. and Moridis, G.: TOUGH2 user's guide, version 2.0. LBNL-43134, Lawrence Berkeley National Laboratory, Berkeley, California (1999).

Yeh, A., Croucher, A. and O'Sullivan, M.J.: Recent developments in the AUTOUGH2 simulator. *Proc. TOUGH Symposium 2012*, Berkeley, California, September 17-19 (2012).