11th World Congress on Computational Mechanics (WCCM XI)
5th European Conference on Computational Mechanics (ECCM V)
6th European Conference on Computational Fluid Dynamics (ECFD VI)
E. Oñate, J. Oliver and A. Huerta (Eds)

FURTHER IMPROVEMENTS IN THE CONVERGENCE OF TOUGH2 SIMULATIONS

JOHN O'SULLIVAN*, ADRIAN CROUCHER*, ANGUS YEH* AND MIKE O'SULLIVAN*

* Department of Engineering Science, University of Auckland, Private Bag 92019, Auckland 1142, New Zealand, jp.osullivan@auckland.ac.nz, www.des.auckland.ac.nz

Key Words: *TOUGH2, convergence, inverse modelling, geothermal reservoir modelling.*

Abstract. Numerical modelling has become an important tool in managing geothermal systems and planning their exploitation for renewable energy production. The TOUGH2 simulator has been the industry standard tool for developing numerical models for many years. It includes several different equation-of-state modules and thus can be used for modelling many different kinds of geothermal fields as well as other complex sub-surface flow problems such as carbon sequestration and nuclear waste storage. Under certain conditions TOUGH2 simulations stall at a relatively small time step size and are unable run up to the very large times required for a natural state simulation. This behaviour leads to slow model development and also poses a significant obstacle to inverse modelling using iTOUGH2 or PEST as forward simulations are more computationally expensive and may not finish. In previous work the authors identified and analysed conditions leading to stalled simulations and proposed corrections to the air-water and CO^2-water equation-of-state modules to eliminate the behaviour.

While these improvements prevent simulations from stalling at small time steps they do not prevent a limiting time step size from occurring as simulations approach a large simulation time. This behaviour is observed often, affects all equation-of-state modules and can add significant computational time to simulations. Improvements to the convergence criteria used by the TOUGH2 algorithm during its internal Newton-Raphson iteration are proposed following an analysis of the algorithm's behaviour as the time step size becomes very large. The results presented show that this improvement can dramatically reduce the total simulation time for all equation-of-state modules. In particular the improvement can reduce the computational time required for calculating numerical derivatives during inverse modelling by an order of magnitude. This has a profound effect upon our ability to carry out inverse modelling and uncertainty analysis of numerical models of large-scale, real geothermal systems.

1 INTRODUCTION

Numerical simulation is an important tool for planning and managing the development of geothermal systems [1]. Since its development in the 1980s TOUGH2 [2] has become the industry standard simulator and is now widely used. At the same time TOUGH2 has been extended to include several different equation-of-state modules allowing it to be used for

modelling many different kinds of geothermal fields, as well as other complex sub-surface flow problems.

Many TOUGH2 simulations must be carried out in order to obtain a well-calibrated model of a geothermal system [3]. During the calibration process, model parameters such as permeabilities and heat and mass inputs are adjusted in order to match measured observations. When TOUGH2 simulations converge slowly or do not converge the calibration process becomes difficult, time-consuming and ultimately impossible. This is especially true when inverse modelling tools such as iTOUGH [4,5] or PEST [6] are used because they depend upon simulations converging without manual intervention.

Most convergence issues encountered during TOUGH2 simulations occur due to two different types of behaviour. The first type of behaviour occurs when a specific set of circumstances cause the TOUGH2 simulator to stall at a relatively small time step size. This makes it difficult or impossible for the simulation to reach the large simulation time required for a natural state model. This behaviour is well known to geothermal reservoir modellers and has also been reported by [7] in CO_2 sequestration simulations using TOUGH2. This issue was solved in our previous work [8] and is summarised here for completeness.

The second type of behaviour does not prevent the simulation from reaching the natural state simulation time but does limit the maximum time step size that can be achieved during the simulation. This significantly increases the real time required to carry out natural state simulations and therefore has a large impact on model development. In particular, inverse modelling and uncertainty quantification become very computationally expensive when this behaviour occurs as many thousands of forward simulations may be required.

The objective of the present work was to improve the convergence properties of TOUGH2 simulations to allow efficient model development and enable the use of inverse modelling and uncertainty quantification tools. This required analysing the TOUGH2 solution algorithm, identifying the causes of the two types of behaviour and then developing changes within the AUTOUGH2 code, the University of Auckland version of TOUGH2 [9].

To test the changes proposed a number of simulations of real geothermal systems that displayed poor convergence behaviour were collected over several months. The significant improvements in their convergence behaviour as a result of the changes are discussed in the following sections. A brief discussion regarding the impact of the improvements on the computational cost of inverse modelling and sensitivity analysis is also presented.

2 THE TOUGH2 SIMULATOR

The TOUGH2 simulator solves mass and energy-balance equations to determine the properties of non-isothermal flows of multiphase, multicomponent fluids in porous and fractured media [2].

2.1 Governing equations

The mass-balance and energy-balance equations solved by TOUGH2 can be written in the following form [10]:

$$\frac{d}{dt}\int_V A_\kappa dV = -\int_A \mathbf{n}.\mathbf{F}_\kappa dA + \int_V q_\kappa dV \tag{1}$$

where V is the volume of integration, A_κ is the amount of each quantity κ within the volume, A is the surface of the volume, $\mathbf{F}_\kappa$ is the flux of quantity κ across the surface A, $\mathbf{n}$ is the normal vector to the surface A and q_κ represents any sources or sinks in the volume. The number of quantities for which balance equations must be solved is determined by the equation of state (EOS) that has been selected to simulate the geothermal system. The EOS is selected by considering both the important properties of the system being simulated and the nature of the issues that being investigated. In this work four common types of EOS have been used. Their details are given in Table 1:

Table 1: Types of Equations of State

Equation of State	Quantities Simulated
EOS1	mass of water, energy
EOS2	mass of water, mass of CO2, energy
EOS3	mass of water, mass of air, energy
EWASG	mass of water, mass of non-condensable gas, mass of chloride, energy

The amount of each component per unit volume is calculated as the sum of the contributions from each phase as shown in Equation (2):

$$A_\kappa = \varphi(\rho_l S_l X_{\kappa l} + \rho_g S_g X_{\kappa g}) \tag{2}$$

Here φ is the porosity and for each phase, β, the density is given by ρ_β, the saturation by S_β and the mass fraction by $X_{\kappa\beta}$. The liquid phase is indicated by the subscript l and the gas phase by the subscript g. For the amount of energy in the volume the definition includes an additional term for the contribution of the rock:

$$A_e = (1-\varphi)\rho_r c_r T + \varphi(\rho_l u_l S_l + \rho_g u_g S_g) \tag{3}$$

T is the temperature, ρ_r the density of the rock, c_r its heat capacity and u_β the internal energy of phase β. The flux of each component $\mathbf{F}_\kappa$ in Equation (1) is calculated using the contribution of each phase $\mathbf{F}_\beta$ weighted by the mass fraction:

$$\mathbf{F}_\kappa = X_{\kappa l}\mathbf{F}_l + X_{\kappa g}\mathbf{F}_g \tag{4}$$

In some equations of state for TOUGH2 a dispersion term can be added to (4) but in most geothermal systems the effects of diffusion and hydrodynamic dispersion are small. For the energy flux a conductive term is also included where K is the thermal conductivity and the enthalpy of each phase β must be taken into account as shown in Equation (5).

$$\mathbf{F}_e = h_l\mathbf{F}_l + h_g\mathbf{F}_g - K\nabla T \tag{5}$$

The flux of each phase is given by the two-phase form of Darcy's Law:

$$\mathbf{F}_\beta = -\frac{\mathbf{k}\,k_{r\beta}}{\nu_\beta}(\nabla p + \rho_\beta \mathbf{g}) \tag{6}$$

Here $\mathbf{k}$ is the permeability tensor (usually assumed to be diagonal), $k_{r\beta}$ the relative permeability of the phase β, ν_β its viscosity, p the pressure and $\mathbf{g}$ is gravity. Note that for this work the effect of capillary pressure was not considered.

Discretising in space and applying implicit time stepping reduces Equation (1) to a set of coupled non-linear equations which can be written as:

$$\frac{V_i}{\Delta t^{n+1}}(A_{\kappa i}^{n+1} - A_{\kappa i}^{n}) = -\sum_j a_{ij} F_{\kappa ij}^{n+1} + \sum_p q_{\kappa ip}^{n+1} \tag{7}$$

In Equation (7) the superscript refers to the time step at which the quantity is calculated. The term a_{ij} is the area of the interface between block i and block j and $F_{\kappa ij}^{n+1}$ is the flux of each quantity κ across the same interface. The term $q_{\kappa ip}^{n+1}$ represents p separate source terms of quantity κ in block i and finally the time step size is given by Δt^{n+1}.

The discrete form of Equation (6) is used to calculate the fluxes of each quantity:

$$F_{\beta ij}^{n+1} = -\left(\frac{k\, k_{r\beta}}{\nu_\beta}\right)_{ij}^{n+1} \left[\frac{p_j^{n+1} - p_i^{n+1}}{d_{ij}} - \rho_{\beta ij}^{n+1} g_{ij}\right] \tag{8}$$

To calculate the relative permeability and viscosity term in Equation (8) upwind differencing is used:

$$\left(\frac{k_{r\beta}}{\nu_\beta}\right)_{ij}^{n+1} = \begin{cases} \left(\frac{k_{r\beta}}{\nu_\beta}\right)_i^{n+1} & \text{phase } \beta \text{ flows } i \text{ to } j \\ \left(\frac{k_{r\beta}}{\nu_\beta}\right)_j^{n+1} & \text{phase } \beta \text{ flows } j \text{ to } i \end{cases} \tag{9}$$

The direction in which each phase flows is determined by the pressure gradient and the effects of gravity as defined by the following equation:

$$G_{\beta ij}^{n+1} = \frac{p_j^{n+1} - p_i^{n+1}}{d_{ij}} - \rho_{\beta ij}^{n+1} g_{ij} \tag{10}$$

Now $G_{\beta ij}^{n+1} > 0$ means that phase β flows from block j to block i and $G_{\beta ij}^{n+1} < 0$ means the opposite.

In both Equation (8) and Equation (10) g_{ij} is the component of gravity acting normally to the interface between block i and block j and the density of the phase at the interface $\rho_{\beta ij}^{n+1}$ is calculated simply from the average for the two blocks:

$$\rho_{\beta ij}^{n+1} = \frac{\rho_{\beta i}^{n+1} + \rho_{\beta j}^{n+1}}{2} \tag{11}$$

In Equation (8) the permeability and block distance terms are combined and calculated as a harmonic weight:

$$\frac{d_{ij}}{k_{ij}} = \frac{d_i}{k_i} + \frac{d_j}{k_j} \tag{12}$$

In Equation (10) the block distance term is simply the sum of the distance from each block centre to their shared interface:

$$d_{ij} = d_i + d_j \tag{13}$$

The values of the secondary parameters h_β, u_β, ρ_β and v_β are calculated using the primary variables and steam table data supplied through the EOS modules.

2.2 Solution process

Gathering all the terms in Equation (7) to the left hand side gives a coupled system of nonlinear equations, which must be solved at each time step to calculate the new values of the primary variables in each block:

$$\mathbf{r}(\mathbf{x}) = \mathbf{0} \tag{14}$$

This system of equations is solved iteratively using the Newton-Raphson (NR) method which can be written in the following form [2]:

$$-\mathbf{J}_k(\mathbf{x}_{k+1} - \mathbf{x}_k) = \mathbf{r}(\mathbf{x}_k) \tag{15}$$

where k represents the NR iteration number and **J** is the Jacobian matrix of the system of equations defined as:

$$\mathbf{J} = \left[\frac{\partial \mathbf{r}}{\partial \mathbf{x}}\right] \tag{16}$$

Note that in TOUGH2 the Jacobian is calculated numerically by evaluating the change in **r** corresponding to a small change in **x**.

The NR method is complete and the values of the primary variables at the new time step have been calculated once $\mathbf{r}(\mathbf{x}_k)$ has been reduced sufficiently to meet one of two convergence criteria [2]. The first is that the absolute value of the ratio of $\mathbf{r}(\mathbf{x}_k)$ to $\mathbf{A}_\kappa$ has been reduced to less than a specified value ε_1:

$$\left|\frac{\mathbf{r}(\mathbf{x}_k)}{\mathbf{A}_\kappa}\right| \leq \varepsilon_1 \tag{17}$$

The second criterion is used when $\mathbf{A}_\kappa$ is smaller than the convergence parameter ε_2. In this case the following relationship is used:

$$|\mathbf{r}(\mathbf{x}_k)| \leq \varepsilon_1 \varepsilon_2 \tag{18}$$

TOUGH2 uses adaptive time stepping which means that if the NR method successfully converges within a specified number of iterations then the time step size is increased. If it successfully converges, but only after more than the specified number of iterations, the time step remains the same. When the NR does not converge within a specified iteration number limit then the time step is reduced and another attempt is made at solving for the primary variables.

The simulation is complete for a natural state model once the total simulated time reaches the required target, usually set at ~1.0×10^{16} seconds. This large target time does not reflect the real time that the system took to develop over geological time but rather is a computational technique for ensuring that the system is at equilibrium.

3 TWO TYPES OF CONVERGENCE ISSUES IN TOUGH2

In general the solution process described in the previous section is robust and relatively efficient. The highly nonlinear nature of the equations governing flow in geothermal systems mean that the adaptive time stepping often reduces and increases the time step size many times during the solution process. This behaviour is expected and doesn't prevent successful simulations. However, in certain circumstances two types of behaviour occur that prevent

TOUGH2 simulations from converging efficiently. The first is when simulations stall at a relatively small time step size and the simulation is unable to progress. The second is when simulations approach the natural state simulation time but reach a limiting time step size causing additional computational cost.

The plots in Figure (1) show two examples of each type of behaviour from four simulations of different geothermal systems. In plot (a) results are shown from the simulations of the Wyang Windu system in Indonesia and the Lihir Island system in Papau New Guinea. The Wyang Windu simulation uses EOS2 while the Lihir Island simulation uses EOS3. It is clear that both simulations have stalled at relatively small time steps that cycle around values less than $1x10^7$ seconds. Plot (b) shows examples of the convergence issues experienced as simulations approach the natural state simulation time. These results are taken from an EOS2 simulation of the Ohaaki system and an EWASG simulation of the Rotorua system, both of which are in New Zealand. The Ohaaki simulation was started from a good set of initial conditions and so quickly achieves a large time step size before reaching a limit around $2x10^{13}$ seconds. The Rotorua simulation takes many more time steps for the solution to evolve but also eventually reaches a limiting step size of around $5x10^{13}$ seconds. In the following sections the circumstances leading to each of these behaviours is analysed and described before changes to the TOUGH2 algorithm are given that correct the behaviour.

3.1 Stalled simulations due to a moving unsaturated zone

When the solution process stalls with the time step unable to increase above a relatively small value the simulation is prevented from reaching the specified total time and hence it does not complete. This behaviour is often manifested by a cycle of time step reductions followed by a single time step increase as can be seen in Figure (1a).

This situation occurs when the NR method is able to converge within the maximum number of iterations at one time step size, but cannot converge at a larger time step size. TOUGH2 can be set to output the maximum residual, $\mathbf{r}(\mathbf{x}_k)$, at each iteration throughout the simulation which reveals that the problem is usually caused by the residual at the larger time step not falling below the convergence criterion for some particular block.

a) b)

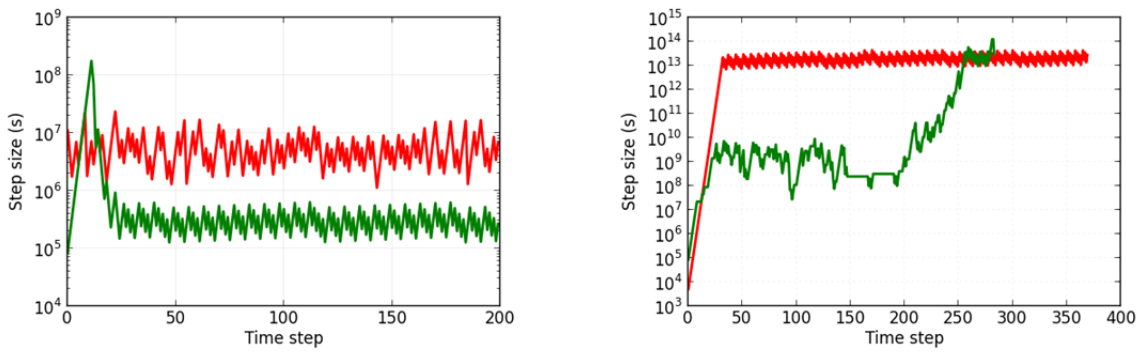

Figure 1: Two types of convergence problems experienced in TOUGH2 simulations. Plot (a) shows the time step size evolution for two stalled simulations; the Wayang Windu model using EOS2 (-) and the Lihir Island model using EOS3 (-). Plot (b) shows the time step size evolution for two simulations with slow convergence at large time steps; the Ohaaki model using EOS2 (-) and the Rotorua model using EWASG (-).

Often the block in question is close to the interface between a saturated and an unsaturated zone. The interface is usually moving and hence two-phase conditions are either evolving or disappearing within blocks in the local area. By examining the conditions in the area, it is possible to determine how the interface is moving and estimate the new primary variables. In the past modellers have attempted to implement this fix by intervening in the simulation either manually or by using scripts [11] to adjust the primary variables in the block or in its neighbours.

In some cases this allows the simulation to progress and complete, but in others it does not, in which case it may be necessary to take the more drastic step of adjusting local permeabilities. This is unsatisfactory not only because there may be no physical basis for the permeability change, but also because it is not guaranteed to solve the problem. A much more satisfactory approach is to analyse and understand the problem and develop a solution within the AUTOUGH2 code itself [8].

In Figure (2) plots show the gas saturation and the gas density calculated in the blocks that caused the Wyang Windu and the Lihir Island simulations to stall. In the case of the Lihir Island system the problem occurs in a block which is two-phase initially while the block below it is single phase liquid. Plot (a) shows that the gas phase in the block disappears slowly but as soon as it has disappeared completely it evolves again immediately creating the cycling behaviour and time step stall shown in Figure (1a). Physically this corresponds to the water table rising and falling within the block. Plot (b) reveals that the problem causing the behaviour is the discontinuity in the value of the gas density ρ_g calculated using Equation (11). This in turn causes discontinuities in the calculations of $\mathbf{r}(\mathbf{x})$ and $\mathbf{J}$ during the NR solve which causes the time step to stall.

For the Wyang Windu simulation the problem occurs in a single phase liquid block that is above a two-phase block. Plot (c) in Figure (1) shows that the gas evolves in the liquid block but the quickly disappears before evolving once again. Physically this corresponds to a boiling zone that is moving up and down on the edge of the problem block. The gas density in the block is shown in plot (d) which reveals that although it does not drop to zero when only liquid is present it is also discontinuous.

As described previously [8] the discontinuities occur in the calculations of $\mathbf{r}(\mathbf{x})$ and $\mathbf{J}$ because of the upwind differencing used by TOUGH2 in Equation (9). A situation can exist (and occurs surprisingly often) where the flow direction for each phase, calculated by Equation (10), can be from a block where the phase is present to one where it is not. From Equation (9) the upwind differencing means that a non-zero relative permeability will be calculated at the interface between the blocks and subsequently a non-zero flux will be calculated for the phase across that interface using Equation (8). Because Equation (8) also contains a contribution from gravity and the phase density, if the value of the density of the phase is discontinuous then the calculated flux is also discontinuous.

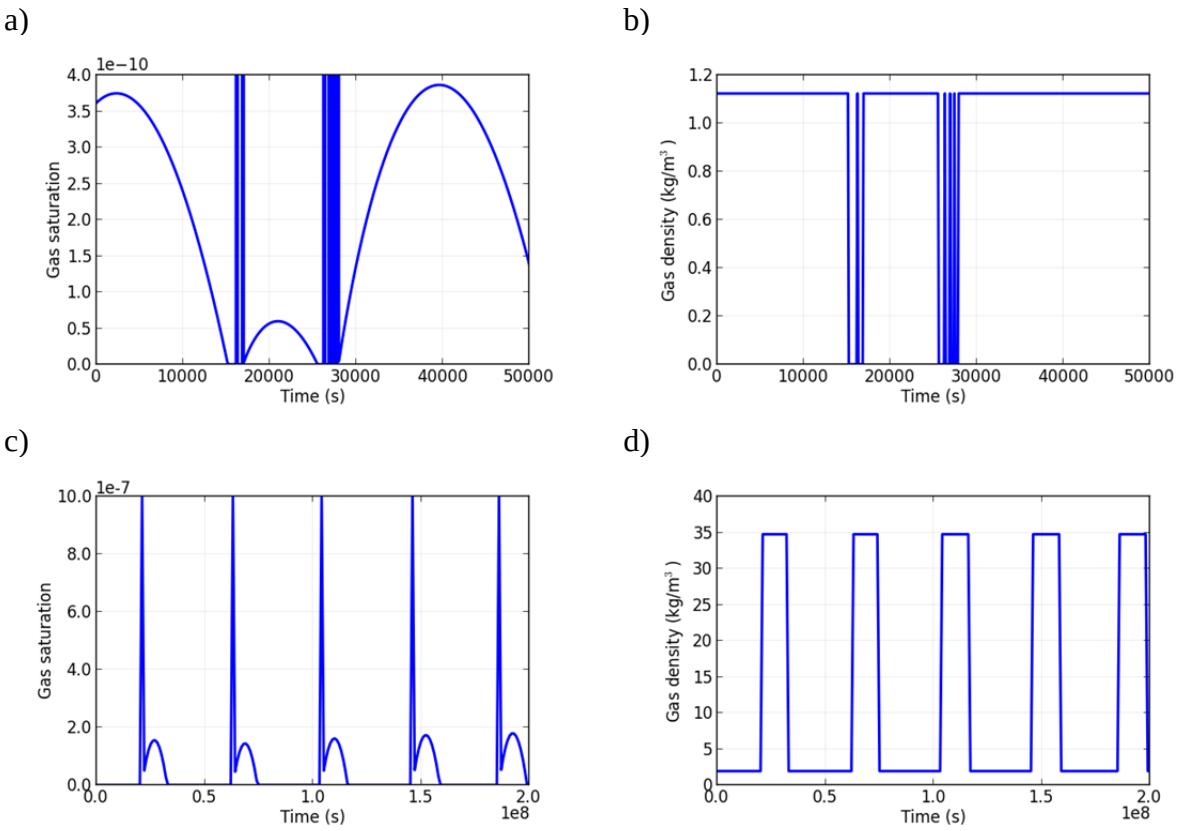

Figure 2: Results for (a) gas saturation, (b) gas density in the problem block in a stalled simulation of the Lihir Island system. The results for (c) gas saturation, (d) gas density in the problem block in a stalled simulation of the Wyang Windu system

The solution to the problem is to ensure that the value of a phase density is continuous prior to a phase evolving or after it disappears. This is achieved by calculating the phase density in the same way as it would be for the two-phase mixture regardless of whether or not the phase is present in the block. The corrections are given in Table 2 and the time step behaviour as a result of the corrections is shown in Figure (3) for both examples. In plot (a) the Lihir Island simulation immediately progresses to a large time step once the correction has been applied and achieves the natural state simulation time within 70 time steps. The Wyang Windu simulation had clearly stalled further from the natural state solution and required over a 100 time steps before the natural state solution was achieved. The maximum time step sizes in both corrected simulations were around $4x10^{13}$ seconds. A further improvement to the TOUGH2 algorithm is proposed in the next section that eliminates this limiting time step size.

Table 2: Corrections to the gas density calculations

Equation of State	Previous Calculation	New Calculation
EOS2	$\rho_g = 0$	$\rho_g = \rho_{vapour} + \rho_{air}$
EOS3	$\rho_g = \rho_{vapour}$	$\rho_g = \rho_{vapour} + \rho_{CO2}$

a) b)

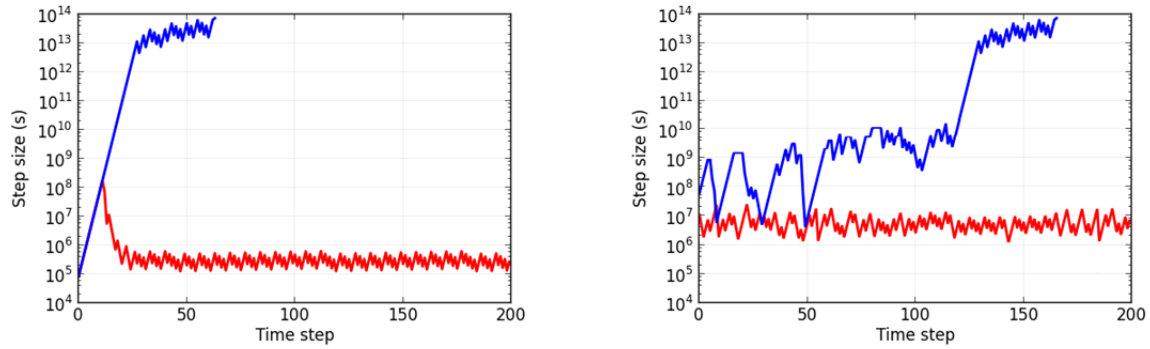

Figure 3: The convergence behaviour before (▪) and after (▪) the gas density calculation correction for (a) the Lihir Island model using EOS3 and (b) the Wayang Windu model using EOS2.

3.2 Time step size limitation approaching large simulation times

Since corrections to the density calculations described in the previous section have been made, many models have been tested and no longer exhibit the stalling behaviour during the main part of the simulation. However, as simulations approach large simulation times they often still display poor convergence behaviour due to the development of a limiting time step size. The results for the corrected simulations in Figure (3) show a limiting time steps of approximately $4x10^{13}$ seconds. For most simulations a natural state is deemed to be obtained once a total time of $1x10^{16}$ seconds is reached, meaning that at least 250 time steps would be required for these simulations. In some models the limiting time step at can be much smaller, hence requiring hundreds or thousands of time steps to reach a natural state simulation time.

Investigations have shown that the limiting time step size does not affect all simulations. In particular, pure water simulations (EOS1) are rarely affected. This information allowed us to focus our efforts on certain areas of the TOUGH2 algorithm to identify the cause of the problem. These efforts revealed that the time step size limitation is caused by loss off significance when updating the primary variables during the NR iteration.

Once the residuals $\mathbf{r}(\mathbf{x}_k)$ have been calculated and neither of the convergence criteria in Equations (17-18) have been met, then the primary variables $\mathbf{x}_k$ are updated to their new values $\mathbf{x}_{k+1}$ using Equation (15). As the simulation approaches the natural state simulation time the system approaches a steady state and these updates become very small. In Equation (7) this equates to the fluxes, sources and sinks being balanced and the accumulation terms being divided by the large time step size Δt^{n+1}. Because TOUGH2 uses dimensional variables, very small updates can lose significance when applied to large values such as the pressure. Also for very small variables such as the mass of air deep in a geothermal system, the updates can become so small that they are affected by machine zero. When the significance of the updates to the primary variables is lost, they remain unchanged from one NR iteration to the next and hence the residual does not reduce. This creates a limit to the time step size for the simulation.

The solution proposed in this work is to include the relative change in the primary variables as well as the residuals when determining convergence. This means that when the updates to the primary variables that are calculated by the NR method are below a prescribed level the block in question is considered to be converged and the NR method is allowed to proceed.

This additional convergence criterion can be written as:

$$\left|\frac{\mathbf{x}_{k+1} - \mathbf{x}_k}{\mathbf{x}_k}\right| \leq \boldsymbol{\varepsilon_3} \tag{19}$$

For this work a value of $\boldsymbol{\varepsilon_3} = 1\text{x}10^{-10}$ was found to be sufficient to allow simulations to progress without reaching a limiting time step size without altering the simulation behaviour. Detailed analysis of many different simulations was carried out to ensure that this additional criterion gave exactly the same simulation results and altered only their convergence behaviour.

Figure (4) shows the convergence behaviour for four different simulations with and without the additional convergence criterion. The examples are from simulations using a different equation of states in order to demonstrate the importance of the improvement for all TOUGH2 simulations. Note that the simulations in plots (a) and (c) started from initial conditions close to the natural state where as simulations (b) and (d) did not.

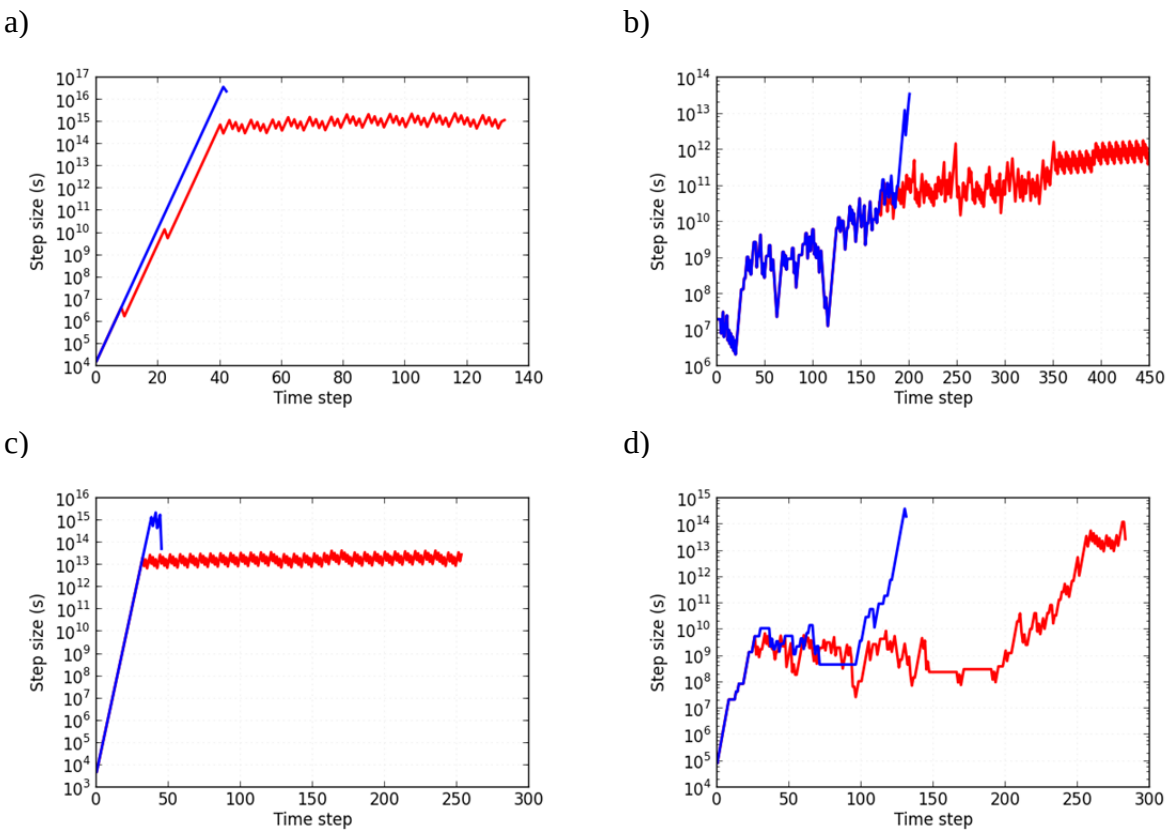

Figure 4: The convergence behaviour of simulations using the previous convergence criteria (-) and using the proposed convergence criteria (-). For simulations using EOS1 (a), EOS3 (b), EOS2 (c) and EWASG (d).

Plot (a) shows that the time step limitation also affects pure water simulations but only at very large time steps (above $5x10^{14}$ seconds). This is because these simulations do not include the mass of a non-condensable gas as a primary variable and hence do not experience the same loss of significance when this mass approaches zero in certain parts of the geothermal system. In plot (b) the simulation using the new convergence criterion appears to suffer from a new time step limitation around $1x10^{15}$ seconds. This behaviour occurs because at very high time step sizes the system of equations solved by the NR become more ill-conditioned and in some case the linear solver can fail to converge. Including improvements to TOUGH2's linear solver is an area of on-going research.

3.3 Implications for inverse modelling and uncertainty quantification

Both inverse modelling and uncertainty quantification rely upon sensitivity analysis of models to their parameters. This may require hundreds or thousands of forward model runs often from initial conditions close to the final solution and with only small changes to a single parameter. Using the previous convergence criteria even small changes to a model parameter could lead to simulations requiring several hundred time steps. Using the proposed convergence criteria this is often reduced to less than fifty and significant reductions in the number of NR iterations required at each time step.

To demonstrate the impact of the speed up achieved tests were carried out using the Lihir Island model. It consists of 9683 blocks, covers a physical area of 103 km^2 and is solved using EOS3. For more details of the model refer to O'Sullivan et al. [12]. Table 3 show a comparison of the real time required for carrying out an inverse modelling analysis using the previous convergence criteria and the new criteria. Both sets of results were generated on Intel Xeon 3.20 GHz processors.

Table 3: Simulation speed-up during inverse modelling analysis

	Number of Simulations	Total Simulation Time
Previous Convergence Criteria	4000	~ 2 x 10^5 s
New Convergence Criteria	4000	~ 1 x 10^4 s

4 CONCLUSION

Corrections to density calculations in the TOUGH2 algorithm have been developed and tested that prevent simulations from stalling at relatively small steps. This ensures that simulations converge to a large natural state total time. A new convergence criterion has also been introduced to the Newton-Raphson method used in the TOUGH2 algorithm that prevents simulations from reaching a limiting time step size. The new criterion has been tested and shown to significantly reduce real simulation time while producing identical simulation results. These two improvements to the TOUGH2 algorithm allow for more efficient model development and enable effective use of inverse modelling and uncertainty quantification tools.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the New Zealand government MBIE Geothermal Supermodels grant which partially supported this research.

REFERENCES

[1] J. Burnell, E. Clearwater, A. Croucher, W. Kissling, J.P. O'Sullivan, M. O'Sullivan and A. Yeh, Future directions in geothermal modelling, *Proceedings (electronic) 34rd New Zealand Geothermal Workshop*, University of Auckland, Auckland, New Zealand, November 19-21, 2012.

[2] K. Pruess, K. Oldenburg and G. Moridis, *TOUGH2 user's guide, version 2.0*, Lawrence Berkeley National Laboratory, University of California, Berkeley, 1999.

[3] M.J. O'Sullivan, K. Pruess and M.J. Lippmann, State of the art of geothermal reservoir simulation, *Geothermics*, 30(4), 395-429, 2001.

[4] Finsterle, S., K. Pruess, D.P. Bullivant and M.J. O'Sullivan, Application of inverse modeling to geothermal reservoir simulation, *Proc. 22nd Workshop on Geothermal Reservoir Engineering*, Stanford University, Stanford, California, 1997.

[5] Finsterle, S., iTOUGH2 User's Guide, Report LBNL-40040, Updated Reprint, Lawrence Berkeley National Laboratory, Berkeley, Calif., 2007a.

[6] J. Doherty, *PEST: model-independent parameter estimation. User manual, 5th edition*, Watermark Numerical Computing, Corinda, Australia, http://www.sspa.com/pest, 2010.

[7] D. Noy, S. Holloway, R. Chadwick, J. Williams, S. Hannis, and R. Lahann, Modelling large-scale carbon dioxide injection into the Bunter Sandstone in the UK Southern North Sea, *International Journal of Greenhouse Gas Control*, 9, 220-233, 2012

[8] J. O'Sullivan, A. Croucher, A. Yeh and M. O'Sullivan, Improved convergence for air-water and CO2-water TOUGH2 simulations, *Proceedings (electronic) 35th New Zealand Geothermal Workshop*, Rotorua, New Zealand, November 18-20, 2013.

[9] A. Yeh, A. Croucher, and M.J. O'Sullivan, Recent developments in the AUTOUGH2 simulator, *Proceedings TOUGH Symposium 2012*, Berkeley, California, September 17-19, 2012.

[10] M.J. O'Sullivan, Geothermal reservoir modelling theory, Lecture delivered at University of Auckland, Auckland, New Zealand, unpublished, 2012.

[11] J. O'Sullivan, D. Dempsey, A. Croucher, A. Yeh, and M. O'Sullivan, Controlling complex geothermal simulations using PyTOUGH, *Proceedings Thirty-Eighth Workshop on Geothermal Reservoir Engineering*, Stanford University, Stanford, California, February 11-13, 2013.

[12] J.P. O'Sullivan, A. Croucher, M.J O'Sullivan, L. Stevens, and M. Esberto, Modelling the evolution of a mine pit in a geothermal field at Lihir Island, Papua, New Guinea, *Proceedings (electronic) 33rd New Zealand Geothermal Workshop*, University of Auckland, Auckland, New Zealand, 21-23 November, 2011.