Proceedings World Geothermal Congress 2015
Melbourne, Australia, 19-25 April 2015

Resource Assessment, a Review, with Reference to the Australian Code

Malcolm A Grant

14A Rewi Rd, Auckland 1023, New Zealand

malcolm@grant.net.nz

Keywords: Stored heat, resource assessment, Australian Geothermal Code, numerical simulation, EGS

ABSTRACT

Resource assessment has in the past relied on stored heat or volumetric methods, and more recently on numerical simulations. Attempts to codify the process have been spectacularly unsuccessful, as shown by the example of the Australian Code which is biased high, sometimes by a large multiple. Attempts to compensate by using Monte Carlo methods have been at times comical failures. Numerical simulation requires both a natural state and a history match. However in the assessment of a new field only the natural state is available. Simulations based upon only a natural state match can be in significant error, but still a considerable improvement on stored heat methods.

1. INTRODUCTION

Stored heat or volumetric assessments of geothermal resources are appealingly simple: The resource being exploited is heat. A stored heat calculation simply computes the amount of heat in the resource, similarly to computing the amount of ore in an ore body. The method has theoretical support in numerical simulations of resource production. While there are significant unknowns in any resource, some of these may be covered by probabilistic approaches, notably a Monte Carlo method. The Australian Geothermal Code represents one specification of such stored heat assessments.

However the experience of recent decades, with the development of significant numbers of geothermal resources, has shown that the method is highly unreliable, and usually biased high. The tendency to overestimates, in particular, has led to the reduced credibility of the method. An example is quoted where simple application of the apparently simple rules gives a ridiculous result. Much of the problem lies in the "recovery factor", the proportion of the resource that can actually be exploited, where comparison with actual performance shows past values have been in all cases too high, as is the current version of the Australian Code.

2. STORED HEAT DEFINITIONS

The amount of heat stored in a geothermal reservoir is given, in simple form, by:

$$Q = \int \rho_f C_f (T - T_{ref})\, dV \tag{1}$$

Where $\rho_f C_f$ is the heat capacity of the rock (including pore fluid), T is the rock temperature and T_{ref} is a cutoff or base temperature. This formulation omits dependence on porosity to which the result is only slightly sensitive. The integral is taken over the entire volume of the reservoir which is being assessed. The practicalities of geological complexity mean that only a fraction η of this heat will in fact be recoverable, and the amount is reduced by this "recovery factor". Finally, the produced heat will be converted to electricity by plant with thermal efficiency η'. This gives the amount of electrical energy recoverable from the resource as:

$$E = \eta\eta' Q = \eta\eta' \int \rho_f C_f (T - T_{ref})\, dV \tag{2}$$

The method is appealingly simple: draw isotherms, evaluate the integral over the volume, multiply by the recovery factor and thermal efficiency to get a total quantum of electricity. Then dividing by the planned lifetime gives a power capacity in MW.

There are some parameters in the formula. The base temperature T_{ref} is normally taken as the reject temperature T_r of the turbine, on the basis that it is the heat supplied to the plant, above this temperature, that is being used. The thermal efficiency η' is determined by the laws of thermodynamics and plant design. For binary plant that is commonly used, T_{ref} is usually around 80-90°C, and η' is frequently around 10%. For the recovery factor η a value must be assumed.

The stored heat method was first developed by USGS (White & Williams, 1975), for a national assessment of geothermal reserves. Simple models of flow in fractured porous medium were used to estimate an average value of the recovery factor of 25% (Nathenson, 1975). This value of 25% was widely copied and has been used in most applications of the method since then. Sometimes the inevitable uncertainty is knowledge about the resource is reflected by probabilistic results – usually a Monte Carlo approach, with a range of values being assumed for various parameters, sometimes including the recovery factor. It is hoped that by doing this a more robust final answer is obtained. More recently the stored heat method has been codified in the Australian Geothermal Reporting Code (AGEA AGEG 2010). The method is not changed from the USGS original.

In all reserve assessments there is an assumption about the technology and economics – proven reserves use current technology and current economics. Reserves dependent upon an anticipated change in either technology or economics form a contingent reserve. It is quite common, when a field has been over-assessed, and subsequent development makes this apparent by failure to support the station, to appeal to deeper drilling, or stimulation, or new exploration. All of these are possible responses to a fuel shortfall but, if successful, do not make the original assessment valid: it must have been produced on the basis of the technology and field knowledge then available.

3. STORED HEAT IN PRACTICE

3.1. Observed results

The stored heat method is purely theoretical, lacking any support from observation. Ideally such a method would be validated by an estimate being made, and then later compared with actual results when the field was subsequently developed. There were many geothermal developments of hydrothermal systems through the late 20th century, and although there were no explicit comparisons it became apparent that past estimates had been too high. This was first reported by Grant (2000), and later by Sanyal et al (2002, 2004), Stefansson (2005) and Benoit (2013). It was observed that past estimates had in some cases been several-fold overestimates. Because it is an heuristic fudge factor, effort has tended to focus on the recovery factor, with these observations generally being regarded as showing that past values were too high.

Sanyal et al. (2002) compared past stored heat assessments against actual performance and numerical simulation, and found that factors of 5-10% were "a more reasonable range of values". Sanyal et al. (2004) similarly reviewed the USGS assessments (Muffler, 1978), and found that the total resource was one-third of the original estimate, and that recovery factors should lie in the range 3-17% with a mean of 11%. Williams (2004) similarly reviewed performance in three US fields and found recovery factors closer to 10% than 25%, and that recovery varied strongly between fields. Note that in all three studies the recovery factor is in every case under 25%, so the implication is that *every* field reviewed was initially overestimated. These results are supported by the observation that US geothermal operates at around 60% of capacity (Lund et al. 2010), arguing for much past oversizing. The same proportion applies to the state of Nevada, so this result is not over-weighted by The Geysers.

These results represent the only published evaluations of the stored heat method compared against observation. They clearly establish that 25% is too high a recovery factor, and that a mean value of around 11% corresponds to observed results. Beyond establishing the correct average recovery factor, there are a wide range or recovery factors – 3-17% covers the entire range of observed results. This indicates that any result is subject to an error of ±70% about the median of 10%.

No method was proposed for discriminating between fields with low and high recovery. In a different context, Wilmarth & Stimac (2014) indicate differences in power density (which imply similar differences in recovery factor) depending on tectonic environment, with fields in extensional tectonic environments having higher recovery than those in compressive environments.

3.2. The Australian code

The Australian Code comes in two parts, the "Code" and the "Lexicon" (AGEA AGEG 2010a,b). The Code stipulates a minimum level of detail for reporting. It is non-prescriptive about methodology, except that it says "defer to the Lexicon if in doubt". In the Code's first draft, the Lexicon recommended a recovery factor of 25%. In the second edition, in the light of Williams' work, this was revised to 14-17.5%. It is difficult to see any justification for this value which is still too high, and in consequence the Code, on average, still overstates reserves by up to 75%. Further, the Lexicon claims an error of ±50%, but this understates the range of observed results.

4. OTHER PARAMETERS

4.1. Base temperature

The focus on the recovery factor has diverted attention from other uncertainty within the stored heat model. For example, Sarmiento & Björnsson (2007) reviewed past reserve estimates in The Philippines and found that they were roughly in line with subsequent performance. The author can also confirm this latter observation from experience in The Philippines. However, it is critical to note that the stored heat method in use in The Philippines is different from the USGS method. The base temperature, T_{ref}, normally used is not the reject temperature of the turbine, but the temperature T_{Pmin} at which wells cease to operate. This is typically 180-200°C. The practical consequence is that stored heat, computed using the Philippine method, for typical Philippine resources, is 35-40% of the value that would be computed using the USGS method. The author has seen simultaneous evaluations of the same resource, using the Philippine and the USGS method. The latter gave a result 2.5 times the former. Thus in fact the results of Sarmiento & Björnsson, while supporting the Philippine results with a recovery factor of 25%, also indicate that had the USGS method been used, the recovery factor should have been 10% rather than 25%.

Which choice for the base temperature is appropriate depends upon an assumption about the state of the reservoir at abandonment (Clotworthy et al. 2010). The heat that is extracted from the reservoir is the difference between the initial and abandonment states, integrated over the entire reservoir. Consider contrasting examples, of a reservoir produced by cold sweep and a reservoir produced by boiling down or mixing. Assume that the production and injection wells are all ideally placed so that production continues from all wells, until all fail simultaneously as their temperature falls too low.

If there is a homogeneous reservoir containing liquid, with a smooth inward sweep of water from the injection wells, the production area remains near original temperature until the thermal front arrives. This will be at the injection temperature, which is the turbine reject temperature. So in this case all the heat above the reject temperature is produced, and the reject temperature is the proper choice for the base temperature. Conversely, consider a high-temperature reservoir containing two-phase fluid which boils continuously under production; or a reservoir subject to rapid mixing of injectate into the production area. In both cases the reservoir cools steadily under production, and production ceases when temperature falls to the minimum producible temperature. That temperature is the temperature of the reservoir at abandonment, and so only the heat above this temperature has been produced.

The difference can be shown clearly by considering a hypothetical liquid reservoir at 180°C. The minimum production temperature T_{Pmin} is also 180°C. A binary turbine using the produced fluid has a reject temperature T_r of 80°C. If the latter is used as the cutoff temperature there is a substantial resource. If T_{Pmin} is used as cutoff there is no resource at all. Now consider actual exploitation of this resource. If waste water is injected at a distant site, from which it returns, sweeping water at reservoir temperature ahead of it, the wells will continue to produce for an extended period and there is a resource to exploit. On the other hand, if injectate returns immediately into the production area, mixing with and cooling the reservoir, the reservoir temperature quickly drops below 180°C, and the wells and the project cease operation. In the first case T_r is the correct choice for cutoff; in the second case T_{Pmin} is the correct choice and there is very little exploitable resource present.

Of course in the latter case the use of downhole pumps would make the resource exploitable. This is a change in the assumed technology, as the presumed T_{Pmin} would no longer apply. Such a technological change does not modify the prior assessment, it introduces a new assessment dependent upon different assumptions.

Actual field behaviour normally lies between these extremes, so there is some argument for either choice. It is important that the choice made be clear, and consistently applied. In particular, if some field experience is being used to justify the choice of recovery factor, it is important that the same method be used throughout. In this the recommendation of Atkinson (2012) for the use of analogue fields is appropriate.

Garg & Combs (2011) give further elaboration on reference temperature, conversion efficiency and how these apparently simple parameters depend on the specific case. Franz et al (2015) likewise elaborate on parameter problems, in particular the reject temperature.

4.2. Volume

Further assumptions arise in the definition of the reservoir volume. The problem is that many geothermal fields contain regions of very poor permeability, such that they do not constitute a reservoir at all. For a volume of hot rock to contribute to production, and hence to the field's reserves, there must be sufficient permeability within it that fluid flow, and pressure changes, propagate through the rock during the production lifetime. Often part of the field fails to reach this minimal level of permeability.

One approach has been to ignore this problem, and assume this is one of the many complications swept into the recovery factor: there is within the recovery factor an implicit allowance for part of the reservoir being impermeable. A contrasting approach has been to consider that proven reserves derive only from regions known to be permeable. At whatever the state of exploration, contours are drawn around the productive wells and this region only is counted. This second approach tends to err on the low side, because of the probable reserves that may exist in regions not yet drilled.

Either approach can be used provided it is used consistently. For undrilled fields where the only information on area is geophysical, there is little alternative to using the entire area outlined by geophysical data. It would seem that the recovery factor for such an approach should be lower than when the area considered is restricted to that drilled successfully.

A further issue arises in respect of the depth of the reservoir. Consider a hypothetical field in which permeability decreases markedly with increasing depth (as is often the case). An array of wells has already been drilled to mid depths, supporting some level of production. A new well is drilled to significantly greater depth. It finds increasing temperature but no significant permeability below that already found. Has it added to the field reserves by adding more volume at the reservoir bottom? There is at this point no known productive capacity at greater depth. The hot rock and fluid at this greater depth only increase the field's long-term production if, over field life, fluid circulates to these depths – the most likely scenario would be injection at depth. If the permeability is too low the additional hot rock is worthless for production. The same issues apply to the definition of reservoir depth as to the area.

4.3. Range of variation of different estimates

The stored heat method is not particularly well-constrained, and results obtained vary quite widely depending on the details of the calculation and the policies of the author. This can be seen in the results reviewed above. Sanyal et al. (2002) reviewed a set of past assessments and found a recovery factor of 5-10%. Sanyal et al. (2004) reviewed past USGS assessments and found an average recovery factor of 11%, or an overestimate by a factor of 2. The first group would have a larger overestimate. Presumably this first group came from different author(s). A third contrast is the Philippine results discussed above, where 25% was actually appropriate, but using a different definition of stored heat. Such wide variations in practice are quite normal, and simply illustrate how poorly constrained this method is. Wide (several-fold) overestimates are common, although underestimates are seldom encountered.

5. PROBABILITY

Monte Carlo methods are frequently used to provide some accounting for the inevitable uncertainties in a resource at best only partly-explored. Table 1 reproduces two results from a national assessment of New Zealand by Lawless & Lovelock (2002), using the method of the Australian Code. In the data columns (Area, thickness, mean temperature) the data are given as minimum/mode/maximum, and the capacity is given as 10th percentile/median/90th percentile. The percent void space is also included because the calculation uses a recovery factor proportional to void space (2.5xvoid).

Table 1. New Zealand national assessment (Lawless & Lovelock 2002), selected fields

Field	Area km^2	Thickness m	Mean Temp °C	Void %	Capacity, MWe
Mangakino (1)	8/9.5/17	1500/1700/2200	220/230/250	8/10/12	65/85/120
Mangakino (2)	0/8/10	1500/1700/2200	220/230/250	8/10/12	20/47/70
Ohaaki	5/11/12	1800/2100/2500	260/275/280	6/8/10	100/135/175

There were two versions of the paper, with different estimates for Mangakino. Mangakino was drilled after this assessment. After four deep wells, the field was abandoned. In a result unique in New Zealand geothermal exploration, the thermal anomaly was present but proved to be conductively heated in rock of very low permeability. There is no ability to produce this resource and so the actual capacity is 0 MWe. This result lies far outside the range of the assessed probability distribution. In any undrilled field, there is some risk of failure. The second Mangakino estimate extends the area range down to zero, but this still produces a vanishingly small probability of a zero outcome, and so this modification still does not cover the actual realisation – it would be necessary to add a discrete quantum at 0 MWe. Rather than so modifying the probability distribution, it is probably better to assign a prior probability of failure (~5-10%), and then consider the current distribution as applying if this contingency does not eventuate.

Ohaaki power station came on line in 1988, with installed capacity 112MWe. There was significant surplus production capacity at startup. The station ran at full load for 5 years, then ran down rapidly and production since has been about half the installed level, despite significant drilling of deeper wells (Clotworthy 2008, Clearwater et al 2011). Levelising the actual generation would give a capacity near 80MW, as originally assessed by Grant (1979). Again, the actual result lies well outside the range of possible outcomes from the Monte Carlo assessment. In this particular case source of the error is clear. Ohaaki has a permeable area of only 3 km^2, despite the much larger field area. It is an extreme case of the issue discussed above, of hot but impermeable regions within the geothermal anomaly.

A more detailed examination of the risks involved in Monte Carlo assessments is given by Garg & Combs (2010), who examine a the field evaluation of Silver Peak, previously published by GeothermEx (2004). By making reasonable variations on the assumed parameters, widely different outcomes are produced. It is clear that there is far greater uncertainty in the proven capacity than is represented in the original assessment. Garg & Combs also give a detailed discussion of the sensitivity to reject temperature and the effect of changes in this parameter on other assumptions, reinforcing that the stored heat method is not nearly as simple as it first appears.

The problem in all these cases is that the concept of the Monte Carlo method has been misapplied, by using overly restrictive estimates of parameter ranges. "Reasonable" assumptions are made about the possible ranges of parameters. But, a probability distribution must include not just "reasonable" outcomes. It must include *all* outcomes that are possible in the given state of knowledge. If there is information unknown, which may produce a particular outcome, that outcome must be included within the probability distribution. The parameter ranges must not be "reasonable", but inclusive of *all* possibilities. This is a very common fault in the use of the Monte Carlo method in geothermal resources. If new information produces a new distribution, this new distribution must always be contained within the range of the prior distribution. If new information might extend the range of a distribution, then it must be so extended initially.

In the case of Mangakino above, the assessment was made before deep drilling (there had been one shallow well earlier). It is an observational fact that some fields are drilled and abandoned. Usually the reason is adverse chemistry, but lack of permeability is another reason sometimes seen, and realised here. Before drilling, there is some possibility that the field will be abandoned, and hence some possibility of a capacity of zero. This possibility must be contained within the range of probability. The revision to produce the second estimate for Mangakino also shows the failure to include all possibilities. In the light of additional information, the area range was extended outside the first range. If the first range were properly constructed, additional information would only produce a restriction – anything that could result from additional information should already be included. Also note in Table 1 the quite narrow range of possible reservoir thicknesses. Geothermal fields vary greatly. Some have essentially only one permeable interval, and so are simply hosted within this one aquifer. There are other fields where drilling beyond 3km depth has continued to find permeability and production. So, the range of thicknesses actually realised ranges from a few hundred metres to a few kilometres. Absent any more knowledge, it is very surprising to so restrict the range of possible thicknesses.

In the case of Ohaaki, Lawless & Lovelock argue: *"Some of resource has demonstrated low permeability, but only to a certain depth, and so is all included in the higher areal estimate."* That is, knowing that the resource as then drilled has only a very limited permeable area, because there might be permeability at depth over a wider area, the greater area is included. This argument deliberately excludes from the range of possible areas, the area as then known; and consequently incorrectly modifies the probability distribution. In the absence of knowledge about permeability at greater depth, the possible outcomes must include both good and bad results of deeper drilling, and the argument advanced by Lawless & Lovelock is an argument for extending the range of possible areas, not for excluding part of that range.

The Silver Peak case similarly shows too-restrictive assumptions about possible outcomes of further exploration. All the cases presented by Garg & Combs can be supported by argument, and so, in the present state of knowledge, a probability distribution describing this knowledge must include them all.

It can be further observed that the use of the Monte Carlo method does not provide protection against over-estimation. If the parameters are assumed with values too high, the resultant probability distribution will also average too high. Ohaaki provides a clear example – actual performance lies well below the 10%-ile. Having a distribution on the inputs does not correct for, in this case, the area and recovery factors being assumed too high over their respective entire assumed ranges. When the mean values of the parameter distributions assumed are biased high, the result will be similarly biased.

More elaborate statistical methods are discussed by Onur et al. (2010)

6. EGS

6.1 Theoretical estimates

The conclusions above apply only to hydrothermal systems, as they are the source of all data used. There is very little data for any other systems. Initial estimates of heat recovery from EGS use fractured medium models, similar to the reasoning of Nathenson

(1975). Sanyal & Butler (2005), using such modelling, found a recovery factor to be typically 40-50%, which was used as a starting point for MIT (2006).

6.2 Observations

Williams (2010) concluded that "Field observations and modelling studies indicate that values for the recovery factor less than [10%] are likely to be representative". There is one published analysis for EGS based upon field data (Grant & Garg 2012), which indicates an even lower recovery factor, 2%. It would be expected that EGS would have lower recovery than naturally fractured systems, as the fractures will be less pervasive. Although microseismicity indicates significant reservoir volumes stimulated by fracturing, the reservoir developed, through which there is significant fluid flow, appears to be markedly smaller (Williams 2010, Grant & Garg 2012). As with hydrothermal systems, it appears that modelling studies have erred on the optimistic side, but for EGS by an even greater amount. Recovery factors for EGS should be at best a few percent.

The theoretical estimates have assumed uniform or random fracturing. It can be observed that in practice the fracture system is neither uniform nor random. Instead there is a single plane of fracturing, with permeability that is effectively swept by fluid extending only a small distance away from the plane. If anything, the reality is nearer to the original HDR concept of 'flow in a penny-shaped crack'.

6.3 The Australian code

The Lexicon recommends a default recovery factor of 14%. This is an overstatement of the observed result by a factor of 7.

7. RESOURCE ASSESSMENT

7.1. As practiced in New Zealand

A desk survey was made of geothermal resource decisions (to commit to or to approve either development or expansion) made in the last 20 years in New Zealand. This was a simple tabulation of decision points and the resource assessment method used to support the decision. Within the New Zealand legal system, developments require "resource consents", which are public processes. There may be consents for exploration or development, or both. Exploration consents were not considered as these typically require relatively little supporting information. The decisions that were considered were: developer's commitment to proceed; granting of development consent; and bank decision to finance.

25 such decisions were identified. There were often special features to any individual case. At Kawerau for example, a field model has been developed, and has been approved by the consent authority, on condition that it is available to all developers. All later consent applications use it, and therefore there was no consideration of what method to use. At Wairakei a field model has long existed but is proprietary. Competing applications require an independent model. There is also a variable amount of legacy information from the government geothermal exploration program of the 1960s-1980s.

Among the 25 decisions identified, reservoir simulation was used 21 times, and in 3 cases it was not used at the first decision but was used later. Power density and stored heat were each used 6 times, but always in support of simulation, which was either produced at the same time or promised for the next stage. There was one case where total well flow was used. This high-level survey establishes that the dominant means of resource assessment in New Zealand is numerical simulation, at all stages of development. This includes undeveloped fields, see for example Ngatamariki (Grant & Bixley 2011).

7.2. Simulation of an undeveloped resource – the Australian code

The Australian Code Lexicon claims "*At the exploration stage, when limited data are available, numerical simulation is unlikely to give a more realistic estimate of long term capacity than simpler volumetric methods.*" This sweeping claim is at variance with industry practice and experience. It is industry practice to use simulation of undeveloped resources, and has been for some years - about 30% of the simulations reported in the survey of O'Sullivan et al (2001) are of undeveloped resources. The process of reservoir simulation is that accuracy increases with increasing amounts of data – more constraints improve model calibration. A simulation of an undeveloped field involves significant calibration against the natural state pressure and temperature distribution plus possibly interference tes ts. The idea that a volumetric estimate, ie a simple average over the entire volume, is better determined is nonsense.

8. CONCLUSIONS

The best estimate of recovery factor for hydrothermal systems, based on observed results, is 10%. The resulting resource estimates are subject to an error of ±70%. This is not reduced by the used of Monte Carlo methods.

The Australian Geothermal Code & Lexicon, even in its revised form, overstates recovery factor for hydrothermal systems by up to 75%, and for EGS by a factor of 7. Even with these corrected, error is understated compared to the range of observed results. A claim that volumetric estimates are as good as simulation in undeveloped fields is contrary to established industry practice.

The Lexicon is neither an accurate reflection of observed field performance nor of current industry practice.

REFERENCES

AGEA AGEG 2010a "The geothermal reporting code Second edition"

AGEA AGEG 2010b "Geothermal lexicon for resources and reserves definition and reporting"

Atkinson, P., 2012 "Proved Geothermal Reserves - A Framework and Methodology" Proceedings, New Zealand geothermal workshop.

Benoit, 2013 “An empirical injection limitation in fault-hosted Basin and Range geothermal systems” Transactions, Geothermal Resources Council, v37.

Clearwater, E.K., O’Sullivan, M.J., & Brockbank, K., 2011 “An update on modelling the Ohaaki geothermal system” NZ Geothermal Workshop

Clotworthy, A.W., 2008 “Ohaaki geothermal field – field management” presented at NZGA seminar

Clotworthy, A.W., Lawless, J.V., & Ussher, G., 2010 “What is the end point for geothermal developments: modeling depletion of geothermal fields” World Geothermal Congree paper 2239

Franz, P., Neville-Lamb, M., Azwar, L., &Quinã0, J., 2015 “Calculation of geothermal stored heat from a numerical model for reserve estimation” These proceedings

Garg, S.K., & Combs, J., 2010 “Appropriate use of USGS volumetric “heat in place” method and Monte Carlo calculations” Proc, 35th Workshop on geothermal reservoir engineering, Stanford University

Garg, S.K., & Combs, J., 2011 “A re-examination of USGS volumetric “heat in place” method” Proc, 36th Workshop on geothermal reservoir engineering, Stanford University

GeothermEx (2004) “New geothermal site identification and qualification” Consultant report for California Energy Commission PIER report no. P500-04-051

Grant, M.A. & Garg, S.K., 2012 “Recovery factor for EGS” Proc, 37th Workshop on Geothermal Reservoir Engineering, Stanford University

Grant, M.A., 1980 “Broadlands optimization” Transactions, Geothermal Resources Council, v3, pp269-270

Grant, M.A., 2000 “Geothermal resource proving criteria” Proceedings, World Geothermal Congress, pp2581-2584

Grant, M.A., & Bixley, P.F., “Geothermal reservoir engineering, 2nd edition” Academic Press

Lawless, J., & Lovelock, B., 2002 “New Zealand’s geothermal resource” paper presented at NZ Geothermal Association

Lund, J.W., Gawell, K., Boyd, T.L., & Jennejohn, D., 2010 “The United States of America country update 2010” World Geothermal Congress paper 0102

MIT, 2006 “The future of geothermal energy”

Muffler, L.P.J., 1978 “Assessment of geothermal resources of the United States – 1978” U.S. Geological Survey, Circular 790, 163p.

Nathenson, M., 1975 “Physical factors determining the fraction of stored energy recoverable from hydrothermal convection systems and conduction-dominated areas” USGS Open-file report 75-525

Onur, M., Sarak, Hülya, and Türeyen, Ö.,İ., 2010 “Probabilistic resource estimation of stored and recoverable thermal energy for geothermal systems by volumetric methods” World Geothermal Congress paper 2219

O’Sullivan, M.J., Pruess, K., & Lippmann, M.J., 2001 “State of the art of geothermal reservoir simulation” Geothermics 30, pp395-429

Sanyal, S.K., Henneberger, R.C., Klein, C.W., & Decker, R.W., 2002 “A methodology for the assessment of geothermal energy reserves associated with volcanic systems” Transactions, Geothermal Resources Council v26, pp59-64

Sanyal, S.K., Klein, C.W., Lovekin, J.W., & Henneberger, R.C., 2004 “National assessment of U.S. geothermal resources – a perspective” Transactions, Geothermal Resources Council v28, pp355-362

Sanyal, S.K., & Butler, S.J., 2005 “An analysis of power generation prospects from Enhanced geothermal systems” Transactions, Geothermal Resources Council, v29, pp131-137

Sarmiento, Z.F., & Björnsson, G., 2007 “Reliability of early modelling studies for high-temperature reservoirs in Iceland and The Philippines” Proceedings, 32nd Workshop on Geothermal Reservoir Engineering, Stanford University

Stefansson, V., 2005 “World geothermal assessment” Proceedings, World Geothermal Congress paper 0001

White, D.E., & Williams, D.L., 1975 “Assessment of geothermal resources of the United States – 1975” Geological Survey Circular 726, USGS

Williams, C.F., 2004 “Development of revised techniques for assessing geothermal resources” Proceedings, 29th Workshop on Geothermal Reservoir Engineering, Stanford University, pp276-280

Williams, C.F., 2007 “Updated methods for estimating recovery factors for geothermal resources” Proceedings, 32nd Workshop on geothermal reservoir engineering, Stanford University

Williams, C.F., 2010, “Thermal energy recovery from enhanced geothermal systems – evaluating the potential from deep, high-temperature resources” Proceedings, 35th Workshop on geothermal reservoir engineering, Stanford University

Wilmarth, M., & Stimac, J., 2014 “Worldwide power density review” Proc., 39th Workshop on geothermal reservoir engineering, Stanford University.