



Stored-heat assessments: a review in the light of field experience

M. A. Grant

MAGAK, 14A Rewi Rd, Auckland 1023, New Zealand

Correspondence to: M. A. Grant (malcolm@grant.net.nz)

Received: 22 June 2014 – Revised: 21 October 2014 – Accepted: 5 November 2014 – Published: 17 December 2014

Abstract. 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 can be covered by probabilistic approaches, notably a Monte Carlo method. The Australian Geothermal Reporting 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.

There are further problems, usually overlooked, in the way that the reservoir volume and “cutoff temperature” are defined. Differing approaches mean that results between different reports are not comparable. The different approaches also imply unrecognised assumptions about the physical processes controlling reservoir depletion. The failure of Monte Carlo methods is similarly due to unrecognised violation of logical consistency in the use of probabilities.

The net effect of these problems is that the method is not a simple means to generate a rough resource estimate, and it often generates faulty results. Usually, such results are overestimates. Monte Carlo methods do not provide a protection against these errors.

The Australian Geothermal Reporting Code should be used for hydrothermal systems with an average recovery factor of 10%. With this average, results are subject to an error of $\pm 70\%$. For enhanced geothermal systems (EGS), the recovery factor should be a few percent.

1 Stored-heat definitions

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

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

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

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

The method is appealingly simple: draw isotherms, evaluate the integral over the volume and 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 megawatts.

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 the 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 the U.S. Geological Survey (USGS; White and 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 % has been used in most applications of the method since then. Sometimes the inevitable uncertainty 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 will be 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.

2 Stored heat in practice

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 is 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) and Stefansson (2005). 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 %: the implication is that every field reviewed was initially overestimated. These results are supported by the observation that US geothermal generation 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 evaluation of the stored-heat method compared against observation. They clearly establish that 25 % is too high a recovery factor and that an average value of around 10 % corresponds to observed results. Beyond establishing the correct average recovery factor, there are a wide range of recovery factors: 3–17 % covers the entire range of observed results. This indicates that any result is subject to an error of at least a factor of 2, or alternatively ± 70 %.

No method was proposed for discriminating between fields with low and high recovery. In a different context, Wilmarth and 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.

The Australian code, in its first draft, recommended 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 %.

3 Other parameters

3.1 Base temperature

The focus on the recovery factor has diverted attention from other uncertainty within the stored-heat model. For example, Sarmiento and 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, derived from Icelandic practice (Pálmason et al., 1985), is different from the USGS method. The base temperature, T_{ref} , normally used is not the reject temperature of the turbine, but the temperature $T_{\text{P,min}}$ 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 larger. Thus in fact the results of Sarmiento and 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. The heat that is extracted from the reservoir is the difference between the initial and abandonment states, integrated over the entire reservoir. Consider the 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 the 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.

Garg and Combs (2011) give further elaboration on reference temperature, conversion efficiency and how these apparently simple parameters depend on the specific case.

3.2 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 for fluid flow, and pressure changes, to 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 with respect to 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 have 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 pro-

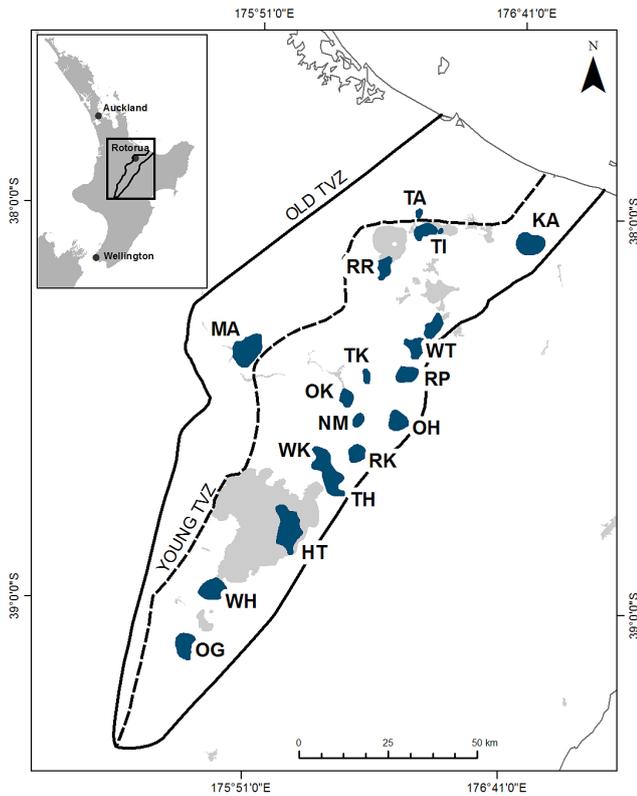


Figure 1. Geothermal fields in the Taupo Volcanic Zone of the North Island, New Zealand. MA = Mangakino, OH = Ohaaki (GNS Science, personal communication, 2014).

duction. The same issues apply to the definition of reservoir depth as to the area.

4 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 and 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 ($2.5 \times \text{void}$). Figure 1 shows the location of the fields discussed below.

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 proba-

bility 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 ($\sim 5\text{--}10\%$) and then consider the current distribution as applying if this first contingency does not eventuate.

Ohaaki Power Station came online in 1988, with installed capacity of 112 MWe. 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 (Clearwater et al., 2011). Again, the actual result lies well outside the range of possible outcomes from the Monte Carlo assessment. In this particular case the error is clear. Ohaaki has a permeable area of only 3 km², 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 and Combs (2010), from which Table 2 is reproduced. The tabulated rows give the capacity in megawatts at the 90th, 50th and 10th percentile, and the ratio of the 10th and 90th percentile values.

Case 1 is the field evaluation previously published by GeothermEx (2004). The other four cases are simply variations on the original assessment, making changes that appear to lie within the range of possibility. The upper (10%) capacity varies somewhat, but the median varies more and the 90th percentile varies by a factor of 4. As this percentile is normally taken as the proven capacity, it is clear that there is far greater uncertainty in the proven capacity than is represented in the original assessment (Case 1). Garg and 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.

In the case of Mangakino above, the assessment was made before deep drilling (there had been one shallow well ear-

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

Field	Area	Thickness	Mean temp	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

Table 2. Megawatt capacity for Silver Peak, Nevada (Garg and Combs, 2010).

	Case 1	Case 2	Case 3	Case 4	Case 5
90 %	41	41	13	9	8
50 %	82	82	64	44	44
10 %	146	150	140	102	137
10 % MW/90 % MW	3.6	3.7	10.8	11.3	17.1

lier). 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 0. 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. Thus, 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 and 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 and 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 and 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 10th percentile. 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.

5 Enhanced geothermal systems (EGS)

The conclusions above apply only to hydrothermal systems, as they are the source of all data used. There are very few data for any other systems. Initial estimates of heat recovery from enhanced geothermal systems (EGS) use fractured medium models, similar to the reasoning of Nathenson (1975). Sanyal and Butler (2005), using such modelling, found a recovery factor to be typically 40–50 %, which was used as a starting point for MIT (2006).

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 and Garg 2012), which indicates even lower recovery factors: 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 volumes stimulated by fracturing, the reservoir that developed, through which there is significant fluid flow, appears to be markedly smaller (Williams, 2010; Grant and Garg, 2012). As with hydrothermal systems, it appears that modelling studies have erred on the optimistic side, and for EGS by an even greater amount. Recovery factors for EGS should be at best a few percent.

6 Conclusions

One conclusion is immediate: past recovery factors have been too high, and comparison with actual performance shows that an average value of 10 % should be used. For this reason the Australian code, even in its second edition, is biased high. With the recovery factor corrected, the results are subject to an error of $\pm 70\%$. For EGS, smaller recovery factors – only a few percent – are indicated and the Australian code is again biased high.

There are significant implicit assumptions about the underlying physical processes in the use of the stored-heat method, which means in practice that there is great uncertainty in the calculation – this may or may not relate to the variation in outcomes just noted. Monte Carlo methods do not provide protection against these uncertainties.

Acknowledgements. Thanks are due to the reviewers, and in particular to E. Juliusson for drawing attention to the “Icelandic method”.

Edited by: G. Beardsmore

Reviewed by: L. C. A. Gutiérrez-Negrín and E. Juliusson

References

- AGEA AGEG: The geothermal reporting code Second edition, 2010a.
- AGEA AGEG: Geothermal lexicon for resources and reserves definition and reporting, 2010b.
- Clearwater, E. K., O’Sullivan, M. J., and Brockbank, K.: An update on modelling the Ohaaki geothermal system, NZ Geothermal Workshop, 2011.
- Garg, S. K. and Combs, J.: Appropriate use of USGS volumetric heat in place method and Monte Carlo calculations, Proc, 35th Workshop on geothermal reservoir engineering, Stanford University, 2010.
- Garg, S. K. and Combs, J.: A re-examination of USGS volumetric heat in place, method, Proc, 36th Workshop on geothermal reservoir engineering, Stanford University, 2011.
- GeothermEx: New geothermal site identification and qualification, Consultant report for California Energy Commission PIER report no. P500-04-051, 2004.
- Grant, M. A. and Garg, S. K.: Recovery factor for EGS, Proc, 37th Workshop on Geothermal Reservoir Engineering, Stanford University, 2012.
- Grant, M. A.: Geothermal resource proving criteria, Proceedings, World Geothermal Congress, 2581–2584, 2000.
- Lawless, J. and Lovelock, B.: New Zealand’s geothermal resource, paper presented at NZ Geothermal Association, 2002.
- Lund, J. W., Gawell, K., Boyd, T. L., and Jennejohn, D.: The United States of America country update 2010, World Geothermal Congress paper 0102, 2010.
- MIT: The future of geothermal energy, 2006.
- Muffler, L. P. J.: Assessment of geothermal resources of the United States – 1978, U.S. Geological Survey, Circular 790, 163p., 1978.
- Nathenson, M.: Physical factors determining the fraction of stored energy recoverable from hydrothermal convection systems and conduction-dominated areas, USGS Open-file report 75–525, 1975.
- Pálmason, G., Johnsen, G. V., Torfason, H., Sæmundsson, Ragnars, K., Haraldsson, G. I., and Halldórsson, G. K.: Mat á Jarðvarma Íslands, (in Icelandic) Orkustofnun report OS-85076/JHD-10, 1985.
- Sanyal, S. K., Henneberger, R. C., Klein, C. W., and Decker, R. W.: A methodology for the assessment of geothermal energy reserves associated with volcanic systems, Transactions, Geothermal Resources Council 26, 59–64, 2002.
- Sanyal, S. K., Klein, C. W., Lovekin, J. W., and Henneberger, R. C.: National assessment of U.S. geothermal resources – a perspective, Transactions, Geothermal Resources Council 28, 355–362, 2004.
- Sanyal, S. K. and Butler, S. J.: An analysis of power generation prospects from Enhanced geothermal systems Transactions, Geothermal Resources Council, 29, 131–137, 2005.
- Sarmiento, Z. F. and Björnsson, G.: Reliability of early modelling studies for high-temperature reservoirs in Iceland and The Philippines, Proceedings, 32nd Workshop on Geothermal Reservoir Engineering, Stanford University, 2007.
- Stefansson, V.: World geothermal assessment, Proceedings, World Geothermal Congress paper 0001, 2005.
- White, D. E. and Williams, D. L.: Assessment of geothermal resources of the United States – 1975, Geological Survey Circular 726, USGS, 1975.
- Williams, C. F.: Development of revised techniques for assessing geothermal resources, Proceedings, 29th Workshop on Geothermal Reservoir Engineering, Stanford University, 276–280, 2004.
- Williams, C. F.: Updated methods for estimating recovery factors for geothermal resources, Proceedings, 32nd Workshop on geothermal reservoir engineering, Stanford University, 2007.
- Williams, C. F.: Thermal energy recovery from enhanced geothermal systems – evaluating the potential from deep, high-temperature resources Proceedings, 35th Workshop on geothermal reservoir engineering, Stanford University, 2010.
- Wilmarth, M. and Stimac, J.: Worldwide power density review Proc., 39th Workshop on geothermal reservoir engineering, Stanford University, 2014.