Geothermal Resources Council *Transactions*, Vol 22, September 20-23, 1998

Sustainable Geothermal Power: The Life Cycle of a Geothermal Field

James W. Lovekin
Geothermex, Inc.

ABSTRACT

Some proponents of geothermal energy have described this energy source as renewable, but many geothermal fields show declines in output as exploitation proceeds. On the other hand, those who would call geothermal a depletable energy source have to explain how some mature fields are able to produce with negligible declines and no apparent limit to the amount of recoverable energy.

The confusion arises out of the attempt to describe geothermal resources using inappropriate conceptual models of how resources behave. This paper presents a conceptual model that is better suited to describe the expected performance of a geothermal field over its entire life cycle. The paper also describes a set of terms that can be used to quantify geothermal resources and to guide plans for development.

Resource Life Cycles

People expect different things from different types of resources. Figure 1 shows simplified conceptual models of life cycles for three resource types that have been used as analogs for geothermal resources in the past.

Figure 1. Life Cycles of Three Resource Types

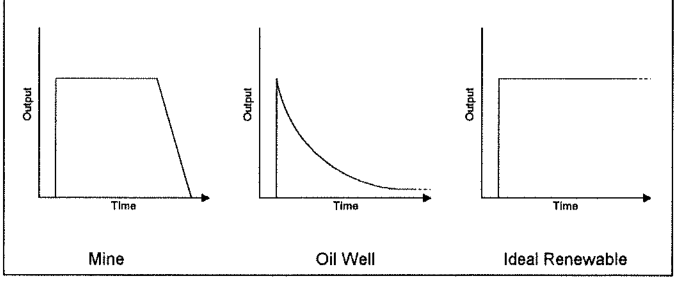

In the mine model, the developer delineates a body of ore in advance and builds an infrastructure to extract the ore at a steady rate over the life of the project. When the ore is exhausted, production drops off rapidly, and the mine shuts down. The key elements of this conceptual model are: (1) the resource is considered from the outset to be finite (although it may not be fully delineated when production begins); (2) there is a high threshold of capital investment up front (including a mill to process the ore) before any of the ore can be brought to market; (3) once underway, production proceeds at a steady pace (assuming no dramatic change in market conditions); and (4) the project has a definite and predictable end.

In the model for an oil well, the developer drills the well and produces it at the maximum possible rate. The well's output declines until it reaches an economic limit, and the well is abandoned. Like the mine model, the model for an oil well presumes from the start that the amount of recoverable resource is finite. The models differ somewhat in the threshold amount of initial capital: while a discovery well may be expensive (including the costs of exploration and earlier dry holes), the ratio of investment to revenue is lower for subsequent wells, and the developer usually does not have to build a new refinery to bring the oil to market. The most striking difference between the two models is the tapering decline in the oil well's output, driven by the response of the reservoir itself rather than the developer's specification of an optimal rate. The end of the life cycle for an oil well is a little fuzzier: if operating costs are low enough and the price of oil is favorable, the well may sustain flow at a low rate ("stripper production") for many years. But ultimately, the developer expects that the well will be shut in and there will not be enough oil left in the field to justify drilling new wells.

The model for an ideal renewable resource is simple: the developer builds a facility which operates at a steady rate forever. People commonly idealize hydroelectric, wind, and solar facilities in this way. The key elements of this conceptual model are: (1) the energy recoverable from the resource is infinite; (2) the threshold investment varies with the resource type, but is manageable; (3) there are no declines; and (4) the project has no end. Of course, no facility exists that fully meets this ideal. Dams silt up, wind turbines wear out, solar panels degrade. Still, the renewable model captures something special about these resources.

They draw on vast and elemental forces: the rain, the wind, the sun. With routine maintenance of the facilities to exploit them, these resources can be considered practically renewable on the time scale of human endeavors.

The Life Cycle Of A Geothermal Field

None of the foregoing conceptual models is quite adequate to describe geothermal resources. The best conceptual model for geothermal resources is a hybrid of all three. Figure 2 illustrates the life cycle of a geothermal field. This idealization assumes an unlimited market for geothermal energy at a profitable price -- a condition that has probably never yet occurred. Because of market limitations and relatively short operating histories, most existing fields have shown only portions of this overall cycle. But since planning for geothermal development typically assumes profitable prices over some time span, a model that describes expected field performance under such conditions should be useful. At least, it can provide a terminology to assist in formulating plans -- and in reacting to market changes.

Figure 2. Life Cycle of a Geothermal Field

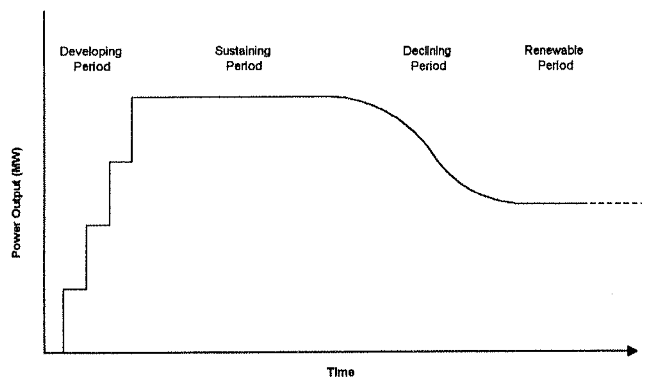

One can divide the life cycle of a geothermal field into four periods: developing, sustaining, declining, and renewable. In the developing period, increments of plant capacity come on line in steps. This is not an explicit feature of any of the three resource types discussed above (though it may be present), but it is the norm for a geothermal development. Like a mine, a geothermal development has a high threshold of initial investment: it needs not only wells but a power plant or a direct-use facility to generate revenue. However, the degree of certainty that can be achieved in defining the resource before going on line is typically lower for a geothermal development than for a mine. Even when a geothermal developer has mapped out a heat anomaly by drilling and has demonstrated production from some of the wells, the long-term output of the project will depend on how the wells interact with each other and with the surrounding hydrologic system. These interactions will determine such critical characteristics as production decline rates, injection breakthrough, and changes in non-condensible gas content. One can conduct tests in advance to estimate these characteristics, but at some point developers typically find it prudent (and more profitable) to install some increment of the plant capacity they are ultimately hoping for. This allows them to see how the reservoir actually responds and greatly enhances the degree of certainty in planning for subsequent increments of plant capacity.

In the sustaining period, the output of the geothermal field remains steady over an extended period of time. The level of output is fixed by the capacity of the plant, which generally corresponds to contractual commitments. This portion of the cycle is very similar to the sustained output of a mine once its infrastructure is in place. It is also similar (possibly deceptively so) to the start of output from a renewable resource. Even under a condition of unlimited demand for geothermal energy, one would not plan to produce a geothermal field on the model of an oil well (that is, starting at the maximum possible rate and beginning an immediate decline). The high threshold of initial investment provides a natural incentive to the geothermal developer to install only the plant capacity that can be sustained -- one would not normally want to pay for additional plant capacity that could be utilized only briefly.

Typically, a developer will drill enough wells at the start to have a surplus of capacity at the wellhead. This allows for anticipated declines in well productivity and provides operating flexibility in case individual wells need workovers. Initially, the constant output of the field is sustained by this surplus capacity at the wellhead. Later, when the initial wells have declined, constant output is sustained by drilling make-up wells. Unless decline rates prove to be very low, drilling make-up wells is an expected part of geothermal development and is normally provided for in the initial financing of the project.

Figure 3. Performance of an Individual Geothermal Well

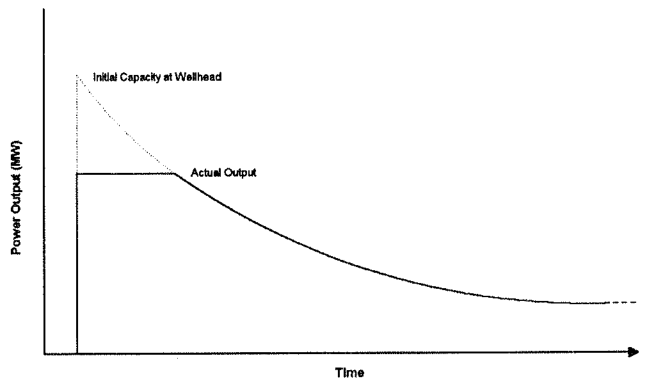

Figure 3 shows the expected performance of an individual geothermal well. It is very similar to the conceptual model of an oil well, except that the geothermal well initially produces at less than its maximum capacity because it shares the load at the plant with other wells. The situation is analogous to production from an oil well constrained by a regulatory limit (an "allowable") or by limited pipeline capacity. Even though the full capacity of the geothermal well is not used initially, this capacity usually declines with time as pressures in the reservoir are drawn down. Eventually, the capacity of the well declines to a point at which the well's full output is needed. The capacity and the actual output of the well become the same, and both follow the path of tapering decline.

Figure 4 illustrates how an individual well fits in with other wells in a geothermal wellfield. The actual output of the wellfield is fixed by the plant capacity, substantially less than the combined initial capacity at the wellhead. In addition, there is often a level of surplus capacity required at the wellhead by contractual commitment. When the combined capacity of the wellfield declines to the level of the required surplus, the developer drills make-up wells. The actual output of the plant is sustained at a constant level until the combined capacity of the initial wells plus the make-up wells declines to that level.

Figure 4. Performance of Geothermal Wellfield

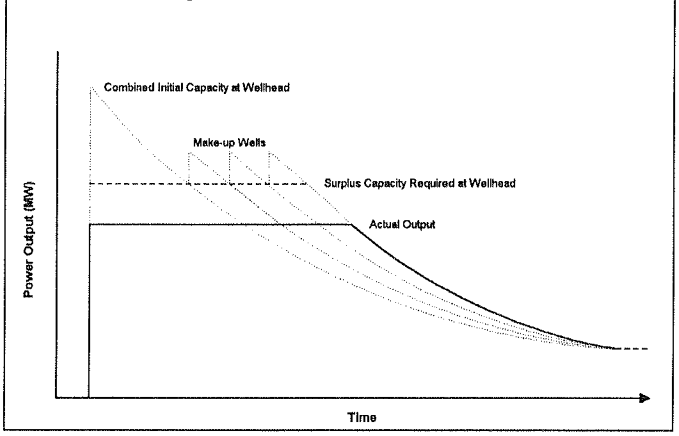

Then starts the declining period in the life cycle of the field. Even under the condition of unlimited demand for geothermal energy, it eventually becomes uneconomic to drill additional make-up wells because of interference between the wells. (If initial declines are very low and no make-up wells are needed, then in a condition of unlimited demand the developer would install more plant capacity -- to a point at which declines are significant and make-up wells become necessary.) The declining period begins when the underlying declines in the capacity of individual wells become evident in the actual output of the entire facility. The field makes a transition from the steady output of a mine to the tapering decline of an oil well. Those who might have prematurely applied the conceptual model of a renewable resource (and perhaps a few over-eager developers) are naturally disappointed. But the onset of declining output does not represent a failure of the project. It is an expected and predictable phase in the geothermal life cycle. A developer who understands this can structure the project's finances so that, by the time the declines begin, the initial investment is already repaid and the project is running profitably with just routine costs for operations and maintenance (O&M). If the declines are gradual and O&M costs do not escalate too rapidly, the field can continue to turn a profit for many years.

As declines continue, the net withdrawal of mass from the reservoir (production minus injection) decreases. Eventually, the field may reach a state in which net mass withdrawals are matched by natural recharge and the amount of extractable heat within the reservoir is still abundant. If operations are still profitable at this level of output, the field enters the fourth phase of its life cycle: the renewable period. At last, the field fits the model of a renewable resource. Given enough time, there is virtually no limit to the amount of energy it can produce. Declines are negligible. Production can continue at this level indefinitely.

Unfortunately, not all fields make it to the renewable period. The condition of unlimited demand for geothermal energy at a profitable price does not mean that the developer of a particular field can raise prices above a level competitive with other fields. A price that is profitable early in a field's life may not be adequate when output falls, due to the loss of economies of scale. The level of output supported by natural recharge may be below the economic limit. But the concept of a renewable period allows for the possibility of long-term, steady production and accommodates the observed longevity of some actual fields.

Quantification of Geothermal Resources

The conceptual model of a geothermal life cycle provides a frame of reference for comparing different fields. It helps illustrate the meaning of several terms that are commonly used (and sometimes misused) to describe how big geothermal fields are.

For a depletable resource such as a mine or an oil well, the term **reserves** describes the amount of a resource one expects to recover through the end of a project's life. On a forecast of output vs time, reserves are the area under the curve. In the example of a mine,

$$\begin{matrix}\text{tons of ore}\\ \text{per year}\end{matrix} \;X\; \begin{matrix}\text{projected years}\\ \text{of operation}\end{matrix} \;=\; \begin{matrix}\text{reserves}\\ \text{in tons.}\end{matrix}$$

For an open-ended, renewable resource, the concept of reserves can be misleading. Given enough time, a photovoltaic cell and a hydroelectric dam can both produce an unlimited amount of energy. To compare reserves in a meaningful way, one has to specify a **project life**. For geothermal fields, the project life for each increment of plant capacity has traditionally been 30 years -- for reasons that have more to do with the expected life of the plant equipment than with the projected performance of the reservoir. Thus, a geothermal facility that is forecast to have a steady output of 100 megawatts (100 MW) for 30 years has reserves of 3,000 MW-years.

The usage of the term **reserves** described above is generally compatible with the terminology set forth by Muffler and Cataldi (1978), which describes "reserve" as that portion of a geothermal resource that has been identified and "that can be extracted legally at a cost competitive with other energy sources at the time of determination." In the terminology of Muffler and Cataldi, recovery of geothermal energy is contemplated to take place "in an industrial time frame (10-100 yr)." The usage described here makes explicit the idea that recovery is to take place within a specified project life.

Another important term in the quantification of geothermal resources is **capacity**. This expresses the rate at which energy is produced, and it is expressed in units of power (for example, MW). The size of a geothermal power plant is specified in terms of its capacity. Similarly, the term capacity can describe the potential output of any other power-generating component of a geothermal facility -- such as an individual well or an entire reservoir.

The units of reserves and capacity are worth stressing, because they are often a source of confusion. A consumer buying electricity pays for a certain amount of energy each month, typically expressed in kilowatt-hours. Energy can be expressed in units of power X time. Since reserves for a geothermal field represent an amount of energy, they are expressed in these terms -- for example, kilowatt-hours or MW-years. (The SI unit for energy -- the joule -- could also be used, but units of power X time are customary.) Since capacity is a rate of energy production, it is expressed in units of power -- for example, kilowatt-hours per hour, or simply, kilowatts.

The developer of a geothermal power plant is concerned with both reserves and capacity, because contracts for the sale of electricity customarily have both energy and capacity components; that is, the developer is paid not just for the amount of energy which a facility produces but also for the capacity of the facility (the rate at which it produces energy). Similarly, the developer of a direct-use facility is typically paid for the amount of thermal energy provided and must maintain enough capacity to meet the thermal load. However, neither reserves nor capacity explicitly address how long a geothermal field can sustain a given level of output. A reserve number by itself says nothing about the timing of production -- a field that starts at 150 MW and ramps down evenly to 50 MW over 30 years has identical reserves (but very different economics) compared to a field that operates for 30 years at a steady output of 100 MW. A capacity number by itself is just a snap-shot -- it expresses the power output at a single point in time. When a developer is deciding how big a geothermal facility should be, it is crucial to know the **sustainable capacity** of the field -- that is, the capacity which the field can sustain for a specified period of time.

Categories Of Reserves

As discussed earlier, the steady output of a geothermal field during the sustaining period results first from using the spare capacity of initial wells and later from drilling make-up wells. Because the make-up wells represent future decisions to expend significant amounts of capital to maintain output, the reserves they develop are sometimes distinguished from the reserves developed by initial wells. As shown in Figure 5, total reserves can be represented as the area under the output curve from the start of a project to the end of the specified project life. These reserves can all be considered **proved** if they can be forecast with sufficient certainty (customarily, with a probability of 90% or greater). By analogy to reserve definitions adopted by the Society of Petroleum Engineers and the World Petroleum Congress (1997), one can classify reserves supported by existing wells as **proved developed reserves** and reserves supported by make-up wells yet to be drilled as **proved undeveloped reserves**.

Figure 5. Geothermal Reserves Determined by Specified Project Life

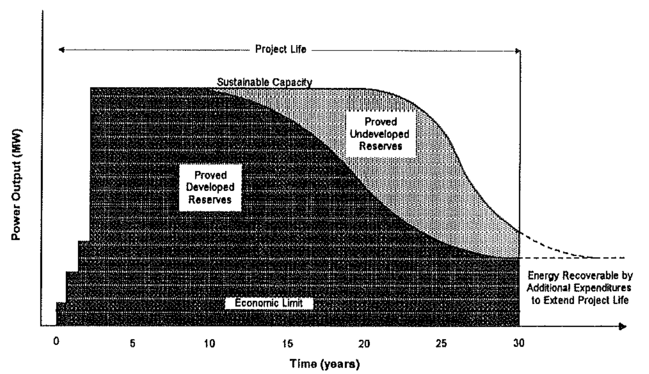

Figure 5 illustrates two other points as well. First, the declining period may start before the end of the specified project life. In a 30-year project, the sustainable capacity for which a power plant is sized may only be intended to apply for, say, the first 20 years. Of course, a developer may decide to size the plant to maintain constant output for the entire project life. But it is quite common at the time of initial financing for a developer to schedule no make-up wells in the project's latter years. Because of the time-value of money, the present value of such wells and the reserves they might develop have a relatively small impact on the project's economics. The developer can always re-visit the decision on whether to maintain constant production by additional drilling or by other capital expenditures. This decision would reflect the market conditions at the time and a better understanding of reservoir performance based on years of operating history. If the developer decides to allow output to decline, the project would still be considered successful as long as it has provided an acceptable return on investment.

Second, Figure 5 illustrates that, beyond the project life, there may be production of additional energy that would not at present be considered reserves. Recovering this energy may eventually entail just continuing routine O&M, with no additional capital expenditures. Still, energy beyond the project life is not considered reserves because of the inherent uncertainties of market conditions and reservoir performance. As these things become better known over time, the developer typically revises the forecast of project life and updates the estimate of reserves.

In Figure 5, the economic limit (that is, the minimum output at which the facility can operate profitably) remains consistently low throughout the project life. Figure 6 illustrates the potential impact of a change in market conditions; in this case, a step upward in the economic limit,

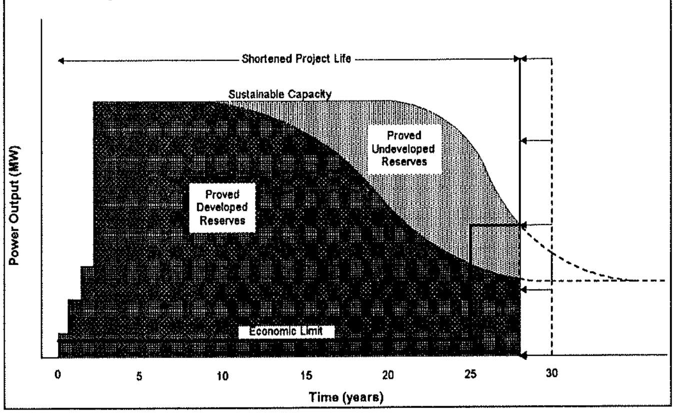

Figure 6. Geothermal Reserves Determined by Economic Limit

either due to a decrease in the price of energy produced or an increase in operating costs. If the developer has foreseen this change, it should be reflected in the original estimate of project life, and reserves should be unaffected. If a rise in the economic limit occurs for unexpected reasons, the developer may have to shorten the estimate of project life and decrease the estimate of reserves.

Conclusion

Even though few geothermal fields have gone through an entire life cycle, the conceptual model of this cycle provides a frame of reference for comparisons between fields. With this model in mind, we can avoid the pitfall of comparing, say, the sustainable capacity of one field with the instantaneous initial capacity of another. Knowing that declines are expected can help avoid overblown claims and consequent disappointments. Knowing that sustainable (even renewable) output is indeed possible can help promote geothermal energy development and maintain our enthusiasm in the face of market adversity.

Acknowledgments

The ideas in this paper have evolved in numerous spirited discussions with my colleagues at CalEnergy and GeothermEx, as well as in the industry at large. A workshop on sustainability convened by the GRC in March 1994 helped me to gel the ideas. The members of the GRC Technical Standards Committee have helped me to hone the definitions of terms. I express my thanks and appreciation to the many individuals who have participated in these various forums -- you know who you are.

References

Muffler, P., and Cataldi, R. (1978), "Methods for regional assessment of geothermal resources," *Geothermics*, Vol. 7, pp.53-89.

Society of Petroleum Engineers and World Petroleum Congress (1997), "Petroleum reserve definitions," *SPE Publication*, March 1997.