

Optimization of Well Positions for SAGD

Jason A. McLennan, Weishan Ren, Oy Leuangthong, and C. V. Deutsch

Centre for Computational Geostatistics
Department of Civil & Environmental Engineering
University of Alberta

The application of steam-assisted gravity drainage (SAGD) to recover heavy oilsands is becoming increasingly important in the northern Alberta McMurray Formation because of the vast resources/reserves accessible with this production mechanism. Choosing the vertical locations of SAGD well pairs is a vital decision to be made for reservoir evaluation and planning. The inherent uncertainty in the distribution of geological variables should be an integral part of this decision. Geostatistical simulation is used to capture geological uncertainty. This geological uncertainty is used to determine a distribution of the best possible well pair locations.

A simulated annealing algorithm is developed and described for determining the optimum stratigraphic locations of SAGD well pairs. The objective is to maximize recovery R. There are three basic steps to the methodology: (1) model the uncertainty in the top continuous bitumen (TCB) and bottom continuous bitumen (BCB) surfaces, (2) formulate the optimization problem and constraints within a simulated annealing framework, and (3) solve the optimization problem using simulated annealing. The methodology is described and implemented on a subset of data from the Athabasca Oilsands in Fort McMurray, Alberta, Canada.

Background

Located in Northern Alberta, the Athabasca oilsands deposit spans 40,000 square kilometers and contains over one trillion barrels of original bitumen-in-place (Denbina, 1998). In the past thirty years, oilsands have increased from 2% to 30% of Canada's total annual oil production. Surface mining up to depths of approximately 130 meters, are currently extracting and producing approximately one quarter of this production; however, a mere 10% of the Athabasca oil reserve is located sufficiently close to the surface to allow the continued use of economical surface mining methods. The demand for innovative in-situ oil sands extraction technology to recover the deeper oil sands is high (Edmunds, 1999).

In 1978, Dr. Roger Butler, introduced the concept of Steam assisted gravity drainage. When SAGD was accepted by industry and government as a commercial in-situ recovery process, the Alberta Energy and Utilities Board (AEUB) reported a 22.7 billion cubic meter or 4000% increase in their Alberta bitumen reserves statement.

Figure 1 illustrates the SAGD concept. The procedure is applied to horizontal, parallel and vertically aligned well pairs. The well pair length and vertical separation are on the order of 1-2 kilometers and 5 meters, respectively. The upper well is known as the "injection well" and the lower well is known as the "production well". The SAGD process begins by circulating steam in both wells to establish a thermally connected steam chamber anchor. When the bitumen in between the well pair is heated, steam circulation through the production well is stopped and

steam is injected into the upper injection well only. A cone shaped steam chamber, anchored at the production well, then begins to develop upwards from the injection well encountering more and colder reservoir material. As new bitumen surfaces are heated, the oil lowers in viscosity and flows downward along the steam chamber boundary into the production well by way of gravity.

Problem

There are several reservoir parameters and operational decisions that affect the overall connected reservoir recovery. However, the spatial distribution of geological reservoir variables, reservoir connectivity, and positioning of SAGD well pairs has the most impact on recovery efficiency (McLennan, 2004). For example, even if all other reservoir and operational parameters are optimal, a non-optimal SAGD well pair position according to the geological heterogeneity can result in devastatingly low R , which could ultimately lead to total well pair failure.

Geological uncertainty is an unavoidable reality for any reservoir recovery project. And since production performance is significantly related to reservoir geology, SAGD production performance is also always uncertain. Geostatistical simulation provides a model of geological uncertainty through multiple realizations of geological variables such as facies type, porosity, water saturation, and permeability. These geological realizations can be used to calculate various production performance measures by way of transfer functions such as flow simulation and the optimization algorithm proposed in this work.

The top and bottom continuous bitumen (TCB and BCB) surfaces are useful summaries of geological heterogeneity for SAGD reservoirs. These surfaces enclose the maximum amount of potential recoverable SAGD reserves. Multiple realization pairs of these surfaces can be created based on the previously constructed model of geological uncertainty by applying reservoir specific cutoffs to combinations of geological variables. The suite of TCB / BCB surface pairs summarizes the geological uncertainty and is all that is required for optimizing the SAGD well elevations.

This work describes a SAGD well pair optimization algorithm given several realizations of the TCB / BCB surface. In particular, we answer the question: *At what stratigraphic elevation should one drill the horizontal production well of a SAGD producer in order to maximize the ultimate recovery R of the SAGD process?* In practice, the aerial SAGD well pair locations are fixed. That is, multiple horizontal well pairs are drilled from a drilling pad in a pattern so that the suite of multiple SAGD well pairs will be able to deliver steam to the entire reservoir volume, see Figure 2. The well pairs are usually offset by a constant aerial distance d . The proposed problem does not involve optimizing the aerial positions of well pairs.

Although the areal SAGD well pair locations are fixed, their vertical locations are not – we can optimize them. Consider the A-A' cross-sectional view in Figure 2 through one of the SAGD well length trajectories. The TCB and BCB surfaces and the horizontal producer well path are shown. At stratigraphic elevations in between the TCB and BCB surfaces, bitumen in the reservoir is connected, which means the cumulative steam chamber can reach and produce the bitumen within this region. If the producer well is located at an elevation in between the TCB and BCB surface, bitumen above the well will be produced and bitumen below the well will not. The latter reserves contribution is referred to as the producible reserves from the effective producer well length. On the other hand, if the producer well is located below the BCB surface, only a portion of the bitumen in between the TCB and BCB surfaces can be produced (from the cumulative effect of adjacent SAGD well pairs). This reserves contribution is referred to as the

producible reserves from the non-effective producer well length. The optimum producer well elevation then is the one that maximizes the sum (over the well length trajectory) of producible reserves from the effective and non-effective well lengths. Once the optimum producer well elevation is established, the injection well location must be 5m above and parallel along the trajectory of the producer well.

The stratigraphic optimization process is repeated for each realization of the TCB / BCB surfaces pair. The uncertainty in this optimum elevation can then be assessed by, for example, the elevations that correspond to a low, medium, and high R case. The optimization algorithm used can be classified as a simulated annealing method.

Simulated Annealing

The technique of simulated annealing is based on an analogy with the physical process of annealing and is typically applied to global optimization problems. Annealing is the process by which a material undergoes extended heating and is slowly cooled. During cooling, thermal vibrations permit a reordering of the molecules to a highly ordered or low energy state lattice (Kirkpatrick, 1989). The *annealing* process can be simulated using the following steps:

1. Generate a SAGD producer well elevation (analogous to the initial alloy) by drawing a random elevation within the range of BCB elevations along the well length trajectory.
2. Define an energy or objective function (analogous to the Gibbs free energy) as a measure of the difference between the maximum recovery possible and the actual recovery.
3. Perturb the model (analogous to the thermal vibrations in true annealing) by drawing another random producer well elevation.
4. Accept the perturbation (thermal vibration) if the objective function decreases; reject it if the “energy” has increased.
5. Continue perturbing the model until a low objective function (energy) is achieved.

Some details of the proposed methodology are now given.

Methodology

There are three basic steps to the methodology: (1) model the TCB and BCB surfaces by applying cutoff criteria to multiple geostatistical realizations, (2) formulate the optimization problem and constraints within a simulated annealing framework, and then (3) solve the optimization problem using simulated annealing. The details for the last two steps are presented in this work.

The first step is to create a 3-D numerical uncertainty model representing the heterogeneity of several geological variables such as facies type, porosity, permeability, and water saturation. Geostatistical simulation algorithms are almost always used for this purpose. These methods honor the original conditioning well data and allow access to local and joint uncertainty (Deutsch, 2002). From the 3-D geological realizations of geological properties, multiple TCB / BCB surface pairs can be calculated according to some cutoff criteria. For example, at a particular aerial location, the BCB elevation could be taken as the lowest elevation within the corresponding column of model cells, above which, the permeability exceeds a minimum permeability cutoff 10mD along a minimum connected bitumen thickness of 8m. Certainly, several combinations of

petrophysical property cutoff criteria could be formulated depending on the practical application. The result is multiple realizations of the 2-D TCB and BCB surface elevations:

$$TCB_{(x_i, y_j)}^l \quad x_i = 1, \dots, X; y_j = 1, \dots, Y; l = 1, \dots, L$$

$$BCB_{(x_i, y_j)}^l \quad x_i = 1, \dots, X; y_j = 1, \dots, Y; l = 1, \dots, L$$

There $l = 1, \dots, L$ TCB / BCB surface realization pairs. The locations $x_i, i = 1, \dots, X$ and $y_j, j = 1, \dots, Y$ represent the aerial easting and northing grid cell locations of the reservoir. At stratigraphic elevations in between the BCB and TCB surfaces, bitumen in the reservoir is connected. This means that the cumulative steam chamber from a SAGD horizontal well pair will reach and produce the bitumen within this region of the reservoir. Outside either of these surfaces, there is no bitumen or the bitumen is interrupted by some impermeable facies such as shale or mudstone and the steam chamber is unable to sweep these regions of the reservoir.

There are four basic components to formulating the optimization problem: (1) identify the set of data, (2) identify the set of variables involved in the problem, together with their domains of definition, (3) identify the set of constraints that defines the set of feasible solutions, and (4) identify the function to be optimized (Castillo, 2002). We now describe each of these essential components of an optimization problem as they apply to optimizing a SAGD producer well elevation location.

Given the drilling pad location and orientation, the azimuth α (positive angle from true north) angle and length l of a SAGD well pair is essentially fixed. For example, in Figure 2, the south-east SAGD well pair is orientated at an azimuth of 90° and has a length approximately 50 m less than half the reservoir width. The α and l data define the trajectory of the SAGD well pair. The trajectory is then used to extract the L TCB and BCB surfaces along the well path:

$$TCB_{(u_i, v_j)}^l \quad u_i = 1, \dots, U; v_j = 1, \dots, V; l = 1, \dots, L$$

$$BCB_{(u_i, v_j)}^l \quad u_i = 1, \dots, U; v_j = 1, \dots, V; l = 1, \dots, L$$

Here, the locations $u_i, i = 1, \dots, U$ and $v_j, j = 1, \dots, V$ represent the easting and northing aerial grid cell locations along the well trajectory. Obviously, the TCB surface must always be greater than the BCB surface. Note, Figure 2 shows schematically the TCB and BCB surfaces along the south-east SAGD well pair trajectory. Later, when we formulate the constraints, we will see that the mean and variance of the BCB surface elevations are required. These are straightforward to calculate as:

$$m_{BCB}^l = \frac{\sum_{i=1}^U \sum_{j=1}^V BCB_{(u_i, v_j)}^l}{U \cdot V} \quad l = 1, \dots, L$$

$$\sigma_{BCB}^{l2} = \frac{\sum_{i=1}^U \sum_{j=1}^V \left(BCB_{(u_i, v_j)}^l \right)^2}{U \cdot V} - \left(m_{BCB}^l \right)^2 \quad l = 1, \dots, L$$

A single constraint is used in the model to limit the feasible stratigraphic domain for the optimum Z_{WELL} location. First, we assume that the BCB surface along the SAGD well trajectory behaves as

a normal distribution. We can fully define this Gaussian distribution with the BCB mean m_{BCB} and variance σ_{BCB}^2 (defined above). The constraint imposed on the Z_{WELL} optimization variable is that it must lie within 3 standard deviations of the mean value:

$$\left(m_{BCB}^l - 3 \cdot \sigma_{BCB}^l\right) \leq Z_{WELL} \leq \left(m_{BCB}^l + 3 \cdot \sigma_{BCB}^l\right)$$

Figure 3 illustrates this constraint for the BCB surface from Figure 2. In the figure, the particular Z_{WELL} value shown is coincident with the mean m_{BCB} BCB surface elevation. By formulating this single constraint, several other constraints are avoided. All the practical variable constraints are all included in the above constraint equation. For example, we do not need to state that Z_{WELL} must be less than the TCB surface along the well trajectory.

Within a simulated annealing framework, the objective function is the minimization of the difference between the maximum potential recovery and the actual recovery R given a candidate producer well elevation Z_{WELL} . The maximum potential thickness recovery is simply the sum (over the well length trajectory) of the TCB less the BCB surface elevations. The actual recovery is the sum of the TCB surface elevations less the candidate producer well elevation Z_{WELL} over the effective well length) and 80% of the maximum potential thickness recovery (for the non-effective well length).

We can see from the objective function that we require two indicator functions. The utility of the first indicator a_{EFF} is to subtract 100% of the producible bitumen from the effective well length and the utility of the second indicator b_{N-EFF} is to subtract 80% of the producible bitumen from the non-effective well length. Also, the objective function, as formulated within the framework of a simulated annealing process above, is equivalent to maximizing SAGD recovery R , which was the original problem. The following is a summary statement of the optimization problem:

Minimize (the objective function):

$$O = \sum_{i=1}^U \sum_{j=1}^V \left[TCB_{(u_i, v_j)}^l - BCB_{(u_i, v_j)}^l - a_{EFF} \left(TCB_{(u_i, v_j)}^l - Z_{WELL} \right) - b_{N-EFF} \left(TCB_{(u_i, v_j)}^l - BCB_{(u_i, v_j)}^l \right) \right]$$

$$O = \sum_{i=1}^U \sum_{j=1}^V \left[(1 - b_{N-EFF}) \left(TCB_{(u_i, v_j)}^l - BCB_{(u_i, v_j)}^l \right) - a_{EFF} \left(TCB_{(u_i, v_j)}^l - Z_{WELL} \right) \right]$$

$$\text{where... } a_{EFF} = \begin{cases} 1 & \text{if } Z_{WELL} \geq BCB_{(u_i, v_j)}^l \\ 0 & \text{otherwise} \end{cases} \quad ; \quad \text{and } b_{N-EFF} = \begin{cases} 0.8 & \text{if } Z_{WELL} \leq BCB_{(u_i, v_j)}^l \\ 0 & \text{otherwise} \end{cases}$$

[by perturbing] The Variable: Z_{well}

[subject to] Constraints:

$$\left(m_{BCB}^l - 3 \cdot \sigma_{BCB}^l\right) \leq Z_{WELL} \leq \left(m_{BCB}^l + 3 \cdot \sigma_{BCB}^l\right)$$

$$\text{where } m_{BCB}^l = \frac{\sum_{i=1}^U \sum_{j=1}^V BCB_{(u_i, v_j)}^l}{U \cdot V} \quad ; \quad \sigma_{BCB}^{l2} = \frac{\sum_{i=1}^U \sum_{j=1}^V \left(BCB_{(u_i, v_j)}^l \right)^2}{U \cdot V} - \left(m_{BCB}^l \right)^2$$

[given] Data:

$$TCB_{(u_i, v_j)} \quad u_i = 1, \dots, U; v_j = 1, \dots, V$$

$$BCB_{(u_i, v_j)} \quad u_i = 1, \dots, U; v_j = 1, \dots, V$$

The steps required to solve the above optimization problem within a simulated annealing framework are now presented:

1. Draw a uniform random number and standardize it to represent a z -score value between -3 and +3 on a standard normal distribution, z .
2. Calculate the non-standard normal SAGD producer well elevation Z_{WELL} as:

$$Z_{WELL} = m_{BCB} + z \cdot \sigma_{BCB}^2$$

3. Calculate an initial value of the objective function O .
4. Repeat 1, 2 for a different Z_{WELL} (perturbation) and calculate the associated value of the objective function O .
5. Accept the new Z_{WELL} perturbation if the objective function value O decreases; reject the new Z_{WELL} otherwise.
6. Repeat steps 1, 2, and 5 until a maximum number of iterations it_{MAX} have been performed or the objective function value O has stabilized at a minimum within some tolerance ϵ .
7. Repeat for L realizations

A simple example is now presented to illustrate the procedure.

Implementation

We now apply the methodology to calculate the optimum SAGD producer well elevation for a particular aerial SAGD well pair. The data used in this example are a subset of coreholes extracted and re-formatted from a real dataset within the McMurray Formation. Twenty vertical exploration coreholes are positioned onto a 400 x 400m grid, see Figure 4. There are 4 coreholes in the easting direction and 5 coreholes in the northing direction spanning an area 1600m east-west by 2000m north-south.

A 3D geological model of porosity, permeability and water saturation uncertainty is already available and was created using geostatistical simulation. There are 100 realizations of the geology. Figure 5 shows the 3D model of the first realization of porosity. The BCB bitumen surface is taken as the lowest elevation where a column of at least 8m of continuous net bitumen resides immediately above. Net bitumen is calculated as being above a porosity cutoff of 12%, above a permeability cutoff of 10 mD, and below a water saturation cutoff of 85%. And the TCB surface is calculated as the highest elevation where a column of at least 4m of continuous net bitumen resides immediately below. Figures 6 and 7 show a histogram and aerial color-scale map of the first TCB and BCB realizations.

The aerial location of the SAGD well pair is fixed by an azimuth angle $\alpha = 140^\circ$ and length $l = 1500$ m. Figure 8 shows the aerial location and trajectory of the SAGD well pair with the α and l parameters superimposed on the average BCB surface. A realistic well pad orientation is also

included. We now extract the first realization TCB and BCB surface pair. Figure 9 illustrates these surfaces along the 1500 m well trajectory from Figure 8. Note the vertical scale is the same in both plots. There are 751 TCB and BCB elevations along the well trajectory.

We now look at the distribution of the first BCB surface realization elevations in particular. Figure 10 is a histogram of the BCB realization surface. The mean m_{BCB} and variance σ_{BCB}^2 of this distribution is 202.33 and 9.0, respectively. To find the optimum SAGD producer well elevation Z_{WELL} we assume this distribution is normal and constrain the stratigraphic feasible region of Z_{WELL} to be within 3 standard deviations of the mean:

$$\begin{aligned} (202 - 3(9.0)) \leq Z_{WELL} \leq (202 + 3(9.0)) \\ 175 \leq Z_{WELL} \leq 229 \end{aligned}$$

The simulated annealing process is now implemented to find the optimum Z_{WELL} position that minimizes the difference between total possible recovery and actual recovery, or equivalently, maximizes the actual recovery R . The objective function has units of thickness (m). The initial perturbation is set at $Z_{WELL} = m_{BCB} = 202.33$. The initial value of the objective function is 5,207m. The annealing process continues until 1000 perturbations have been imposed. Figure 11 shows all 1000 objective function tries corresponding to all 1000 Z_{WELL} perturbations. As well, Figure 11 shows the objective function value at each iteration. We see that the objective function reaches its minimum energy state at a value of 4,094 m in under 400 iterations. This objective function value corresponds to a SAGD producer well elevation of $Z_{WELL} = 197.5$ m. For confirmation that we have maximized actual recovery R , a plot of the recovery in percent is plotted vs. elevation in Figure 12. Indeed, a maximum recovery of $R = 88\%$ occurs at the 197.5 m Z_{WELL} elevation. Figure 13 shows the optimum SAGD producer well location superimposed on the BCB surface along the well trajectory.

The optimum SAGD well elevation is found over 100 TCB / BCB surfaces created from the 100 geological realizations. Figure 14 shows the uncertainty of the 100 optimum elevations as a histogram. The optimum elevation with the lowest, medium, and highest recoveries are all 88%. In this case, there is little uncertainty in the maximum recovery; however, other reservoirs may involve higher uncertainty.

Discussion and Conclusion

The simulated annealing optimization algorithm is a brute-force and greedy way of getting an optimal solution. Usually, a fairly complex objective function with more dimensionality and variables is used with a high associated CPU time and resource cost. In our case, optimizing the SAGD producer well locations is a simple 1-D optimization problem. Nevertheless, this allows us to integrate more complexity into the optimization problem for more reservoir-specific scenarios.

We could consider several variations or modifications of the optimization problem. For example, we may want to simultaneously optimize several pairs of SAGD producer well pairs, each with a different orientation α and length l . In this case, we would need to add an additional constraint on α so that no well pair trajectories intersect at any length along their trajectory. We could also consider optimizing the aerial well pair positions. A dynamic constraint mechanism would need to be added so that there is no overlapping.

We optimized according to percent recovery only; however, there are other fundamental SAGD performance parameters, such as oil rate and steam-oil-ratio SOR. In a practical study, the

objective function may be to simultaneously maximize oil rate and minimize SOR. Definitely, this will affect the optimum producer well elevation results. For example, we assumed that we could get 80% production from non-effective well lengths due to the impact from adjacent horizontal wells. Although 80% may be an accurate figure, the oil rate and steam-oil ratio may be adversely affected by producing bitumen in this manner. In this case, the optimum Z_{WELL} position would increase above 197.5 m in Figure 13.

References

Akgun, F., 2002. *Drilling Stable Horizontal Wells – A Key Problem in Modern Oil Well Drilling*, Canadian International Petroleum Conference, June 2002, Calgary, Alberta, Canada.

Alberta Research Council (ARC) and Alberta Energy Research Institute (AERI), 2001. *Thermal Gravity Processes*, Power-point, AERI/ARC Core Industry Research Program.

Card, C.C., Woo, J. S., Wang, C. and Hu, Z., 1998. *CNPC Liaohe Dual Well SAGD – The Journey from Vision to Reality*, Canadian Institute of Mining and Metallurgy (CIM), no. 134.

Castillo, E., Conejo, A., Pedregal, P., Garcia, R., and Alguacil, N., 2002. *Building and Solving Mathematical Programming Models in Engineering and Science*. Wiley.

Denbina, E.S., 1998. *SAGD Comes of AGE!*, Journal of Canadian Petroleum Technology, 37(7):9-12.

Deutsch, C. V., Dembicki, E. and Yeung, K. C., 2002. *Geostatistical Determination of Production Uncertainty: Application to Firebag Project*, Center for Computational Geostatistics (CCG), University of Alberta, Edmonton, Alberta, Canada.

Edmunds, N.R. and Sugget, J.C., 1995. *Design of a Commercial SAGD Heavy Oil Project*, International Heavy Oil Symposium, June 1995, Calgary, Alberta, Canada.

Edmunds, N., 1999. *SAGD: Present and Future Challenges*, APEGGA Professional Development Conference, April 1999, Calgary, Alberta, Canada.

Kirkpatrick, S., Gelatt, C., and Vecchi, M., 1989. *Optimization by Simulated Annealing*. Science, Vol. 220, No. 4598, pp. 671-680.

Llaguno, P.E., Moreno, F., Garcia, R., Mendez, Z., Escobar, E., 2002. *A Reservoir Screening Methodology for SAGD Applications*, Canadian International Petroleum Conference, June 2002, Calgary, Alberta, Canada.

McLennan, J.A. and Deutsch, C.V., 2003. *Guide to SAGD (Steam Assisted Gravity Drainage) Reservoir Characterization Using Geostatistics*. Centre for Computational Geostatistics (CCG), Edmonton, Alberta.

Miller, R., 2000. *Optimization: Foundations and Applications*, Wiley.

Shangqi, L., Yongrong, G., Zhimian, H., Naiqun, Y., Liping, Z. and Suning, H., 1998. *Study on Steam Assisted Gravity Drainage with Horizontal Wells for Super-Heavy Crude Reservoir*, Canadian Institute of Mining and Metallurgy (CIM), no. 217.

Venkataraman, P., 1996. *Applied Optimization with Matlab Programming*




Figure 1: The Steam Assisted Gravity Drainage (SAGD) Process. A schematic illustration of the SAGD process. Steam concentration is depicted on the grayscale shown in the top left; the setting, including the horizontal well dimensions, is shown in the top right. The middle illustrates the initial stage of SAGD when steam is injected into both wells and the steam chamber begins to form. The bottom shows the overriding SAGD mechanics and geometry when steam is injected into the injection well only.




Figure 2: Optimizing the Vertical SAGD Producer Well Location. The top shows a typical aerial SAGD well pair pattern. These aerial well pair locations are fixed in practice. The bottom shows a cross section through the reservoir along the trajectory of one of the SAGD producer wells. The optimum producer well location is the elevation that maximizes the combined production from the effective and non-effective producer well length segments.

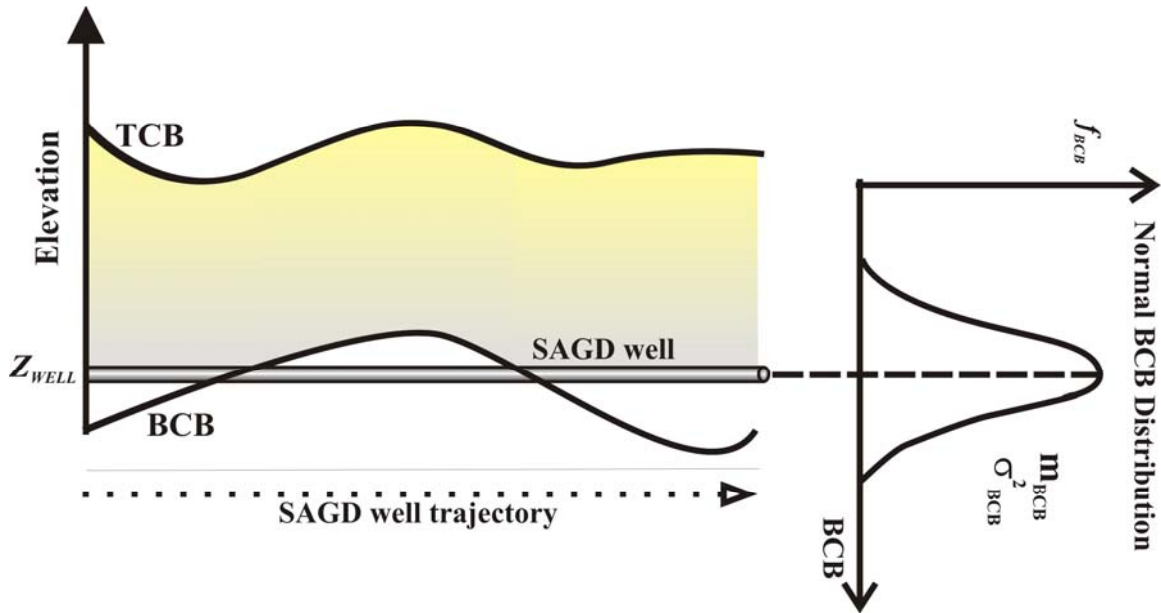


Figure 3: The Z_{WELL} constraint. An illustration of the assumption that the BCB surface elevations are normally distributed along the well trajectory. In the minimization of the objective function, Z_{WELL} is constrained to be within 3 standard deviations of the m_{BCB} value.




Figure 4: The Spatial Domain. An illustration of the corehole data configuration in plan view. There are 20 exactly vertical conditioning coreholes within a 2.5 x 2.0 km aerial reservoir domain.




Figure 5: Porosity Model. An easting, northing and two elevation cross sections through the first realization porosity model. The raw porosity data are also shown along the corehole traces. The porosity units are in %.

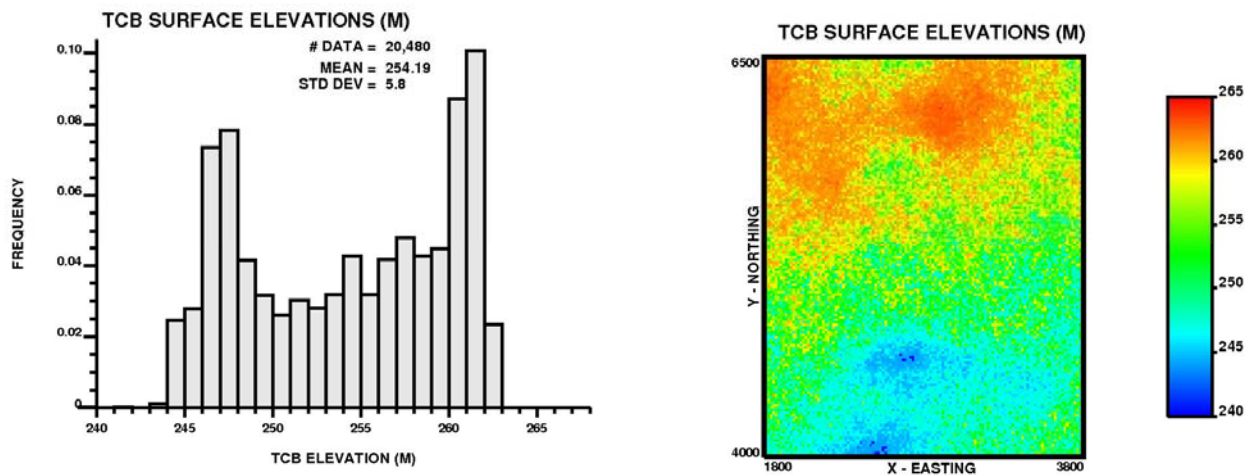


Figure 6: The TCB Surface. A histogram and aerial map of the first realization top continuous bitumen TCB surface elevations.

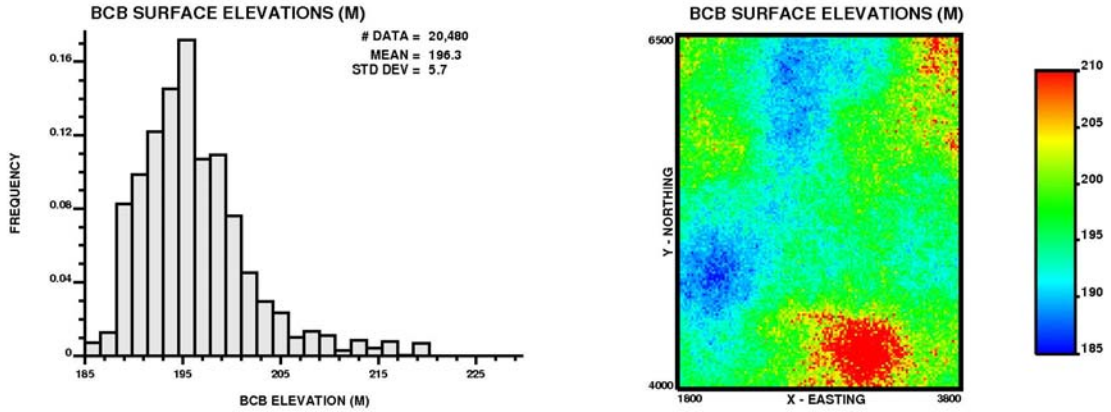


Figure 7: The BCB Surface. A histogram and aerial map of the first realization bottom continuous bitumen BCB surface elevations.

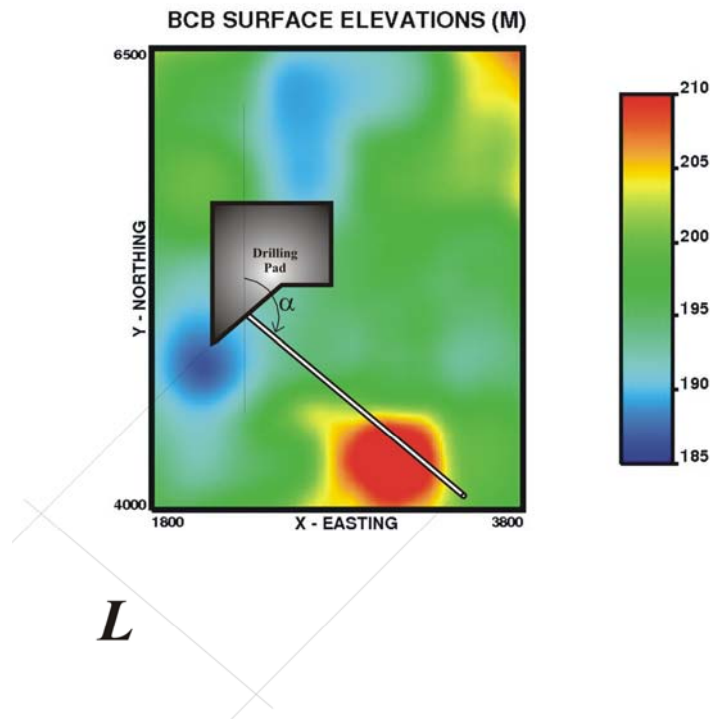


Figure 8: The Aerial SAGD Well Pair Location. An illustration of the aerial SAGD well location (superimposed on the BCB surface) considered in this example. The azimuth angle and length are 140° and 1500m, respectively.




Figure 9: The TCB and BCB Well Trajectory Surface. An illustration of the first TCB and BCB surface realizations along the trajectory of the well.

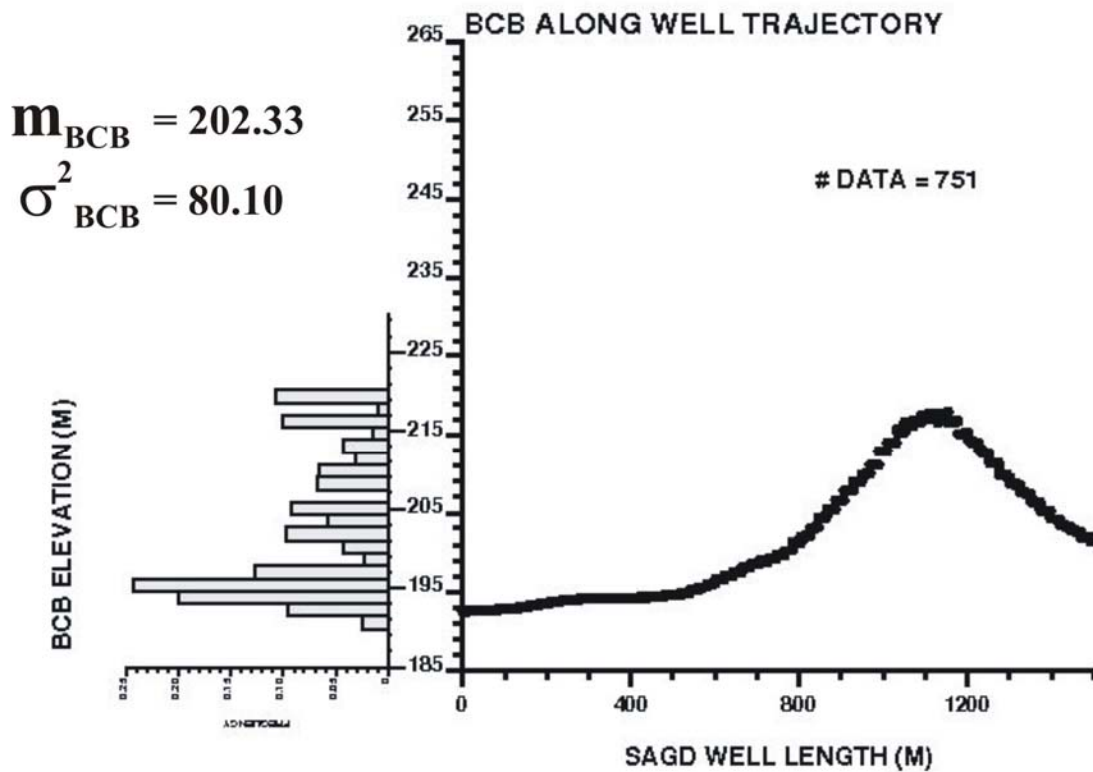


Figure 10: The BCB Well Trajectory. The histogram of the first realization BCB surface elevations encountered along the SAGD well path is shown on the left. The mean m_{BCB} and variance σ^2_{BCB} of this distribution are 202.33 and 80.10, respectively. The average BCB surface is shown on the right.




Figure 11: Simulated Annealing Results. The objective function value tries are shown with respect to elevation on the left; the objective function value by iteration is shown on the right. We observe an optimum SAGD producer well elevation at 197.5 m.

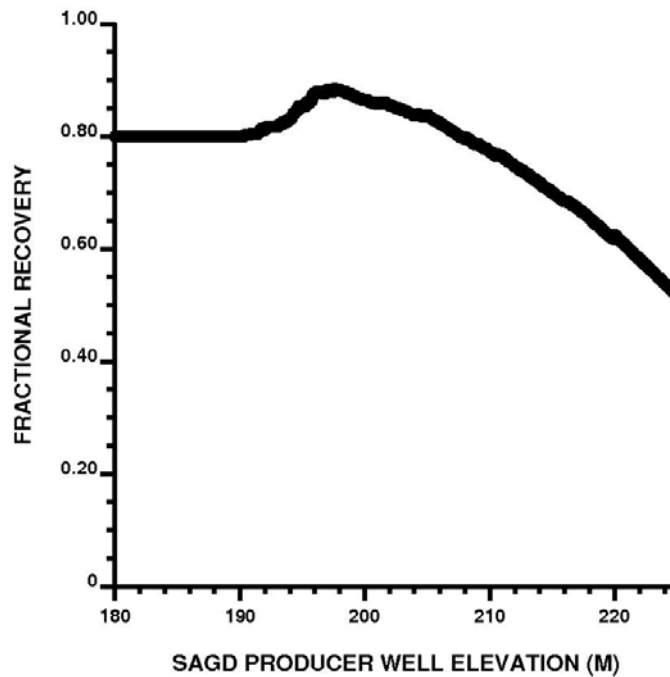


Figure 12: Does the Optimum Maximize Recovery R ? A plot of percent recovery versus SAGD producer well elevation confirms we have maximized recovery at a Z_{WELL} value of 197.5 m. The maximum recovery is $R = 88\%$.

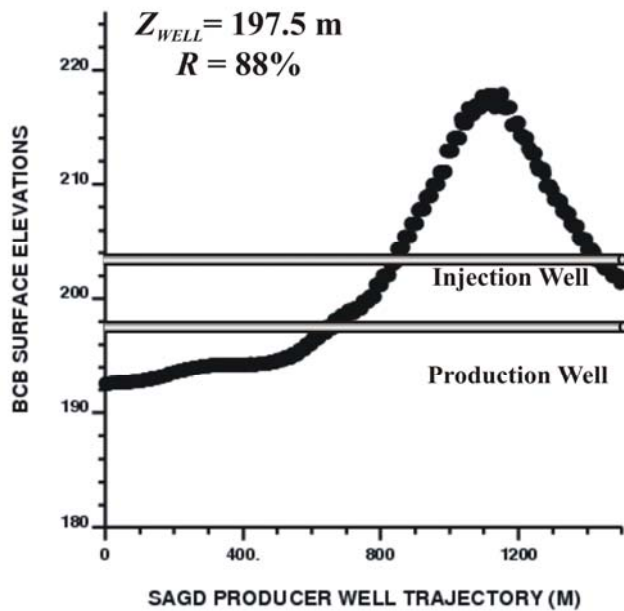


Figure 13: The Optimum Z_{WELL} . The optimum SAGD producer well elevation and trajectory is shown with respect to the BCB surface along the well trajectory. The optimum Z_{WELL} value is 197.5 for a maximum recovery of $R = 88\%$. The injection well is also shown.

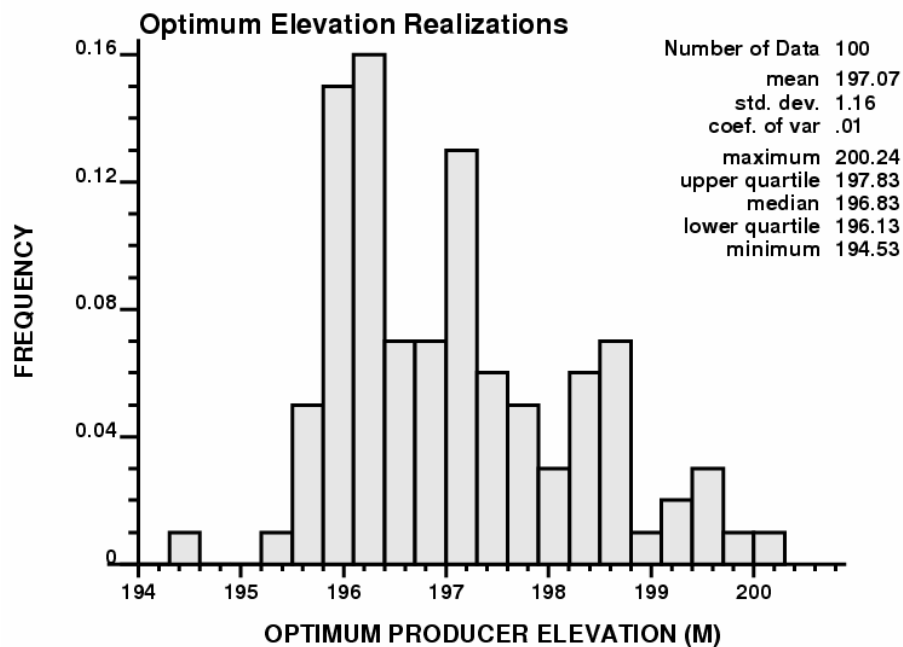


Figure 14: Uncertainty in optimum Z_{WELL} . The histogram of optimum SAGD well elevations over all 100 TCB / BCB realizations pairs.