

SPE/Petroleum Society of CIM/CHOA 78996

Automatic Determination of Well Placement Subject to Geostatistical and Economic Constraints

Karl P. Norrena, University of Alberta; Clayton V. Deutsch, SPE, University of Alberta

Copyright 2002, SPE ITOHOS/ICHWT conference.

This paper was prepared for presentation at the SPE International Thermal Operations and Heavy Oil Symposium and International Horizontal Well Technology Conference held in Calgary, Alberta, Canada, 4–7 November 2002.

This paper was selected for presentation by an SPE ITOHOS/ICHWT Program Committee following review of information contained in an abstract submitted by the author(s). Contents of the paper, as presented, have not been reviewed by the Society of Petroleum Engineers, the Petroleum Society of CIM/CHOA and are subject to correction by the author(s). The material, as presented, does not necessarily reflect any position of the Society of Petroleum Engineers, the Petroleum Society of CIM/CHOA, its officers, or members. Papers presented at SPE ITOHOS/ICHWT meetings are subject to publication review by Editorial Committees of the Society of Petroleum Engineers. Electronic reproduction, distribution, or storage of any part of this paper for commercial purposes without the written consent of the Society of Petroleum Engineers is prohibited. Permission to reproduce in print is restricted to an abstract of not more than 300 words; illustrations may not be copied. The abstract must contain conspicuous acknowledgment of where and by whom the paper was presented. Write Librarian, SPE, P.O. Box 833836, Richardson, TX 75083-3836, U.S.A., fax 01-972-952-9435.

Abstract

Optimal well placement is a complex problem that requires detailed models of the reservoir structure/geometry and the petrophysical properties such as facies, porosity, permeability, and fluid saturation. The reservoir development team attempts to integrate all of these aspects when devising a well plan for optimal reservoir exploitation. Ideally the well locations would be selected with the assistance of a flow simulator; however, this is impractical due to time and CPU requirements. This paper presents a technique for selecting optimal well locations for fine-tuning with a flow simulator. The technique constructs the well placement problem as an optimization problem to be solved with simulated annealing. The global objective function consists of multiple component objective functions. Each component represents a desirable feature or constraint in the problem. Optimality is defined as the best balance among the component objectives. The format of the technique is flexible and can incorporate 3-D geostatistical models of uncertainty and multiple constraints. The proposed method iteratively refines initial well locations and trajectories until the global objective is maximized. Several examples are shown. Optimal well placement in a steam assisted gravity drainage context is illustrated.

Introduction

One task of a reservoir development team is to set a well plan that, given all available information, is reasonable. A well plan is set with the help of a reservoir model. A reservoir model highlights candidate regions for well placement. This initial well plan is *static* because it does not account for the dynamics of fluid flow. The static well plan is adjusted to a dynamic well plan with the aid of a flow simulator. The

process is iterative. This process can be expensive in terms of professional and CPU time. Selecting good static well plans is important because it will reduce the iterations required and lead to better decisions.

Assembling a good static well plan is difficult due to heterogeneities and uncertainties in the subsurface reservoir parameters. Accounting for this information in a decision-making framework is the subject of this paper. Various approaches have been proposed including optimization techniques such as mixed integer programming and neural networks.^{1,2} Integer programming requires the objective function to be expressed as a linear function, and neural networks require training and a library of training images. One common technique for accounting for uncertainty involves selecting the P_{05} , P_{50} , P_{95} realizations and selecting a well plan that is jointly optimal on these using a flow simulator. This approach is impractical in most cases due to the number of possible well locations and the computational expense required for evaluating the realizations.

This paper proposes a technique for selecting a good static well plan. The problem is posed as an optimization problem for simulated annealing. Simulated annealing is an optimization routine particularly well suited to optimizing highly combinatorial problems such as the problem of selecting well locations. Consider the placement of two wells on a 2-D reservoir model on a 50 x 50 grid. If the reservoir development team were to exhaustively evaluate every location on the grid there are many combinations:

$$C_2^{2500} = \frac{2500 \cdot 2499}{2!} = 3123750$$

Selecting well locations with respect to uncertainty would require $L \cdot C_2^{2500}$ evaluations, where L is the number of realizations. In practice a reservoir model may have a grid size of millions of cells and hundreds of realizations. The combinatorial becomes incomprehensible in size.

Simulated annealing is a numerical analogy to the thermodynamic process of annealing. An objective function measures the extent of order in an initial solution and successively perturbs the decision variables in the objective function until the objective function is reduced. Simulated annealing accepts perturbations that decrease the objective function and probabilistically accepts those that increase the

objective function with progressively diminishing probability. Progressively diminishing the probability of selecting perturbations that increase the objective function is analogous to slow cooling in annealing. The probability distribution used is the Boltzmann distribution. There are several advantages to using simulated annealing: (1) the objective function need not be continuous, linear, or differentiable; (2) simulated annealing always converges to an optimal result; (3) simulated annealing can be made to converge quickly. However, the implementation of simulated annealing and requires some tradeoffs. In particular, the annealing schedule is difficult to set. There are empirical rules that are very useful for establishing an efficient and fast annealing schedule. There are many good references on the subject of simulated annealing (SA)^{3,4,5}.

Methodology

The idea is to: (1) construct an objective function that quantifies the performance of a well plan, (2) propose an initial well plan, (3) quantify the performance of the initial well plan, (4) perform a random perturbation to the well plan, (5) apply the simulated annealing decision rule and accept or reject the perturbation, (6) repeat steps 4 and 5 until the well plan optimizes the objective function.

Multiple realizations of porosity (ϕ), water saturation (s_w), facies type (f_i), top of reservoir, and the thickness of the reservoir are used. The top of reservoir and thickness models use a corner point grid topology and the other models use a center point grid topology. The proposed algorithm can select well locations for three different well paths: (1) vertical wells, (2) deviated wells, and segmented wells. Figure 1 shows an example reservoir model with the three well types superimposed.

The decision variables that are perturbed are the coordinates of the segments that make up the well paths in the well plan. The perturbation mechanism randomly selects a segment and moves it by a random $\Delta x, \Delta y, \Delta z$. Figure 2 illustrates the perturbation mechanism for the three well types. In the perturbation of a vertical well a single pair of random $\Delta x, \Delta y$ are drawn and two Δz values are drawn ($\Delta z_{top}, \Delta z_{bottom}$). A perturbation for a deviated well draws six random coordinates, and the segmented well draws $3n+3$ random coordinates, where n is the number of segments. The Perturbation mechanism is important in SA optimization problems. The perturbations must not be too drastic or most perturbations will not be accepted and convergence will be slow. The perturbations must not be too minor or many perturbations will be required to achieve convergence or the algorithm cannot jump out of a local minima. Standard practice is to choose a reasonable mechanism and any inefficiencies will be revealed in slow convergence. The algorithm coded here rarely takes more than a few minutes on a PC for convergence; thus, any inefficiency translates to acceptable CPU time.

The objective function is a measure of performance of the well plan. The aim is to satisfy several component objectives through a weighted linear combination of the component objectives:

$$O_g = \sum_{c=1}^C w_c O_c$$

where w_c is the weight for the component objective O_c . The weights ensure that each component objective contributes appropriately to the global objective and that issues related to scaling and units are avoided. The units for the global objective is set to dollars and simulated annealing will be used to find a well plan that maximizes profit. Any number of component objectives can be included in the global objective function. Three components are considered and one is proposed in this paper: (1) profitability, (2) alignment to a fixed drilling platform, (3) cost per unit well length, and (4) connected pore volume.

The profit component objective function is calculated in expected value over all L realizations. Calculating O_{profit} over all L realizations accounts for uncertainty and yields a well plan that is jointly optimal on all of the L realizations. One cannot determine L individual optimal well plans and use them to infer a single joint optimal well plan. Figure 3 illustrates the notion of the joint optimal solution with a reservoir model consisting of 5 realizations of porosity. The global objective is to select a single well location with high porosity and robust with respect to uncertainty. The stack of realizations on the left show the well location that maximizes porosity over L realizations, no other location yields higher porosity over the entire model of uncertainty. The realizations on the right have optimal well locations on each realization. There is no easy way to reconcile the L optimal locations into a single optimal location.

The profit component objective function is written as the sum of the by the N well profits (pw_i) over the L realizations:

$$O_{profit} = \sum_{l=1}^L \sum_{i=1}^N pw_i^l$$

The well profits are constrained by the intersected geo-body, the drainage radius, and accounting. A geo-body is a collection of communicating blocks that share a specified attribute. Only cells of the same geo-body as intersected by the well path and falling within the drainage radius are included in the by-well profit calculations. Blocks in overlapping drainage radii are accounted for once by assigning them to the closest well. The by-well profits are calculated over the entire reservoir (V) as:

$$pw_i^l = \sum_V i_g^l(\mathbf{u}, \mathbf{u}') \cdot i_d(\mathbf{u}, \mathbf{u}') \cdot i_w(\mathbf{u}) \cdot profit^l(\mathbf{u})$$

where i_g^l is a binary variable indicating whether or not the cell at location \mathbf{u} belongs to the same geobody as the cell intersected by the well path at location \mathbf{u}' on realization l :

$$i_g^l(\mathbf{u}, \mathbf{u}') = \begin{cases} 1, & \text{if } geobody^l(\mathbf{u}) = geobody^l(\mathbf{u}') \\ 0, & \text{otherwise} \end{cases}$$

and i_d is a binary variable indicating whether or not the cell at location \mathbf{u} falls within the drainage radius centered at location \mathbf{u}' :

$$i_d(\mathbf{u}, \mathbf{u}') = \begin{cases} 1, & \text{if } \sqrt{(\mathbf{u} - \mathbf{u}')^2} \leq \text{drainage radius} \\ 0, & \text{otherwise} \end{cases}$$

and i_w is a binary variable indicating whether or not the cell has been accounted for by another well:

$$i_w(\mathbf{u}) = \begin{cases} 1, & \text{if the cell is not taken} \\ 0, & \text{otherwise} \end{cases}$$

and $profit^l$ is:

$$profit^l(\mathbf{u}) = \phi^l(\mathbf{u}) \cdot (1 - s_w^l(\mathbf{u})) \cdot thickness^l \cdot p_h$$

with $thickness$ being the thickness of the cell centered at \mathbf{u} and p_h is the specified price per unit volume of hydrocarbon. This formulation of profit accounts for the interaction between wells and the lithofacies model. Each cell is attributed to one well. The lithofacies constraint is accounted for by only considering cells that have the same geobody type as the cell that is intersected by the well. The lithofacies model could be constructed using any technique. The weight w_{profit} is set to be equal to 1 since the units are in terms of dollars.

A well plan may be constrained by the location of the drilling rig and/or the desire to drill multiple wells from a single location. A well plan that maximizes O_{profit} may be expensive or impossible to drill from a fixed platform or pad location. One goal is to find the well plan that strikes the best balance between O_{align} and O_{profit} . Figure 4 shows wells with good and bad well alignment. The maps show two wells and their associated drainage radii (the dark lines surrounded by a shaded area) in a hypothetical reservoir. The wells in the left map have good alignment to the drilling platform, the wells on the right would require highly deviated wells with tight turning radii. Figure 5 shows one metric for measuring well alignment; the sum of the Euclidean distance (d_e) and perpendicular distance (d_p) between the start of the well path and platform location for the N wells:

$$O_{align} = \sum_{i=1}^N d_{e_i} + d_{p_i}$$

The units for objective O_{align} are not dollars. The weight w_{align} is used to convert to dollars. High w_{align} values put more emphasis on well alignment than low values. Professional judgment is required to selecting w_{align} .

The component objective function O_{length} is the sum all the well lengths (l_w) in the well plan multiplied by a cost per unit length:

$$O_{length} = \sum_{i=1}^N l_{w_i} \cdot p_l$$

The proposed global objective function is flexible enough to permit new component objective functions to be added easily. Permeability is not considered in the present objective function. Permeability is often correlated to porosity and the connectivity of porous reservoir captures a lot of the information contained in permeability.

Examples

Two synthetic reservoirs have been prepared to illustrate the proposed technique. The first example is a small reservoir with a single source of uncertainty. The second example has uncertainty in all input variables. The maps show the well path(s) as a dark line and the drainage radii as shaded areas surrounding the well path. The lithofacies geobodies are shown as shaded objects on a white background.

Example 1

Consider a 50x50 grid with 10 realizations of porosity. Water saturation is held constant at 20%. The top of the reservoir is flat and the thickness is constant. The geobody model was created from a set of normally distributed unconditional realizations of porosity built using sequential Gaussian simulation. The realizations are converted to geobodies by identifying all communicating nodes that fall above a threshold value. An example conversion is shown in Figure 6. The map on the left is a realization generated by sequential Gaussian simulation. The associated geobody map is shown on the right. The net-to-gross ratio is about 0.3

Figure 7 shows a horizontal well in a single realization. The objective is to maximize profit given a single well; the component objectives of well alignment and well length were not used in this first example. The drainage radius was set to 5 units. The maps show the progression of the optimization starting at the initial well plan (top left) to the final optimized well plan (bottom right). The number of perturbations is indicated at the top of each map.

Figure 8 shows optimal well plans for 1 to 3 wells on the same map. The cost per unit length component objective function was used; note that the well paths are of minimal length but maximize profit.

Figure 9 shows an example with 3 wells with the complete global objective function. For comparison the map on the left does not include the component objectives of length and alignment. The map on the right shows the results using the complete global objective function. The location of the drilling platform/pad is shown as a large dot. The wells on the map on the right are well aligned with the locations of the drilling platform at the expense of giving up some of the reservoir on the upper right.

An optimal well plan on 10 realizations for 3 vertical wells is shown in Figure 10. The well locations are shown as black dots surrounded by a shaded drainage radius. The component objectives of length and alignment do not apply here. Note that the well positions are the same on each map. This is analogous to the stack of realizations on the left side of Figure 3. One could find realization-specific optimal well locations that out perform the joint optimal solution, but those locations would not outperform the joint optimal solution over multiple realizations.

The optimal well plan for a single well is not a subset of the well plan for 2 or more wells. This means that we cannot work backwards from the n wells well plan to solve for the $n-1$ well plan nor work forward to solve for the $n+1$ well plan; each well plan represents a unique problem.

None of the examples took more than a few seconds to optimize, even with 10,000 perturbations. The technique also works in 3-D but visualization is difficult.

The initial well plan could be selected automatically or provided by a professional. The initial well plan has minimal impact on the final result, but could affect the number of perturbations required to reach optimality.

Example 2

This reservoir is a 2-D cross section. The reservoir consists of fluvial channels. The sources of uncertainty are the top surface, the thickness of the reservoir, porosity, and the locations and size of the sand bed forms. The stratigraphic correlation is proportional. The crest of the reservoir is, on average, 500m below surface, and the limbs average 650m below surface. The cells in the grid measure 50m in the x direction. The average thickness of the reservoir is 100m. There are 100 nodes in x and z . There are two facies types: channel sand and shale. The distribution of porosity in the channel sand is shown in Figure 11. The average is 20% with a standard deviation of 7.

The water saturation model is constant (15%) above OWC and 100% below. Figure 12 shows the sectional view of a single realization of the reservoir model. The darkest shade is below OWC and has an s_w of 1.0. The shapes with the lightest color represent channel beds.

Figure 13 shows the optimal well plan over 50 realizations for a single vertical well having a drainage radius of 500m and a global objective of maximizing profit only. Below the map is a chart showing the profit for every location in the reservoir. The dashes are the locations that SA sampled. The optimal location is at $x=1875$. Figure 14 is a histogram of the sample locations that SA visited over 10,000 perturbations. Note that every feasible location was sampled and that the region with the optimal location was heavily sampled.

Application to a SAGD Process

The steam assisted gravity drainage (SAGD) process is an important method for recovering bitumen. The SAGD process consists of a pair of parallel wells: a steam injector above a recovery well, see Figure 15. The injected steam increases the mobility of the bitumen within a drainage block. Gravity draws the bitumen down into the recovery well.

Placing the well pairs requires consideration of remnant shale lenses. Due to sparse delineation wells and imperfect seismic data the locations of the shale cannot be known with certainty. Shale above the producer will reduce the performance of the well pair. The optimal vertical well placement is one that minimizes the probability of shale impeding performance. Moreover, the optimal well pair location maximizes recoverable reserves; bitumen below the recovery well cannot be recovered since the process relies on gravity.

A SAGD process consists of a “network” of well pairs drilled from a central pad location, see

Figure 16. The problem of selecting a pad location must also be considered: each well is subject to individual constraints such as those discussed above and joint constraints such as maximizing total production. We do not tackle this problem; we address the problem of selecting an optimal elevation for a single well pair.

Our example considers 20 exploration core holes on a regular 400m x 400m grid. Data were re-gridded to a fine 10cm resolution to preserve facies boundaries. Facies and porosity models are available for a 32x40x190 grid with regular blocks of size 50x50x0.5m. A subset of this model was used to construct a high-resolution model in the region of a single well pair, see Figure 17. The high-resolution model has 151x50x95 nodes measuring 1x20x1m.

The objective is to select an optimal vertical position for a well pair located in the middle ($x=76$) of the model. 3-D visualization software is very useful to understand the variability of the model and to assist in locating the wells; however, it becomes difficult to apply visual criteria when there are many geostatistical realizations to consider. We will use 20 realizations for our example.

Central XY, XZ, and YZ slices are shown on Figures 18, 19, and 20. Three figures are used to illustrate the variability

(i) the central slice of the first realization, (ii) the average of the central slice over 20 realizations, and (iii) the average of all slices over all realizations. Low porosity is shown by light gray scale; high porosity is dark gray scale.

The objective function is to maximize connected pore volume above the producer well location. The pore volume above the producer will be greater when the well is positioned low in the reservoir; however, the McMurray formation has more shale at the bottom and a producer positioned too low will recover less on account of the greater fraction of shale.

The connected pore volume is calculated with an ad hoc procedure, see Figure 21. Each 2-D XZ slice of each realization is considered independently. Slices are only accounted for if the well location considered intersects net reservoir. The connectivity of each block to the producer well is based on the amount of net reservoir that is encountered in a straight line of sight. The greater the fraction of net reservoir the more likely the bitumen will be produced. A multiplicative factor is considered as a function of that net fraction, see Figure 22. There is no discount when the net fraction is 1 and the factor is zero when there is no net. Three different penalty functions are used to assess the sensitivity of the final placement to this input parameter.

Figure 23 shows the results for the three penalty curves. The weighted pore volume must decrease as the penalty becomes more severe. Note, however, that the optimal location of the horizontal well is about the same in all three cases: Linear = 24m, Moderate = 24m, and Severe = 23m. The curves on Figure 23 are averages over all 20 realizations being considered. The curve denoted "Total Above" indicates the total pore volume (not penalized) above a well location whether or not non-net reservoir is intersected.

Figure 24 shows the resource above the producer well (the light line) and the total resource above the well for the first three realizations. Note that each realization yields a different optimal well location: realization 3 gives an optimal well location of 20m (38177.5 m³), realization 12 gives an optimal location of 29m (33282.7m³), and realization 16 gives an optimal location of 24m (32749.9m³). There is no clear way to reconcile the three realization specific optimal well locations into a single optimal well location, we must determine the optimal jointly.

This overall optimal location is at 25 m from the base of the model. A histogram of the 20 connected pore volumes for this location is shown at the center of Figure 25. We also show the histograms for placing the 5m higher and 5m lower. By construction, the optimal well location is the best on average although any particular realization is different.

In practice, the true reservoir is unknown and we can only determine the best producer location *in expected value*. The range of results on Figure 25 tells us the uncertainty in the

connected resource. Selected realizations could be processed through a flow simulator to evaluate the variability in production. The connected pore volume we calculate for placement could be used as a ranking measure to choose a limited number of realizations for flow simulation.

Concluding Remarks

Making the best possible decisions in the face of uncertainty is an important problem. The overall idea is to determine the best positions *in expected value* over realizations that quantify geological uncertainty. Clearly, the reasonableness of the results depends on how well the realizations reproduce all available geological, geophysical, and engineering data.

The well positions we generate are good starting points to be refined with flow simulation and consideration of other engineering constraints. It would be interesting to incorporate a fast flow simulator in the optimization problem.

Acknowledgements

We thank Jason McLennan and other researchers at the Centre for Computational Geostatistics for the input SAGD example which is part of a manual on SAGD reservoir characterization using geostatistics.

References

1. Vasantharajan, S. and Cullick, A. S.: "Well Site Selection Using Integer Programming," *Proceedings of IAMG'97 CIMNE* (1993) 1 421
2. Doraisamy, D.: "Methods of Neuro-Simulations for Field Development," *Society of Petroleum Engineers* (1997) SPE 39962
3. Kirkpatrick, S., Gelatt, C. D. and Vecchi, M. P.: "Optimization by Simulated Annealing" *Science*, (May 1983) 220(4598) 671
4. Ouenes, A., Bhagavan, S., Bunge, P. H. and Travis, B. J.: "Application of Simulated Annealing and Other Global Optimization Methods to Reservoir Description: Myths and Realities," SPE 28415 (September 1994)
5. Deutsch, C. V. and Cockerham, P. W.: "Practical Considerations in the Application of Simulated Annealing to Stochastic Simulation," *Mathematical Geology* (1994) 26(1)
6. Deutsch, C. V. and Journel, A. G.: *GSLIB: Geostatistical Software Library and User's Guide*, second edition Oxford University Press, NY (1998) 369
7. Beckner, B. L. and Song, X.: "Field Development Planning Using Simulated Annealing - Optimal Economic Well Scheduling and Placement," SPE 30650 (1995)
8. Dembicki, E.A., Yeung, K.C., and Deutsch, C.V., "Application of 3-D Geostatistical Models to Characterize and Simulate an Oil Sands Reservoir using the Steam Assisted Gravity Drainage (SAGD) recovery process in northeastern Alberta, Canada," AAPG/SPE Conference, California, June 2000

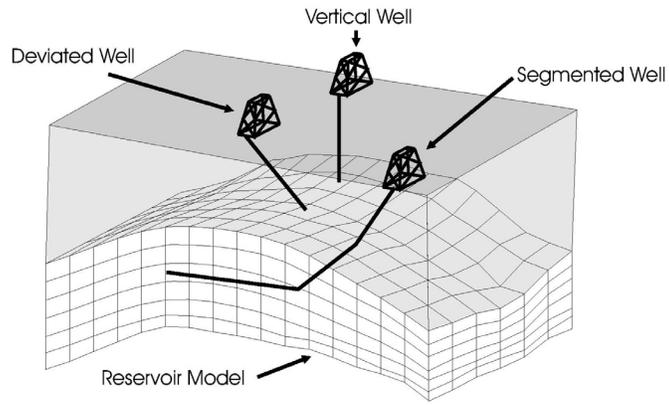


Figure 1: An example reservoir model showing the three different well types considered in the paper.

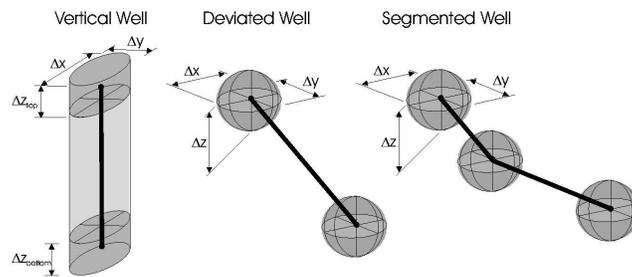


Figure 2: The perturbation mechanism for each of the 3 well types.

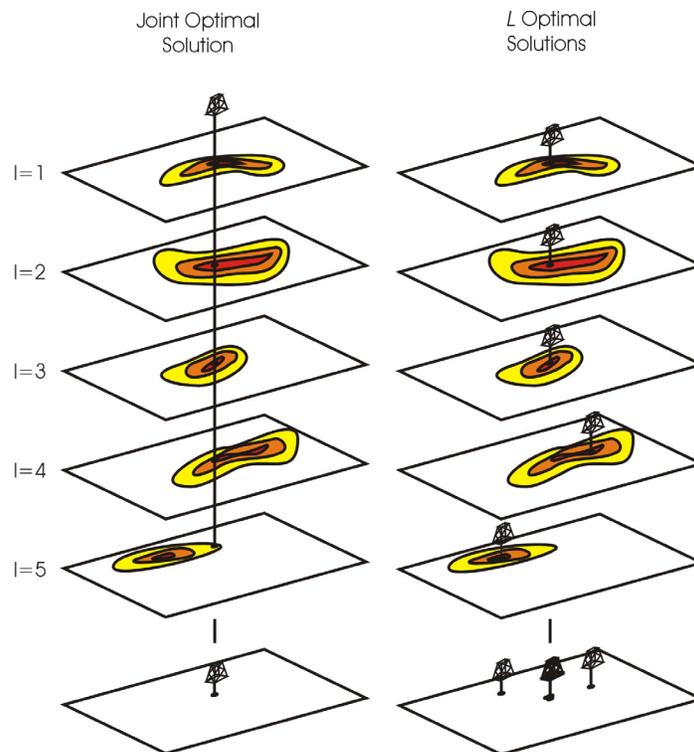


Figure 3: Optimization over the entire model of uncertainty. Note that there are 5 wells on the bottom right map: 3 are in the middle and nearly overlap.

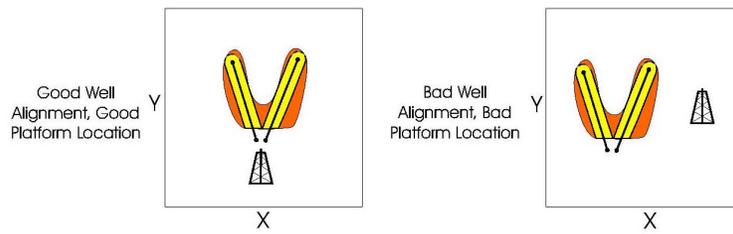


Figure 4: The map on the left shows two wells that have good alignment to the drilling platform, the map on the right shows two wells with bad well alignment.

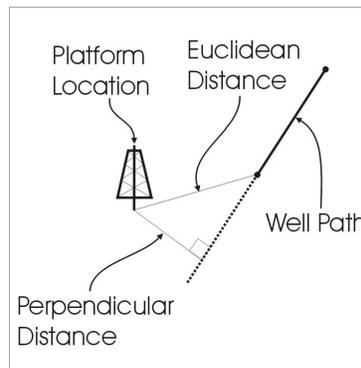


Figure 5: One metric of well alignment is the sum of the perpendicular and Euclidean distance between the location of the drilling platform and the well path.

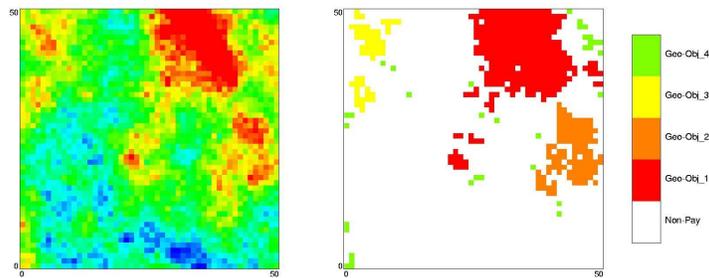


Figure 6: An example geobody conversion.

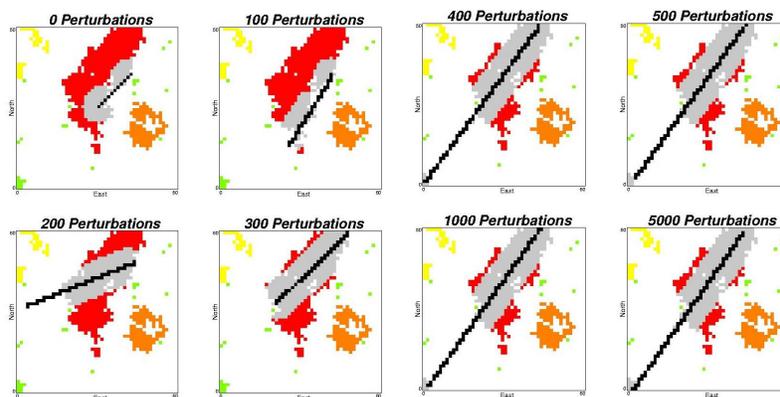


Figure 7: A 2-D example showing the progression from the initial guess (top left) to the final optimized horizontal well path (bottom right).

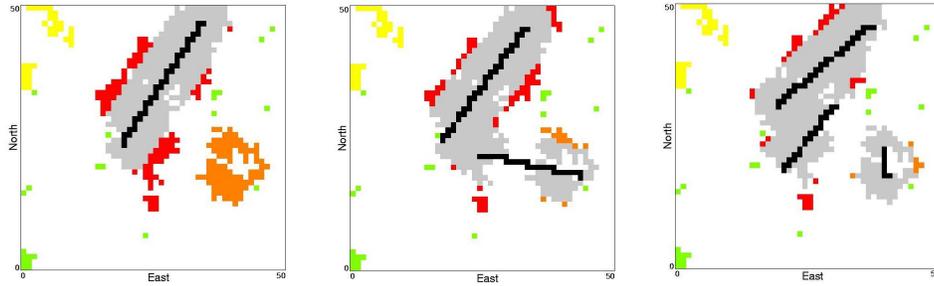


Figure 8: An example showing 1, 2, and 3 wells using the cost per unit length component objective function.

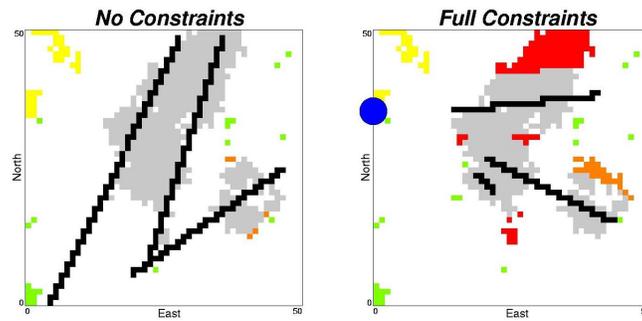


Figure 9: An example well plan with three wells. The map on the left has no constraints, the map on the right has full constraints. The drilling platform location is denoted by a large dot on the left side of the map.

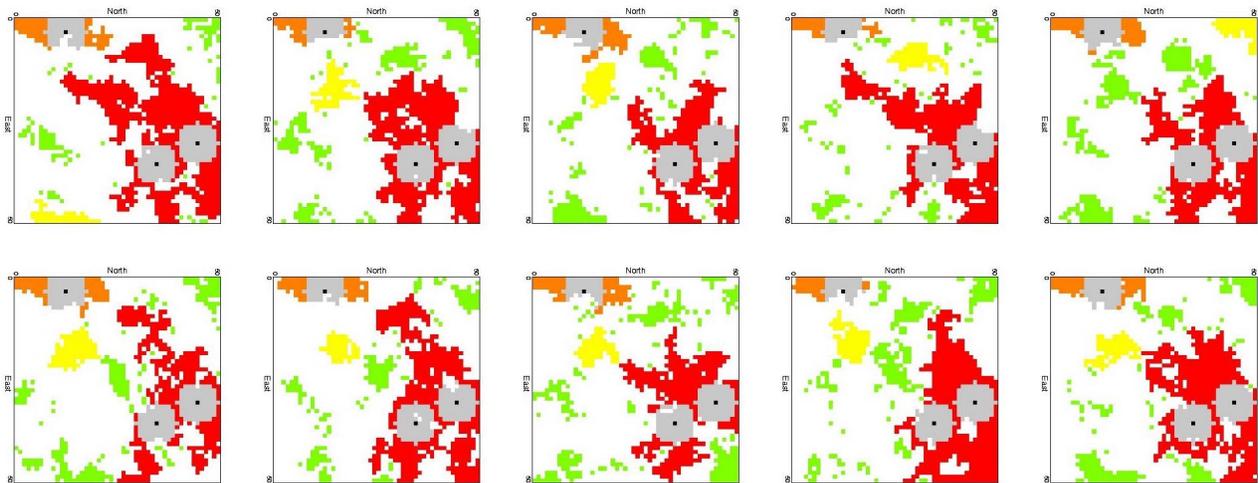


Figure 10: An optimized 2-D well plan consisting of 3 optimal well locations on a model of uncertainty with 10 realizations.

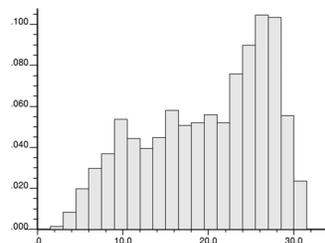


Figure 11: The distribution of porosity in the channel sand for example 2.

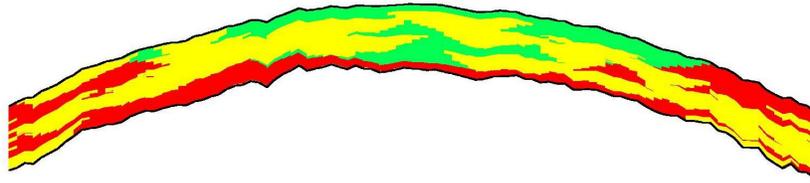


Figure 12: The reservoir model for Example 2. The geobodies are the lightest shade, the darkest shade is below OWC.

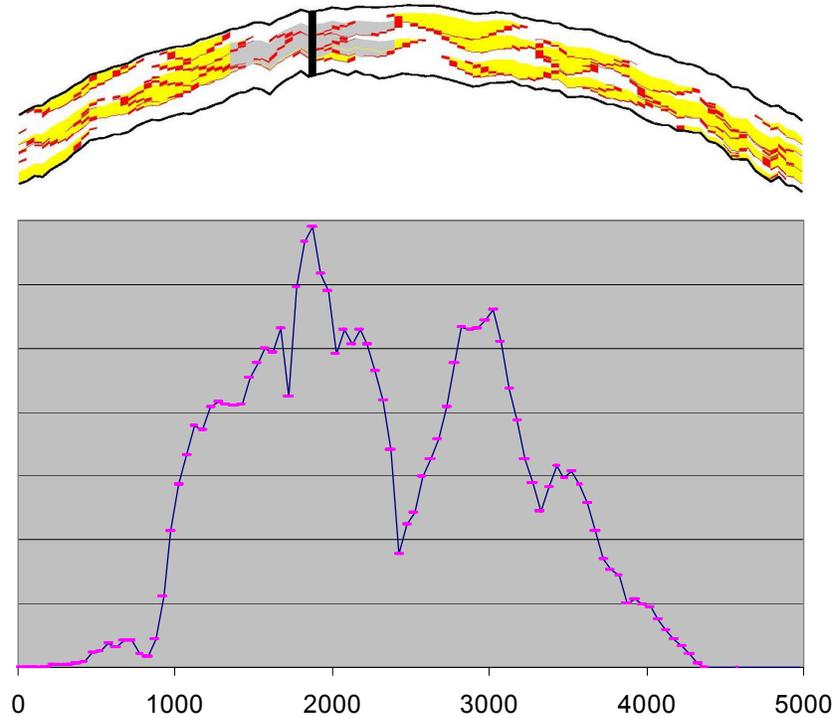


Figure 13: The optimal well location for a single well on one realization is $x=1875$. Shown below is a chart of the exhaustive profit at every possible location on the reservoir (the line on the chart) and the locations that SA sampled before reaching convergence (the squares).

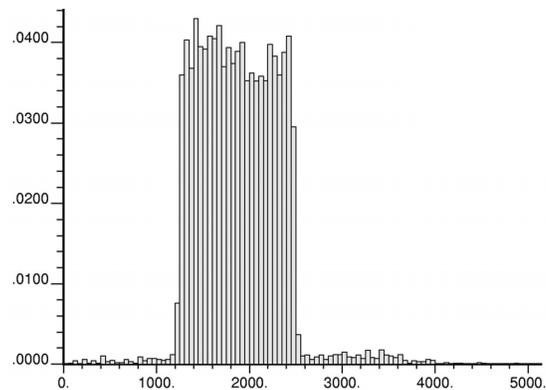


Figure 14: The distribution of SA sample locations in the x direction.

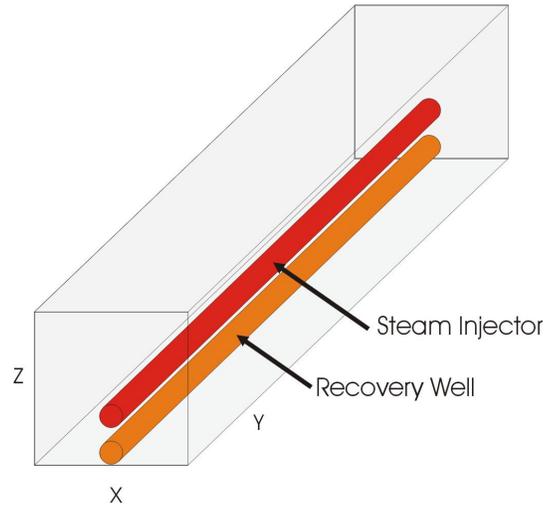


Figure 15: A SAGD well pair within a drainage block.

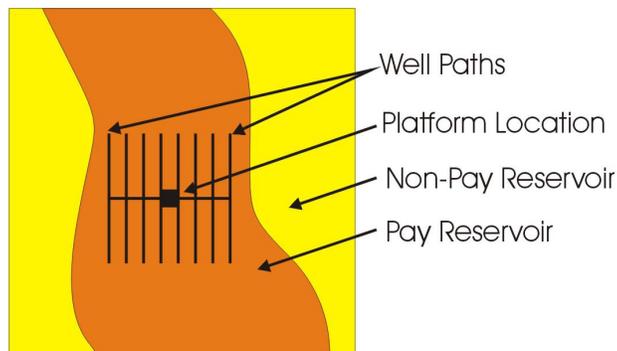


Figure 16: A central platform servicing several well pairs.

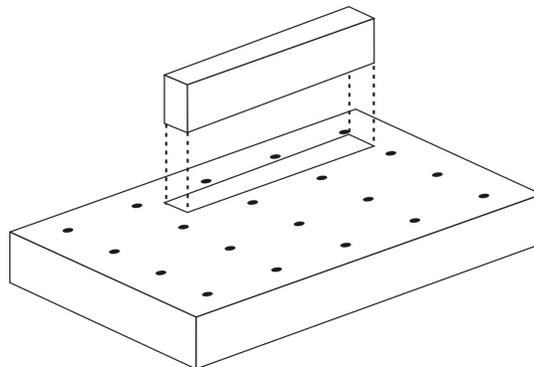


Figure 17: The high-resolution example model is a subset of a larger coarse model.

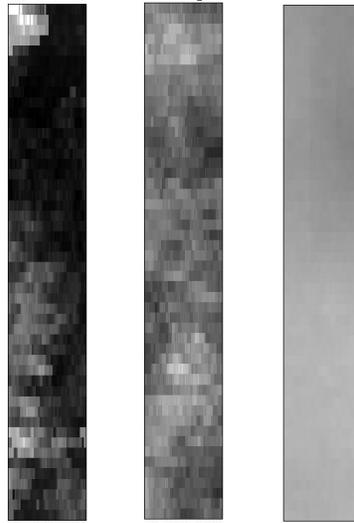


Figure 18: Central xy slice of first realization, same slice but averaged over 20 realizations, “projection” of averaged values.

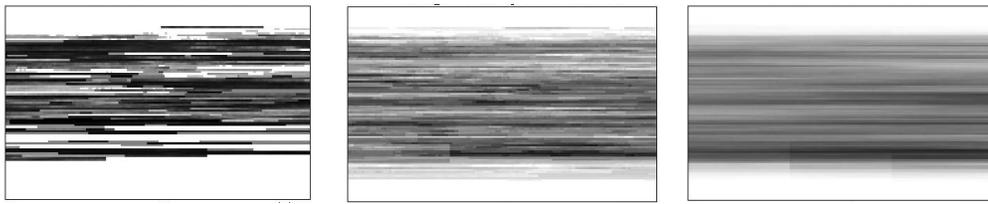


Figure 19: Central xz slice of first realization, same slice but averaged over 20 realizations, “projection” of averaged values.

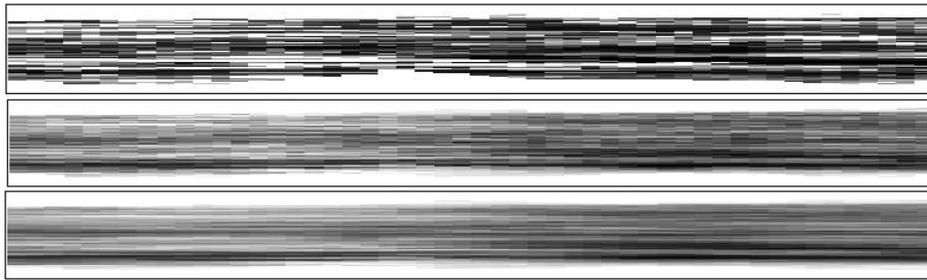


Figure 20: Central zy slice of first realization, same slice but averaged over 20 realizations, “projection” of averaged values.

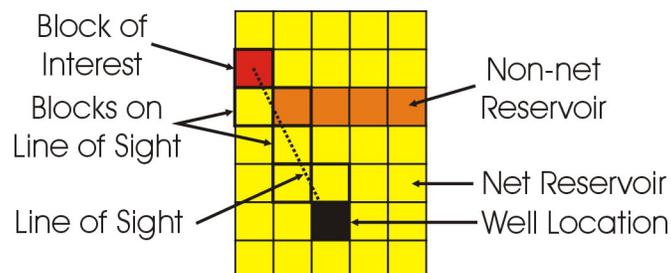


Figure 21: Connected pore volume is calculated by evaluating the fraction of non-net blocks on a line of site path from a block of interest to the well location. The example has 20% non-net on the line of site path and would be lightly penalized. If the fraction of non-net were higher then the penalty would increase.

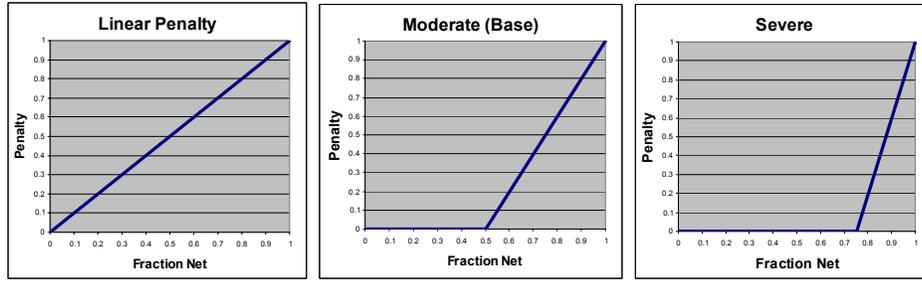


Figure 22: The three penalty functions used to calculate the expected resource.

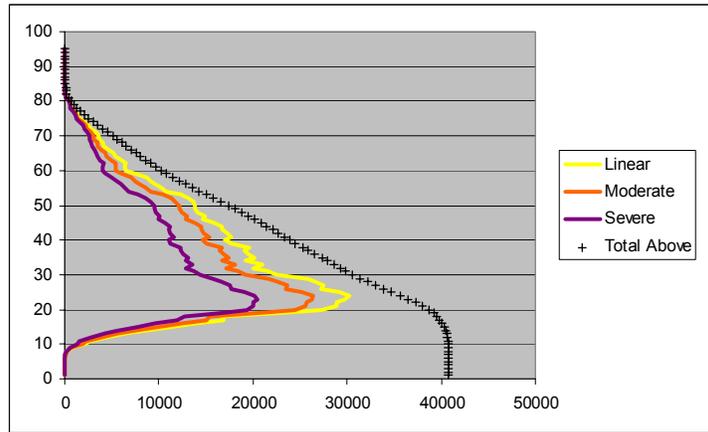


Figure 23: The x axis is pore volume (blocksize x porosity), the y axis is well location. The solid lines are expected resource above the well calculated using four penalty functions. The dashed line is the expected total remaining pore volume above the well. The no penalty line and the total above line are not the same since the no penalty line does not include pore volume at well locations where the well does not intersect reservoir whereas the total above calculation does.

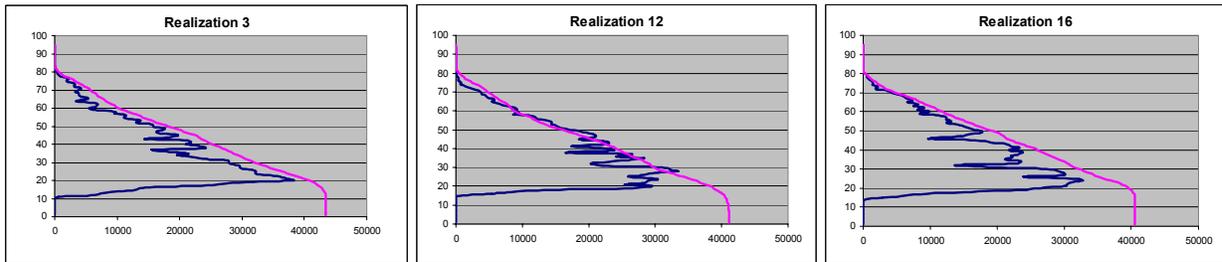


Figure 24: The resource above the well (the light line) and the total resource above the well for three realizations. Note that each realization yields a different optimal well location: realization 3 gives an optimal well location of 20m (38177.5 m³), realization 12 gives an optimal location of 29m (33282.7m³), and realization 16 gives an optimal location of 24m (32749.9m³)

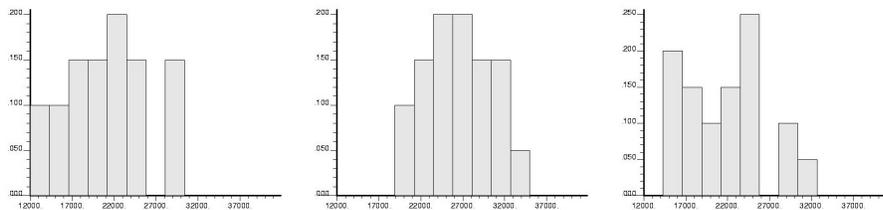


Figure 25: The histograms of pore volume for a well located 5m above the optimal well location of 24m (mean = 29282.98), at the optimal well location (mean = 32177.5), and 5m below the optimal well location (mean = 31471.62). Note that in expected value the optimal well location yields the highest expected resource.