

University of Alberta

DECISION MAKING USING GEOSTATISTICAL MODELS OF UNCERTAINTY

by

Karl Patrick Norrena

A thesis submitted to the Faculty of Graduate Studies and Research in partial fulfillment
of the requirements for the degree of **Doctor of Philosophy**.

in

Mining Engineering

Department of School of Mining and Petroleum Engineering

Edmonton, Alberta
Spring 2007



Library and
Archives Canada

Bibliothèque et
Archives Canada

Published Heritage
Branch

Direction du
Patrimoine de l'édition

395 Wellington Street
Ottawa ON K1A 0N4
Canada

395, rue Wellington
Ottawa ON K1A 0N4
Canada

Your file *Votre référence*
ISBN: 978-0-494-29721-6
Our file *Notre référence*
ISBN: 978-0-494-29721-6

NOTICE:

The author has granted a non-exclusive license allowing Library and Archives Canada to reproduce, publish, archive, preserve, conserve, communicate to the public by telecommunication or on the Internet, loan, distribute and sell theses worldwide, for commercial or non-commercial purposes, in microform, paper, electronic and/or any other formats.

The author retains copyright ownership and moral rights in this thesis. Neither the thesis nor substantial extracts from it may be printed or otherwise reproduced without the author's permission.

AVIS:

L'auteur a accordé une licence non exclusive permettant à la Bibliothèque et Archives Canada de reproduire, publier, archiver, sauvegarder, conserver, transmettre au public par télécommunication ou par l'Internet, prêter, distribuer et vendre des thèses partout dans le monde, à des fins commerciales ou autres, sur support microforme, papier, électronique et/ou autres formats.

L'auteur conserve la propriété du droit d'auteur et des droits moraux qui protègent cette thèse. Ni la thèse ni des extraits substantiels de celle-ci ne doivent être imprimés ou autrement reproduits sans son autorisation.

In compliance with the Canadian Privacy Act some supporting forms may have been removed from this thesis.

Conformément à la loi canadienne sur la protection de la vie privée, quelques formulaires secondaires ont été enlevés de cette thèse.

While these forms may be included in the document page count, their removal does not represent any loss of content from the thesis.

Bien que ces formulaires aient inclus dans la pagination, il n'y aura aucun contenu manquant.


Canada

The devil is in the details.

Lutz Röhrich

There is no free lunch.

Milton Friedmond

Deadlines are your friend.

Dr. C. V. Deutsch

To my wife Leanne, my son Cole,
and my dog Boca, with you, all is possible

Abstract

The selection of dig limits and well locations in the mining and petroleum industries have enormous economic consequences. Sometimes these decisions are easy, but usually the decision making process is obscured by uncertainty and an abundance of plausible alternative decisions. Geostatistical tools can be employed to account for uncertainty, but incorporating a model of uncertainty into a decision making framework adds another layer of complexity. Common practice for decision making under uncertainty is to construct a global objective function that quantifies loss and to use an optimization algorithm to find the decision that minimizes loss. Many optimization algorithms require that the global objective function meet certain criteria for use. The dig limit and well location selection problems cannot be set up in a way that satisfies the constraints for many optimization algorithms. The algorithm known as Simulated Annealing has very few constraints and can be used to select dig limits and well locations that minimize loss, given a well constructed global objective function.

This dissertation develops techniques that semiautomatically select dig limits and well locations that account for subsurface uncertainty. The dig limit selection technique has the additional feature of selecting dig limits that account for the limitations of the mining equipment. The techniques are referred to as semiautomatic because the user must select seed dig limits or well locations.

For dig limit selection, applications to hypothetical and real mines are explored. As well, hand drawn dig limits and semiautomatically selected dig limits are compared. In experiments, semiautomatic dig limits always outperform the hand drawn dig limits. When the semiautomatic dig limit selection technique was applied at the Bingham Canyon mine improvements of up to 1.5% are observed. The semiautomatic well location selection technique is applied to hypothetical reservoirs and a real data set: the Smiley Buffalo waterflood project. The semiautomatic well location selection outperformed well locations selected by the Asset Team by more than 19%.

Acknowledgements

I would like to thank my wife Leanne for her love, patience, and friendship. I would like to thank my son, Cole, who unknowingly inspired me to finish. I thank my thesis supervisor, Dr. Deutsch, for his time, boundless knowledge, patience, and generosity. I also gratefully acknowledge the support that I received from my colleagues. In particular, Julian Maxilliano Ortiz Caberra, Oy Leuanthong, and Steve Brown, deserve special thanks. These folks listened to my gripes and complaints after I had exhausted my wife. Their friendship and patient ears helped me stay the course. I would also like to thank NSERC for funding two years of studies, and Nexen Inc. for supporting this effort.

Contents

1	Introduction	1
1.1	Problem Setting	1
1.1.1	Selection of Dig Limits	2
1.1.2	Selection of Well Locations	11
1.2	Dissertation Outline	16
2	Theoretical Foundations	19
2.1	Geostatistics	19
2.1.1	Random Function Concept	19
2.1.2	Moments of Random Variables	20
2.1.3	Stationarity	20
2.1.4	Spatial Correlation	21
2.1.5	Kriging	22
2.1.6	Geostatistical Simulation	24
2.1.7	Spaces of Uncertainty	26
2.2	Simulated Annealing For Decision Making	27
2.2.1	Motivation for Using Simulated Annealing	27
2.2.2	Simulated Annealing Background	28
2.2.3	Thermodynamic Annealing	29
2.2.4	Thermodynamic and Simulated Annealing	29
2.2.5	Simulated Annealing Algorithm	30
2.2.6	Annealing Schedule	31
2.2.7	The Global Objective Function	33
2.2.8	Perturbation Mechanism	34
2.2.9	Simulated Annealing and Optimality	34
2.2.10	Simulated Annealing in the Literature	35
3	Semi-Automatic Dig Limit Selection	37
3.1	Theory	37
3.1.1	Expected Profit Map	40
3.1.2	Global Objective Function	44
3.2	Implementation Details	47
3.2.1	Expected Profit Map Implementation Details	49
3.2.2	Dig Limit Selection Details	51
3.2.3	Comments	59
4	Dig Limit Selection Case Studies	62
4.1	Considerations for Semi-Automatic Dig Limit Selection	62
4.1.1	Required Number of Perturbations	62
4.1.2	Effect of the Annealing Schedule	68
4.1.3	Repeatability of the Algorithm	68
4.2	Comparing Hand Drawn and Semi-Automatic Dig Limits	70
4.3	Copper Mine Case Study	75

4.4	Bingham Canyon Mine	80
4.4.1	Data and Geostatistical Models	82
4.4.2	Comments	87
5	Semi-Automatic Well Location Selection Theory	95
5.1	Jointly Optimal Well Locations	95
5.2	Global Objective Function	96
5.2.1	Component Objective Functions	98
5.2.2	Calibration Procedure	106
5.3	Implementation Details	107
5.3.1	Well Topology	107
5.3.2	Grid Topology	107
5.3.3	Initial Well Plan	108
5.3.4	Identification of Intersected Grid Blocks	108
5.3.5	Perturbation Mechanism	110
5.3.6	Well Selection Algorithm	112
5.3.7	Comments	114
6	Well Location Selection Case Studies	116
6.1	Considerations for Selecting Well Locations	116
6.1.1	Sequential Well Planning	116
6.2	Sequential Well Location Selection With Incremental Information	118
6.2.1	Methodology 1	121
6.2.2	Methodology 2	128
6.2.3	Comments	135
6.3	Well Selection on the Smiley Buffalo Field	135
6.3.1	Basis for Experimentation	135
6.3.2	Geology Review	137
6.3.3	Static Model	138
6.3.4	Dynamic Model	142
6.3.5	Asset Team Flow Simulation Results	149
6.3.6	Well Location Selection	150
6.3.7	Drainage Radius and Injection Radius Selection	154
6.3.8	Mini Model Results	154
6.3.9	Structural Penalty	162
6.3.10	COBJ Construction	168
6.3.11	Calibration to Flow Simulation Results	176
6.3.12	Experimental Results and Comparison	178
6.3.13	Comments	184
7	Concluding Comments	186
7.1	Dig Limit Selection	186
7.1.1	Dig Limit Selection Application Comments and Limitations	187
7.2	Well Location Selection	187
7.2.1	Well Location Selection Application Comments and Limitations	188
7.3	Future Work	189
7.3.1	Dig Limit Selection	189
7.3.2	Well Location Selection	190
	Bibliography	192

List of Figures

1.1	Map <i>a</i> shows dig limits that might be too difficult to mine.	3
1.2	This map shows the sample locations for the dig limit selection examlpe. . . .	5
1.3	The map shows the expected profit values used in the preliminary example. . .	6
1.4	The polygon shown in the map is the initial dig limit polygon.	6
1.5	The maps show the dig limit polygon during the optimization.	8
1.6	This map shows the final dig limits after 57601 perturbations.	9
1.7	The two areas identified as Area 1 and Area 2 will be examined	9
1.8	The two maps shown present the expected profit values. for the example . .	10
1.9	The map shows the fraction of the blocks falling inside the dig limits.	10
1.10	The top of structure maps for the well location example.	14
1.11	The five well plans shown are used in the calibration exercize.	15
1.12	The scatterplot shows NPV on the <i>y</i> -axis.	16
1.13	The map shows the final optimal well pattern.	16
1.14	The three maps show the three different 5 spot spacings.	17
2.1	An example experimental semivariogram is shown.	21
2.2	Four semivariogram models are commonly used to model.	22
2.3	This is a flowchart for the simulated annealing algorithm.	32
2.4	Shown are the results from two optimization runs.	35
3.1	The Figure shows two dig limit cases.	39
3.2	This schematic shows the transformation of the distribution.	41
3.3	An illustration of the effect of accounting for the mining equipment.	41
3.4	A hypothetical recovery curve.	42
3.5	Distributions of uncertainty in grade for two blocks of material.	43
3.6	An illustration of the notion of an optimal dig limit.	44
3.7	Shown in the Figure is an example penalty function.	46
3.8	The penalty function penalizes for difficult to dig polygons.	46
3.9	The digability factor interacts with the penalty function.	47
3.10	The map shows the expected profit values.	49
3.11	This map shows the optimal classification of blocks.	50
3.12	This map shows the expected profit values for ore.	50
3.13	The map shown posts dig limits for both ore types.	51
3.14	This flow charts illustrates the steps the dig limit selection.	52
3.15	The decomposition approach.	54
3.16	An alternative decomposition approach.	54
3.17	The perturbation mechanism for perturbing the polygon vertices.	55
3.18	Summary information from the dig limit selection.	56
3.19	The automatic initial dig limit selection algorithm steps are shown.	59
3.20	An illustration of how the automatic initial dig limit fails.	60
4.1	Three synthetic maps for exploring dig limits.	62
4.2	The chart shows the fraction of total revenue versus number of perturbations. .	63

4.3	The chart shows the fraction of total revenue versus number of perturbations.	63
4.4	The chart shows the fraction of total revenue versus number of perturbations.	64
4.5	The initial dig limits for each shape are shown at the top.	66
4.6	Plotted on this chart are the results for three different annealing schedules.	68
4.7	These maps show the dig limits resulting from the fast, moderate, and slow schedules.	68
4.8	A chart of the global objective function value for the four random seed.	69
4.9	The map shows four dig limits superimposed on a single map of expected profit.	70
4.10	This map of expected profit was provided to several mining professionals.	71
4.11	The map for the dig limit challenge.	71
4.12	The maps shown compare the hand drawn dig limits and selected dig limits.	74
4.13	The map shows the sample locations and grades for a copper mine in Chile.	75
4.14	The histogram of copper grades for the blasthole samples.	75
4.15	The experimental and modeled semi-variogram.	76
4.16	The map shows kriged estimates of percent copper.	76
4.17	The map is of a single realization of percent copper from the model.	77
4.18	This map shows a block averaged percent copper map.	77
4.19	The expected profit map used to select semi-automatic dig limits.	78
4.20	None of the blocks had an expected grade greater than the cutoff grade.	78
4.21	The image on the left shows grade to profit conversion.	79
4.22	Shown in this Figure is the seed dig limit.	79
4.23	This Figure shows a digability catalogue.	80
4.24	The maps show the sample locations for the two benches being considered.	82
4.25	The maps show rock types for the 4940 (left) and 5640 (right) benches.	83
4.26	From the top of the Figure are the Datum, ID3, and Simulation EWOR maps.	85
4.27	From the top of the Figure, are the Datum, ID3, and Simulation EWOR maps.	86
4.28	These maps are for the 4940 region results.	88
4.29	These maps show the 5640 region results.	89
4.30	These maps are for the 4940 region results.	90
4.31	These maps are for the 5640 region results.	91
4.32	These maps are for the 4940 region results.	92
4.33	These maps are for the 5640 region results.	93
5.1	The jointly optimal solution intersects regions of high reservoir quality.	95
5.2	The chargeable volume is the sum of the blocks falling in the drainage.	101
5.3	The sectional view of the reservoir marked <i>a</i> has three geobody numbers.	102
5.4	Three different well types are considered: vertical, deviated, and segmented.	106
5.5	The corner point grid system is used to model reservoir structure.	107
5.6	The well path is discretized into a number of points.	109
5.7	The process for determining the <i>z</i> -index for the discretized point.	109
5.8	This Figure illustrates the perturbation mechanism.	110
5.9	The flow chart for the well selection algorithm.	113
6.1	An unconditional realization is shown on the left.	115
6.2	The maps shown depict the progress of the algorithm.	117
6.3	The map shows the three horizontal well solution.	118
6.4	These maps show an areal perspective and of the optimal solutions.	118
6.5	There is more than one optimal one well solution.	119
6.6	This schematic illustration shows the histograms of recoverable volume.	120
6.7	This schematic illustrates the procedure to assess value.	122
6.8	The map shows the positions of the four conditioning data.	123
6.9	This histogram shows the distribution of transformed porosity.	123
6.10	Shown is a single porosity realization.	123
6.11	The map shows the geobodies and the four well locations.	124

6.12	Together these maps represent one leg from Figure 6.7.	124
6.13	The red, blue and green histograms.	125
6.14	Here is a schematic illustration of the histograms of uncertainty.	126
6.15	The map on the left is the secondary information.	128
6.16	The histogram for the 20 realizations is shown on the left.	128
6.17	These results use the incremental data from the Red case.	129
6.18	These results are for the Orange case.	129
6.19	These results use the incremental information from window averaging.	129
6.20	These results use the incremental information from the bottom histogram.	130
6.21	These are the final well locations for the green case.	130
6.22	The top left histogram show the Red case results.	131
6.23	The map shows the basal surface of the Smiley Buffalo Section 15 reservoir.	133
6.24	The map shows the isochore thickness of the reservoir zone.	135
6.25	These histograms show the raw and upscaled facies proportions.	136
6.26	The raw and upscaled porosity distributions for Smiley Buffalo	137
6.27	The raw and upscaled permeability distributions for Smiley Buffalo.	138
6.28	The raw and upscaled water saturations for Smiley Buffalo.	138
6.29	The permeability/porosity relationship for Smiley Buffalo.	140
6.30	The permeability/porosity relationship for Smiley Buffalo.	140
6.31	The map shows a facies realization from the Smiley Buffalo reservoir model.	141
6.32	The map shows a porosity realization from the Smiley Buffalo reservoir model.	142
6.33	The map shows a permeability realization from the Smiley Buffalo reservoir model.	143
6.34	The chart compares the average permeability by layer.	143
6.35	The PVT properties for Smiley Buffalo.	144
6.36	The relative permeability curves used for reservoir simulation.	145
6.37	This chart shows the field oil production versus time for the Asset Team wells.	146
6.38	This chart shows the field water production versus time for the Asset Team wells.	146
6.39	This chart shows the field water injected versus time for the Asset Team Selected wells.	147
6.40	This chart shows the barrel oil equivalent value for the field under the Asset Team well plan.	147
6.41	These charts are the cumulative distribution functions for the Asset Team well plan.	148
6.42	The chart shows the hydrocarbon production rate over time for all realizations and the	150
6.43	The map is the vertically averaged pressure for realization 3 after 475 days.	151
6.44	The map is the vertically averaged pressure for realization 5 after 475 days.	152
6.45	The map is the vertically averaged pressure averaged over all realizations after 475 days.	153
6.46	The map shows the water saturation on layer 18 (the bottom of the reservoir) after 20 years for realization 5.	154
6.47	The x -axis shows separation distance, and the y -axis shows the cumulative oil.	155
6.48	The x -axis shows separation distance, and the y -axis shows the cumulative oil.	156
6.49	The x -axis shows separation distance, and the y -axis shows the cumulative water.	157
6.50	The x -axis shows separation distance, and the y -axis shows log cumulative water.	157
6.51	The x -axis shows separation distance, and the y -axis shows the cumulative oil.	158
6.52	The x -axis shows separation distance, and the y -axis shows the COV.	159
6.53	The x -axis shows separation distance, and the y -axis shows the surface area.	160
6.54	The x -axis shows separation distance, and the y -axis shows the geobody size.	160
6.55	The x -axis shows separation distance, and the y -axis shows the HCPV.	161

6.56	The x -axis shows separation distance, and the y -axis shows the permeability.	161
6.57	A map of the azimuth of strike for the Smiley Buffalo reservoir.	162
6.58	This map shows the dip at each location on the basal surface.	163
6.59	The 6 charts show how field oil production changes with respect to separation.	165
6.60	The x -axis shows the EOPT values, and the y -axis shows the predicted EOPT.	166
6.61	The x -axis shows the predicted EOPT values, and the y -axis shows the standardized residuals.	167
6.62	The 6 charts show how field water production changes with respect to separation distance,	168
6.63	The x -axis shows the EWPT values, and the y -axis shows the predicted EWPT values.	169
6.64	The x -axis shows the predicted EWPT values, and the y -axis shows the standardized residuals.	169
6.65	The 6 charts show how field water injected and produced changes with respect to separation distance	170
6.66	This map shows the optimal well locations (solid dots), the Asset Team locations.	173
6.67	This map shows the recoverable volume template for the Asset Team well plan (top)	174
6.68	This map shows the sweep pattern for the Asset Team (top map) selected well plan	175
6.69	This chart shows the field oil production versus time for the Asset Team Selected wells.	176
6.70	This chart shows the field water production versus time for the Asset Team Selected wells.	177
6.71	This chart shows the field water injected versus time for the Asset Team Selected wells.	177
6.72	This chart shows the barrel oil equivalent value for the field under the Asset Team well plan.	178
6.73	These charts are the cumulative distribution functions for the optimal well plan (dark line),	179

List of Tables

1.1	This Table presents the numerical results for the second comparison.	7
1.2	The Table shows the tonnes of ore and waste for the preliminary example . .	7
3.1	The Table is a summary table of results for the grade control experiment. . .	42
4.1	This Table compares the revenues (Rev.) and digability (DIG.)	65
4.2	These values are the annealing schedules used.	67
4.3	These results show the final profit, penalty, and global objective function values.	67
4.4	This table shows the results for comparing the hand drawn and semi-automatic dig limits.	72
4.5	This table shows the results for comparing the hand drawn and semi-automatic dig limits.	72
4.6	Here are the tabulated values for prediction of total improvement.	87
4.7	These results show the predicted improvement using simulation based EWOR. 87	
4.8	These results show the prediction of improvement using semi-automatic dig limit selection alone.	89
6.1	This Table documents the results for 100 trials in the second experiment. . .	128
6.2	The facies semivariograms.	137
6.3	The porosity semivariograms.	138
6.4	The permeability semivariograms.	139
6.5	The water saturation semivariograms.	139
6.6	The tabulated flow simulation results for the Asset Team well plan.	149
6.7	This Table lists the individual penalties that comprise the structural penalty.	160
6.8	This Table shows some fitment statistic for the regressed EOPT component objective function.	164
6.9	This Table tabulates the fitment statistics for the coefficients of the EOPT component objective function.	166
6.10	This Table tabulates the fitment statistics for the coefficients of the EWPT component objective function.	168
6.11	This Table of values compares the component objective function values. . . .	171
6.12	The Table of values compares the pre- and post-calibration global objective function.	172
6.13	This Table enumerates the flow simulation results for the optimal well plan. .	178
6.14	This Table compares the summary statistics for the Asset Team well plan. .	179
6.15	This Table compares the cumulative produced oil volumes for each producer.	180
6.16	This Table compares the cumulative produced water volumes for each producer.	180

Chapter 1

Introduction

1.1 Problem Setting

Making decisions in the mineral resources industries is difficult due to inherent uncertainty, uncertainty that is a consequence of limited knowledge of the subsurface. In the mining industry, mining engineers must classify material as ore and waste. Ore is material that is economic to ship to the mill for processing and extraction of the mineral of interest. Waste is material not classified as ore and shipped to the waste dump. Ore misclassified as waste is missed opportunity to profit. Waste processed as ore does not recover the cost of processing. The mining engineer draws polygons around volumes of ore and waste. The polygons are called dig limits. There is limited information for the selection of dig limits and the decision has high economic importance. The collected data have large spacing and represent a very small fraction of the volume of material to be classified. There is uncertainty in the grade of material between the sample points and it is not clear what dig limits will give the greatest value.

In the petroleum industry, the selection of well locations by reservoir engineers is also burdened with uncertainty. Reservoirs are sampled far less than mineral deposits. There is significant uncertainty in the rock properties between sample locations. An important hydrocarbon recovery technique is injecting water into the reservoir. The injected water displaces hydrocarbon. A good well plan has the water injection well(s) displacing hydrocarbon towards the producing well(s). A poorly selected well plan can have production wells producing only the injected water, or have the injected water displaced away from the producer. The rock properties, as well as the reservoir structure affect the hydrocarbon displacement process. Wells that produce high fractions of water may fail to payback the cost of engineering the location, drilling the well, outfitting the well, or the cost of pumping fluids to or from the well.

In both the mining and petroleum industries, geostatistical tools can be used to quantify uncertainty in the subsurface. Uncertainty is quantified in the form of multiple realizations. Each realization is an alternative model of the subsurface. Each realization integrates and honors all available information. The collection of realizations is called a model of uncertainty. The realizations could be used to enumerate the risk associated with selecting a particular dig limit or well plan; however, there are few tools that an engineer can use to select the decision that maximizes value over the model of uncertainty. The problem of selecting the optimal decision on a model of uncertainty is not trivial. In the case of the dig limit selection problem, the number of possible dig limits cannot be counted, and for the problem of selecting well locations, the number of possible solutions is a combinatorial function of the number of wells and the number of possible locations the wells could take if the wells are vertical wells, and uncountable if the wells are not vertical.

The typical approach to solving decision making problems like the dig limit and well location selection problems is to pose the problem in the form of a mathematical equation

called a global objective function and use some sort of optimization technique to find the solution, that is the dig limits or well locations, that minimizes the function. However, it can be prohibitively expensive to exhaustively enumerate the set of plausible solutions in the search for an optimal solution. The global objective function used for evaluating the goodness a particular decision may be difficult to evaluate, and the global objective function might be poorly behaved and limit types of optimization algorithms that can be used. The selection of the appropriate optimization algorithm and the formulation of the global objective function are essential for selecting dig limits and well locations that strikes the best balance among multiple objectives housed by the global objective function.

The following Sections, Section 1.1.1 and Section 1.1.2, elaborate on the problem specific challenges in selecting dig limits and well locations under uncertainty. The Sections also include preliminary examples of the decision making techniques developed in the dissertation for selection of dig limits and well locations. The Sections following the discussions of the dig limit and well location section problems discuss the layout of the dissertation.

1.1.1 Selection of Dig Limits

In an open pit mine, the mining engineer uses maps of grade information to delineate regions that are economic to process (ore) and regions that are uneconomic to process (waste) with dig limits. Usually the dig limits are selected on a short term basis, that is, they contain one to ten days worth of digging.

Dig limits are used in one of two ways. A survey crew surveys the dig limit coordinates in the pit using painted stakes. Each side of the stake has a different color. One color always faces ore and the other faces waste. When ore and waste cannot be discerned by eye, the shovel operator relies on the stakes. The operator digs to the stakes until they fall into the bucket or onto the operating bench. The shovel operator keeps track of the material as it is loaded in haul trucks and tells the haul truck driver, either by radio or by horn blasts, if the load is ore or waste. The dig limit selection can also be executed is digitally. The engineer or geologist selects dig limits using geomodelling software. The software keeps track of the progress of the shovel using GPS sensors on the shovel. The software indicates to the shovel operator if the scoop is in ore or waste. The software might also indicate to the haul truck operator if the load is ore or waste otherwise, the shovel operator provides that information.

There are economic consequences if material is misclassified. Material that is classified as ore is shipped to the mill for processing and mineral extraction. Material that is really waste and misclassified as ore incurs loss because the material that can be recovered does not cover the cost of shipping, processing, and mineral extraction. Material that is misclassified as waste is shipped to a waste dump and most mines never revisit the waste dump to recover misclassified material. Ore misclassified as waste amounts to a lost opportunity to profit. Loss due to misclassification on a few days worth of digging may not be significant with respect to the total profit capacity of the mine; however, a mine's life may be in the order of decades and systematic loss can accumulate to millions of dollars.

The grade control program is the protocol that a mine uses to ensure that the mill feed is of a minimum quality. A typical grade control program might consist of the following steps:

1. Obtain sample information from drill, blast hole, and geologic data;
2. Assay the drill and/or blasthole data to obtain grade information. Interpret geological data in the context of a geological conceptual model.
3. Map the grade information to the bench of interest.
4. Select dig limits.

We are interested in the last three steps, with most of our interest concentrating on the last two.

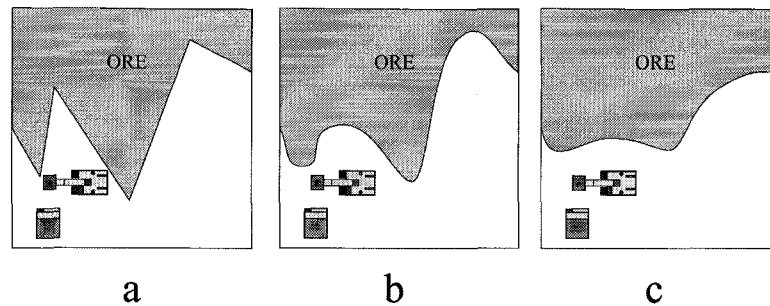


Figure 1.1: Map *a* shows dig limits that might be too difficult to mine. Map *c* shows dig limits that unnecessarily dilute and lose ore. Map *b* shows dig limits that strike the best balance between ease of mining and maximum grade.

Prior to the introduction of geostatistics and computers, mining engineers and geologists mapped by hand the sample data and hand drew the dig limits [5, 6, 54, 62]. There are some drawbacks to hand mapping and hand drawn dig limits:

1. Uncertainty in grade at unsampled locations is not accounted for in a quantitative manner, that is, no provision is made for assessing the impact of uncertainty and misclassification.
2. Grade information from previously mined benches and exploration drilling is not easy to consider.
3. Mining equipment limitations are not systematically accounted for, that is, the limits may be unrealistically complex or overly simplistic. Dig limits that are too complex are difficult to mine and the mining equipment cannot perform efficiently. Simplistic dig limits might dilute or waste ore.
4. Hand contoured dig limits are subjective, that is, there is neither a reproducible procedure, or an objective measure of optimality.

Figure 1.1 shows three mining scenarios. The scenario in map *a* shows dig limits that are too complex for the mining equipment to mine efficiently; the angles between the vertices are too acute and there may not be enough room for the shovel to operate efficiently or the haul trucks to manoeuvre. The operators may spend too much time attempting to mine the dig limits precisely. The scenario in map *c* shows dig limits that are too smooth. Ore is either diluted or wasted, so that the dig limits can be easily excavated. Optimal dig limits, that is, dig limits that can be shown to minimize loss, minimize diluted and lost ore and are not too complex. Map *b* shows dig limits that strike the best balance between ease of mining and minimizing lost and diluted ore.

The advent of geostatistics made it possible to construct models of uncertainty in grade information. Drill and blasthole samples provide grade information, but only at sampled locations. Similarly, geologic information is available only at the mine face or wall, not at locations where mining will occur in the future. The model of uncertainty provides a set of realizations of possible grade values at unsampled locations. This information is used to help select dig limits.

The challenge that the mining engineer or geologist faces is how to use the set of realizations to select dig limits that minimize loss. For each realization, there is a corresponding dig limit that minimizes loss. This is not particularly useful information for the shovel operator, who is faced with the reality of having to load haul trucks with ore or waste; the shovel operator can mine only to a single dig limit, not to a set. Within the literature there

is documentation of how other researchers have used models of grade uncertainty to identify regions of ore and waste.

Deutsch et. al. [23] present a geostatistical simulation based technique for identifying ore and waste blocks that maximizes profit. They call the technique Maximum Profit Selection (MPS). The method starts with a model of uncertainty in grade. The realizations are processed to provide a distribution of profit if the block were classified as ore and as waste. The block is classified as ore if the expected profit as ore is higher than the expected profit of the block as waste. The profit of a block depends on the mineral commodities present, the price of the commodities, the cost of recovery, the ore and waste mining cost, and the lost profit due to wasting ore. This technique is repeated over all blocks in the orebody model. The dig limits are interpreted as the interface between ore and waste blocks. This type of dig limit will be referred to as “blockwise” dig limits.

Srivastava et. al. [77] use simulation to propose dig limits that incorporate the swing radius of a shovel into dig limit selection. They start with an initial data set and construct a model of uncertainty. Then, on each realization, a swing radius is superimposed on each node. If all the blocks with centers in the radius can be mined profitably, the block centered on the node is called ore; otherwise, it is waste. Next, the algorithm counts the fraction of times a node is called ore. If a node shows too much uncertainty, the authors recommend that a sample be obtained at this location. Using the new information, the model of uncertainty is rebuilt and the process of superimposing a swing radius is repeated. The whole process is repeated until all nodes can be identified as ore or waste within an acceptable level of uncertainty. The final dig limits are interpreted as the blockwise representation of dig limits.

Glacken [30] develops a grade control technique that replaces the traditional approach for block selection with a probabilistic approach. The probability of a block to exceed a deterministic cutoff grade is compared to a threshold probability that is calculated in a way that accounts for the mine’s risk profile. Blocks are classified as ore if they exceed the probability threshold and as waste otherwise.

Schofield et. al. [74] develop a profit based grade control system. They construct a distribution of grades at unsampled locations and transform subsequently transform the distribution into profit. They use the profit values to classify blocks such that profit is maximized. The criteria used to classify blocks is not discussed in the work. The authors claim to have implemented the technique successfully at several Australian mines.

Remark

While the techniques discussed above account, to varying degrees, for uncertainty and misclassification, they do not account for the limitations of the mining equipment. The final dig limits are represented by the interface between ore and waste blocks and assume that the mining equipment can perfectly mine blocks of material. There is a need for an algorithm that optimally selects dig limits according to the selectivity of the mining equipment.

The proposed dig limits need to account for the cost of misclassification, the limitations of the mining equipment, and ultimately minimize loss. Accounting for the limitations of the mining equipment and minimizing loss are competing objectives. Dig limits that account for the limitations of the mining equipment may not minimize loss, and dig limits that minimize losses in equipment efficiency may lose or dilute too much ore. A large number of plausible solutions must be evaluated to find the optimal solution. The set of plausible solutions is a function of the number of vertices defining the dig limit and their locations and are constrained by the limitations of the mining equipment, and the location of the ore. The “goodness” of a dig limit polygon might be poorly behaved with changes in the location or number of vertices.

Dig Limit Selection Preliminary Example

This Section presents a preliminary example of the dig limit selection algorithm. The general approach for selecting optimal well locations consists of the following steps:




Figure 1.2: This map shows the sample locations and the corresponding grades used in the preliminary example.

1. Construct a model of uncertainty on the mineral grade(s).
2. Transform the model(s) of uncertainty to an expected profit map.
3. Use the expected profit map to select initial dig limits by hand.
4. Use a global objective function that quantifies the goodness of the dig limits and simulated annealing to select optimal dig limits.

The data are from a copper mine in Chile. Figure 1.2 shows the locations and grades of the blasthole samples. The model of grade uncertainty was built using Sequential Gaussian simulation, and the grid has 24x24 blocks that measure 4x4x5m. The model of grade uncertainty is transformed into the map of expected profit shown in Figure 1.3. The expected profit map summarizes the profitability of a block of material and accounts for the risk of misclassification.

The legend for the expected profit map shows a grey scale gradation of profit values from waste to ore. The darkest blocks have the highest expected profit and the lightest have the least profit. The dark areas represent regions that would be profitable to ship to the mill (ore), and the lighter regions would not be as profitable (waste). The large connected shaded region in the top right corner will be used to select dig limits and the other regions will be ignored.

Initial dig limits are shown in Figure 1.4 as a polygon that outlines a high profit region. The enclosed region contains an expected profit of \$407,437. By observation, the initial dig limit polygon unnecessarily dilutes and wastes ore. Optimal dig limits minimize dilution and still account for the limitations of the mining equipment.

Figure 1.5 shows the dig limits at selected intervals during the optimization. The first interval is shown in map *a* and is for 500 accepted perturbations. Map *b* through *h* are after 1000, 2000, 4000, 8000, 16000, 32000, and 64000 accepted perturbations respectively.

The maps show some interesting features:

- The dig limits in Map *a* are noisier than the initial dig limits.
- Map *b* shows smoother dig limits that have begun to close in on an ore region.
- The dig limits in maps *c* through *f* eliminate the noisy areas of low values and get to the “core” of the ore body.




Figure 1.3: The map shows the expected profit values used in the preliminary example.




Figure 1.4: The polygon shown in the map is the initial dig limit polygon.

Perfect Selection			
	Tonnes	Initial Profit (\$)	Final Profit (\$)
Ore	66720	443047	443047
Waste	25440	-25440	-25440
Total	92160	417607	417607
Inside Dig Limits			
	Tonnes	Initial Profit (\$)	Final Profit (\$)
Ore	52884	410072	411948
Waste	1	-2635	-460
Total	52885	407437	411488
Outside Dig Limits			
	Tonnes	Initial Profit (\$)	Final Profit (\$)
Ore	13835	32975	31098
Waste	25438	-22805	-24979
Total	39274	10170	6119

Table 1.1: This Table presents the numerical results for the second comparison.

	Perfect Selection		Inside Dig Limits		Outside Dig Limits	
	Tonnes	Profit	Tonnes	Profit	Tonnes	Profit
Ore	90200000	7534.06	70375096	7079.78	19824904.00	454.28
Waste	35192000	-603.92	2255850	-8.29	32936150	-595.63
Total	125392000	6930.14	72630946	7071.49	52761054	-141.35

Table 1.2: The Table shows the tonnes of ore and waste for the preliminary example

- Maps g and h work out the fine details of the dig limits. The coarse features give the biggest improvement on the dig limits. The fine details give smaller improvements and require many perturbations to resolve.

The final dig limits are shown in Figure 1.6. The dig limits are smooth, but not arbitrarily smooth. If a block has high profit the dig limits are more tortuous. Consider the areas outlined as Area 1 and Area 2 in Figure 1.7. The expected profit values for each area are shown in Figure 1.8. The dig limits for the areas are shown as grey lines. The dig limits for Area 1 are more tortuous than those for Area 2, so that the block valued at 2.003 in Area 1 is almost fully within the dig limits. Area 1 also has more dilution than Area 2. Waste is mined as ore to incorporate the high grade block in Area 1.

A summary table of results is shown in Table 1.1. The fraction of ore increases in the final dig limits. The overall increase comes to about 1%. This may not seem like a big improvement, but profit accumulates over a large number of dig limits.

Figure 1.9 shows the fraction of each block contained in the dig limits for the preliminary example. The fraction of blocks map can be used to estimate tonnes of ore and waste. With an assumed average density of ore of 2750 kg/m^3 and 2650 kg/m^3 for waste, a 5m bench height, and a price for copper of $\$2.75/\text{kg}$ the results for the dig limits are shown in Table 1.2. The results in Table 1.2 could be used for plant scheduling.

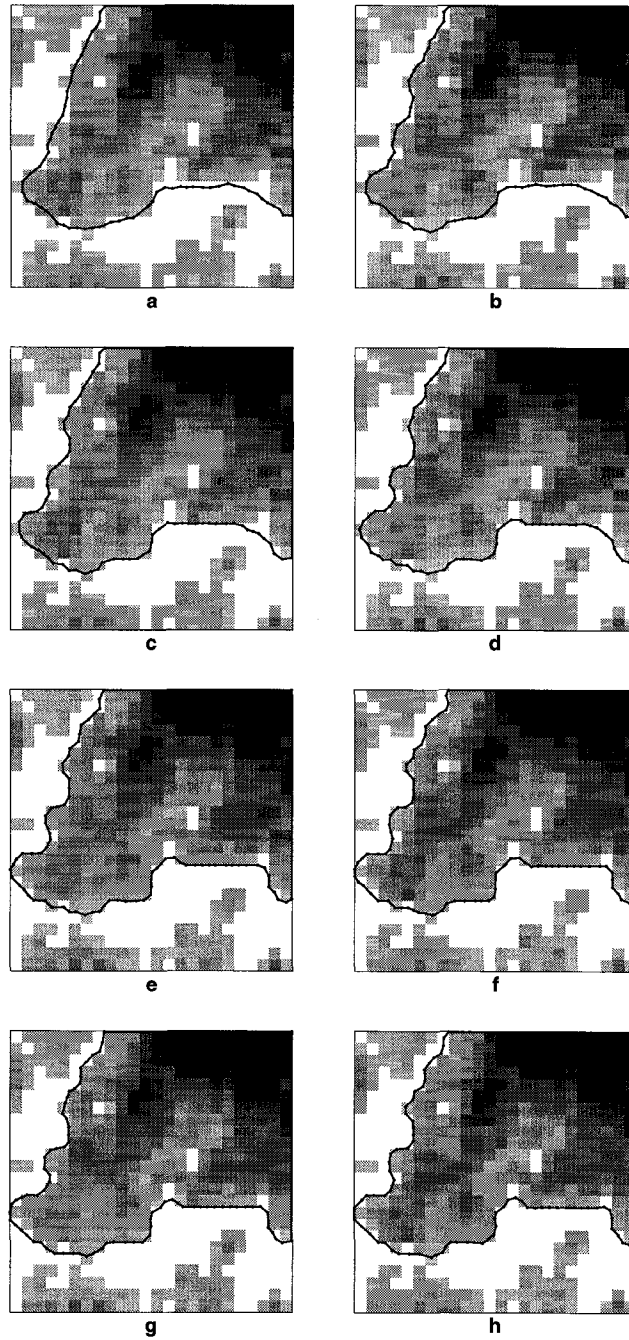


Figure 1.5: The maps show the dig limit polygon during the optimization. Maps *a* shows the dig limits after 500 perturbations. Subsequent maps double the number of perturbations, i.e. map *b* shows dig limits after 1000 perturbations.

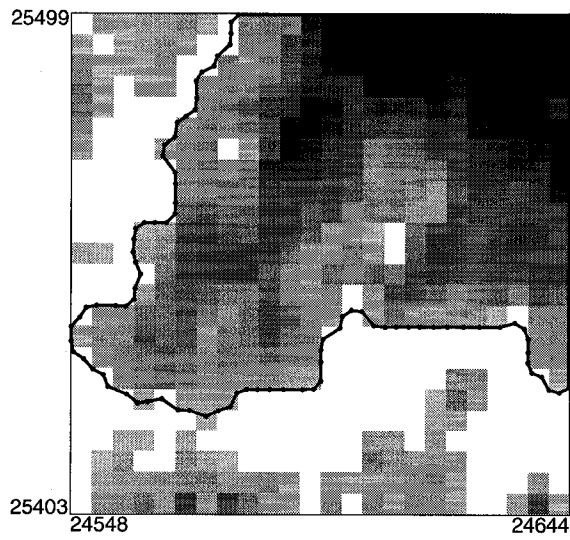


Figure 1.6: This map shows the final dig limits after 57601 perturbations.

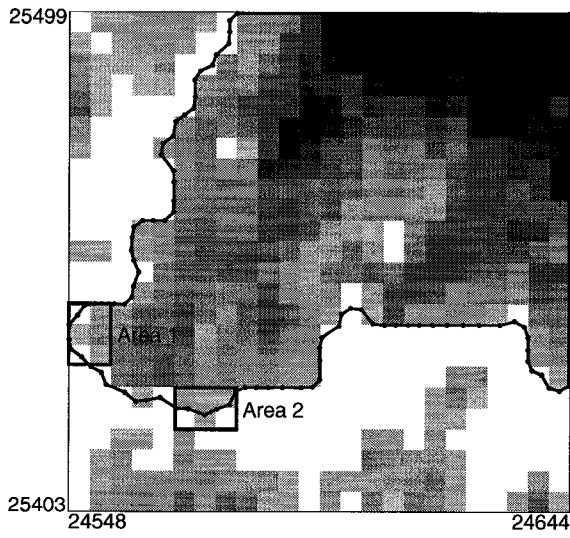


Figure 1.7: The two areas identified as Area 1 and Area 2 will be examined more closely.

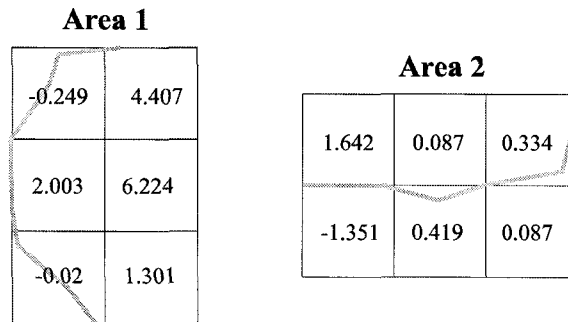


Figure 1.8: The two maps shown present the expected profit values for the two areas identified in Figure 1.7. The map on the left is for Area 1 and the map on the right is for Area 2. The dig limit polygons are superimposed on the areas under consideration.




Figure 1.9: The map shows the fraction of the blocks falling inside the dig limits. The black blocks fall completely inside the dig limits and the white ones do not fall in the dig limits at all.

Comments

The preliminary example demonstrates the dig limit selection technique developed in Chapter 3. One advantage that the dig limit selection technique enjoys is consistency. The technique provides dig limits that consistently account for the limitations of the mining equipment, are reproducible despite the initial dig limits or the random number seed.

1.1.2 Selection of Well Locations

Well locations have significant economic importance. Petroleum companies rely on the well locations for demonstration of reserves and for generation of income. A multidisciplinary team called an asset team uses a reservoir model to highlight candidate well locations. The reservoir model might be as simple as a structure map or as complex as a detailed 3D digital reservoir model. The candidate locations might be exploration or appraisal wells, infill wells, or injectors.

Ideally, candidate well locations are evaluated using the reservoir model and flow simulation. Flow simulation provides estimates of fluids produced, and if injectors are present, the volume of fluid injected. When used together, a reservoir model of uncertainty and flow simulation can give distributions of produced and injected fluids. Candidate locations that maximize the asset teams objective(s) make up the proposed well plan.

Identification of well locations that maximize the asset teams objectives is difficult. The asset team must consider many variables when selecting a well plan and some of the variables may interact. If the well plan is selected by hand, some objectives may be overlooked, over-, or understated, and the results may be impossible to reproduce or justify.

A large volume of literature exists for techniques that select well locations. The literature can be categorized as techniques that use static information only, techniques that use static and dynamic information, and hybrid techniques that use both static information and proxies of dynamic information.

Well Placement Techniques That Use Static Information

Vasanthrajan and Cullick [81] use mixed integer programming to select a static well plan based on a reservoir quality map. The technique incorporates a number of constraints: the number of wells, minimum well spacing, and a specified drainage radius. They acknowledge that even for small grids the problem can have prohibitively large CPU and memory requirements. They reduce the complexity of the problem by decomposing the problem into smaller regions and they only consider solutions that meet or exceed specific criteria, that is, they do not consider solutions in which the distance between any two wells is less than a threshold value. They use a single deterministic reservoir model rather than a model of uncertainty.

Ierapetritou et al. [38] also use mixed integer programming. Their method is similar to that of Vasanthrajan and Cullick, but they tackle larger reservoir models and more constraints by reducing the problem into a series of smaller mixed integer linear programming problems. They also acknowledge CPU and memory limitations for their technique.

Seifert et. al. [75] show a method that uses a reservoir model of uncertainty. They rank the well locations statistically and select the set that performs best. The set of well locations they consider are preselected by hand. The well trajectories are also hand selected: the well trajectories form an array similar to a pincushion of possible locations emanating from a central pad location. The trajectories are equally spaced about the pad. The disadvantage of this technique is that it does not consider the entire solution set; it only considers the hand selected locations, and these may offer only sub-optimal locations.

Goggin et. al. [31] demonstrate a method for horizontal well placement. The method starts with a reservoir model and maximizes recovery with a minimum number of wells. The platform locations are fixed and hand selected. The idea is to propose a large set of candidate well trajectories from the two platforms, and select from the set the fewest

number of wells required to maximize production. The method requires that the candidate well locations be selected by hand.

Wang [85] uses simulated annealing to select optimal well locations over a reservoir model of uncertainty. The technique starts with an initial well configuration that is perturbed until a configuration is found that maximizes either the statically connected pore volume or the pore volume that falls within a specified drainage radius. The well paths can be vertical, deviated, or horizontal. Combinations of well types are not permitted.

Well Placement Techniques That Use Dynamic Information

Rosenwald and Green [72] use integer programming to select a static well plan. The stated objective is to select the appropriate well timing to achieve a desired time/production relationship. The algorithm starts with a set of wells with known locations and productions and a desired time/production relationship. The objective is to select the well timing that minimizes the difference between the observed time/production relationship and the desired time/production relationship. It is difficult to solve the problem in the continuous case, so the problem is discretized into two time steps. The authors use the solution for the first time step as the seed for the solution of the second time step. The authors constrain the solution such that the observed production does not exceed the desired production and only permit up to 3 producing wells. The method involves running a numerical flow simulation to obtain a series of potential coefficients relating the flow at a given location to the pressure drop in the remaining candidate well locations. The simulation is run for each of the candidate well locations, so for n candidate well locations there would be n production/pressure drop profiles. The combinatorial of production/pressure drop profiles for combinations of wells is solved by assuming superposition over the n relations.

Wagenhofer and Hatzignatiou [84] present a method that maximizes gas and water breakthrough times. The method starts by placement of a well in the horizontal middle of the reservoir and calculation of the predicted breakthrough times for gas and water. If gas breaks through first the algorithm moves the well further away from the gas/oil contact. If water breaks through first the well is moved away from the oil/water contact. The technique does not account for uncertainty.

Beckner and Song [9] present a method for field development planning and well placement in which they use simulated annealing to propose scheduling that maximizes net present value (NPV) with the use of a flow simulator. They construct a reservoir model consisting of 36 cells, the center point of each being a candidate well location. The reservoir model is set to have seven different scenarios. They randomly select a scenario and subject it to a flow simulator. The algorithm uses SA to find an optimal solution constrained by well cost, the maximum number of wells, permeability, water saturation, and relative permeabilities of oil/water and oil/gas. The authors only discuss a 2-D application.

Hybrid Techniques for Selecting Well Locations

Da Cruz et. al. [12] build a reservoir quality map by performing flow simulation with a single well at several active locations on the grid. The quality of each location is evaluated by its cumulative hydrocarbon production or its net present value. Locations not simulated are populated using kriging. Although only a single well location is simulated, the authors conjecture that the quality map is a good proxy for flow simulation under a multiple well scenario, since the interactions between wells and reservoir heterogeneity are accounted for.

Gyagler [34] uses flow simulation and a few strategic well locations to construct a data set and then uses kriging and the production data to construct a proxy map of cumulative hydrocarbon. The proxy map is searched for the optimal solution.

Pan [66] uses a similar approach to Gyagler's, but shows results from least squares and kriging to construct the proxy surface.

Comments

The techniques in the literature use a diverse range of tools to select well locations. Some reasons for the diversity of work is discussed below:

- The number of solutions may be very large. Consider selecting the locations of an injector and a producer on a reservoir model of uncertainty consisting of 100 realizations on a 50x50 grid. The exhaustive solution set has $\frac{2500 \cdot 2499}{2!} = 3,123,750$ solutions. Construction of the distributions of fluid production for each solution requires $3,123,750 \cdot 100$ flow simulation runs. Flow simulation is computationally expensive and a single flow simulation run can take hours.
- Some of the objectives in the well plan may compete. Consider the distance between injectors and producers in a waterflood. Depending on the mobility ratio, if the injector is too close to a producer, early water breakthrough can occur. If the injector is too far away from the producer hydrocarbon can be displaced in unfavorable directions.
- The global objective function for well locations may be poorly behaved over the solution space; small changes in well locations may result in dramatic changes in the goodness of the well locations.

Well Selection Preliminary Example

This Section presents a preliminary example of the well location selection algorithm developed in Chapter 5. The algorithm selects optimal waterflood well plans. The general approach for selecting optimal well locations consists of the following steps:

1. Construct a reservoir model of uncertainty and a dynamic reservoir model.
2. Construct a global objective function to quantifies the objectives the asset team has set for the well plan.
3. Select a seed candidate well plan.
4. Obtain flow simulation results for the candidate locations.
5. Calibrate the global objective function to the flow simulation results.
6. Use the calibrated global objective function and simulated annealing to select optimal well locations.

The example is based on the synthetic reservoir model used in a technical report [73]. The reservoir is situated at a depth of around 3000m and has a thickness of 60m. The grid dimensions are 50x50x20 with blocks that measure 75x75x3m. The model of uncertainty has 4 attributes: porosity (ϕ), horizontal permeability (kh), horizontal/vertical permeability ratio (kv/kh), and top of reservoir. Porosity, kh, and kv/kh are modeled as constant properties with porosity at 0.1%, 0.15%, 0.20%, kh at 500mD, 1000mD, 1500mD, and kv/kh at 0.4,0.5,0.6. The top surfaces are made using kriging. The surfaces have a structural high near the middle. The corners of the top surfaces are held constant and the elevation of the high is varied. Maps of the tops surfaces are shown in Figure 1.10.

In this example the full space of uncertainty is examined and consists of 81 possible combinations of tops, porosity, kh, and kv/kh ratio, which are submitted to flow simulation.

The flow simulation model is kept simple. The initial reservoir pressure is 315 barsa at a reference depth of 3015 m. The oil/water contact is always below the base of the reservoir at a depth of 3080 m, but the larger the top surface curvature, the closer this contact is to the base. Well control parameters have been summarized in Table 1.3. A constant productivity index (PI) of $4 \text{ m}^3/\text{day}/\text{kgf}/\text{cm}^2$ is used for all runs. All the other required parameters

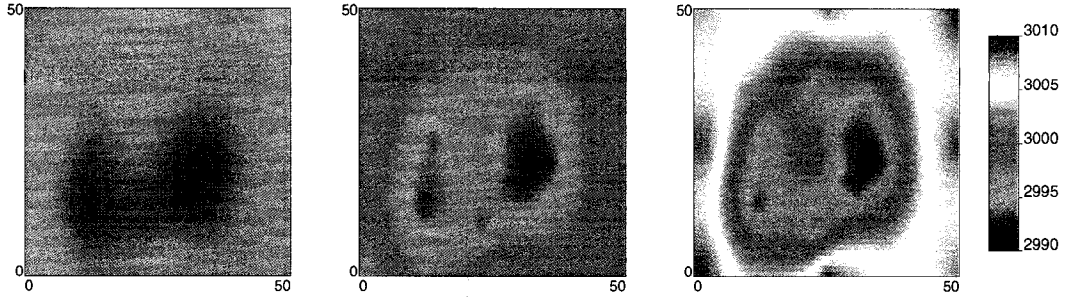


Figure 1.10: Starting from the left are the top of structure maps for the structural model of uncertainty. The mean elevations for each, starting from the left are 2994, 2995, and 3005, respectively.

Well Type	Control	Perf. Range (layer)	Min. BHP (barsa)	Max. BHP (barsa)	Max. Rate m^3
Producer	BHP	1 to 10	125		1500
Injector	BHP	1 to 20		360	2000

Table 1.3: The flow simulation model was kept simple. This Table specifies some of the initialization parameters.

needed to perform the flow simulation are based on a medium/heavy oil scenario in an Albertan basin.

Flow simulation is performed using ECLIPSE 100 (Schlumberger, 2003a) on a 2.83 GHz IBM compatible desktop computer with 2.5 GB of RAM. The time requirement for each flow simulation ranges from about 5 to 7 minutes. As reservoir model complexity increases and more input variables are used, the computation requirements will increase. In most cases enumeration of the full space of uncertainty is impractical.

Five well patterns were used to calibrate the global objective function for a total of 405 flow simulation runs, see Figure 1.11. The dashed lines are reservoir delineation boundaries.

The global objective function aims to maximize hydrocarbon pore volume within a drainage radius and minimize the distance between injectors and producers. The drainage radius is set to 500m.

The calibration procedure modifies the global objective function to be a good proxy for flow simulation. The criteria for calibration is barrels of oil equivalent well plan performance (WPP). The flow simulation results are used to estimate the NPV over the simulation life of the reservoir. The calibration procedure modifies the global objective function such that the correlation coefficient between the global objective function and the flow simulation based NPV values is maximized. The average WPP for the five well plans, starting with well plan *a* are $1612 * 10^3 m^3$, $1726 * 10^3 m^3$, $1099 * 10^3 m^3$, $1450 * 10^3 m^3$, and $1431 * 10^3 m^3$ respectively.

The scatterplot of the calibrated global objective function and the NPV values is shown in 1.12. The correlation coefficient is 0.927. The strange appearance of the scatterplot - the streaks and clusters of points - is due to the model of uncertainty having multiple realizations of constant properties rather than distributions of properties; if the properties were populated using geostatistical tools, the scatterplot would appear more cloudlike.

The final well pattern is shown in Figure 1.13. The optimal result is close to a regular five-spot well plan. The configuration component objective function does not aim to provide symmetrical results.

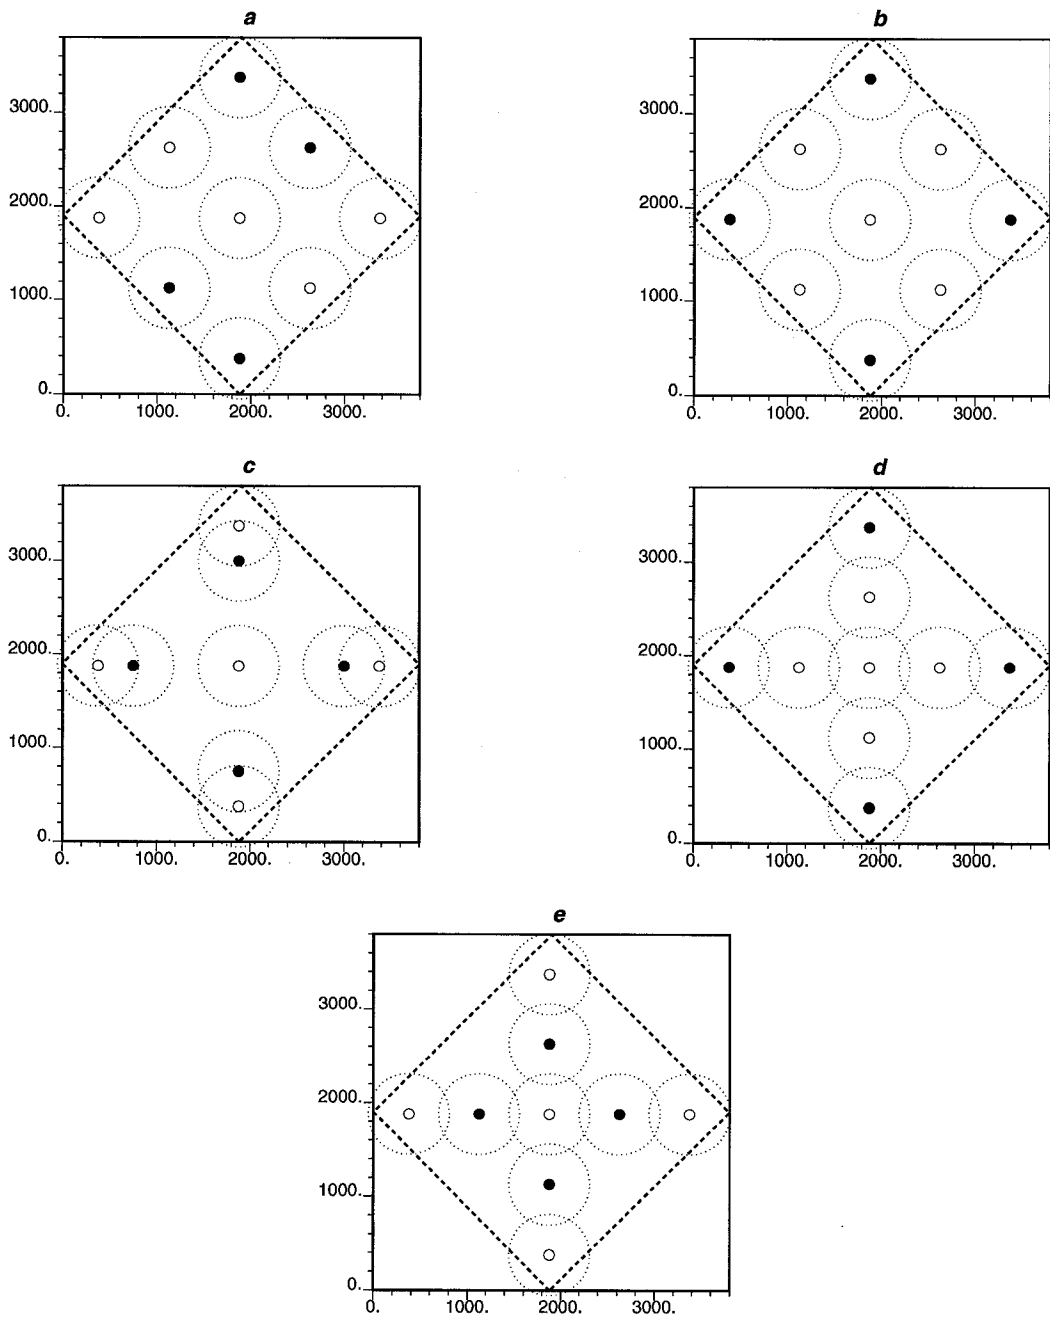


Figure 1.11: The five well plans shown are used to calibrate the global objective function. Rings represent producers and discs represent injectors.




Figure 1.12: The scatterplot shows NPV on the y -axis and the static global objective function on the x -axis. The correlation coefficient is 0.927.

The average WPP for this well plan is $1746m^3 * 10^3$. Compare this to the best WPP of the calibration well plans of $1726m^3 * 10^3$.

It is impractical to exhaustively examine all possible well locations to identify the global optimal well plan. The number of possible configurations to be evaluated is too large. There are however, works that demonstrate that a regular five-spot pattern is the best pattern for a homogenous reservoir [13, 35, 68, 86].

The spacing of the five-spot pattern can affect WPP. If the spacing is too close or too far, WPP is compromised. The patterns shown in Figure 1.14 are examined using flow simulation.

Patterns a , b , c give the respective following average WPPs: $1722m^3 * 10^3$, $1735m^3 * 10^3$, and $1752m^3 * 10^3$. Well spacing c gives the best results. A comparison of these results with the results from the well location optimizer shows that the algorithm did not select the optimal dynamic location. There are a couple of reasons for this. The problem of selecting a well location in a homogenous reservoir is particularly difficult because the properties do not vary in space and no location picked on the static information would outperform any other. The global objective function is merely a proxy to actual flow simulation results. An optimization algorithm that used flow simulation directly might place the wells in their optimal dynamic locations.

1.2 Dissertation Outline

Chapter 2 discusses the theoretical material required to develop the techniques for selecting dig limits and well locations. Geostatistics is used to construct the models of uncertainty used in the decision making. Categorical and continuous data are discussed. The concept of a random function is introduced. Kriging, the basis for many geostatistical simulation techniques, is introduced. Essential simulation concepts for constructing models of uncertainty are discussed. Simulated annealing and the theoretical issues surrounding the application of simulated annealing are covered. A technique for speeding up simulated annealing is also presented.




Figure 1.13: The map shows the final optimal well pattern.




Figure 1.14: The three maps show the three different 5 spot spacings considered.

Chapter 3 presents a technique for selecting optimal dig limits. A grade to profit transform is presented. The transform avoids the limitations associated with using grade information for classification. The transform permits easy accounting for multiple minerals and multiple ore types. The theory and implementation details of a dig limit selection algorithm are discussed.

Chapter 4 is devoted to exploring the dig limit selection technique through a number of synthetic and real data case studies. A dig limit challenge is presented in which hand drawn dig limits are compared to automatic dig limits.

Chapter 5 presents a technique for automatically selecting optimal well plans. The concept of the joint optimal solution is presented. The theory and implementation details of well location selection are covered.

Chapter 6 presents a number of case studies that use synthetic and real reservoirs are presented. Synthetic examples are used to explore the functionality and parameters of the algorithm. The value of incremental information is presented. A real example based on a water flood project in Saskatchewan is used to demonstrate the applicability of the algorithm.

Chapter 7 presents some remarks on the developed decision making techniques, including suggestions on parameter selection and future research directions.

Chapter 2

Theoretical Foundations

This chapter presents theoretical aspects of geostatistics, Section 2.1, and simulated annealing, Section 2.2. Geostatistics is the application of probability and statistical theory to spatial processes. Simulated annealing is a general optimization algorithm that has analogies to thermodynamics.

2.1 Geostatistics

Geostatistics is based on the notion that spatial variables can be described by a spatial probability law. Sample data are considered location-specific realizations of a regionalized variable that is characterized by a spatial random function. There are a number of excellent references for geostatistics: [11, 14, 21, 43, 50].

2.1.1 Random Function Concept

Geologic phenomenon have deterministic and random features. The random part is not a simple superposition of random noise on a deterministic model of geology. The random part is usually spatially correlated. Modeling the random part requires the application of the tools of probability and statistics. The probabilistic theory of random functions is widely used to describe properties of geologic phenomena [57] because it synthesizes the deterministic features and spatially correlated random features.

A random variable (RV) is a variable that can take a series of outcomes. A RV is denoted by an upper case letter, Z , and an outcome of Z is denoted by a lower case letter z . The location dependence of an RV is denoted as $Z(\mathbf{u})$, where \mathbf{u} is a vector specifying a location in space within the study area A .

A RV that can assume a continuum of values is called a continuous random variable; otherwise it is a discrete or categorical RV. A mineral grade is an example of a continuous RV and a rocktype is an example of a categorical RV.

A categorical RV has K categories, $s_k = 1, \dots, K$, and the location \mathbf{u} must belong to one and only one of the K categories. If \mathbf{u} belongs to the category s_k it is set 1, $i(\mathbf{u}; s_k) = 1$, or 0 otherwise:

$$i(\mathbf{u}; s_k) = \begin{cases} 1, & \text{if location } \mathbf{u} \text{ is in category } s_k \\ 0, & \text{otherwise} \end{cases} \quad (2.1)$$

A continuous RV can be transformed to a categorical RV. The continuous data are categorized according to K threshold values z_k , $k = 1, 2, \dots, K$:

$$i(\mathbf{u}; z_k) = \begin{cases} 1, & \text{if } Z(\mathbf{u}) \leq z_k \\ 0, & \text{otherwise} \end{cases} \quad (2.2)$$

A RV can be fully characterized by its cumulative distribution function (cdf):

$$F_Z(\mathbf{u}; z) = Prob\{Z(\mathbf{u}) \leq z\} \quad (2.3)$$

The derivative of the cdf, if it exists, is the probability density function (pdf) $f(\mathbf{u}, z) = F'(\mathbf{u}, z)$ and gives the probability for the outcome to fall in the interval dz .

A random function (RF) is a collection of dependent RVs, $\{Z(\mathbf{u}), \mathbf{u} \in A\}$. The RF is fully characterized by the set of N -variate cdfs for any number N and any choice of locations $\mathbf{u}_i, i = 1 \dots, N$, within the study area A :

$$F_Z(\mathbf{u}_1, \dots, \mathbf{u}_N; z_1, \dots, z_N) = Prob\{Z(\mathbf{u}_1) \leq z_1, \dots, Z(\mathbf{u}_N) \leq z_N\}, \forall \mathbf{u}_\alpha \in A, \alpha = 1, \dots, N \quad (2.4)$$

2.1.2 Moments of Random Variables

The first moment of a RV is the expected value or mean. It is a measure central tendency:

$$E\{Z(\mathbf{u})\} = m(\mathbf{u}) \quad (2.5)$$

where $E\{\cdot\}$ is the expected value operator. The expected value is a linear mathematical operator that weights the possible values the RV can take by their probability of occurrence. For example, the expected value of the function $g(Z)$, Z a continuous random variable, is:

$$E\{g(Z)\} = \int_{-\infty}^{\infty} g(z)f_Z dz \quad (2.6)$$

The variance measures the average variability of the data about the mean:

$$Var\{Z(\mathbf{u})\} = E\{(Z(\mathbf{u}) - m(\mathbf{u}))^2\} = \sigma^2(\mathbf{u}) \quad (2.7)$$

The variance is the difference between the squared first moment and the second moment, $E\{Z^2\}$. The square root of the variance gives the standard deviation (σ). Geostatistical studies rarely go beyond second moment statistics.

The covariance is a second moment two point statistic that measures the correspondence between RVs at two different locations, say, \mathbf{u}_1 and \mathbf{u}_2 :

$$Cov\{\mathbf{u}_1, \mathbf{u}_2\} = E\{(Z(\mathbf{u}_1) - m(\mathbf{u}_1)) \cdot (Z(\mathbf{u}_2) - m(\mathbf{u}_2))\} \quad (2.8)$$

The variogram is another second moment statistic. It measures the dissimilarity of two RVs located at two different locations:

$$2\gamma\{\mathbf{u}_1, \mathbf{u}_2\} = E\{(Z(\mathbf{u}_1) - Z(\mathbf{u}_2))^2\} \quad (2.9)$$

2.1.3 Stationarity

An RF is said to be stationary within the domain A if its multivariate cdf is invariant with any translation \mathbf{h} :

$$F(\mathbf{u}_1, \dots, \mathbf{u}_N; z_1, \dots, z_N) = F(\mathbf{u}_1 + \mathbf{h}, \dots, \mathbf{u}_N + \mathbf{h}; z_1, \dots, z_N) \quad (2.10) \\ \forall \mathbf{u} \in A, \forall \mathbf{u} + \mathbf{h} \in A$$

Stationarity of the multivariate cdf is not a property of the RF; it is a decision made by the practitioner. A decision of stationarity allows data to be pooled over the domain A and is essential for inference. Rock types and facies are sometimes used to pool data according to geologic criteria, statistical criteria, or both.

First order stationarity implies invariance of the mean, $E\{Z(\mathbf{u})\} = m$. Second order stationarity implies invariance of second order moments such as the covariance, $Cov\{\mathbf{u}, \mathbf{u} + \mathbf{h}\} = E\{(Z(\mathbf{u}) - m) \cdot (Z(\mathbf{u} + \mathbf{h}) - m)\}$.




Figure 2.1: An example experimental semivariogram is shown. The x -axis is distance and the y -axis is semi-variance. The experimental semivariogram values are shown as dots.

2.1.4 Spatial Correlation

A hallmark of geostatistics is the use of measures of spatial correlation in estimation and simulation. The two most common measures are the covariance and the semivariogram.

Under a decision of stationarity, and if the variance exists, the covariance in Equation 2.8 can be written to consider two RVs at different locations to give a measure of spatial correlation:

$$C(\mathbf{h}) = E \{Z(\mathbf{u}) \cdot Z(\mathbf{u} + \mathbf{h})\} - m^2 \quad (2.11)$$

where \mathbf{h} is a displacement vector. If $\mathbf{h} = 0$, then Equation 2.11 becomes the variance of the data:

$$C(\mathbf{h}) = C(0) = \sigma^2 \quad (2.12)$$

For this reason, the variance of the data is sometimes denoted $C(0)$. The variogram is more commonly used as a measure of spatial variability:

$$2\gamma(\mathbf{h}) = E \{(Z(\mathbf{u}) - Z(\mathbf{u} + \mathbf{h}))^2\} \quad (2.13)$$

A decision of stationarity yields a relationship between the covariance and the variogram:

$$C(\mathbf{h}) = C(0) - \gamma(\mathbf{h}) \quad (2.14)$$

In practice, the experimental *semi*-variogram, literally one half the variogram, is calculated from the data as the average squared difference between data:

$$\gamma(\mathbf{h}) = \frac{1}{2N(\mathbf{h})} \sum_{i=1}^{N(\mathbf{h})} (Z(\mathbf{u}_i) - Z(\mathbf{u}_i + \mathbf{h}))^2 \quad (2.15)$$

where $N(\mathbf{h})$ is the total number of pairs of data separated by the lag vector \mathbf{h} . The semi-variogram values are calculated over a number of locations and a number of lags over the area of interest A . The calculated values are then plotted on a chart as shown in Figure 2.1. The semivariogram values are on the y -axis and the corresponding lag vector distances on the x -axis. The dots are experimental semivariogram values at a particular lag distance.

Experimental semivariograms are calculated in the directions of minimal and maximal spatial continuity to characterize spatial correlation in 2D, and perpendicular to this plane




Figure 2.2: Four semivariogram models are commonly used to model the experimental semivariogram. Starting from the top left and rotating clockwise, is the nugget effect are the spherical model, the Gaussian model, and the exponential model.

for 3D. The experimental values are not used in subsequent geostatistical algorithms. Experimental semivariograms are modeled with mathematical functions. The mathematical models ensure licit measures of spatial variability in all directions and for all distances.

Four common semivariogram models are shown in Figure 2.2 [15]. The four models can be used alone or combined to characterize the spatial variability of the data. The nugget effect semivariogram represents the absence of spatial correlation and is the limit case for spatial correlation. Note that $\gamma(0) = 0$, so in the nugget effect model there is a discontinuity at a very small distance from $\mathbf{h} = 0$ that honors this constraint. There are many good references that discuss calculating and modeling the covariance and semivariogram. [4, 7, 22, 32, 33, 64].

2.1.5 Kriging

Kriging has two main roles in geostatistics: (1) the calculation of optimal estimates $z^*(\mathbf{u})$ of the variable Z with an associated kriging variance at an unsampled location \mathbf{u} , and (2) the construction of a posterior conditional probability distribution for the purpose of geostatistical simulation. Different kriging algorithms have become available since its formalization by Matheron [56]. All versions of kriging are variations of the basic linear regression algorithm and corresponding estimator known as non-stationary simple kriging:

$$[z^*(\mathbf{u}) - m_Z(\mathbf{u})] = \sum_{\alpha=1}^n \lambda_{\alpha} \cdot [z(\mathbf{u}_{\alpha}) - m_Z(\mathbf{u}_{\alpha})] \quad (2.16)$$

where $z^*(\mathbf{u})$ is the estimate at the location of interest, $m_Z(\mathbf{u})$ is the mean at the location to be estimated, λ_α are the weights for the α nearby data $z(\mathbf{u}_\alpha)$, and $m_Z(\mathbf{u}_\alpha)$ is the mean at the data locations. The different variations arise from different assumptions on stationarity and the application of constraints.

The most basic form of kriging is simple kriging. In simple kriging the mean is assumed to be stationary, that is, $m_Z(\mathbf{u}) = m$. The simple kriging form of the kriging estimator in Equation 2.16 is written as follows:

$$z^*(\mathbf{u}) = \sum_{\alpha=1}^n \lambda_\alpha \cdot z(\mathbf{u}_\alpha) + \left[1 - \sum_{\alpha=1}^n \lambda_\alpha(\mathbf{u}) \right] m \quad (2.17)$$

The weights, λ_α , for the data are selected such that the error variance is minimized. The error variance σ_E^2 is written as follows:

$$\begin{aligned} \sigma_E^2(\mathbf{u}) &= E \{ Z^*(\mathbf{u}) - Z(\mathbf{u}) \} \\ &= E \{ [Z^*(\mathbf{u})]^2 \} - 2 \cdot E \{ Z^*(\mathbf{u}) \cdot Z(\mathbf{u}) \} + E \{ [Z(\mathbf{u})]^2 \} \\ &= \sum_{\alpha=1}^n \sum_{\beta=1}^n \lambda_\alpha \lambda_\beta E \{ Z(\mathbf{u}_\alpha) \cdot Z(\mathbf{u}_\beta) \} - 2 \cdot \sum_{\alpha=1}^n \lambda_\alpha E \{ Z(\mathbf{u}) \cdot Z(\mathbf{u}_\alpha) \} + C(0) \\ &= C(0) - 2 \cdot \sum_{\alpha=1}^n \lambda_\alpha C(\mathbf{u}, \mathbf{u}_\alpha) + \sum_{\alpha=1}^n \sum_{\beta=1}^n \lambda_\alpha \lambda_\beta C(\mathbf{u}_\alpha, \mathbf{u}_\beta) \end{aligned} \quad (2.18)$$

The expansion shows some important features of the kriging algorithm. The first part is the global variance of the data and accounts for the variability of the data. The middle part accounts for the closeness of the data to the location to be estimated by using the covariance $C(\mathbf{u}, \mathbf{u}_\alpha)$. The last part accounts for the redundancy of the data by using the covariance between data at locations \mathbf{u}_α and \mathbf{u}_β , $C(\mathbf{u}_\alpha, \mathbf{u}_\beta)$. The error variance accounts for the configuration of the data but does not account for the local variability of the data values. That is because the data values are not incorporated into the calculation, only the stationary covariance is used.

A classic technique for optimization of a function is to take the partial derivative with respect to the variables of interest and find values for the variable that make the partial derivative equal 0. The first derivative of the error variance expansion is:

$$\frac{\partial [\sigma_E^2(\mathbf{u})]}{\partial \lambda_\alpha} = 2 \cdot \sum_{\beta=1}^n \lambda_\beta Cov \{ \mathbf{u}_\alpha, \mathbf{u}_\beta \} - 2 \cdot Cov \{ \mathbf{u}, \mathbf{u}_\alpha \}, \quad \alpha = 1, \dots, n \quad (2.19)$$

The equations above can be simplified to the equations below to give what are known as the simple kriging equations:

$$\sum_{\beta=1}^n \lambda_\beta Cov \{ \mathbf{u}_\alpha, \mathbf{u}_\beta \} = Cov \{ \mathbf{u}, \mathbf{u}_\alpha \}, \quad \alpha = 1, \dots, n \quad (2.20)$$

Linear algebra is used to solve for the weights λ_α . In matrix form, the kriging equations are written as follows:

$$\begin{bmatrix} C(\mathbf{u}_1, \mathbf{u}_1) & \dots & C(\mathbf{u}_1, \mathbf{u}_n) \\ \vdots & \ddots & \vdots \\ C(\mathbf{u}_n, \mathbf{u}_1) & \dots & C(\mathbf{u}_n, \mathbf{u}_n) \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_n \end{bmatrix} = \begin{bmatrix} C(0, 1) \\ \vdots \\ C(0, n) \end{bmatrix} \quad (2.21)$$

The matrix on the left contains the data-to-data vector covariances. The vector in the middle of the equation contains the weights that are to be solved for. The vector on the

right of the equation contains the covariances between the data values and the points to be estimated. The weights, λ_n , are found by solving the n equations simultaneously. The system of equations must be positive definite – this is one reason that the covariances are modeled – and no two data can be collocated: $\mathbf{u}_\alpha \neq \mathbf{u}_\beta$ for all $\alpha \neq \beta$. If these two criteria are met, then the kriging solution is unique and the kriging variance is non-negative.

The simple kriging error variance can be simplified to:

$$\sigma_{SK}^2 = C(0) - \sum_{\alpha=1}^n \lambda_\alpha C(\mathbf{u}, \mathbf{u}_\alpha) \quad (2.22)$$

The kriging estimate and the kriging error variance are found by substituting the weights into the kriging estimator in Equation 2.17 and Equation 2.22, respectively.

Other Variations of Kriging

Two important variations of the general kriging estimator are simple indicator kriging (SIK) and ordinary kriging (OK). SIK is the categorical version of the simple kriging algorithm. SIK predicts the probability of a categorical random variable to prevail at an unsampled location. The primary application of indicator kriging is in indicator simulation. The following references offer more complete discussions on indicators and indicator kriging [11, 32]. OK is considered to be the most robust form of kriging [11, 32, 43]. OK assumes that the mean is unknown, but can be estimated from data. The estimation of the mean is performed inside the OK system of equations; however, a decision of stationarity is still required for the variance, and covariance over the area of interest. OK constrains the kriging weights such that they sum to one: $\sum_{\alpha=1}^n \lambda_\alpha = 1$. There many other forms of kriging. The following are good references on the variety of kriging algorithms: [11, 32, 50, 83].

2.1.6 Geostatistical Simulation

Kriging gives maps that have less spatial variability than the variable it intends to model. Geostatistical simulation avoids the smoothing effect of kriging; simulation reproduces the spatial variability, histogram, and data.

The aim of geostatistical simulation is to construct a set of L alternative realizations of the spatial distribution of a variable. Each realization is equally likely to be drawn. Taken together, the set of L alternative realizations represents a model of uncertainty that reproduces the spatial fluctuations of a variable over the area of interest as well as the histogram and the data at their locations. [18, 44, 50].

The realizations are used as input to engineering simulation studies that observe the relationship between input and output uncertainty for a process or production scenario. The observed output uncertainty is used in decision making. An example is processing a reservoir model of uncertainty using a flow simulator to observe fluid production for a particular well configuration. Each realization of the reservoir model yields a different response from the flow simulator. The set of responses constitutes an output space of uncertainty. The distribution of reservoir responses for each production scenario is used to select the best configuration with respect to uncertainty. In a mining setting, the model of uncertainty could be used to observe the fluctuation in grade from the mine to the mill for the purpose of developing an ore blending strategy.

Simulation algorithms can be partitioned into continuous and categorical variable algorithms. The categorical simulation algorithms can be further subdivided into pixel based and object based simulation techniques; see below and [19, 55, 76].

Continuous Variable Simulation

Sequential Gaussian simulation (SGS) is a kriging based simulation technique for constructing continuous variable realizations that can be directly conditioned to data. According to

Chilès and Delfiner [11], one of the first applications of SGS was by Alabert [3]. Currently, SGS prevails as the preferred geostatistical simulation technique [11, 18, 76]. SGS relies on sequentially visiting unsampled nodes and simulating a value that is to be considered data for simulation at other unsampled locations under a multiGaussian model. The sequential path is selected as a random path over the area of interest to avoid artifacts [22, 32, 42]. SGS proceeds as follows:

1. Transform data to a standard normal distribution. The simulated values are drawn from Gaussian distributions.
2. Go to a location \mathbf{u} and perform kriging to get the mean and corresponding kriging variance.
3. Draw a random realization from the Gaussian distribution with the mean estimated by $z^*(\mathbf{u})$ and variance of $\sigma_{SK}^2(\mathbf{u})$ to get the simulated value $z_s(\mathbf{u})$.
4. Add $z_s(\mathbf{u})$ to the set of data. Using kriging and previously simulated values forces the covariance to be reproduced over the grid [18].
5. Visit all locations in random order and repeat steps 2-4.
6. Back-transform all data values and simulated values after all locations have been visited.
7. Create another realization by repeating the previous steps.

The data distribution is transformed to a standard normal distribution using the normal scores transform $y = G^{-1}(F_Z(z))$. The transform permits the application of the multi-Gaussian model and the use of some of its congenial properties. Under the multiGaussian model, the mean and variance completely specify a distribution. The mean and variance of the conditional cdf at location \mathbf{u} are identical to the kriged estimate $z^*(\mathbf{u})$ and variance σ_{SK}^2 obtained from the $n(\mathbf{u})$ data $z(\mathbf{u}_\alpha)$. The standard normal distribution is preserved (in expected value) over the area of study. This provides a consistent platform for backtransformation of the simulated values back into the original data space.

Some references that discuss the application and theory of SGS include [11, 18, 32, 55]. Other continuous variable simulation algorithms include turning bands [44], simulated annealing [20] and LU simulation [2].

Categorical Variable Simulation

An important application of categorical variable simulation is the construction of facies or rock type models of uncertainty. The categorical simulators can be classified as pixel based and object based simulators. Pixel based categorical simulators regard each pixel as belonging to a category. The categorical values correspond to facies types and are assigned on a pixel by pixel basis. Sequential indicator simulation is a popular pixel based technique. Object based facies modeling represents the spatial distribution of facies as a collage of objects.

Sequential Indicator Simulation

SIS uses simple indicator kriging to estimate the probability for category s_k to prevail at location \mathbf{u} . The formulation of the simple indicator kriging algorithm is the same as the simple kriging approach presented in Section 2.1.5 except that the categorical information $i(\mathbf{u}; s_k)$ replaces the continuous information z .

SIS visits the nodes in the grid sequentially, as in SGS. The basic idea is to use SIK to construct the local probabilities for each of the k categories to prevail and randomly draw a category [46, 47, 48, 49, 61]. One of the first applications of SIS appears in [51]. SIS can be used for both categorical and continuous data [17, 32].

Object Based Facies Modeling

Object based modeling is founded on the notion that facies or rock types can be modeled as a collage of objects or shapes within a matrix [18, 26]. For example, a facies model might be represented by a collage of ellipses of different sizes, channels, or channels with levees of varying widths. Logical rules can be used for placing the objects: they can overlap, erode, or combine. Logical rules can also be applied that prevent particular facies types from existing adjacently [11, 25, 26].

Object based simulation requires many parameters to specify the shapes and dimensions. Seismic information and analogue information derived from outcrops or modern processes can be used to infer these parameters.

The objects must be placed so that they appear realistic and honor available data. Two approaches are used. The first draws objects and keeps only those that satisfy some constraints. The second directly conditions the objects by randomly drawing an object and later conditioning it by iteratively modifying its parameters until the data are honored.

2.1.7 Spaces of Uncertainty

The goal of this Section is to demonstrate that a model of uncertainty occupies a space of uncertainty that can easily become ungainly and difficult to consider in the context of decision making.

The root purpose of building a model of uncertainty is to observe how input uncertainty affects a process or mathematical model. The process or mathematical model is referred to as a transfer function because the function transfers input uncertainty to output uncertainty. A transfer function may be complex and provide a nonlinear mapping from the input space of uncertainty to the output space of uncertainty. Examples of transfer functions include flow simulators that numerically simulate the flow of fluids through porous media and mining simulators that mimic the process of excavating and shipping material to the mill for processing. The output space of uncertainty is built by feeding the input space of uncertainty, one realization at a time, into the transfer function. The output space of uncertainty is used for decision making.

Most times, the input space of uncertainty is larger than the output space of uncertainty. Consider for example the input space of uncertainty for a hypothetical reservoir. A guiding philosophy for the construction of a model of uncertainty is to build the model hierarchically according to data scale. The conceptual geologic model is commonly the largest scale source of information. Among other things, the conceptual geologic model specifies the depositional environment and provides insight into the internal architecture of the reservoir. Suppose the conceptual geological model for the hypothetical reservoir specifies two scenarios: (1) A turbidite reservoir deposited in a confined basin that is composed of several channelized flows, or (2) A turbidite reservoir deposited in a poorly confined basin and composed of several nearly lobe-like flows. Next on the hierarchy might be structure. Structural information is usually provided by seismic data. Suppose that the hypothetical reservoir has a salt lense above it. Salt propagates sound differently than sandstone. The geophysicist might have a model of uncertainty on velocity, a basal surface, and optimistic and pessimistic structural models. Next on the hierarchy is the facies model. There is uncertainty in the facies proportions and their spatial correlation. Last in the hierarchy is the petrophysical properties and these will have uncertainty in their distributions and spatial correlation as well. This concludes the characterization of uncertainty on the static parameters, that is, uncertainty not related to the fluid properties in the reservoir. For most reservoirs there is uncertainty in the PVT tables, the relative permeabilities of the fluids, and fluid density. Uncertainty in fluid properties are also included in the input space of uncertainty.

For the decision making exercise, suppose that for the hypothetical reservoir, water and hydrocarbon rates are required so that appropriate facilities can be priced. In this case the output space of uncertainty is far smaller than the input space of uncertainty. Clearly, it is

challenging for an engineer to simultaneously consider both the input space of uncertainty and the output space of uncertainty and with rigor make optimal decisions.

The hypothetical example discussed above shows how large the space of uncertainty can become. Adding input uncertainties can cause the space of uncertainty to grow very large very fast. A challenge in decision making under uncertainty is selecting a decision given a large space of uncertainty. The space of uncertainty can easily become so complex that it is impossible to visualize or comprehend. Optimization tools can be used to select the best decision in light of the model of uncertainty and do so consistently and with rigor. Later in this thesis simulated annealing will be used to select dig limits and well locations for a given model of uncertainty.

2.2 Simulated Annealing For Decision Making

Engineers are often faced with the problem of selecting the best decision from a set of alternatives. The best decision is the one that minimizes loss as quantified by a global objective function. Often the set of alternative decisions is large and an optimal decision cannot be clearly identified because there are many viable alternatives or because there are only subtle differences among the alternatives, or for both reasons. Simulated annealing is a general optimization algorithm that is good at identifying an optimal solution in the presence of a large solution space with multiple viable solutions.

2.2.1 Motivation for Using Simulated Annealing

Three classic deterministic optimization tools are the gradient based techniques, the simplex search algorithm, and linear programming. These examples are sensitive to the formulation and behavior of the global objective function. The global objective functions for the dig limit and well location selection problems are not well behaved and their behavior precludes the use of classic optimization routines. There are two reasons that the gradient optimizers cannot be used. First, the global objective functions have multiple local minima and gradient optimizers have no facility to climb out of local optima and continue searching for global optima. Second, gradient type optimizers require differentiable global objective functions; however, the dig limit and well location global objective functions are not differentiable since they are functions of non-parametric RVs. The simplex search algorithm cannot be used because it has no facility to climb out of local optima. Implementation of linear programming to solve the dig limit and well location selection problems under a full model of uncertainty maybe impractical. There is an upper limit to the size of the linear programming optimization model that is dictated by computing hardware and software. For example, a linear programming model with 10,000 variables and 10,000 constraints could have a "coefficient matrix" with $10,000 \times 10,000 = 100$ million elements. Arrays this size might be too large or cumbersome to handle. Linear programming optimization performs many matrix inversions. Inversions can be give numerically unstable results on large arrays when array components are very large or small values. Furthermore, the global objective function for well placement may not always be linear.

Deterministic optimizers are not good choices for selecting dig limits and well locations. The dig limit and well location selection problems share some similarity to the class of problems known as nondeterministic polynomial time complete (NP-complete) problems. A NP-complete problem is a decision problem that is solvable only in a non-deterministic way. The term nondeterministic implies that a trial and error approach for generating solutions is required, rather than a deterministic or rule based approach. This means that it is possible to guess the solution (by some nondeterministic algorithm) and then check it. The guesses could be purely random or established by some strategy.

The term "polynomial time complete" refers to the predicted period of time or number of steps required to arrive at an *optimal* solution. The number of steps to arrive at a solution increases with the number of solutions N as $exp(const. \times N)$. The term *optimal*

is emphasized because for some of the algorithms, the solution cannot be shown to be the global optimal solution, but can be shown to quickly provide a solution that is within a range of optimality. Not all NP-complete problems have good optimization algorithms, and for some problems, finding a good and quick approximately optimal solution is enough.

A classic example of an NP-complete problem is the traveling salesman problem: a salesman must visit n cities only once and return to the starting city in the shortest total distance. The problem can be constructed to have multiple constraints [69]. The postman problem and the Hamiltonian circuit are two other examples. An entertaining NP-complete problem can be found on most PCs and is known as Minesweeper. Some practical NP-complete problems that need solutions are the placement of computing chips on a printed circuit board, the routing of garbage trucks, and the sequencing of traffic lights.

Genetic algorithms and simulated annealing have been demonstrated to be good at solving NP-complete problems. Both algorithms have the feature of being able to climb out of local minima.

Genetic algorithms are inspired by Darwin's theory of evolution. The central idea is that, given an initial solution, an optimal solution can be found by iterative evolution from solution to solution. The genetic algorithm was first introduced by Holland [37].

Genetic algorithms are popular. They are easy to code and are easy to use in parallel computing environments. There are several drawbacks to the application of genetic algorithms. Genetic algorithms are not designed to ergodically sample the solution space so there is no statistical guarantee of global convergence [41, 67]. Genetic algorithms are CPU intensive and have large memory requirements. The perturbation mechanism can be too aggressive for some problems. The perturbation mechanism swaps out large portions of the solution and this leads to potentially large fluctuations in the global objective function. The large fluctuations may require that a very large number of perturbations be used to find an optimal solution. There are proofs of global optimality for certain applications of genetic algorithms, but these rely on simulated annealing-like assumptions and parallelisms and require elitism to be part of the algorithm [36]. Elitism is the case where the new population always carries the best parents from the previous population into the new population.

2.2.2 Simulated Annealing Background

Two essential components led to the development of simulated annealing as an optimization algorithm: (1) Monte Carlo simulation, and (2) the numerical modeling of the thermodynamic process of annealing. Monte Carlo simulation is a term that S. Ulam coined for a statistical sampling procedure that was later formalized by N. Metropolis [59]. The numerical model for the annealing process was provided by Metropolis et al. [58]. The algorithm was developed for observing the properties of alloys that avoided having to physically make them and is sometimes referred to in the literature as the Metropolis Algorithm.

Kirkpatrick et al. developed the first simulated annealing optimizer in 1982 and published the work in 1983 [52]. They used the term *simulated annealing* to describe their use of the Metropolis Algorithm for the purpose of optimization. Černý independently developed an optimizer that is based on thermodynamics and has also been called simulated annealing; however, his work was not published until 1984 [82].

Geman et al. were the first to prove the necessary and sufficient conditions for the convergence of the algorithm [29]. The first application of SA to the geosciences was by Farmer [28]. Deutsch applied SA to the problem of generating models of uncertainty [17, 20].

Since its development, SA has solved many optimization problems once thought intractable. The problems that SA is particularly suited to solving share three characteristic features: (1) the decision space is very large; (2) the global objective function is not well behaved; and (3) global objective function consists of multiple objectives that sometimes conflict. The problems of dig limit selection and well location selection have these features.

2.2.3 Thermodynamic Annealing

Thermodynamic annealing is a process used to modify material at a molecular level. Thermodynamic annealing involves heating a material to a temperature just below its melting point and slowly cooling it. At high temperatures, the molecules have high mobility and are poorly organized. The molecules move into configurations of increasing order as the material is slowly cooled. If the material is cooled quickly, the molecules may not have sufficient opportunity to find highly organized configurations, but instead take on poorly organized configurations that have high energy. Slow cooling keeps the whole of the material in thermodynamic equilibrium and ensures that all the molecules have a chance to find highly organized low energy configurations. If the material is cooled under optimal conditions, the final result is a crystal with highly organized configurations of molecules and minimal energy.

The molecules sample all possible configurations during cooling and molecules will always exist in high and low energy configurations at any temperature. Thus, at any given temperature the molecules move between energy states and will possess distribution of energy states. Molecular movement between energy states is probabilistic and is specified by the Boltzmann distribution.

The probability of a molecule to change energy states, $P\{\Delta E\}$, where ΔE is the corresponding change in energy and is expressed as:

$$P\{\Delta E\} = \begin{cases} 1, & \text{if } \Delta E \leq 0 \\ e^{(-\frac{\Delta E}{kT})}, & \text{if } \Delta E > 0 \end{cases} \quad (2.23)$$

where k is the Boltzmann coefficient that relates energy to temperature, and T is the absolute temperature in Kelvin. Molecules will always move from a high energy configuration to a low energy configuration when it can: $P\{\Delta E\} = 1$. Sometimes molecules will take on high energy configurations.

The initial temperature and rate of cooling are dictated by an annealing schedule and are critical to the final state of organization in the material. If the initial temperature is too low or if the material is cooled too fast, the material may have defects.

Every material has its own annealing schedule that achieves a high level of organization in a minimum cooling time. The key to this fast promotion of organization is the critical temperature. The critical temperature is the temperature at which molecules have enough energy to move but not so much energy that they are dislodged from low energy configurations. A well designed annealing schedule heats the material just above the critical temperature, holds the material at the critical temperature sufficiently long enough to get all the molecules mobile, cools very slowly over the critical temperature, and then cools quickly to the final resting temperature.

2.2.4 Thermodynamic and Simulated Annealing

SA is a numerical analogy to the thermodynamic process of annealing. SA numerically simulates the effect of heating a material and slowly cooling it. The following discussion lists the correspondence between thermodynamic annealing and simulated annealing.

The energy of the material corresponds to the global objective function: The global objective function is a measure of loss, given a configuration or set of decision values. In SA, loss corresponds to energy. An optimal solution is a least loss solution and corresponds to a highly organized low energy molecular configuration. Mathematically, we let $\Delta E = \Delta O_g$, where O_g is the calculated value of the global objective function.

Molecular configuration corresponds to decision variable values: The values the decision variables take correspond to a configuration of molecules. High global objec-

tive function values have decision variable values far from values that represent the optimal solution.

The interaction of particles corresponds to the SA perturbation mechanism: In thermodynamics, molecular mobility is provided by interactions with adjacent molecules. In SA, the values the decision variables take are randomly modified by a perturbation mechanism. The perturbations randomly sample different solutions and are analogous to molecular sampling of different configurations.

Thermodynamic temperature corresponds to simulated temperature: In thermodynamics temperature controls the probability of a molecule to move from one configuration to another via the Boltzmann distribution. With SA, temperature controls the probability to move from one configuration of decision variables to another.

The thermodynamic annealing schedule corresponds to a simulated annealing schedule: In both thermodynamical and simulated annealing, the annealing schedule specifies the initial temperature, T_0 , and the rate of cooling. The simulated annealing schedule specifies how to reduce the temperature T as a function of the perturbations. Time is analogous to the number of perturbations. Holding T constant over a large number of perturbations is analogous to holding a material at a constant temperature for a long time.

The thermodynamic Boltzmann distribution corresponds to the simulated Boltzmann distribution: In thermodynamics, the Boltzmann distribution specifies the probability of a molecule to move from one configuration to another, whereas in simulated annealing, it specifies the probability for moving from one solution to another. The probability is a function of the magnitude of change in the global objective function, $\Delta O_g = O_p - O_i$, where O_i is the value of O_g before the perturbation and O_p is the value after the perturbation, and the current temperature. If the perturbation results in a decrease in O_g , then $P_{\Delta O_g} = 1$. This leads to the SA probabilistic acceptance/rejection rule:

$$P\{\Delta O_g\} = \begin{cases} 1, & \text{if } \Delta O_g \leq 0 \\ e^{\frac{-\Delta O_g}{T}}, & \text{otherwise} \end{cases} \quad (2.24)$$

Acceptance or rejection of loss-increasing perturbations is performed by Monte Carlo simulation: a random number in the set $[0, 1]$ is drawn and compared to $P\{\Delta O_g\}$. If the random number is less than $P\{\Delta O_g\}$, the perturbation is accepted. Note that the only modification to the Boltzmann distribution in SA is to group or combine the T and k parameters.

2.2.5 Simulated Annealing Algorithm

Figure 2.3 shows a flowchart for the simulated annealing algorithm. The following text describes the flowchart.

- Use an initial guess to calculate an initial value for the global objective function.
- Perform a perturbation: randomly draw a decision variable and randomly change it.
- Compute ΔO_g , the change in the global objective function due to the perturbation.
- Use the SA probabilistic acceptance/rejection rule in Equation 2.24 to determine if the decision variable values will be saved as the current solution or not.
- Consider the annealing schedule and determine if T needs to be reduced.

- This step in the algorithm contains tests to determine if the algorithm should be stopped. If any of the stopping criteria have been satisfied, the algorithm is stopped. If the none of the stopping criteria have been satisfied, the SA algorithm performs another perturbation.

2.2.6 Annealing Schedule

The annealing schedule specifies the initial temperature and how the temperature will be reduced. The annealing schedule affects the final solution.

The initial temperature must be high enough to permit enough perturbations to occur to find an optimal solution. A high initial temperature gives a correspondingly high probability of accepting global objective function increasing perturbations from the Boltzmann distribution. This feature facilitates climbing out of locally optimal solutions, and permits broad sampling of more of the solution space. If the initial temperature is too high, it may take a large number of perturbations to impose order again. The initial temperature must be low enough that too many productive changes to the decision variables are not undone by subsequent perturbations. The technique outlined in [63] identifies the critical temperature for helping to select an initial temperature; otherwise, selecting an initial temperature requires an iterative approach. One iterative approach for selecting T_0 would be to start with a high value of T_0 , quickly cool, plot a chart of the value of the global objective function versus the number of perturbations, and compare the plots for a number of initial temperatures. A good initial temperature would be one that is slightly higher than the temperature that gives the fastest decline in the global objective function.

Reducing the temperature as the algorithm proceeds systematically reduces the probability of accepting loss-increasing perturbations; eventually, only loss-reducing perturbations are accepted. If the temperature is reduced too quickly, the final solution may be sub-optimal because SA will not have performed enough perturbations to climb out of a local solution. If temperature is lowered too slowly, the annealing schedule may not be of practical use.

There are three common approaches for specifying a reduction schedule for T : (1) use a continuous function to reduce T with each perturbation; (2) reduce T after a number of perturbations; and (3) accept only global objective function reducing perturbations, that is, set $T = 0$ to give a procedure known as quenching. Schedules 1 and 2 rely on a good previously selected initial temperature.

The temperature reduction schedule proposed by Kirkpatrick et al. systematically reduces T at each perturbation according to the following annealing schedule [53]:

$$T_k = \lambda T_{k-1} = \lambda^k T_0 \quad (2.25)$$

where T_k is the temperature for the k^{th} perturbation, T_0 is the initial temperature, and λ is a multiplicative factor less than 1, but greater than 0, called the reduction factor. Reported useful values of λ are in the range $0.8 < \lambda < 0.99$.

Geman et al. [29] used a continuous inverse log function to set a temperature reduction schedule:

$$T_k = \frac{T_0}{\ln(k+1)}, \quad k = 1, \dots, \infty \quad (2.26)$$

Geman et al. use the above annealing schedule to prove the convergence property of SA. This annealing schedule takes an exceptionally large number of perturbations and is often impractical.

A stepwise annealing schedule, as developed in Numerical Recipes in Fortran [69], is used in this thesis. The approach is as follows. Select a good initial temperature and multiply T by a reduction factor λ if: (1) a number of attempted perturbations ($k_{attempted}$) have been carried out, or (2) a number of accepted perturbations (k_{accept}) have been performed. This

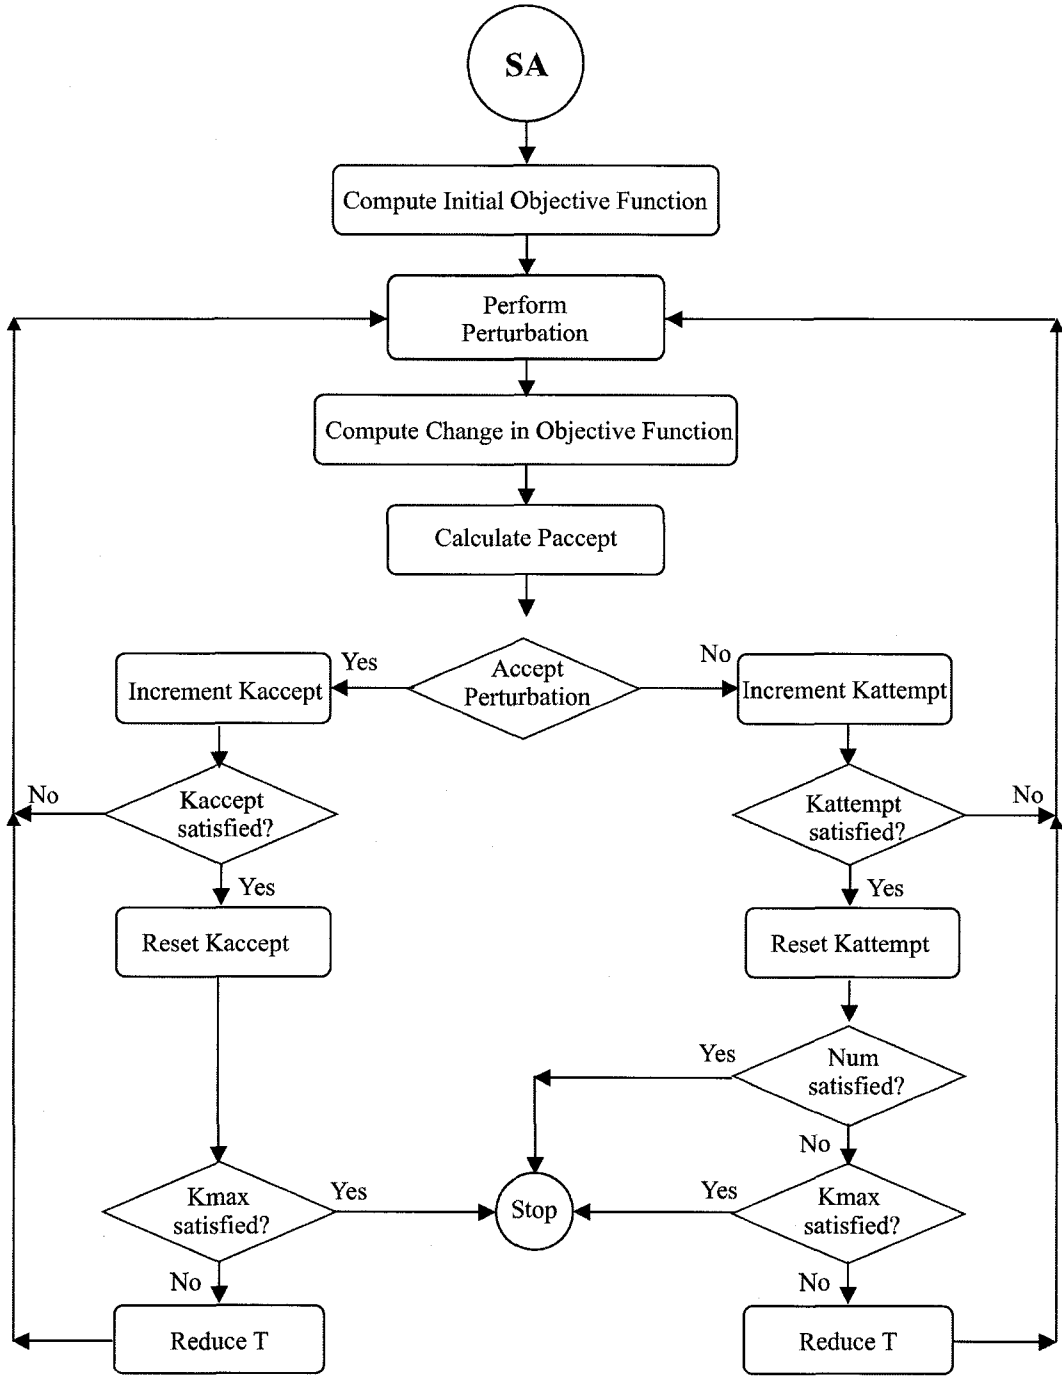


Figure 2.3: This is a flowchart for the simulated annealing algorithm.

approach permits a number of perturbations to occur for a specified T and allows SA to achieve equilibrium with respect to T_k .

The selected approach is the most common approach in the geostatistical literature [18, 22, 32]. One drawback to this approach is that a number of parameters (T_0 , $k_{attempted}$, k_{accept} , and λ) need to be specified. Optimal parameter selection can be difficult. $k_{attempted}$ should be roughly 10 times larger than k_{accept} to permit the algorithm to sample the solution space well. k_{accept} should be selected according to the size of the solution space, i.e., if the grid for selecting dig limits or well locations is small, k_{accept} should be small values as well. It is difficult to provide more precise recommendations because each dig limit and well placement problem is unique. An iterative approach may have to be implemented that compares the value of the final global objective function to the value of the parameters. The following references discuss techniques on selecting initial simulated annealing parameters [1, 65].

Stopping Criteria

There is no way to determine if an optimal solution has been found. A stopping criteria must be invoked to stop the algorithm. Three stopping criteria are used in this thesis. The algorithm stop if a specified number of perturbations have been performed (k_{max}) or if too many $k_{attempted}$ have been carried out without reaching k_{accept} . The algorithm also stops if the percent change in O_g over a number of $k_{attempted}$ perturbations is less than ΔO_{min} .

Comments on the Annealing Schedule

In the early stages of SA, large scale reductions of the global objective function are made. Progressively smaller reductions are made as the algorithm proceeds. At high T values, large scale improvements are easily undone due to the probability of accepting unfavorable perturbations. This means that perturbations that result in large reduction in O_g persist whereas perturbations that have small scale improvements are more often overwritten. Reducing T reduces the probability of accepting unfavorable perturbations and SA starts to converge. If T is reduced too quickly the algorithm will not have sampled a sufficient portion of the solution space and will provide a local optimal solution.

There are a number of vague statements about the annealing schedule being "too fast" or "too slow". The selection of these parameters is problem-specific thus parameter selection requires some experimentation. Each dig limit and well location selection problem is unique. One obvious feature that makes each problem unique is the spatial distribution of properties. There is no way to predict what the annealing schedule should be thus a trial and error approach must be applied. This is a disadvantage of SA, but it is shared by many of the algorithms capable of solving combinatorial problems.

2.2.7 The Global Objective Function

The global objective in most decision making problems is a composite of several smaller component objectives. For example, optimal dig limits should include as much ore as possible, yet minimize the cost of mining. Similarly, the selection of optimal static well locations should maximize recoverable petroleum and minimize the number of wells. The global objective function combines multiple objectives into a single global objective function as a weighted sum of component objective functions:

$$O_g = \sum_{i=1}^N \lambda_i \cdot O_i \quad (2.27)$$

where λ_i is the weight for the i^{th} component objective function, and O_i is the i^{th} component objective function.

Selecting weights can be challenging because the weights have two effects on the optimization. First, the weights calibrate the component objectives to have the same units. Second, the weights control the influence a component objective has on the final solution. Selecting weights may be an iterative procedure. For example, the weights for selecting dig limits with large mining equipment may not be applicable to dig limit selection with small equipment.

The global objective function in this thesis is posed as a proxy to profit. The weights convert the component objectives to values that are metrics of profit. It is acknowledged that the global objective function is not measuring profit in the accounting or economic sense. There are advantages to the proxy of profit approach. Profit is a universally understood unit, most objectives can be easily converted to profit-like units, and maximization of profit is a goal for most mining and petroleum companies.

Component objective functions measure loss with different magnitudes. Component objectives that have high magnitudes tend to control the optimization and the final solution because satisfying them leads to the large reductions in the global objective function. Scaling component objective functions gives the added flexibility of suppressing or stimulating individual component objectives.

The formulation of the global objective function for the dig limit selection problem is discussed in Chapter 3, and the formulation of the global objective function for the well locations selection problem is covered in Chapter 5.

2.2.8 Perturbation Mechanism

The perturbation mechanism is important because a perturbation mechanism that performs changes that are too large or too small will take longer to find an optimal solution. If the perturbations are too large, then the perturbation mechanism may not be able to access the entire solution space or may require a large number of perturbations to sufficiently sample the solution space. If the perturbations are too small, a large number of perturbations maybe required to sufficiently sample the solution space or a large number of perturbations may be required to climb out a local optima.

The perturbation mechanism used here is to randomly select a dig limit vertex or a well and randomly modify its location. There are no theoretical constraints on the magnitude of the change, but experience has shown that for dig limit selection, maximum perturbation distances of 20% of the block size gives good results. For the well location selection problem, maximum perturbation distances of 20% of the field size gives good results.

2.2.9 Simulated Annealing and Optimality

Depending on the annealing schedule, SA is statistically guaranteed to converge on an optimal solution [1, 29, 40]. This is known as the SA property of (weak) ergodicity. The property implies that given enough perturbations SA will converge to an optimal solution. The property is classified as weak because it depends on the annealing schedule. The required annealing schedules that guarantee convergence may take too long to be of practical use, and the lack of speed is one of the leading criticisms of SA. There have been many different ideas for speeding up SA [1, 8, 40, 63]. The trade-off is a reduction in the effectiveness of the convergence property and hence the possibility of a suboptimal solution. The assumption is that the fast optimal solution is close enough to being the global optimal.

The annealing schedule parameters require fine tuning to obtain the best results. Fine tuning can be performed by comparing plots of the number of perturbations versus the global objective function for a variety of parameter selections. An example plot of the number of perturbations versus the global objective function is shown in Figure 2.4. The dashed line shows an optimization that finished with a higher final global objective function value relative to the optimization shown in by the solid line output. It can be inferred that the optimization shown by the dashed output provides a sub-optimal solution.

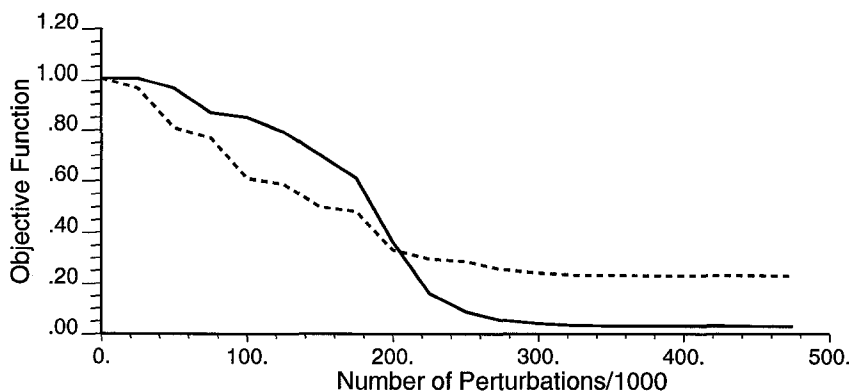


Figure 2.4: Shown are the results from two optimization runs. The dashed results are sub-optimal when compared to the results shown by the solid line.

2.2.10 Simulated Annealing in the Literature

The simulated annealing algorithm is amenable to extensive modification. SA has three important components that can be modified: the acceptance/rejection rule, the annealing schedule, and the perturbation mechanism. Researchers have modified each of these components to spin off new versions of simulated annealing.

The SA description above is a classic application of SA to optimization. Szu et al. [80] show a simulated annealing technique called Fast Simulated Annealing (FSA) that can also be proven to find an optimal solution.

Ingber presents Very Fast Simulated Reannealing (VFSR) that uses a designed distribution rather than the Boltzmann distribution for calculating the probability of accepting/rejecting perturbations. Using this designed distribution he was able to propose an even faster annealing schedule that also can be proven to converge [39]:

$$T_k = T_0 \exp(-ck^{1/n}) \quad (2.28)$$

Ingber's VFSR paper [39] presents an alternative to systematically reducing T throughout the course of the optimization. He sets T as a RV and draws T values from a distribution that is conditional to the k^{th} perturbation. The effect is that the algorithm cycles through periods of convergence and randomization randomly. The author calls this feature re-annealing because higher temperatures can be revisited.

The perturbation mechanism can be easily modified to reduce the required number of perturbations. Press et al. [69] discuss a novel perturbation mechanism that uses a simplex search type hyperplane to identify the direction of global objective function descent. The perturbation mechanism preferentially draws perturbations in the direction of global objective function descent.

These findings are important but not applicable to this thesis. The distribution of the decision variables are x, y, z coordinates. In the case of the dig limit selection problem, the coordinates specify the locations of the vertices of the dig limit polygon. In the case of the well location selection problem, the coordinates specify the trajectories of the wells. The coordinates do not follow a distribution that has a documented proof of convergence in SA.

Tailoring the perturbation mechanism to search in directions of optimality was not considered. A tailored perturbation mechanism is useful for global objective functions that can be quickly evaluated, because solving for the hyperplane requires more objective function evaluations. With the dig limit and well location selection problems the most CPU-intensive part of the algorithm is the global objective function evaluation. An algorithm that requires

more evaluations may actually take more CPU time to converge than one that does not have a tailored perturbation mechanism.

Chapter 3

Semi-Automatic Dig Limit Selection

This Chapter presents the theoretical and implementation details for semi-automatic dig limit selection. The technique selects dig limits that maximize profit under uncertainty and under the constraints of mining equipment.

3.1 Theory

A dig limit is a polygonal boundary used in open pit mining to delineate regions of ore and waste for the shovel operator. The operator depends on the dig limits when ore and waste cannot be discerned visually. Dig limits are selected by the mining engineer or production geologist. They use sample data and geologic information to help select dig limits. They weigh many considerations when selecting dig limits:

- **Uncertainty:** Relatively sparse samples are collected from blastholes or dedicated grade control drilling and there is uncertainty in grade in the intersample regions. Uncertainty makes it impossible to know the true location of the ore/waste contact. The consequence is the possibility of misclassifying material and compromising profit.
- **Mining Equipment:** The mining equipment needs room to operate and different pieces of equipment have varying abilities to follow dig limits precisely. Dig limits that do not account for the limitations of the mining equipment compromise profit because there is unnecessary ore loss or dilution of ore, or the mining equipment spend too much time attempting to follow the dig limits.
- **Maximize Profit:** The mining engineer or production geologist aims to select dig limits that maximize profit.
- **The Presence of Multiple Minerals:** Many mines have multiple minerals. A volume of material may have subeconomic grade for the primary mineral of interest but contain enough secondary or tertiary mineral for the volume to be called ore. To calculate profit, the mine must consider total mineral content and how mineral content might effect recovery or the sale of product.
- **The Presence of Contaminants:** Contaminants may adversely affect the recovery process or the final selling price. Contaminants can catalyze or suppress the recovery process and as a result additives may be required to maintain control of the recovery process.

- **Production Scheduling:** A mine is driven by its production schedule. The production schedule determines the mine life and the cash flow, including capital requirements, operational costs, and revenues. A mine's production schedule specifies tonnage, grade, or concentrate quality targets over a time period. The mining engineer sets short term plans to satisfy the production schedule. If the targets are not met, the mining engineer needs to set short term plans that satisfy the shortfall. This could involve moving mining equipment, reclassifying material, selecting dig limits that are easier to mine so that the equipment can mine more easily for improved production, or directing efforts on waste removal to expose more ore.
- **Mining Sequence:** The mining sequence refers to the order that material is mined. There may be good reasons to preferentially mine waste over ore, when the mill is shut down for example. Alternatively, economic performance might be improved by prioritizing ore production. The need to meet production targets is another. A mill shutdown might force the mine to mine waste for a period of time.
- **Blending Strategy:** Some mines employ a blending strategy to ensure the mill is fed with material of uniform quality. Blending strategies may force the mine to preferentially mine ore, protore, or waste to maintain stockpiles and control the mill feed quality. Blending strategies are particularly important in mines where mineral or contaminant quality is highly variable. Some blending can be implemented in the pit through strategic selection of dig limits.
- **Direction of Mining:** The direction of mining can be important when the mineral deposit is highly anisotropic. If mining proceeds oblique to the continuity of the ore, dilution may be increased.
- **Blasting:** Blasting is used to fragment rock so that mining equipment can easily excavate it. Blasting mixes and moves material. Blasting can change the classification of material from ore to waste.
- **Equipment Positioning:** The dig limits must provide adequate access and working room for the mining equipment. The shovel should not have to wait on haul trucks and haul trucks should not spend too much time waiting on the shovel.

Selecting dig limits under all the above considerations requires compromises. The satisfaction of some considerations may violate others. Some considerations may be very difficult to account for and simplifications may be required. Increasing sample density may decrease uncertainty, but sampling is expensive and cuts into profit. Dig limits that maximize grade may be too tortuous for the mining equipment to mine efficiently and increase the cost of mining. Conversely, dig limits that are smooth enough for the mining equipment to mine with optimal efficiency might dilute and/or lose ore and compromise profit. A lack of information prior to mining makes it impossible to specify the optimal location of mining equipment or select the best direction for mining. The effect of blasting is difficult to account for because many factors, such as the loading, the blast pattern, and geology influence the blast. Evaluation of the bench after blasting is difficult because the material no longer provides stable footing.

An alternative to selecting dig limits is to classify material by the dipperful. Dipper-by-dipper classification would avoid having to take assay information from boreholes, the need to account for movement of material due to blasting, and even avoid having to select dig limits. Unfortunately, dipper-by-dipper classification is impractical. Current technology cannot support dipper-by-dipper sampling and classification. Dipper-by-dipper classification may not even be practical because there would have to be ore and waste haul trucks available to accept a dipperful of either classification as mining progressed, and some mines may not have enough room on a bench for a shovel and two haul trucks. Positioning of two haul trucks and a shovel would be suboptimal when maintaining a toe or crest because the

shovel might have to spend too much time manoeuvring. An alternative to dipper-by-dipper classification is classification of material by a haul truck load. This approach is also impractical: the shovel operator needs to know the classification of each dipperful prior to loading, and must be able to stay in a region of ore or waste long enough to completely load a haul truck.

If some compromises are made, the dig limit selection problem can be posed as a tractable problem having a global objective function that can be optimized with the use of SA.

Although the geostatistical tools covered in Section 2.1 are the best way to manage uncertainty, invoking geostatistics involves compromises. The model of uncertainty represents the mineral deposit as a grid of points. For convenience, the blocks are usually regular and at a coarser resolution than the geology. It is impossible to reproduce heterogeneity at scales smaller than the blocks in the grid and some detail is lost.

Another compromise that is a consequence of using geostatistics is that the grid cannot be used directly to select dig limits. Ideally, the mining engineer would use the grid to identify blocks of ore and waste and mine the blocks accordingly. In practice, using the model this way compromises profit for the following reasons: (1) Mining does not proceed block by block, it proceeds haul truck load by haul truck load. (2) Too much time might be spent mining low value blocks precisely. (3) Material sloughs as it is being mined which destroys the block boundaries. (4) The footprint of the mining equipment may be too large to mine out individual blocks.

To maximize mining efficiency the dig limits must be a smoothed representation of the blockwise dig limits. Smoothed dig limits potentially lose and/or dilute ore and compromise profit. Strategic smoothing can minimize lost profit. Some smoothing algorithms, like successive erosion/dilation and window averaging, do not account for the value of the attribute being smoothed. Figure 3.1 shows two cases. In Case A, the top ore block is marginal and should be left because dilution makes it uneconomic. In Case B, the top ore block is high grade ore and the dilution is acceptable because the value of the ore outweighs the dilution. These two cases are indistinguishable from a binary image cleaning perspective. Moreover, image cleaning typically works with pixels and not polygons. The high value blocks should be smoothed less than low valued blocks because they are worth more. The dig limits should be smoothed according to the limitations of the mining equipment. Small equipment is more selective than large and can mine tortuous dig limits more easily than large.

The dig limit selection technique is a two step procedure. The first step transforms the geostatistical model of grades into a map of expected profit and accounts for uncertainty, the presence of multiple minerals, contaminants, and blending. Expected profit is a single variable and avoids having to invoke multivariate decision making tools. Dig limits selected on an expected profit map aim to maximize profit instead of aiming to get material above a minimum specified grade to the mill. A profit based grade control program can easily incorporate spatial/temporal variations in milling costs, mining costs, recovery factors, and price. An expected profit map can handle blending of ore types. Expected profit maps improve profitability by accounting for uncertainty in grade and the economic consequences of misclassification. The procedure for constructing an expected profit map is covered in Section 3.1.1. The second step selects optimal dig limits, given the expected profit map and an initial dig limit. The dig limits are smoothed enough to obtain the best balance between lost and diluted ore and mining efficiency. Section 3.1.2, below, develops the dig limit selection global objective function.

The procedure for constructing an expected profit map is covered in Section 3.1.1. The second step selects optimal dig limits, given the expected profit map and an initial dig limit. The dig limits are smoothed enough to obtain the best balance between lost and diluted ore and mining efficiency.

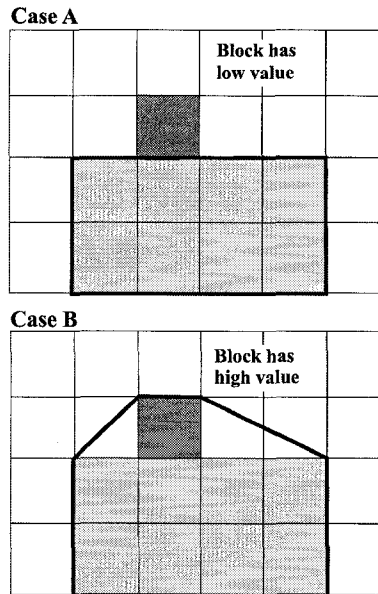


Figure 3.1: The Figure shows two cases. In Case A, the top ore block is marginal and should be identified as waste because dilution incurred by the mining equipment makes it uneconomic. In Case B, the top ore block is high grade ore and dilution is acceptable because the total of the dilution and ore is still economic.

3.1.1 Expected Profit Map

This Section describes the expected profit map transform, and discusses its features and advantages. The expected profit transform is a preprocessing step to dig limit selection. The transform is performed sequentially on each block of the model of grade uncertainty. At each block, the expected profit transform transforms the l grade realizations that characterize the distribution of grade uncertainty to a corresponding distribution of profit. The l profit values are used to calculate expected profit. The expected profit values are then posted on the expected profit map. The expected profit map is used directly in the dig limit selection.

The expected profit transform does not classify material; the dig limit selection algorithm classifies material. The transform considers all blocks, no matter the grade, as ore. The guiding philosophy is that if a block cannot be profitably mined as ore, then it must be waste or low grade stockpile material.

The expected profit transform accounts for lost opportunity costs. A lost opportunity cost is a cost incurred when the opportunity to profit is forfeited. A lost opportunity cost is incurred by the mine when material that could be profited by processing as ore is wasted. A lost opportunity cost can also be incurred when material with mineral content high enough to offset the cost of shipping material but is not ore (protore), is wasted. Although a loss is incurred by processing protore, the loss may be less than the loss incurred by shipping the material as waste. Thus, in some instances, processing protore may result in gain.

The mine will likely never account for lost opportunity cost. It can only be evaluated by sampling the waste dumps. Even with a rigorous waste dump sampling program, lost opportunity cost can only be estimated. Since the expected profit transform considers each block as ore, and for low grade blocks, accounts for the cost of shipping as waste, the probability of misclassification and incurring lost opportunity costs is minimized.

There is no need to consider the cost of misclassifying ore as waste explicitly. Material transformed as ore that does not make gains, including the consideration of lost opportunity

costs, is correctly classified as waste. Furthermore, if it were cheaper to ship barren material as ore plus process it, it would be economic to mine and process all barren material. This is never the case. The cost of shipping material to the waste dump is always less than the cost of shipping material to the mill plus the cost of processing (assuming the material must be processed if it is shipped as ore). Therefore the expected profit gained from shipping and processing material as ore is the essential criteria for classification of material; if a block does not make money as ore it must be waste.

Expected Profit Transform

The grade to expected profit transform has five components:

1. **The grade information** ($z^l(\mathbf{u}), l = 1, \dots, L$): The distribution of grades Z at location \mathbf{u} are represented by the model of uncertainty consisting of multiple realizations.
2. **The cutoff grade** (z_c): The cutoff grade is often defined by management. In general, the cutoff grade is interpreted as the grade at which the mine can operate at an economic threshold. Market conditions and the recovery process control the cutoff grade. Other considerations are the mining cost, administration cost, and contract obligations.
3. **The recovery curve** ($r(z)$): In most mines, the recovery increases with increasing grade (e.g., with constant tailings grade).
4. **The price per unit mineral** (p): This is the price per unit mineral and must consistent with units used for grade.
5. **Cost for processing waste modifier** (cpw): A modifier for increasing or decreasing the value of waste according to mining conditions. The value of this variable can be decreased to make waste more likely to be classified as ore in time of mill feed shortages, or increased in times of ore abundance or mill shut-downs to reduce the likelihood of waste being classified as ore.

Figure 3.2 shows an example grade to expected profit transform. Grade is shown on the x -axis, and profit is shown on the y -axis. The transform function is shown as a curve. The transform function could be a best fit function based on observed recoveries, or a continuous function provided by the metallurgical engineer. The transform function intersects the y -axis at cpw . The vertical line denoted z_c is the cutoff grade threshold. The cutoff grade corresponds to a profit of 0. Material with grade above the cutoff grade gains profit:

$$p_{ore} = (z(\mathbf{u}) \cdot r(z(\mathbf{u})) - z_c \cdot r(z_c)) \cdot price \quad (3.1)$$

The cost of treating material with grade less than the cutoff grade is not constant, because there may be some protore that can be used to offset the cost of treatment. For material with grade less than the cutoff grade, the profit is calculated as:

$$p_{waste} = (z(\mathbf{u}) \cdot r(z(\mathbf{u})) - z_c \cdot r(z_c)) \cdot price - cpw \quad (3.2)$$

The cost of waste, cpw , can be modified to cope with mining issues such as abundance or shortages of ore. Figure 3.3 shows the effect of three different cpw values. The base case is the middle curve. Increasing cpw has the effect of making material marginally below the cutoff grade less profitable and reduces tons classified as ore by the dig limits compared to the base case. Decreasing cpw has the opposite effect: tons are increased because marginal material is more profitable relative to the base case.

The relation used to calculate profit at a location \mathbf{u} for a single realization l is as follows:

$$p^l(\mathbf{u}) = \begin{cases} (z(\mathbf{u}) \cdot r(z(\mathbf{u})) - z_c \cdot r(z_c)) \cdot price, & \text{if } z^l(\mathbf{u}) \geq z_c \\ (z(\mathbf{u}) \cdot r(z(\mathbf{u})) - z_c \cdot r(z_c)) \cdot price - cpw, & \text{if } z^l(\mathbf{u}) < z_c \end{cases} \quad l = 1, \dots, L \quad (3.3)$$




Figure 3.2: This schematic shows the transformation of the distribution of uncertainty in grade to that of profit. The distribution of grade is shown on the x -axis, and the transformed distribution is shown on the y -axis. The transform function is shown as a curve because recovery is assumed to increase with grade.




Figure 3.3: This schematic illustrates the effect of accounting for different cpws on the grade to expected profit transform. A hypothetical profit curve with three cpw scenarios is shown: (1) The cpw ratio is less than 1. (2) The cpw ratio is equal to one. (3) The cpw ratio is greater than one.




Figure 3.4: This Figure shows a hypothetical recovery curve that is used in the small classification example.

Realization	Grade (%)	Recovery (%)	Profit (\$)
1	0.80	0.70	0.00
2	0.60	0.59	-306.65
3	0.71	0.66	-140.97
4	0.67	0.63	-202.24
5	1.30	0.93	973.50
6	0.59	0.59	-321.22
7	1.20	0.90	780.00
8	0.59	0.59	-321.22
9	0.60	0.59	-306.65
10	0.75	0.68	-78.78
Expected Grade:	0.78	Expected Profit:	7.58

Table 3.1: The Table is a summary table of results for the grade control experiment that uses a cutoff grade for classification.

The expected profit value that is posted to the expected profit map is calculated with the use of the l expected profit values from the model of grade uncertainty:

$$\hat{p}(\mathbf{u}) = \frac{1}{L} \sum_{l=1}^L p^l(\mathbf{u}) \quad (3.4)$$

Expected Profit Transform Advantages

The expected profit transform has economic advantages over other techniques such as using a cutoff grade or a probability to be ore cutoff. The sole use of cutoffs may not account for lost opportunity cost and can lead to suboptimal economic performance because material can be misclassified.

The following example illustrates how using cutoff grades can lead to misclassification. Consider the hypothetical nonlinear recovery curve shown in Figure 3.4. Consider a 1 ton block of material that has an average grade of 0.78% copper, and a grade control program that specifies a cutoff grade of 0.8%. The expected grade and profit are calculated with the use of the hypothetical distribution of uncertainty shown in Table 3.1.

The expected profit of the block shows that there is potential for the block to have enough mineral content to offset the cost of shipping as ore and shipping as waste. An expected profit greater than zero shows that the best classification for this block would be ore, yet the block would be waste if classified by a cutoff grade.



Figure 3.5: Distributions of uncertainty in grade for two blocks of material are shown. Note the long tail on the distribution on the left. If a probability of grade greater than the cutoff grade threshold were used for classification then both blocks would have the same probability of being classified as waste, but note that the distribution on the left has a higher probability of being ore.

Classification using a probability to be ore threshold yields the same issues as using a cutoff grade for classification. Figure 3.5 shows models of uncertainty of grade for two blocks of material. Both models have the same probability of exceeding the cutoff grade. However, the model on the left has higher probability for being high grade. If a probability for being ore cutoff is used both blocks are treated identically.

Using a probability to be ore cutoff can lead to misclassification, because the spread of the distribution of uncertainty is not considered in decision making. The expected profit transform does consider the extremes of the distribution by taking the average of the distribution and thus avoids misclassification.

3.1.2 Global Objective Function

The global objective function quantifies profit for a dig limit and is composed of a revenue component objective function ($O_{revenue}$), and a digability component objective function ($O_{digability}$). The global objective function is formulated as follows:

$$O_{global} = OW \cdot [O_{revenue} - \lambda \cdot O_{digability}] \quad (3.5)$$

where OW is a variable used to force the global objective function to comply with the standard practice of minimization of the global objective function in optimization. If the dig limits are intended to enclose ore, $OW = -1$, if waste $OW = 1$.

The Revenue Component Objective Function $O_{revenue}$

$O_{revenue}$ summarizes the value of the material inside the dig limits. $O_{revenue}$ is quantified by superimposing the dig limit polygon on the map of expected profit and taking the sum of all the fractional block values falling within the dig limit:

$$O_{revenue} = \sum_{ix=1}^{nx} \sum_{iy=1}^{ny} frac(ix, iy) \cdot p(ix, iy) \quad (3.6)$$

where $frac(ix, iy)$ is the fractional area of the block indexed at location (ix, iy) within the polygon, and $p(ix, iy)$ is the profit for location (ix, iy) .

The Digability Component Objective Function $O_{digability}$

$O_{digability}$ measures the digability of a dig limit. The weight, λ , is called the digability factor. The digability factor is a variable that is used to specify the how important digability is to the directives of the mine.

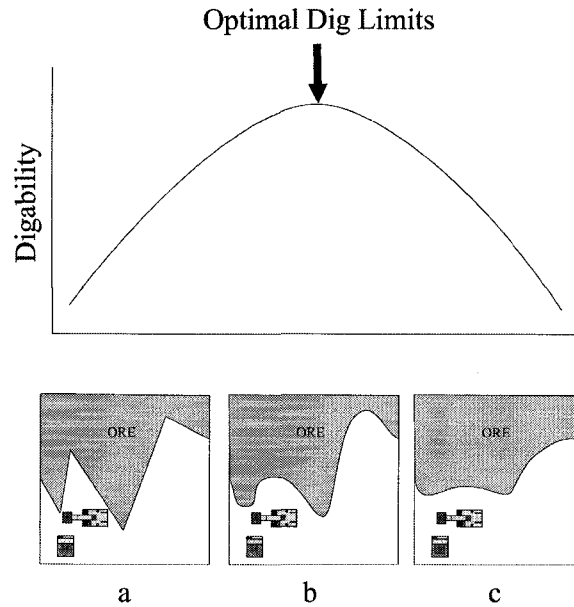


Figure 3.6: This Figure helps to illustrate the notion of an optimal dig limit. At the top of the Figure is a chart with dig limits on the x -axis and digability on the y -axis. The dig limits in maps a and c have low digability. The dig limits in map a are too tortuous and the mining equipment will waste ore trying to mine them. The dig limits in map c are too smooth and unnecessarily waste ore. The dig limits in map b strike the best balance among the objectives of maximizing mining efficiency and minimizing misclassification.

The term digability is used to describe how well the dig limits satisfy the various objectives of the mine, that is, to minimize ore dilution and/or lost ore, and to minimize cost of operations. Dig limits with optimal digability are smooth enough to offer the best compromise between operations costs (time to mine, fuel, wear and tear on equipment) and lost ore. Dig limits with poor digability are either too smooth and unnecessarily dilute and lose ore for the sake of decreasing operations costs or are too tortuous and sacrifice operations costs for reduced dilution and ore loss. Figure 3.6 demonstrates the concept of digability. The dig limits shown on Map a and Map c do not strike an optimal balance on the mines objectives of maximizing efficiency and minimizing lost and/or diluted ore, whereas map b does. The digability curve shown is hypothetical. In practice it may not be smooth or symmetrical.

The value $\lambda \cdot O_{digability}$ should not be considered a cost in the accounting sense. The formulation of $\lambda \cdot O_{digability}$ presented below does not enumerate any of the costs associated with mining, including fuel costs, maintenance, or shut downs. It will be assumed that the operations cost for mining a particular dig limit can be approximated through the measurement of features of the dig limits and transforming this into a metric of digability. The essential feature of the technique is to assign high $O_{digability}$ values to dig limits with poor digability and minimal $O_{digability}$ values for dig limits that have optimal digability, because sub-optimal digability dig limits cost more to mine than optimal digability dig limits.

Quantifying Digability

Quantification of $O_{digability}$ is achieved by measuring the angle of operation (α), then for each vertex on the dig limit ($v, v = 1, \dots, NV$), assigning a corresponding digability penalty

($dig_{pen}(\alpha_v)$). The sum of all the NV digability penalties is used as the as the measure of digability:

$$O_{digability} = \sum_{v=1}^{NV} dig_{pen}(\alpha_v) \quad (3.7)$$

The angle of operation is the angle formed by the vertex under consideration and its two neighboring vertices. A dig limit is a closed polygon with at least 3 vertices, thus, there will always be two neighboring vertices. If $v = 1$, the pair of neighboring vertices are $v = 2$ and $v = NV$; if $v = NV$ the pair vertices are $v = 1$, and $v = (NV - 1)$.

For each α , a corresponding digability penalty, $dig_{pen}(\alpha_v)$, is specified by the digability function. The digability function performs the difficult task of defining the correspondence between the angle of operation and the performance of the mining equipment. Unfortunately, there is no theory to help construct such a function. The digability function could be constructed empirically. Time trials could be performed that measure how long it takes to load a haul truck under a variety of angles of operation. Reconciliation of the in-pit grades delineated by dig limit polygon against the head grade could help provide some insight to equipment induced dilution and ore loss if the modeling practices are sound and do not overestimate or underestimate grade. An alternative approach, which is the one used here, is to infer the digability factor using professional experience.

The following assumptions will be used to guide the construction of a digability function. It is assumed that mining equipment has a range of angles of operation in which digability is optimal – mining efficiency is maximized and dilution and/or lost ore is minimized. This range of angles will be called the angles of optimal operation. When the angle of operation falls within angles of optimal operation, loss is minimized, profit is maximized, and $dig_{pen}(\alpha) = 0$. Of course the mining equipment incur the cost of shipping, but this cost is integrated into the expected profit map and the variable cpw . The loss quantified by $\lambda \cdot O_{digability}$ is solely due to mining sub-optimal dig limits, that is, dig limits that are not calibrated to the mining equipment and are either too smooth or too tortuous. It will be assumed that mining efficiency decreases and the risk of dilution and lost ore increases as the angle of operation decreases because (1) there may not be sufficient room for the mining equipment to maneuver, (2) the ore or waste zone thins making it more difficult to discriminate between ore and waste, and/or (3) more time must be spent to follow the dig limits precisely. It will be assumed that as the angle of operation mining degrades from the range of optimal angles of operation to more acute angles, that mining equipment performance degrades smoothly and digability penalties increase smoothly. As the angle of operation degrades to the point where digability is minimal, that is, dilution and ore loss is maximized and/or mining efficiency is minimized, the digability penalties will asymptotically approach maximum loss. Finally, it will be assumed that the relationships between mining efficiency, dilution and/or lost ore, and the angle of operation are scalable according to the size of the mining equipment. In other words, small mining equipment will have a larger range of angles of optimal operation than large. Due to size, smaller mining equipment has greater ability to mine ore and waste selectively.

An example digability function is shown at the top of Figure 3.7. The angle of operation α is on the x-axis, and the digability penalty is on the y-axis. The digability function is shown as a curved line. The example digability function is not a linear function of the angle of operation. Where α falls in the range of optimal angles of operation, the digability penalty is 0. The digability function shows that the mining equipment gradually degrade in performance as α decreases. As α decreases to $\alpha = 0$ the digability penalty asymptotically approaches maximum penalty.

The area under the digability function in Figure 3.8 has regions with different shades of gray. Each shaded region represents the average digability penalty over a range of angles of operation. The darkest shade represents 100% digability penalty, and the lightest represents 0% digability penalty. Below the digability function is a diagram that shows the same




Figure 3.7: Shown in the Figure is an example penalty function. The penalty function is used to penalize dig limits that have acute angles.

angle of operation and shading combination as shown in the digability function. At small operation angles the digability penalty is at a maximum and the opposite is true when α is large. If mining equipment smaller than the equipment shown in Figure 3.8 were used, the area shaded as 100% penalty would be smaller. The opposite would be true if the mining equipment were larger. The shape of the digability function can be modified to reflect different equipment characteristics. For example, a digability function might show that performance may fall more or less quickly as α decreases.

The y axis on the digability function assigns digability penalties that are scaled from 0-100% of the expected profit blocks that have expected profit greater than 0 within the area of interest. Scaling by the expected value customizes the digability function to the area of interest. This feature ensures that the dig limits in high valued regions are smoothed less and minimizes lost and diluted ore.

The Digability Factor

The digability factor, λ , calibrates $O_{digability}$ to units of cost. With $\lambda = 0$, the emphasis is on $O_{revenue}$. As λ increases the emphasis shifts to $O_{digability}$ and the mining equipment constraints. The digability factor has another function. It translates the digability function along the x -axis. Figure 3.9 shows how the digability factor and the penalty function interact.

The parameter λ and digability function interact in the following ways. If the emphasis is on profit, λ is small, and the digability function is moved to the left and specifies a large range of angles of optimal operation. If the emphasis is on profit, λ is large, and the digability function is moved to the right and specifies a small range of angles of optimal operation.

λ attempts to impart different mining constraints on the dig limits. If a mine selects small λ values, it likely has small mining equipment and tries to mine very selectively, because the ore has high value and dilution is very costly. Mines using small λ s likely have low daily production and the depositional environment is likely not disseminated. If a mine selects large λ values, the ore is likely most profitably mined in bulk by large mining equipment that lacks the ability to mine selectively. This scenario may occur when the deposit is highly disseminated or the ore is of low quality. The above comments are a coarse guide only. The general idea is that the mining engineer would select an appropriate λ value given their experience, and the mining scenario at hand.

3.2 Implementation Details

This section presents the details for implementing the dig limit selection technique. Presented first are the expected profit map construction implementation details. Specifically discussed are the grid definition for an expected profit map, approaches for constructing





Figure 3.8: The penalty function penalizes according to how acute the operation angle is. Shown for illustration is a shovel and the corresponding penalty for a number of different operation angles. Darkly shaded areas are penalized more than light areas.




Figure 3.9: The digability factor interacts with the penalty function by sliding it back and forth on the x -axis. Low digability factors slide the penalty function closer to the y -axis and imply that the equipment is somewhat selective. Large digability factors move the penalty function away from the y -axis and imply that the equipment is not very selective.

expected profit maps in the presence of multiple ore types, and coping with the multiple minerals. Next, the details of applying simulated annealing to optimize the dig limit selection global objective function are discussed. The topics include the dig limit/simulated annealing algorithm, the components that simulated annealing requires to perform optimization on the global objective function, and the setting of a stopping criteria for the algorithm.

3.2.1 Expected Profit Map Implementation Details

A Note on Using Expected Profit for Selecting Dig Limits

Dig limit selection is a classification decision that is repeated many times, perhaps several thousands of times, over the life of a mine. Ideally, the individual classification decisions would maximize profit. In reality, this rarely happens. The individual classification decisions often contain errors and compromise profit. The difference between maximum profit and compromised profit is lost profit. For example, profit is lost when the grade of material classified as ore is mispredicted because the estimate is used for planning and economics.

Expected profit is used to select dig limits because it minimizes expected lost profit, that is, lost profit averages to a minimum over the long term. There is no need to apply subsequent risk qualified decision making techniques, such as a utility function transform to introduce a position on risk; see [70, 71]. The risk of misclassification is accounted for by the grade to expected profit transform. The use of any value other than the expected profit value for classification decisions does not minimize lost profit.

It is known that expected values are good for repeated decisions when probabilities are well known and risk spread over many similar decisions as in the case of the dig limit selection decision. Each of the classification decisions is made, for the most part, under the same conditions. This is not to say that the conditions for decision making are the same at all locations, just that there are no one-off decisions over the life of the mine; for every decision made there exists a decision that has been or will be made under the same conditions at some other location. Consider the data conditions for constructing a local model of uncertainty at location u . For this particular location, and many others, the number and configuration of data are similar due to blasting requirements and symmetries of the blasting plan. Since kriging only relies on the configuration of the data when estimating the variability, and the covariance is the same over the field due to a decision of stationarity, the variance for many of the decisions is the same. Of course the magnitude of the value of a given block varies spatially, but the uncertainty associated is similar over many decisions, so the risk of loss is spread over many decisions.

Expected Profit Map Grid Definition

The grid definition is a regular grid system with an origin at (xmn, ymn) , the number of grid nodes (nx, ny) , and the spacing of the grid nodes $(xsiz, ysiz)$. There is no theoretical requirement for a regular grid; however, many geocellular models in mining are constructed on a regular grid. The y -axis is associated to the north-south direction, and the x -axis is associated to the east-west direction. The grid indices increase from 1 to nx in the positive x direction. The grid indices increase from 1 to ny in the positive y direction. This grid definition follows the GSLIB grid definition [22].

Expected Profit Maps and Multiple Minerals

The presence of multiple minerals can make compiling an expected profit map challenging. There may be multiple minerals that contribute to profit and multiple contaminants that reduce profit. Some minerals are known to catalyze the chemical reactions for recovering minerals of interest, while others suppress the reactions. The precise interaction is subject to numerous parameters, such as temperature, relative concentration, and pressure.




Figure 3.10: The map shows the expected profit values used in the preliminary example.

If the minerals are complementary and not interactive, that is, all of the minerals are minerals of interest and do not affect the recovery process adversely, the solution for building an expected profit map is to construct expected profit maps for each mineral and add the maps to get a merged map of expected profit. This case is rarely encountered, but it might be a satisfactory solution if the interaction between minerals in the mill is insignificant.

More complex relationships might be able to use a multivariate function to characterize the relationship, or a lookup table could be constructed that gives recoveries corresponding to combinations of mineral grades. Such a function or lookup table would require input from the mineral or process engineer.

Expected Profit Maps and Multiple Ore Types

Some mines have grade control programs that use multiple ore types. Multiple ore types can be used, when multiple treatment facilities exist, for more efficient processing of high grade ore, strategic treatment of material high in contaminant concentration, or if there is a need to stockpile material. The dig limit selection algorithm can be used to select multiple ore type dig limits by applying it recursively on expected profit maps prepared for each ore type. The following paragraph describes the procedure for constructing an expected profit map for each ore type.

The first step is to identify which ore type the material most profitably belongs to. This can be done by constructing an expected profit map for each ore type and, for each block, compare the expected profit of a block under each ore type classification. The optimal ore type classification for the block is the classification that maximizes expected profit. The next step is to reprocess the expected profit map for each ore type to show loss due to misclassification. Dig limits are then selected on each ore types reprocessed expected profit map less one; if there are N_{OT} ore types, say high grade, low grade, and waste, one need only select the high and low grade dig limits, because the amalgamation of these two dig limits delineate the waste dig limits.

The following is an example of the multiple ore type procedure. The example is an extension of the preliminary example shown in Chapter 1, Section 1.1.1. The extension deals with the problem of having two ore types (high grade and low grade ore) plus one waste type material.

Figure 3.10 shows the original expected profit map from the preliminary example. Suppose that the mine treats high grade material differently than low grade material, and thus has high grade ($N_{OT} = 1$) and low grade ore types ($N_{OT} = 2$), where N_{OT} is a discrete variable with integers for identification of ore types. The last material type is waste.




Figure 3.11: This map shows the optimal classification of blocks. The darkest blocks are ore type $N_{OT} = 1$. The lighter shaded blocks are ore type $N_{OT} = 2$. The unshaded blocks are waste.

The map shown in Figure 3.11 shows the optimal ore type classification for each block. The dark grey blocks are high grade, the medium grey blocks are low grade and the unshaded blocks are waste. The map was made by selecting the classification that maximized expected profit, and assigning the appropriate integer to reflect the ore type classification.

Figure 3.12 shows the expected profit values for all blocks classified as $N_{OT} = 1$. Comparing the expected profit map shown in Figure 1.3 with the expected profit for $N_{OT} = 1$, shows that much more of the material is "waste" relative to the classification as $N_{OT} = 1$. This map was made by replacing the expected profit values at blocks optimally identified as $N_{OT} = 2$, with the loss that would be incurred if they were treated as $N_{OT} = 2$. The loss due to misclassification is calculated as $N_{OT} = 2: \hat{p}(\mathbf{u}) = \hat{p}(\mathbf{u}, N_{OT} = 1) - \hat{p}(\mathbf{u}, N_{OT} = 2)$.

The dig limits for each ore type are shown in Figure 3.13. There is some material that is misclassified. This is a necessary loss due to mining. There are also a few blocks that are not included in the correct ore type dig limits. Another professional might decide to incorporate these blocks in the dig limits.

The multiple ore type dig limit selection technique can accommodate any number of ore types by repeating the described procedure over each ore type classification.

3.2.2 Dig Limit Selection Details

The Dig Limit Selection Optimization Algorithm

The implementation steps for dig limit selection algorithm are: (1) Transform the model of uncertainty in mineral grade to a map of expected profit. (2) Select an initial dig limit polygon with the aid of the expected profit map. (3) Use simulated annealing to perturb the dig limit polygon until it conforms to a dig limit polygon that maximizes the global objective function. Figure 3.14 shows a flow chart of the dig limit selection algorithm. The text below discusses the nodes in the flow chart:

Read Map of Data and Initial Polygon :

Read in the map of data and save the initial dig limit as $Digi$.

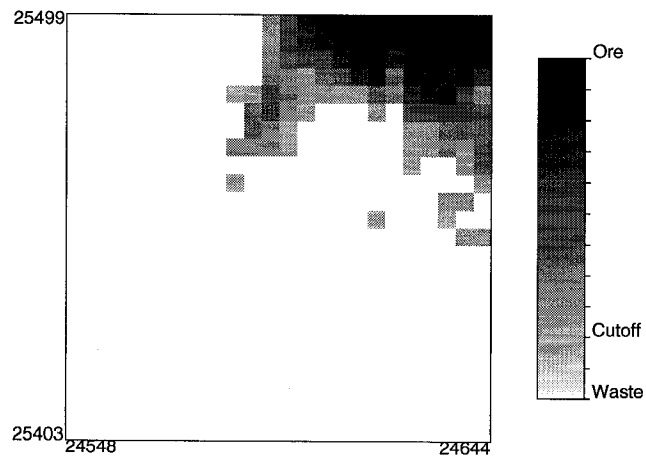


Figure 3.12: This map shows the expected profit values for material identified as high grade ore $N_{OT} = 1$. The unshaded blocks are waste with respect to the high grade ore classified material.




Figure 3.13: The map shown posts dig limits for both ore types.

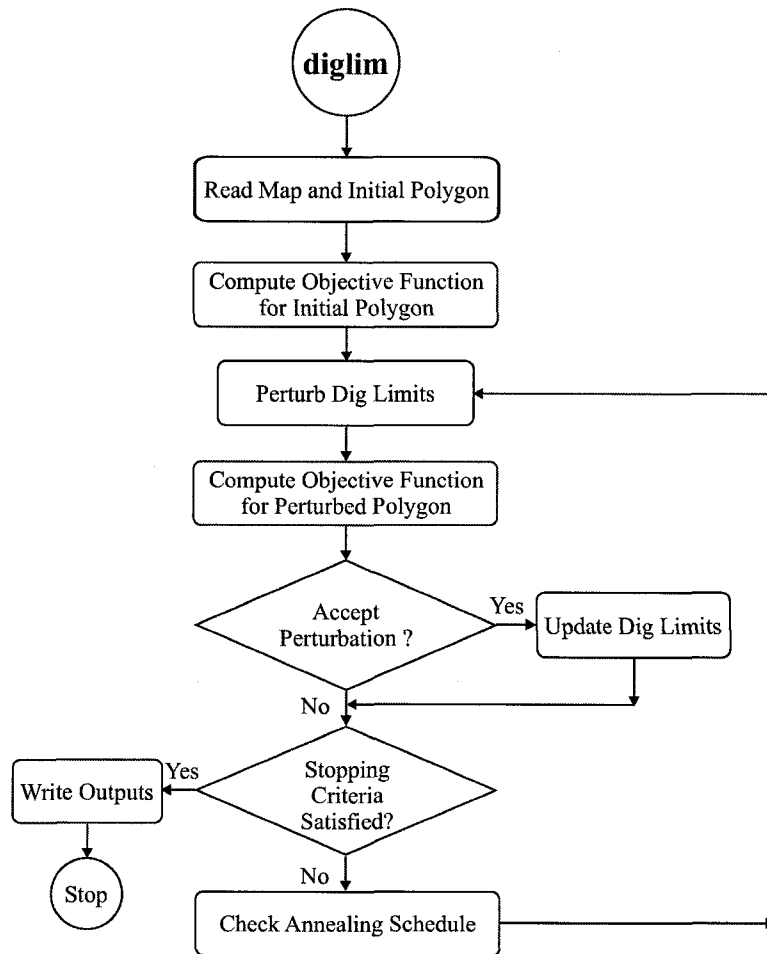


Figure 3.14: This flow charts illustrates the steps the dig limit selection takes during optimization.

Establish the initial value of the global objective function :

The initial dig limit Dig_i is superimposed on the expected profit map. The value of the revenue component objective function, $O_{revenue}$, is calculated. The algorithm scans Dig_i and calculates the digability component objective function $O_{digability}$. The result is the initial global objective function value which is saved as O_i .

Perform a Perturbation :

A perturbation is a random change to the dig limit. The modified dig limit is saved as Dig_p . Section 3.2.2 discusses the perturbation mechanism.

Establish the new value of the objective function :

The perturbed dig limits Dig_p and the mapped data are used to determine the new value of global objective function, O_p .

Accept or reject perturbed dig limit polygon? :

Calculate $\Delta O_g = O_p - O_i$. Use ΔO_g and the current temperature, T , from the annealing schedule to calculate the probability of accepting Dig_p according to:

$$P \{accept\} = \begin{cases} 1, & \text{if } \Delta O_g \leq 0 \\ e^{(-\frac{\Delta O_g}{T})}, & \text{otherwise} \end{cases} \quad (3.8)$$

If ΔO_g is less than 0, set P_{accept} to 1. If ΔO is greater than 0, draw a random number $\in [0, 1]$, and test it against P_{accept} . If the random number is less than or equal to P_{accept} , then the perturbed dig limit is accepted. If the random number is greater than P_{accept} , then the perturbed dig limit polygon is not accepted.

Update the initial global objective function and dig limit polygon: If the perturbation is accepted, save O_p as O_i and save Dig_p as Dig_i . If the perturbation is rejected then restore Dig_i and O_i and go back for another change. If Dig_p is rejected it will be replaced by the next accepted perturbed dig limit. The algorithm only requires the results from the previous and current dig limits. The results from perturbations previous to the most recent perturbation are not saved.

Stopping Criteria Satisfied? : If the stopping criteria from the annealing schedule have been satisfied then stop the algorithm.

Write Outputs If one of the stopping criteria is satisfied then write out the vertex coordinates for the dig limit polygon and a grid specifying the fraction of each block falling within the dig limit polygon.

Check Annealing Schedule : If k_{accept} or $k_{attempt}$ has been satisfied, then apply the reduction factor λ to the temperature T and reset the counters for k_{accept} and $k_{attempt}$. If k_{accept} has not been satisfied and the perturbation was accepted, then increment the k_{accept} counter. If $k_{attempt}$ has not been satisfied and the perturbation was rejected, then increment the $k_{attempt}$ counter.

Loop Until the a Stopping Criteria is Satisfied :

The algorithm perturbs the dig limits until one of the stopping criteria in the annealing schedule is satisfied.




Figure 3.15: For situations where a body of waste is enclosed in a body of ore a decomposition approach can be used to obtain the dig limits. First dig limits for the small waste body dig limits are solved for. Then the larger ore dig limits are solved for.




Figure 3.16: This Figure shows an alternative decomposition approach to the decomposition shown in Figure 3.15.

Selecting Initial Dig Limits

Simulated annealing requires an initial dig limit, which is provided by the mining engineer. The initial dig limit must be a single closed polygon with no crossing segments. The approach used here was to digitize the initial dig limits on a map of expected profit and submit the coordinates of the vertices to the dig limit selection algorithm. For the most part, selecting initial dig limits is a straightforward procedure, but situations such as nested dig limits, multiple ore and waste dig limits and specifying of non-perturbed vertices need further discussion.

Nested Initial Dig Limits

Most benches have regions of commingled ore and waste where a pod of ore resides within a larger region of waste or vice versa. A decomposition approach is adopted to select initial dig limits in this case. A hypothetical example of the decomposition approach is shown in Figure 3.15. The Figure shows a bench with the some ore/waste information. The darker areas represent ore and are demarked with an uppercase *O*. The lighter areas are waste and are demarked with an uppercase *W*. Note that a single dig limit will not partition the bench into ore and waste.

One possible decomposition is shown in Figure 3.15. The first initial dig limit is selected as the small waste pod. The second initial dig limit is the initial dig limit for the orebody. The remaining material is all waste. Another possible decomposition is shown Figure 3.16. In the alternative decomposition, the first initial dig limit is the wastebody. The second initial dig limit is for the imbedded waste. The assemblage of dig limits partition the entire bench.




Figure 3.17: The perturbation mechanism for perturbing the polygon vertices randomly selects a vertex and moves it a random $\Delta x, \Delta y$, as shown in the Figure.

Specifying Non-Perturbed Vertices

The dig limit selection algorithm can be set such that all the vertices are perturbed to conform to optimal dig limits, or some vertices can be set to non-perturbable vertices. This feature is useful when there is a need to constrain a dig limit along the crest or toe of a bench, or when a patchwork of several dig limits are merged over a large area.

The Perturbation Mechanism

A perturbation is the random selection of a vertex in the dig limit polygon and a random change in the location of a vertex. The new location is selected by selecting a random Δx and Δy as show in Figure 3.17. The Figure shows an example dig limit polygon with a candidate vertex highlighted. The range of values for Δx and Δy are indicated by a rectangle centered on the candidate vertex. The perturbed vertex could be anywhere inside this area.

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 find an optimal solution. A practical solution is to choose a reasonable mechanism, and any inefficiencies will be revealed in slow convergence.

The dig limit selection algorithm will not perturb the dig limits beyond the area of interest specified by the expected profit map, or a user selected boundary window. The initial dig limits must fall inside the boundary window. If they do not, the vertices having coordinates outside the boundary window are moved so that they are inside the boundary window. The boundary window can be used to specify a subset of the expected profit map that might represent a shift or days worth of mining. Without such a constraint the dig limit selection algorithm could select optimal dig limits over the entire map of expected profit.

Calculating $O_{revenue}$

Recall that $O_{revenue}$ is calculated as the fraction of the blocks within the dig limits multiplied by the expected profit of the block. The technique used to perform the calculation is adapted from a paper written by C. Deutsch [16]. An alternative is to calculate $O_{revenue}$ as the sum of the blocks with centres falling inside the dig limits. The initial development of the dig limit selection algorithm used this approach, but the results were not satisfactory. The

global objective did not reduce smoothly because perturbations that caused the dig limits to include or preclude a block also caused the global objective function to have strong fluctuations. Due to the fluctuations, more perturbations were required to converge. Also, the dig limits provided by the block centre approach were not optimal because they wrapped around the centre points of the blocks and excluding the rest of the block giving erroneous dig limits. Adapting the algorithm to calculate the fractional area of the block falling inside the dig limits provided a relatively well behaved global objective function and dig limits that did not wrap around the cell centres.

Minimum and Maximum Segment Lengths

The number of vertices in the dig limit polygon is a critical constraint that affects the calculated digability of the dig limit. If there are too few vertices, the dig limit may not have enough flexibility to conform to the ore/waste body or the dig limit will have unavoidably low digability, and may contradict the selectivity of the mining equipment. Dig limits with a large number of vertices conform smoothly the ore body and have high digability, but takes longer to converge, or are difficult to survey in the pit, or may contradict the selectivity of the mining equipment. Areas of high profit may require more vertices to minimize ore loss and dilution than in low valued regions. Unfortunately, there is no way to know a priori the optimal number of vertices required to optimize the dig limit.

A dynamic approach that allows the optimization algorithm to specify the optimal number of dig limit vertices over the course of the optimization. The dynamic approach specifies a maximum and minimum distance between vertices, d_{min} and d_{max} respectively. If two vertices are too close, one is randomly selected for deletion. If two vertices are too far apart, a new vertex is added halfway between the two vertices.

There is no theory or specifying d_{min} and d_{max} . The block size used in the expected profit map serves as a good guide for selecting d_{min} and d_{max} . A good minimum distance is about half of a block, and a good maximum distance is about two times the block size. Another idea is to use the shovel dimensions as a guide to selecting d_{min} and d_{max} .

Stopping Criteria

The dig limit selection algorithm will perturb the dig limits until a stopping criteria is encountered. There is no efficient way to predict the value of the global objective function that leads to optimal dig limits.

The dig limit selection algorithm considers three stopping criteria. Two are imbedded in the annealing schedule: the parameters k_{max} and num can stop the algorithm. The parameter k_{max} is the total number of perturbations and num is the number of times that $k_{attempted}$ is satisfied without an accepted perturbation. The third stopping criteria measures the percent change in the global objective function O_g over $k_{attempted}$ perturbations and stops the algorithm if the change is less than a selected threshold.

Output Information

The output files contain the fraction of tonnes of ore, the average grade of the material inside and outside the dig limits, and the expected profit of the material inside and outside the dig limits. An example summary table is shown in Figure 3.18.

The *Perfect Selection* section lists the profit from ore and waste as well as the tonnes of ore and waste if the material could be mined according to the blockwise dig limits. The *Inside Dig Limits* section lists the profit and tonnes for ore and waste for the material that falls inside the dig limits. The *Outside Dig Limits* section lists profit and tonnes for material outside the dig limits.

Perfect Selection			
Ore Portion =	\$	563540.99	
Waste Portion =	\$	-57154.00	
Tonnes Ore =	501600.0	Average Grade =	.000
Tonnes Waste =	95200.0	Average Grade =	.000
Inside Diglimits \$ -25742.15			
Ore Portion =	\$.32	.00%
Waste Portion =	\$	-25742.46	45.04%
Tonnes Ore =	31.7	Average Grade =	.000 .01%
Tonnes Waste =	25478.2	Average Grade =	.000 26.76%
Total Tonnes =	25509.9	Average Grade =	.000
Outside Diglimits			
Ore Portion =	\$	563540.67	
Waste Portion =	\$	-31411.54	
Tonnes Ore =	501568.3	Average Grade =	.000 99.99%
Tonnes Waste =	69721.8	Average Grade =	.000 73.24%

Figure 3.18: This Table presents summary information from the dig limit selection algorithm. The values are for a waste dig limit.

Selecting a Digability Factor

There is no theory that leads directly to a digability factor that is optimal for a particular mine/mining equipment scenario. Selecting a digability factor is an iterative procedure that involves experimentation and professional judgment. Fortunately, the digability factor only needs to be selected once for each combination of mining equipment.

A useful tool for selecting digability factors is the digability catalog. A digability catalog consists of a number of dig limits selected with the use of a variety of digability factors and the same expected profit map. The mining engineer or production geologist examines the digability catalog and selects the digability factor appropriate for the mining equipment.

If the tonnes per day shipped to the mill increase, but the grade decreases, the digability factor is too high: the dig limits are too smooth and there is too much emphasis on the mining equipment. The solution is to construct a new digability catalog with the current digability factor as the largest digability factor. Selecting a lower digability factor shifts more importance to profitability and dilutes and loses less ore by selecting more tortuous dig limits.

If the reverse occurs, that is, tonnes per day decrease and grade increases, the digability factor is too low and the dig limits are too tortuous. The mining equipment spends too much time mining the dig limits because they are too tortuous and the grade is too high. The solution here is to construct a digability catalog with the current digability factor as the smallest factor and select a larger digability factor.

If the tonnes per day and grade increase, the digability factor is approaching optimality. To investigate fine tuning of the digability factor, construct a digability catalog using the current digability factor as the midpoint, and select digability factors that are incrementally smaller and larger than the current digability factor. Observe the response in the feed tonnes and the average grade and modify as needed.

If the tonnage and grade drop, something other than the digability factor deserves investigation. A decrease in feed grade might be due to the quality of ore. The mining engineer should reconcile the feed tonnes and grade with the expected profit map. If there is consensus, there is no need for investigation. Contradiction between the expected profit map and

the daily tonnage and feed grade are indicators of a problem.

The digability factor is a tuning parameter that depends on the performance characteristics of the mining equipment, the operating strategy of the mine, the skill level of the operators, and other imprecise considerations. A good representative digability factor can usually be selected after a few trials.

3.2.3 Comments

Automatic Selection of Initial Dig Limits

Aside from calibration of the global objective function, which only has to be performed once for each mining equipment combination, the only interaction required by the dig limit selection algorithm is the specification of initial dig limits. Research was conducted on automatically selecting initial dig limits in an effort to make the entire dig limit selection process fully automatic.

Two approaches for automatically selecting initial dig limits were researched. Hand selected initial dig limits were used throughout the development of the dig limit selection technique. Neither of the automatic approaches produced acceptable results and both showed that automatic initial dig limit selection is not straightforward.

The first approach involved automatically identifying the central coordinate of the ore or waste body and planting a circular seed dig limit of a specified radius. The idea was that the dig limit would migrate and engulf the ore or waste zone over the course of the optimization. The approach gave acceptable final results as long as all the ore blocks were corner or edge connected. Disconnected blocks were not incorporated into the final dig limit unless a very aggressive perturbation mechanism was used (maximum perturbation distances in the order of several blocks). If small perturbations were used, the dig limits never experienced a perturbation large enough to include the unconnected blocks in the search for the optimal dig limits. With vary large perturbation distances undesirable blocks were often incorporated in the dig limits. Also, this approach required a large number of perturbations to converge. This was particularly the case when the overall orebody shape differed significantly from a circle.

The second approach converted the expected profit map to a categorical map of ore and waste blocks scanned the area of interest from four directions sequentially and assembled the initial dig limit as it proceeded. The first step transformed the expected profit data to a categorical map where blocks with expected profits greater than 0 are temporarily assigned the value of 1, and 0 otherwise. The second step obtained the minimum and maximum grid indexes for the blocks of value 1 in the area of interest. The initial dig limit selection algorithm considered only this subset of the area of interest. The rest of the procedure is shown diagrammatically in Figure 3.19. The left column shows which face is scanned and the right side shows the subsequent initial dig limits. The scan starts on the bottom left side of the subset of the grid. The scan looks for a block with value 1. When a block with a 1 value is found, the coordinates for the centre of the left face are stored. This is repeated on each side.

The initial dig limit is shown on the bottom right grid. If the blocks with value 1 were all edge and corner connected the technique usually provided acceptable initial dig limits. Unconnected blocks caused the algorithm to generate initial dig limits with crossing segments, as in Figure 3.20.

A dilution/erosion algorithm was added to help remove unconnected blocks, but it did not help much. One might attempt to remove the crossed segments, but reassembly of the polygon with ordered vertices is troublesome. After several experiments, it was concluded that the most practical way to consistently obtain good initial dig limit polygon was to get them from the mining engineer.




Figure 3.19: The automatic initial dig limit selection algorithm steps are shown in the maps above. The algorithm scanned the sides of the area of interest and assembled the initial dig limits as it proceeded.




Figure 3.20: The automatic initial dig limit selection algorithm failed in certain cases that were experienced often. The algorithm produced dig limits that crossed.

Chapter 4

Dig Limit Selection Case Studies

The objectives of this Chapter are to examine some considerations for application of the dig limit selection algorithm, compare the algorithm results to hand drawn dig limits, and show practical applications of dig limit selection on real data.

4.1 Considerations for Semi-Automatic Dig Limit Selection

This Section discusses some key aspects to the application of the dig limit selection algorithm. The number of perturbations required to find optimal results, parameter selection for the annealing schedule and its effect on the final result, the repeatability of the algorithm, and optimality are discussed.

4.1.1 Required Number of Perturbations

As mentioned in Section 2.2.6, SA has no way to determine when or if an optimal solution has been found. The algorithm will continue to perturb the dig limits until some stopping criteria are invoked. The semi-automatic dig limit selection algorithm stops when too many perturbations have been performed, when too many perturbations have been attempted without any being accepted, or when there appears to be no change in the global objective function after a number of perturbations. This Section demonstrates that the number of perturbations required to find an optimal solution depends on the shape of the ore body, the initial dig limits, and the digability factor. These factors must be considered in the selection of the stopping criteria.

Three synthetic expected profit maps are used. The maps were generated in two steps. In the first step, a secondary data set was created by kriging data values in a particular configuration. Next, sequential collocated Gaussian cosimulation was used with the kriged information and a correlation coefficient of 0.95 to construct the final expected profit maps. Each map has 25 by 25 blocks. The maps are contrived, but serve the purpose of demonstrating that the number of required perturbations depends on the problem. The three maps are shown in Figure 4.1. The map on the far left will be called X data; the middle map, C data; and the far right map, Z data. The legend represents synthetic profit and has no ties to anything realistic.

For the first experiment, the fraction of total available revenue, and the initial and final penalty for each shape using three digability factors are compared. The revenue from processing waste (negative revenue) is not used to calculate the fractional revenue because this simplifies the comparison. The cost of milling is kept constant. Each shape has its




Figure 4.1: The three synthetic maps shown will be used to observe the properties of the dig limit algorithm.

own initial polygon and this polygon was used respectively for each experimental digability factor. Three different digability factors will be used: low ($\lambda = 0$), medium ($\lambda = 0.5$), high ($\lambda = 1.0$).

The results for only 20,000 perturbations are shown. More perturbations are required to find the optimal solution in some of the cases. The reason for limiting the number is that the relationships between the shape of the ore body, the initial dig limits, and the digability factor are apparent within 20,000 perturbations. Also, it becomes difficult to observe and examine the results when too many global objective function evaluations are posted.

Figures 4.2, 4.3, and 4.4 show the revenue component objective function evaluation for a given perturbation and digability factor for X data, C data and Z data respectively.

Some highlights of the three Figures are as follows. For X data, the initial revenue starts at just less than 95%. The X data low digability factor case quickly accumulates almost all of the revenue and asymptotically approaches 100% of the revenue. The medium digability factor case does not improve much and plateaus at about 96%. The high digability factor case quickly loses revenue to account for the limitation of the mining equipment and plateaus to about 87% after about 12,500 perturbations. The C data example has the least dramatic results of the three data sets. The initial revenue is about 97% of the total. A large number of perturbations were required to get the low digability factor case dig limits up to about 100%. The medium digability factor case shows almost no change in revenue. Some revenue is lost initially and the curve plateaus quickly to about 97.5%. The high digability factor case plateaus at about 92% after about 7,500 perturbations. The Z data initial revenue is about 94%. The low digability factor Z data case quickly ramps up to 98%, and then requires a large number of perturbations to select the blockwise dig limits. The medium digability factor case settles down quickly and plateaus at about 93%. The high digability factor case falls quickly and takes 16,000 perturbations to plateau at approximately 81%.

All of the plots show some noise. At 20,000 perturbations, the annealing schedule parameters are permitting some global objective function increasing perturbations to occur. The plots show the revenue component objective function only. Adding in the digability component objective function serves to dampen out some of the noise, but has no effect on the observations. The three curves for each shape start at the same fraction of total revenue because for each shape and digability factor, the initial dig limit are the same.

The low digability factor curve is always the top curve. A low digability factor minimally accounts for the constraints of the mining equipment, and amounts to selecting blockwise dig limits. Profit is maximized in the low digability factor case because there is minimal lost or diluted ore.

The high digability factor case is always on the bottom of the plot. High digability factor dig limits are usually very smooth. The smoothing loses and dilutes ore and compromises revenue to account for the limitations of the mining equipment. The high penalty case also




Figure 4.2: The chart shows the fraction of total revenue versus number of perturbations using X data for three digability factors.




Figure 4.3: The chart shows the fraction of total revenue versus number of perturbations using C data for three digability factors.




Figure 4.4: The chart shows the fraction of total revenue versus number of perturbations using Z data for three digability factors.

has the noisiest curve. Some of the noise is due to the annealing schedule, and some of it is due to increasing the importance of the digability component objective function. For high digability factors, digability has high importance. Perturbations that yield improvements in digability may have nearly equivalent adverse effects on revenue causing the revenue, component objective function to fluctuate.

A plateauing curve indicates that accepted perturbations are having increasingly less effect on the global objective function and that the dig limits are approaching optimality. Note that the low and high digability factor cases require more perturbations to plateau than the medium digability factor case. It is a difficult problem to find dig limits that conform to the blockwise dig limits by iteratively perturbing the vertices. It takes many small perturbations to precisely align the dig limits precisely on the block edges. Finding dig limits that minimize a global objective function which incorporates the digability component function is an easier problem, because the vertices do not have to line up on the block edges. Another reason that the blockwise dig limits are more difficult to find is that the only dig limit polygon that minimizes the global objective function is the blockwise dig limits, whereas there may be more than one optimal dig limit polygon that minimizes the global objective function. To summarize, there is only one minimum global objective function value, but there may be more than one optimal dig limit polygon that attains the minimum global objective function value. There are different combinations of the component objective functions that can minimize the global objective function. Section 4.1.3 presents an experiment that explores the concept of multiple dig limit polygons that minimize the global objective function.

The high digability factor case takes more perturbations than the moderate case because small changes in vertex locations cause large changes in the global objective function due to the application of high penalties from the penalty curve. The result is that it becomes necessary to sample more vertex locations to find perturbations that reduce the global objective function.

Of the three levels of digability, the medium digability factor case requires the fewest number of perturbations to plateau. The medium digability factor case initial dig limits have angles of operation that are already well matched to the digability penalty function and include a high fraction of the greater than zero expected profit blocks. This implies that the initial dig limits can affect the number of required perturbations. The Section below compares the dig limit polygons.

	X Data				C Data				Z Data			
	Initial		Final		Initial		Final		Initial		Final	
	Rev	Dig	Rev	Dig	Rev	Dig	Rev	Dig	Rev	Dig	Rev	Dig
Low	0.95	655	0.99	3530	0.99	465	0.99	4585	0.94	825	0.99	7222
Med.	0.95	655	0.96	1740	0.99	465	0.98	393	0.94	825	0.93	433
Hi.	0.95	655	0.87	32	0.99	465	0.92	124	0.94	825	0.85	116

Table 4.1: This Table compares the revenues (Rev.) and digability (DIG.) for the initial and final dig limits on all three ore body shapes.

X, C, Z Dig Limits Comparison

Maps showing the initial and final dig limits for each digability factor are shown in Figure 4.5. The initial dig limits for each shape are on the top row. The next three rows are the final results for the low, medium, and high digability factor cases respectively. The legend is shown on the bottom right and ranges from +10 to -10; however, values below -5 are clipped from the maps.

As expected, the high digability factor dig limits are smoother than the medium digability factor dig limits, which are in turn smoother than the low digability factor dig limits. Note that the low digability factor dig limits do not appear to follow the blocks precisely and on some blocks the dig limits cut the blocks diagonally in half. This is because some blocks have zero or nearly zero expected profit. These blocks, in whole or in part, do not affect the global objective function much at low digability factors. If there is no change in the global objective function, the algorithm cannot tell if it has cut through a block or not. Blocks with clear economic distinction have dig limits that follow the blocks. This feature is interesting in that it demonstrates “confusion” or “indifference” by the algorithm.

The medium digability factor dig limits are similar to the initial dig limits in that none of the angles between segments are less than about 120° . The medium digability factor dig limits are also affected by the near zero blocks. Recall that the medium digability factor dig limits required the fewest perturbations to have the global objective function plateau.

The high digability factor dig limits are much smoother than the medium digability factor dig limits. Some high expected profit blocks are obviously compromised to achieve such smooth dig limits. High digability factor values require compromises and the global objective function cannot be minimized to the same extent as low digability factors.

The fraction of maximum revenue component objective function and maximum digability component function results are shown in Table 4.1. The initial column shows the results from the initial dig limits and the final column shows the optimal results. The *Rev.* column reports the fraction of total revenue and the *Dig.* column the digability component objective function.

The results show that Z data is most sensitive to the digability factor; this likely due to the spatial distribution of expected profit.

Comments

Clearly, each dig limit selection problem has different perturbation requirements. These requirements change according to the shape of the ore/waste body, the digability factor, and the initial dig limits. A recommended approach for selecting the number of perturbations is to construct a plot of the global objective function versus the number of perturbations and look for the plot to plateau. A plot that shows little or no change over a large number of perturbations might indicate that the perturbation distance is too small. If the plot stops before a plateau is reached, it could mean that the stopping criteria need to be modified to allow more perturbations. If the plot plateaus long before the plot ends the stopping

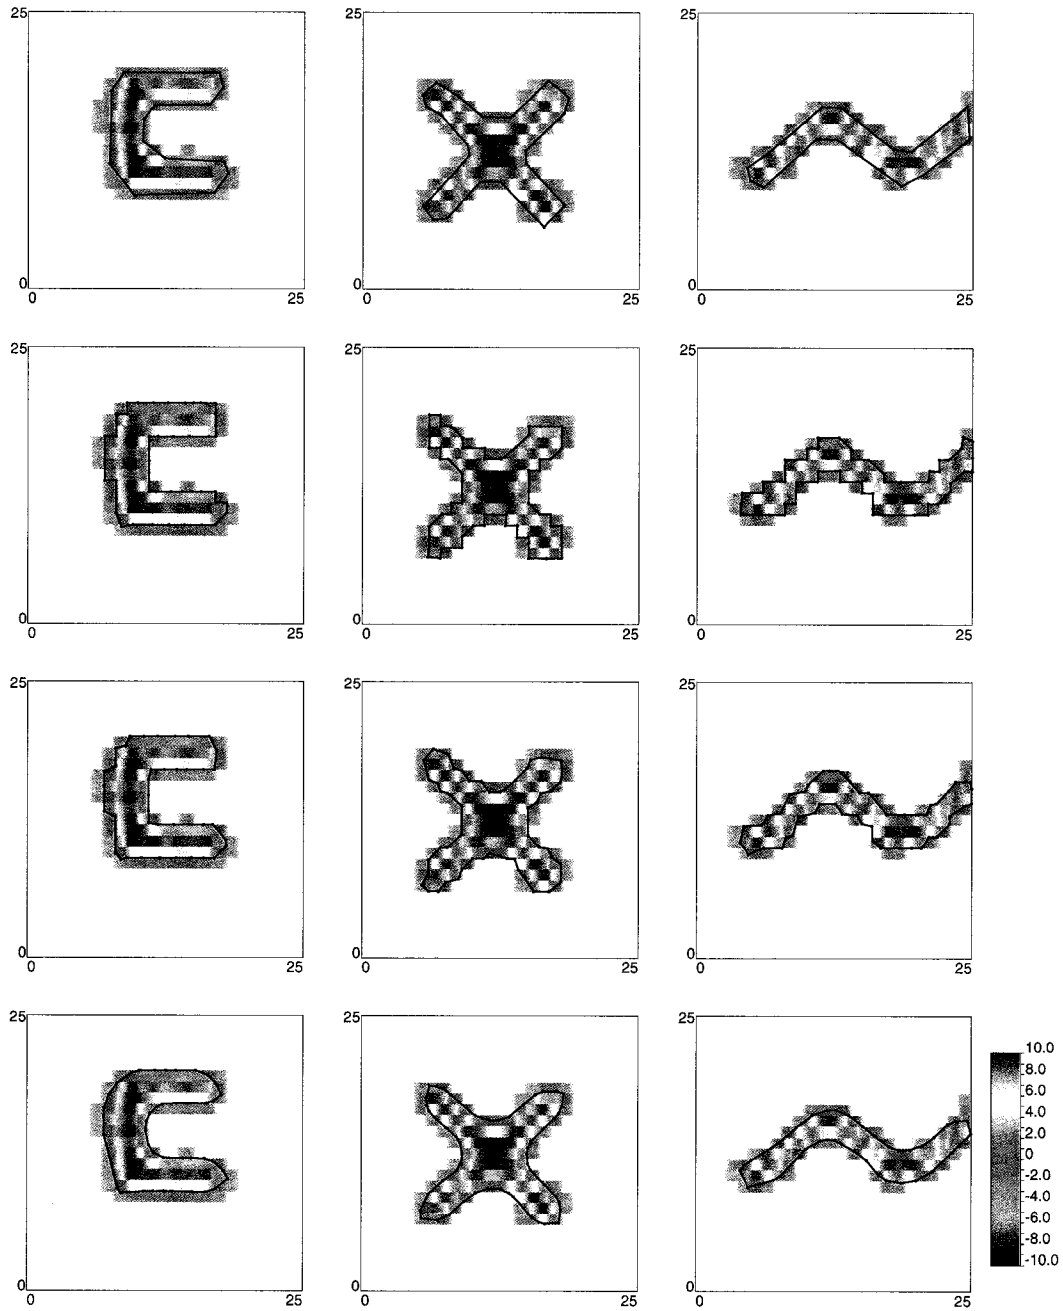


Figure 4.5: The initial dig limits for each shape are shown at the top. The optimal dig limits for each digability factor are shown below, starting with the low digability factor case and with the dig limits that have high digability factor at the bottom.

	Fast	Medium	Slow
t_0	1.0	1.0	1.0
red_{fac}	0.001	0.5	0.9
k_{accept}	5	25	100
k	2	10	40
num	2	10	40

Table 4.2: These values are the annealing schedules used to observe the effect of using different annealing schedules on the optimization.

criteria could be modified to allow fewer perturbations.

4.1.2 Effect of the Annealing Schedule

This Section uses the X data to select dig limits using fast, moderate, and slow annealing schedules to observe how the annealing schedule affects the final solution. If the annealing schedule cools too fast, SA may not be able to sufficiently sample the solution space to find the global optimal solution. If the annealing schedule cools too slow, it may not be of practical use. The parameters used for each schedule are specified in Table 4.2.

Figure 4.6 shows a plot of the global objective function versus the number of perturbations. The global objective function for the fast annealing schedule fell fastest and plateaued with the highest global objective function value (-486.4). The moderate annealing schedule took more perturbations to plateau, and plateaued with a slightly lower global objective function value than the fast annealing schedule (-487.3). The slow annealing schedule took the longest number of perturbations to plateau, but plateaued at the lowest value (-489.5). When the global objective function plateaus it means that the algorithm can no longer perform perturbations that reduce the global objective function. The fast and moderate annealing schedules approach a zero probability of accepting objective function increasing perturbations sooner than the slow annealing schedule and the slow annealing schedule accepts objective function increasing perturbations long after the fast and moderate annealing schedules have stopped. This allows the algorithm to search longer for the final optimal solution.

Figure 4.7 shows the dig limits for each annealing schedule. The fast annealing schedule dig limits are on the left; the moderate, in the middle; and the slow schedule dig limits, on the right. It is difficult to observe any differences between the dig limits. This is because the numeric differences are relatively small.

4.1.3 Repeatability of the Algorithm

The dig limit selection algorithm should give the same dig limits despite the random number seed. Repeatability of the dig limit selection algorithm is demonstrated by selecting 4 different random paths and keeping all other parameters the same.

The numeric results for the four dig limits are tabulated in Table 4.3. The first column in the Table specifies the dig limit polygon number. The second column of the results table shows the fraction of total profit; the third column is the final digability component objective function; and the last column is the final value for the global objective function. The final global objective function values are nearly the same for each dig limit polygon. Also, if a dig limit polygon has a lower fraction of profit than other dig limits, the digability component objective function is correspondingly lower, as shown in a comparison of dig limit polygons number 3 and 4. Figure 4.8 shows a graph of the global objective function versus the number of perturbations.




Figure 4.6: Plotted on this chart are the results for three different annealing schedules on the X data set.




Figure 4.7: These maps show the dig limits resulting from the fast (far left), moderate (middle), and slow (right) annealing schedules.

Random Number	Profit	Final Penalty	Final global objective
1	0.9784	0.26	-479.97
2	0.9782	0.25	-479.98
3	0.9779	0.22	-479.95
4	0.9783	0.26	-479.94

Table 4.3: These results show the final profit, penalty, and global objective function values when four different random seeds are used. The similarity among the results show that the random number seed has no real effect on the final dig limits.




Figure 4.8: The chart shows the global objective function value for the four random seed number runs. The final global objective function values are very close.

Figure 4.9 shows all four dig limits overlaid. The numerical differences between the final dig limit polygons are small. Differences between the dig limit polygons are observable, but may not be significant in the field. For any ore or waste body, there may be more than one dig limit polygon that brings the global objective function to its current value. This result is because the vertices can be arranged in an infinite number of locations to give the same global objective function value whilst still being constrained by the digability penalty function. This is particularly the case when there are a large number of vertices in the dig limit polygon. To summarize, there are a number of ways to construct a dig limit polygon that minimizes the global objective function.

As shown by the tabulated results, all the dig limits converge to nearly the same global objective function value, but not the same dig limit polygon. Any optimal dig limit polygon will be as good as any other polygon. Interestingly, all four of the dig limit polygons have the same number of vertices. The number of vertices is not set deterministically by the user or the algorithm. It is however influenced by the maximum and minimum permissible distances between vertices.

4.2 Comparing Hand Drawn and Semi-Automatic Dig Limits

The Dig Limit Challenge is an experiment designed to demonstrate that the dig limit selection algorithm gives dig limits that equal or better hand drawn dig limits. The idea for the experiment is to have mining professionals select dig limits on an expected profit map; use the hand drawn dig limits to calibrate the dig limit selection algorithm; use the dig limit selection algorithm to select dig limits that have select features of the hand drawn dig limits; and finally, compare the results. This means that for each hand drawn dig limit there is a corresponding semi-automatic dig limit polygon.

The expected profit map used for the Challenge is shown in Figure 4.10. Four mining professionals participated in the experiment and the hand drawn dig limits were scanned and digitized. Most of the parameters for the dig limit selection algorithm were derived from the digitized dig limit polygons and used to calibrate the dig limit selection algorithm. The individual hand drawn dig limit polygons were used as initial polygons. The parameters




Figure 4.9: The map shows four dig limits superimposed on a single map of expected profit. In addition to the fact that the random number seed has no effect on the final dig limits, another notion is conveyed here: there may be more than one optimal dig limit for any dig limit selection problem.




Figure 4.10: This map of expected profit was provided to several mining professionals to draw dig limits on. These hand drawn dig limits were compared to semi-automatic dig limits.




Figure 4.11: This chart shows the 4 different penalty functions calculated from the hand drawn dig limits. The x -axis has the angle of operation and the y -axis has the penalty for a given angle.

d_{min} and d_{max} , the minimum and maximum distance between vertices respectively, were interpreted as the minimum and maximum distances between vertices from the hand drawn polygons. The penalty function for each dig limit polygon was built by calculating the probability of finding an angle less than a specified angle on 10° increments and plotting the results.

Figure 4.11 shows the penalty functions for each professional. The y -axis shows the percent of total penalty for the given operating angle on the x -axis. Professional number 4 gave the penalty function with the lowest selectivity. Penalty function 4 implies that the mining equipment was not very selective because the mining equipment starts gaining efficiency on angles of operation greater than about 100° . Professional number 3 has the most selective equipment. The mining equipment starts gaining efficiency with angles greater than 80° . The other two penalty functions specified selectivities between professionals 3 and 4.

Two comparisons were considered. The first does not constrain the dig limits to the users overall penalty, while the other comparison does constrain the final dig limit polygon to have a similar final value for the digability component function. The second comparison is more fair than the first because it aims to reproduce professional judgment of how selective the mining equipment is by reproducing the final value of the digability component objective function. The results for the first comparison are tabulated in Table 4.4. The units of profit are arbitrary. Under the Hand Drawn column, profit for the hand drawn dig limits was obtained by summing all of the fractions of blocks falling inside the hand drawn dig limits. The penalty column results were obtained by using the individually inferred penalty functions and summing all of the penalties. The semi-automatic column posts the final results from the dig limit selection algorithm for each case. The improvement column posts the percent improvement that the semi-automatic dig limits give over the hand drawn dig limits. The results show that semi-automatic dig limits consistently offer improved profit and digability over every hand drawn dig limit polygon. The greatest improvement in profit is on hand drawn dig limit polygon number 4. Profit is improved by 3.46%, with only a

	Hand Drawn		Semi-automatic		Improvement	
	Profit	Penalty	Profit	Penalty	Profit (%)	Penalty (%)
1	1196888	14529	1210401	11372	1.13	21.73
2	1204202	16347	1210873	10801	0.55	33.93
3	1197635	14487	1212118	7991	1.21	44.84
4	1168998	12111	1209424	12432	3.46	2.65

Table 4.4: This table shows the results for comparing the hand drawn and semi-automatic dig limits. These results do not aim to honor the digability derived from the hand drawn dig limits.

	Hand Limits		Semi-auto Limits		Improvement	
	Profit	Penalty	Profit	Penalty	Profit (%)	Penalty (%)
1	1196888	14529	1214297	14556	1.45	0.19
2	1204202	16347	1216690	16342	1.03	0.03
3	1197635	14487	1216118	14494	1.62	0.05
4	1168998	12111	1210911	12123	3.58	0.11

Table 4.5: This table shows the results for comparing the hand drawn and semi-automatic dig limits. These results do not aim to honor the digability derived from the hand drawn dig limits.

moderate 2.65% decrease in the final value for the digability component objective function. The improvement for Professional number 3's dig limits is 45% in penalty, and 1.21% in profit. Professional number 1's dig limits give similar gains: profit improved by 1.13% and penalty is reduced by 34%. Professional number 2 had the smallest improvement: profit increased by 0.55%, and penalty improved by 34%.

Table 4.4 shows the results for the comparison between hand drawn dig limits and the semi-automatically selected dig limits that are constrained to have the same, or close to the same, final value for the digability component objective function. To constraint the solution another component objective function, $O_{difference}$ is added to the global objective function:

$$O_{difference} = (O_{digability} - O_{hand})^2 \quad (4.1)$$

where $O_{digability}$ is the digability component objective function covered in Section 3.1.2, and O_{hand} are the individual total penalty for the hand drawn dig limits from Table 4.5. $O_{difference}$ is the squared difference between the target total penalty and the final penalty and when the difference between $O_{digability}$ and O_{hand} is large then $O_{difference}$ is large. The effect is that the algorithm aims to minimize the difference between $O_{digability}$ and O_{hand} , but does not precisely honor the total hand drawn penalty. The global objective function for the dig limit challenge becomes

$$O_{global} = O_{profit} + O_{digability} + O_{difference} \quad (4.2)$$

The final values for the digability component objective function match within 0.5% for each case. Also, in each case the semi-automatic dig limits improved profit. Hand drawn dig limit polygon 4 has the biggest improvement in profit of 3.58% and hand drawn dig limit polygon 2 had the smallest improvement of 1.03%. The average improvement is 1.92%.

Figure 4.12 shows the hand drawn and semi-automatically selected dig limits side-by-side. There are only subtle differences between the dig limits; one would not expect to see obvious differences when changes are in the order of 1-3%.




Figure 4.12: The maps shown compare the hand drawn dig limits (on the left) and the semi-automatic dig limits (on the right). The plotted semi-automatic dig limits aim to honor the hand drawn digability.

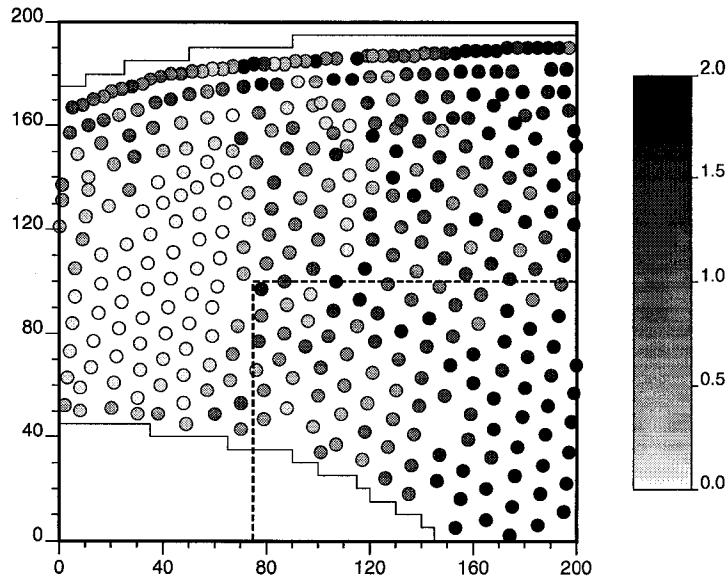


Figure 4.13: The map shows the sample locations and grades for a copper mine in Chile. The drill hole spacing is about 10m. The scale is percent copper. The outlined region will be used to select dig limits.

While the expected profit map used in the Dig Limit Challenge was contrived, it served to show that the dig limit selection algorithm consistently selects dig limits that can recover more profit than hand drawn dig limits which are subjective, inconsistent among professionals, and not reproducible.

4.3 Copper Mine Case Study

This case study shows the selection of dig limits from an undisclosed copper mine in Chile. The case study starts with borehole data, and works through all the steps in a dig limit selection exercise.

Figure 4.13 shows the sample information on a color coded location map. The data values is percent copper assayed from blasthole data. The bottom right corner has most of the high grade material. The area is 200m x 200m subsection from a previously active bench. The histogram of grades is shown in Figure 4.14. There are 372 data with an average copper grade of 1.015 % and a standard deviation of 1.19%. The maximum grade is 11.01%. A selling price of \$ 0.75 per *lb* is used in the example. The cost of shipping of ore and waste are the same and the cost of treatment is \$0.48 per *lb*. These costs give a cutoff grade of 0.8%.

The data were normal scores transformed to calculate and model the semivariogram. The semivariogram in Figure 4.15 shows moderate zonal anisotropy with the direction of maximal continuity about N5°E. A kriged map of grades is shown in Figure 4.16. The kriged map is shown for reference only. The model of uncertainty on grade is was built with Sequential Gaussian simulation. The model of uncertainty has 100 realizations. The block size is 2.5m x 2.5m on an 80x80 grid. A single grade realization is shown in Figure 4.17. The fine scale realizations were block averaged to 5.0m x 5.0m blocks on a 40x40 regular grid. A block averaged grade realization (the same realization already shown) is shown in Figure 4.18.

The block averaged model of uncertainty consisting of 100 realizations is used by the




Figure 4.14: The histogram shows the distribution of copper grades for the blasthole samples. The mean is 1.02 and the cutoff grade is 0.8%




Figure 4.15: The experimental and modeled semi-variogram are shown. The direction of maximal continuity is N5°E (the solid line). There is moderate zonal anisotropy.




Figure 4.16: The map shows kriged estimates of percent copper.




Figure 4.17: The map is of a single realization of percent copper from the model of uncertainty.




Figure 4.18: This map shows a block averaged percent copper map of the fine scale realization shown in Figure 4.17.

expected profit transform to give the map shown in Figure 4.19. The scale is expected profit. The selling price is set at \$0.75/lb. cpw (cost per unit waste) is set to 0 scaled.

The advantage of using the expected profit transform over a cutoff grade is demonstrated by comparing which blocks that the expected profit transform would identify as ore and waste and which blocks a cutoff grade classification would identify as ore. This procedure is used only for this experiment to demonstrate the advantage of the grade to expected profit transform. All other experiments use the dig limit selection algorithm to classify material as ore and waste.

The map in Figure 4.20 identifies blocks that fall into one of the four following scenarios: (1) The blocks that have a cutoff grade classification of ore, but have an expected profit classification of waste ($Waste_{Ez} < zc$). (2) The blocks that both classification schemes called ore (*Ore*). (3) The blocks that both the cutoff grade and expected profit classify as waste (*Waste*). (4) The blocks that the expected profit classification scheme identifies as ore, but the cutoff grade classification identifies as waste ($Ore_{Ez} > zc$). None of the blocks fell into scenario 1. 47.8% of the blocks fell into scenario 2. 48.8% of the blocks fell into scenario 3, and 3.4% of the blocks fell into scenario 4. If the expected profit values are correct, a cutoff grade control system wastes 3.4% more tonnes of ore than the expected profit transform classification technique.

The blocks that would have been classified as waste using a cutoff grade to classify the material, only occur next to regions of waste. Using the cutoff grade for grade control might have sent more ore than necessary to the waste dump and reduced expected profit.

A block misclassified as waste using a cutoff grade and called ore by the ore/waste transform is selected for closer study (block index (21,1)). The corresponding grade to profit transform, histogram of grades, and histogram of profits is shown in Figure 4.21. The 100 points shown in the scatterplot are the 100 profit values from the 100 grade realizations. The expected grade is 0.73%, which is less than the cutoff grade, but the expected profit is \$83.81. Under the expected profit classification criteria, this block would be classified as ore.

The initial dig limits are shown in Figure 4.22. Sometimes there is a need to specify vertices in the dig limit polygon that should not be allowed to be perturbed. This would be the case, for example, when the dig limit polygon is delineates the toe or the crest of a bench. In the Figure, vertices that appear as circles are perturbed; all other vertices are not




Figure 4.19: Here is the expected profit map that will be used to select semi-automatic dig limits. 100 realizations and a selling price of \$0.75/lb were used. The *cpw* (cost per unit waste) was not scaled. The scale is expected profit.




Figure 4.20: None of the blocks had an expected grade greater than the cutoff grade and an expected profit less than 0. Blocks with profit less than zero are classified as waste, all other blocks are ore. Some of the blocks have an expected grade less than the cutoff grade but are classified as ore because the expected profit is greater than 0.




Figure 4.21: The image on the left shows grade to profit conversion for a single realization on the x -axis and profit on the y -axis. The points represent the grade to profit conversion for 100 realizations for a location where the expected grade is less than the cutoff grade. The expected profit is \$83.81. The histogram in the middle, is the histogram of the simulated profit values. The histogram on the right, is the histogram of 100 simulated grades. The expected grade is 0.7318%.

allowed to be perturbed.

A digability factor catalogue is constructed to help select a digability factor and is shown in Figure 4.23. The top left map has a digability factor of 0.1. The rest of the catalogue increments the digability factor by 0.1 from left to right and down to the bottom of the Figure.

Once a digability factor is selected the coordinates for the dig limit are either given to the survey crew for staking, or the material classified by the dig limits is stored in the geocellular model.

4.4 Bingham Canyon Mine

Bingham Canyon Mine undertook a case study to investigate the potential economic advantages of moving from deterministic mapping and hand drawn dig limit grade control to simulation and semi-automatic dig limit selection grade control. The results discussed here centre on dig limit selection and certain aspects of a simulation based approach to grade control.

Bingham Canyon Mine is located near Salt Lake City, Utah, and is operated by Kennecott Utah Copper. The mine is more than 100 years old and is considered to be the world's largest mine. It measures approximately 4km across and is approximately 800m deep and can be discerned from space.

Bingham Canyon is a porphyry copper deposit. The ore body consists of several mineralized stocks and dikes hosted by a regional quartzite. The most important mineral content contributors to the deposit are the monzonite and the quartz monzonite porphyry intrusions. The quartz monzonite sits in the northwest edge of the monzonite, and the monzonite porphyry is inside the quartz monzonite. The primary copper bearing minerals are chalcopyrite, bornite, and chalcocite. Gold and molybdenite are the primary byproduct metals. An endoskarn is present that contains minerals such as talc that can cause problems in the mill.

Prior to the case study, inverse distance cubed interpolation (ID3 model) was used to construct maps of copper, gold, and molybdenum grades. These deterministic models were then jointly transformed from grade information to a proprietary profit indexing criteria called waste/ore ranking (WOR). WOR accounts for all the costs associated with recovering mineral from the material. The grade control engineers used the WOR maps to hand draw dig limits that account for digability and maximize profit.




Figure 4.22: Shown in this Figure is the seed dig limit. The vertices that will be permitted to move are indicated by a small circle.

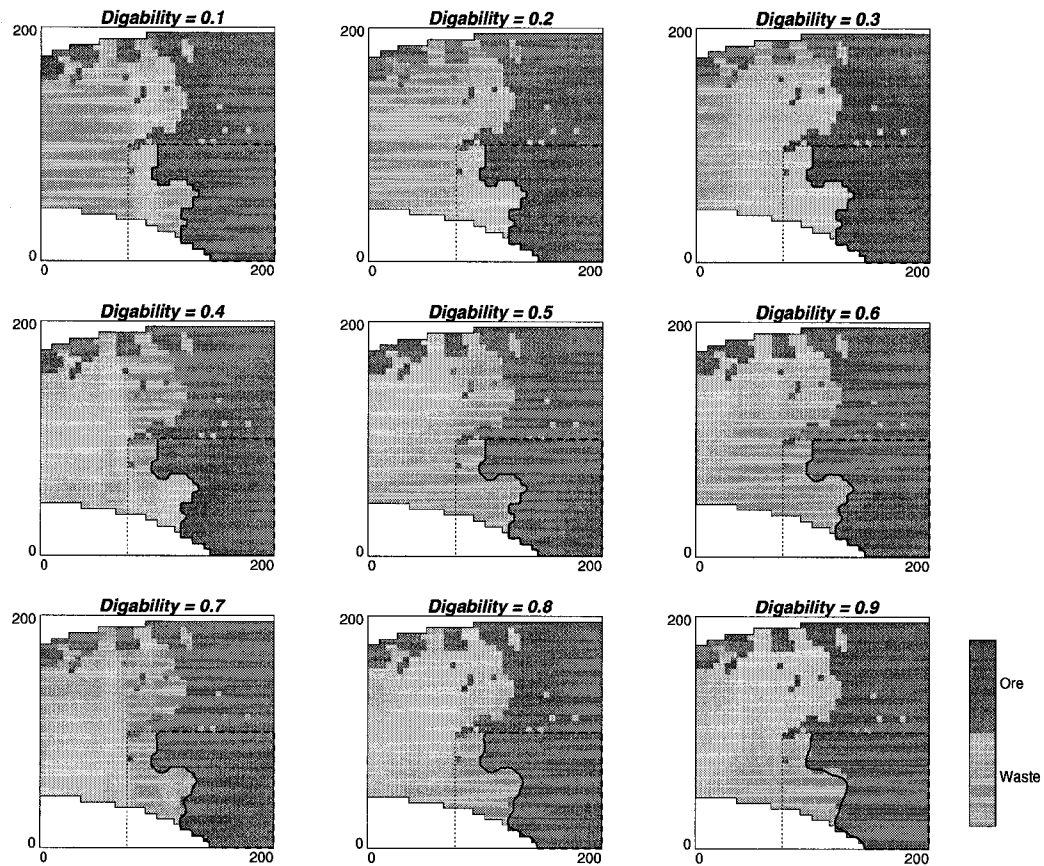


Figure 4.23: This Figure shows a digability catalogue. The digability factors are shown above each of the final dig limits. Dig limits with the lowest digability are shown in the top left corner. Dig limits with the highest digability are shown in the bottom right corner.

The case study involved using sequential Gaussian simulation to construct models of uncertainty on copper, gold, and molybdenum (Simulation model). An expected WOR transform (EWOR) was used to transform the models of uncertainty into maps for dig limit selection.

A datum model was constructed using sequential Gaussian simulation. The datum model was used to compare ID3 and simulation results. The dig limits were compared across all three models.

4.4.1 Data and Geostatistical Models

Two subsections of the 4940 and 5640 benches were considered. The relevant data locations are included inside the bolded polygons shown in Figure 4.24. Blast hole data within the area of interest as well as the data from the bench immediately above were used in the case study. The average data spacing was about 30ft by 30ft.

The final grid block size was 10ftx10ft. The Datum and Simulation models started with 2ftx2ft grid blocks that were averaged up to 10ftx10ft blocks to account for volume variance issues; see [43] and [45]. The ID3 model was modeled to 10ftx10ft resolution directly.

The raw data was the same for the ID3 and Simulation models. As much as possible, the search strategies were kept the same for modeling.

Of special interest is the complex relationship between copper, gold, and molybdenum that varied with location. The types of relationships these minerals share cannot be reproduced using sequential Gaussian simulation alone. The simulation model used a stepwise conditional transformation to remove the correlation between the variables. The stepwise conditional procedure is to perform a normal scores transform to the copper data – copper is transformed first because it is the most important variable – transform the gold data conditional to the normal scores copper data, and transform the molybdenum data conditional to the normal scores gold data. The transformed copper, gold, and molybdenum models can be constructed independently using SGS. The data are backtransformed by reversing the transform steps to restore the original relationship. Details on the stepwise conditional transform can be found in [55].

The metal content models were constrained to a deterministic rock type model. Figure 4.25 shows the rock type maps for the 4940 on the left and the 5640 on the right. In total there are 35 rock types. Not all rock types are represented by the maps. The rock type model was constructed using a coarser grid than the metal content models. The area of interest is shown by the polygon. The influence of the rock type model is apparent on the EWOR maps.

Expected Waste/Ore Ranking Calculation and the Datum Model

The EWOR calculation is similar in principle to the grade to profit transform outlined in 3.1.1, except that there are more and different variables involved in the EWOR transform. The EWOR transform is also a non-linear transform. The operating state of the equipment, the availability of ore, the presence of contaminants and all the costs associated with recovering mineral from the material are incorporated into the EWOR calculation. EWOR is expressed as a rate in terms of dollars per hour:

$$EWOR[\$/hr] = \frac{1}{L} \sum_{l=1}^L (OreValue^l[\$/t] - WasteValue^l[\$/t]) * OreProcessingRate[t/hr] \quad (4.3)$$

The input parameters for the EWOR calculation can change daily. Maximization of profit requires that WOR be re-evaluated when economic criteria change. Constantly changing criteria mean constantly changing dig limits. It is challenging for the mining engineer to select dig limits that consistently account for the limitations of mining equipment and maximize profit.





Figure 4.24: The maps show the sample locations for the two benches being considered. The areas of interest are delineated by polygons. The overlaid grid measures 400x400ft




Figure 4.25: The maps show rock types for the 4940 (left) and 5640 (right) benches. The polygons show the areas of interest. The unshaded portions are waste or previously mined regions. On the 4940 map, the area of interest is extended to the east, but this region is ignored in the dig limit selection because it is known to be a waste region.

The Datum, ID3, Simulation models are all processed by the EWOR transform. The data for the Datum model is obtained by randomly selecting a WOR realization, and extracting samples on a 10ftx10ft spacing. These sample data are used to construct a single realization of WOR on a 2ftx2ft grid. The grid is upscaled to a 10ftx10ft grid and used for selecting dig limits.

Figure 4.26, and 4.27 shows the Datum, ID3, and Simulation EWOR maps for the 4940 and 5640 regions respectively.

The maps all have similar features. The ID3 and Simulation EWOR maps are smoother than the Datum map. ID3 makes smooth maps because the estimates at unsampled locations are averages of nearby samples. The Simulation EWOR map is smooth because the EWOR transform averages the L realizations at every location. The EWOR maps clearly show the effect of the rock type model and that some of the rock types are clearly non-economic.

Basis for Comparison

The object of the comparison is to demonstrate the economic advantage arising from the dig limit selection approach over current mine practices. It is acknowledged that the datum model is not actually the mineralization that would occur sub-surface; however, a basis for comparison is required and the datum model will be treated as a ground truth for comparison. The basis for comparison is as follows:

1. Use the same data to construct a Datum model, an ID3 model, and a Simulation model. The mining engineer at Bingham Canyon constructed the ID3 model. The Datum model and Simulation models were constructed by external consultants Harry Parker and Bruce Davis.
2. Transform the three grade models to EWOR maps.
3. Have the Bingham Canyon mining engineer select dig limits by hand on the ID3 EWOR map using the mines grade control software. This represents current grade control practice in the mine.

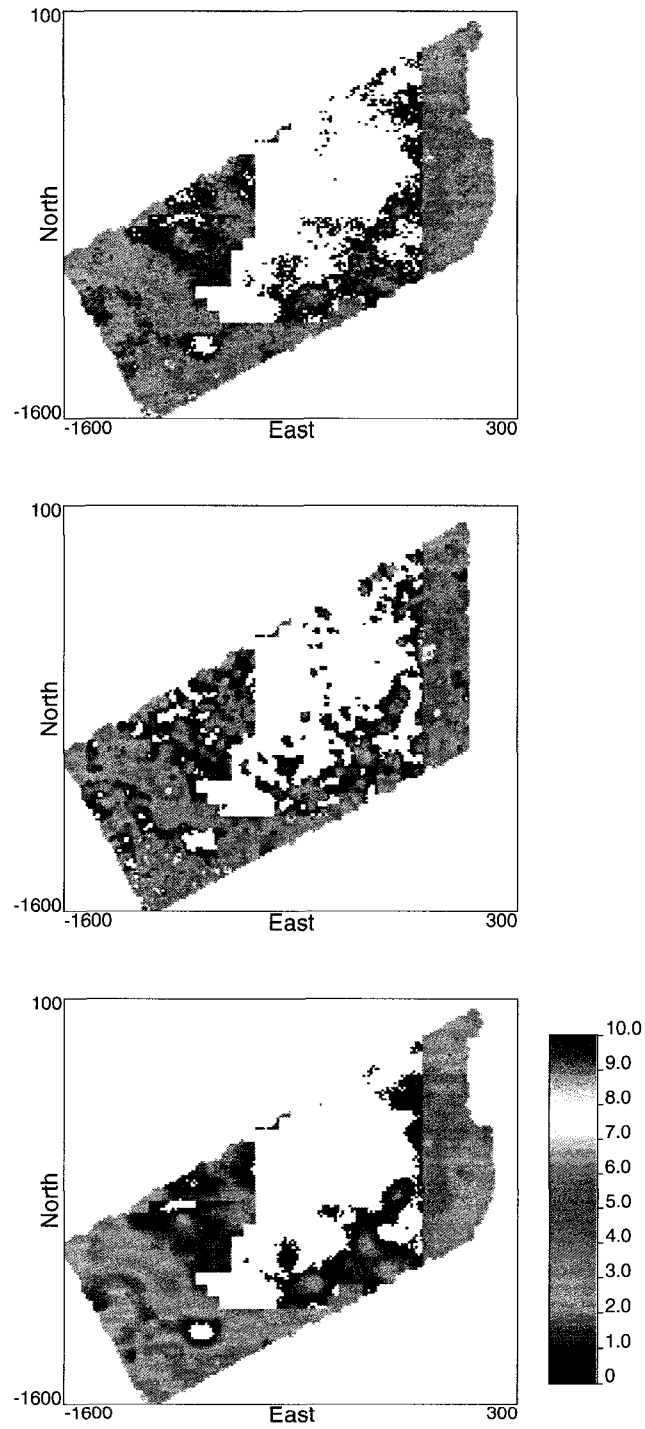


Figure 4.26: From the top of the Figure are the Datum, ID3, and Simulation EWOR maps for the 4940 region. The block size is 10x10. The white areas are waste regions.

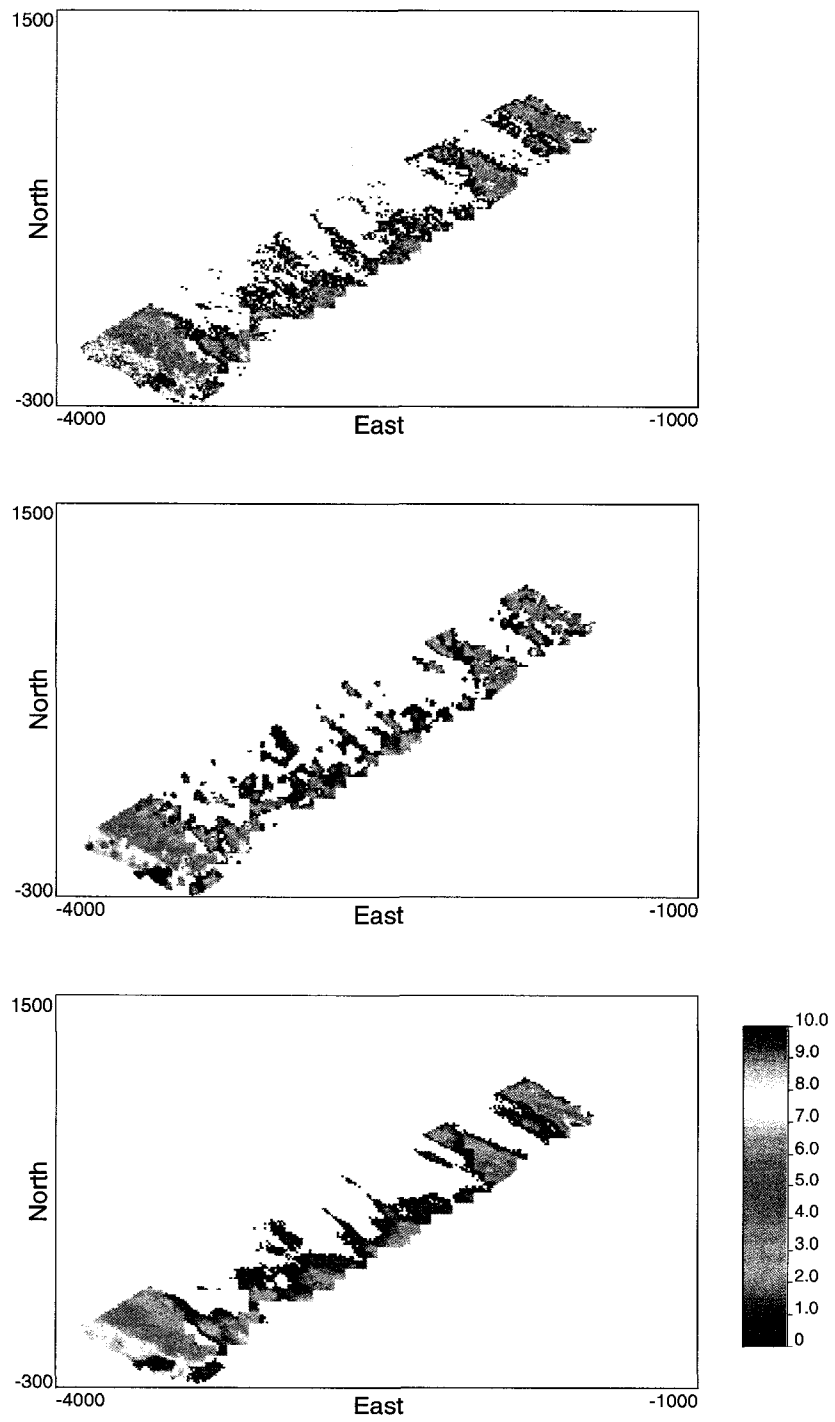


Figure 4.27: From the top of the Figure, are the Datum, ID3, and Simulation EWOR maps for the 5640 region. The block size is 10x10. The white areas are waste regions.

Area	Total EWOR from Hand	Total EWOR from Auto	Total EWOR Datum	Total % Gain
4940	\$34,746,000	\$35,109,000	\$35,343,000	1.03
5640	\$27,467,000	\$27,875,000	\$28,098,000	1.45

Table 4.6: Here are the tabulated values for prediction of total improvement comparing the simulation plus semi-automatic dig limits versus the ID3 method

4. Use the semi-automatic dig limit selection algorithm to select dig limits on the Simulation EWOR map. This represents the potential from the revised grade control practice.
5. Use the semi-automatic dig limit selection algorithm to select dig limits on the ID3 EWOR map. This demonstrates the potential economic advantage derived from using semi-automatic dig limit selection alone.

The benchmark for comparing the different grade control dig limits is the Datum EWOR map. The Datum map will be considered to be the material that the mill will really see, as opposed to what was used to select the dig limits.

The hand drawn dig limits are used as initial dig limits for the semi-automatic dig limits in all cases. The hand drawn dig limits along with some parameters from the mining equipment were used to select parameters for the dig limits. For example, the penalty function was derived from the digitized hand drawn dig limits. A digability catalog was used by the mining engineer to select a digability factor of 0.5. The minimum and maximum distances between vertices in the hand drawn dig limits, along with the size of the shovel, were used to select the minimum (3ft) and maximum (15ft) permissible distances between vertices in the algorithm. The procedure covered in Section 4.1.1 was used to select an annealing schedule. The same annealing schedule was used in all cases.

Semi-automatic Dig Limits on the Simulation Model Versus Hand Drawn Dig Limits on the ID3 Model

The numerical results for total possible improvement, that is, Simulation plus semi-automatic dig limits compared to ID3 with hand drawn dig limits, are shown in Table 4.6. The maps in Figure 4.28, show the dig limits superimposed on the ID3, Simulation, and Datum EWOR maps for the 4940 region. The top left map is the ID3 model and has hand drawn dig limits. Below this are the hand drawn dig limits on the Datum map. On the top right is the Simulation model and the semi-automatic dig limits. On the bottom right is the Datum model with the semi-automatic dig limits. Figure 4.29 shows the same configuration of results, but for the 5640 region.

Table 4.6 shows that the combination of simulation and the semi-automatic dig limit selection improve EWOR by approximately 1% on the 4940 Bench and 1.5% on the 5640 Bench.

4.4.2 Comments

Semi-automatic Dig Limits on the ID3 Model Versus Hand Drawn Dig Limits on the ID3 Model

The effect of adopting a simulation based grade control system on the dig limits can be observed by comparing the semi-automatic dig limits on the ID3 and simulation models and the numerical results tabulated in Table 4.7. The results show that the Simulation approach offered an improvement in the 5640 area (an improvement of approximately 1.5 percent).

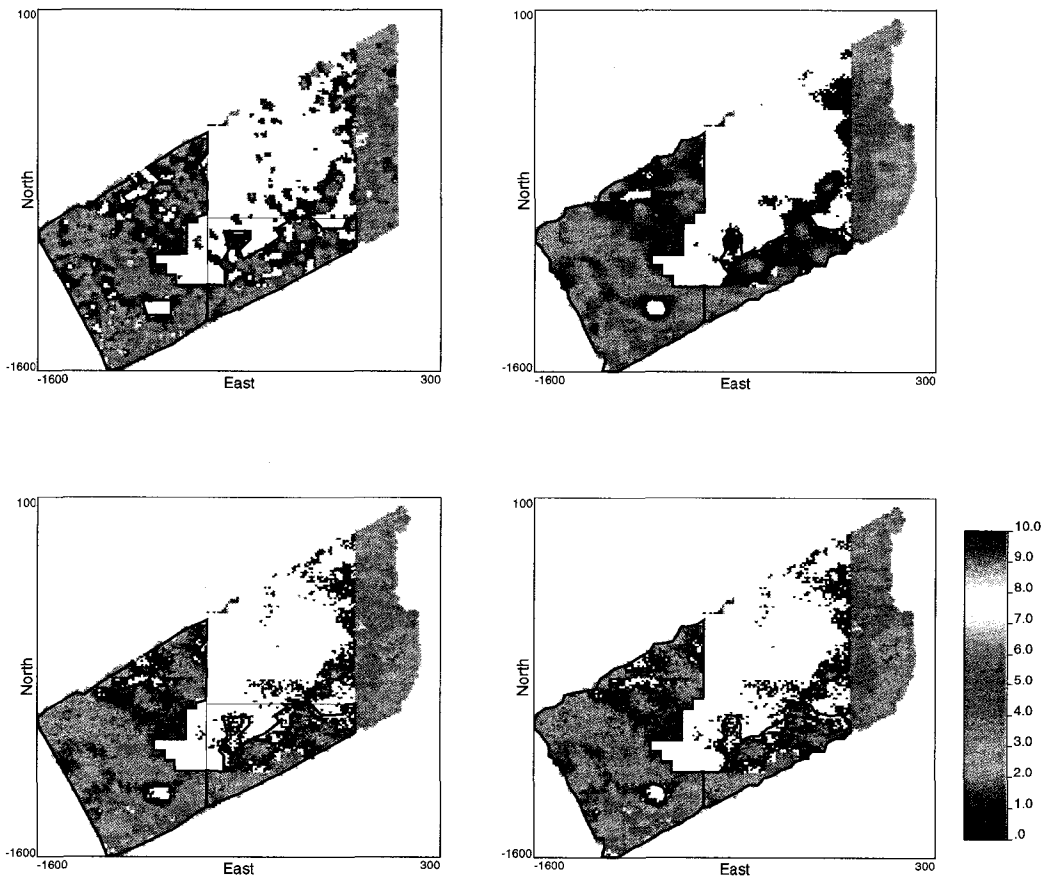


Figure 4.28: These maps are for the 4940 region results. The hand drawn dig limits on superimposed on the ID3 model are shown in the top left map. On the bottom left are the hand drawn dig limits on the Datum model. The semi-automatic dig limits on the Simulation model are shown on the top right, and the semi-automatic dig limits on the Datum model are shown on the bottom right.




Figure 4.29: These maps show the 5640 region results. The hand drawn dig limits on superimposed on the ID3 model are shown in the top left map. On the bottom left are the hand drawn dig limits on the Datum model. The semi-automatic dig limits on the Simulation model are shown on the top right, and the semi-automatic dig limits on the Datum model are shown on the bottom right.

Area	Total EWOR from Hand	Total EWOR from Auto	Total EWOR Datum	Total % Gain
4940	\$35,133,000	\$35,109,000	\$35,343,000	-0.07
5640	\$27,502,000	\$27,875,000	\$28,098,000	1.05

Table 4.7: These results show the predicted improvement using simulation based EWOR.

Area	Total EWOR from Hand	Total EWOR from Auto	Total EWOR Datum	Total % Gain
4940	\$34,746,000	\$35,133,000	\$35,343,000	1.09
5640	\$27,467,000	\$27,502,000	\$28,098,000	0.13

Table 4.8: These results show the prediction of improvement using semi-automatic dig limit selection alone.

One interpretation of this result is that simulation provides little benefit when the bench is mostly ore or waste (e.g., bench 4940) and the waste/ore contact is relatively simple. However, simulation can provide a benefit in areas where the ore/waste contact is complex (e.g., bench 5640) and carries greater uncertainty.

The maps in Figure 4.30 show the resulting dig limits for the 4940 region. The map on the top left shows the ID3 model and semi-automatic dig limits. Below that are the semi-automatic dig limits superimposed on the Datum model. On the top right is the Simulation model and the simulation based semi-automatic dig limits. The semi-automatic Simulation based dig limits are shown superimposed on the Datum model on the bottom right. Figure 4.31, shows the same configuration of results but for the 5640 region.

Semi-automatic Dig Limits on the ID3 Model Versus Hand Drawn Dig Limits on the ID3 Model

Table 4.8 compares the hand drawn dig limits on the ID3 model with the semi-automatic dig limits on the same model. This comparison attempts to assess the impact of using the semi-automatic dig limit program alone. The semi-automatic dig-limit algorithm gave a very slight improvement in the 5640 area and an improvement of approximately 1 percent in the 4940 area and 1.5 percent in the 4940 area. The maps in Figure 4.32 show the hand drawn dig limits on the 4940 EWOR map on the top left and the hand drawn dig limits on the Datum model on the bottom left. The semi-automatic dig limits on the ID3 model are shown on the top right, and the semi-automatic dig limits on the Datum model on the bottom right. Figures 4.33 shows the same configuration of results but for the 5640 region.

Concluding Comments

The Bingham Canyon Case Study showed that in regions where there is high uncertainty in grade a simulation based grade control program prevails and that little improvement is observed where there is little uncertainty. With no loss in improvement, it might be beneficial to adopt the simulation based grade control program. The Case Study also showed that semi-automatic dig limits give improved grade control whether or not a simulation based grade control program is used.

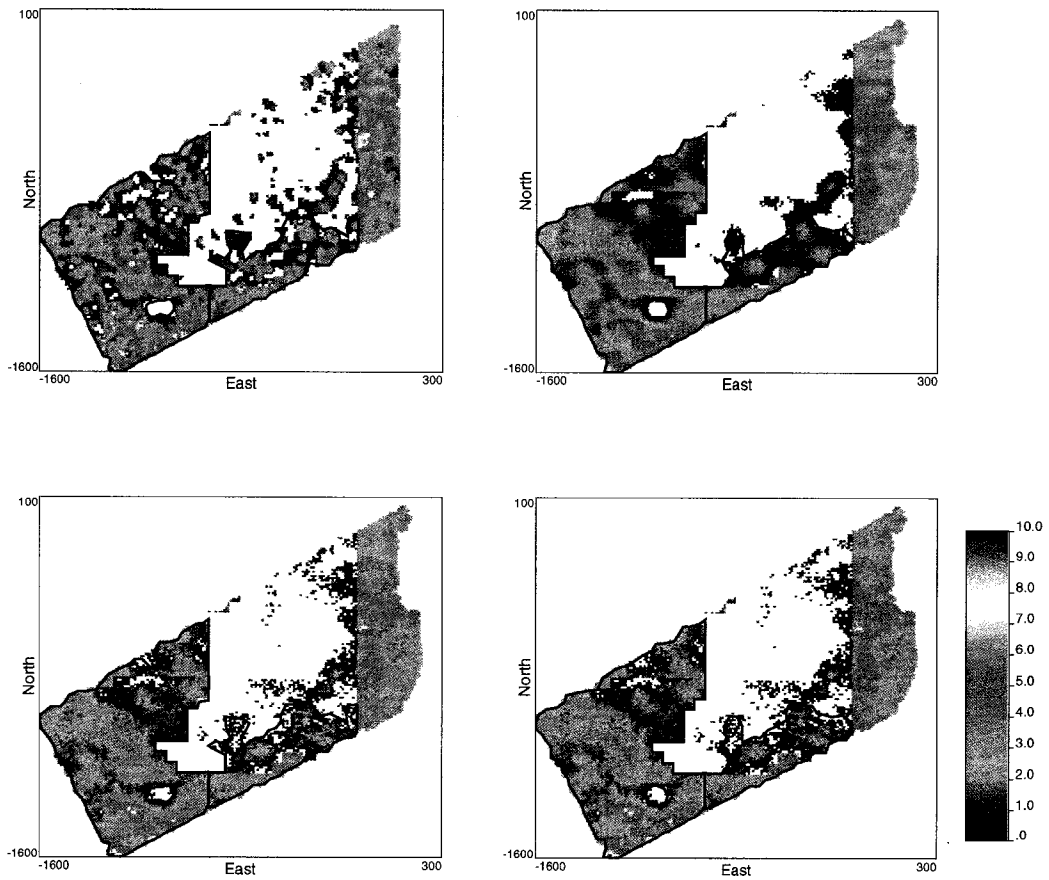


Figure 4.30: These maps are for the 4940 region results. The semi-automatic dig limits are superimposed on the ID3 model and shown in the top left map. On the bottom left are the semi-automatic dig limits on the Datum model. The semi-automatic dig limits on the Simulation model are shown on the top right, and the semi-automatic dig limits on the Datum model are shown on the bottom right.

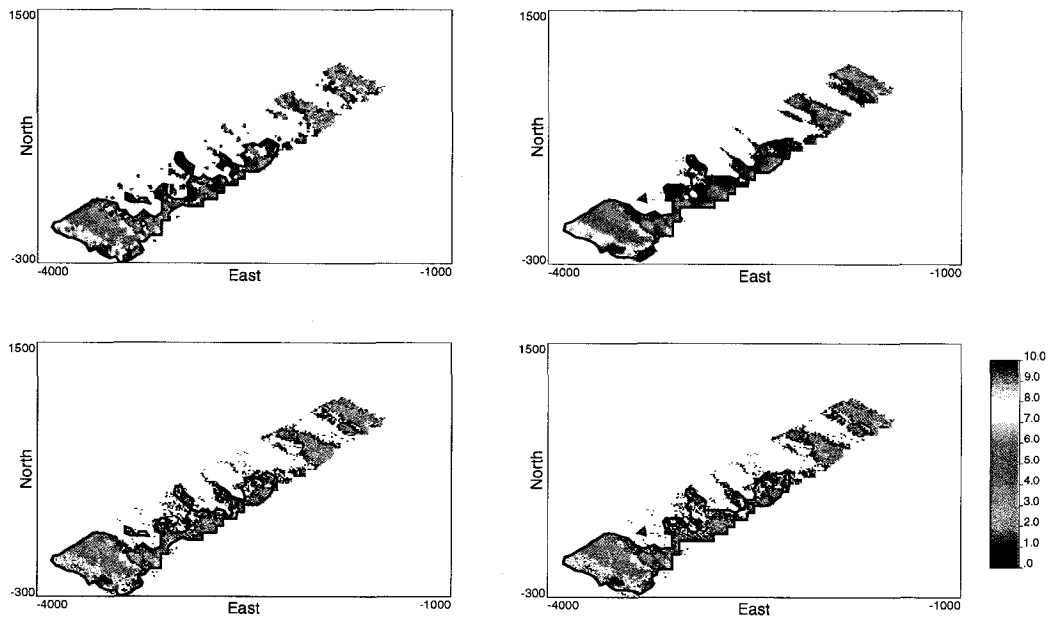


Figure 4.31: These maps are for the 5640 region results. The semi-automatic dig limits are superimposed on the ID3 model and shown in the top left map. On the bottom left are the semi-automatic dig limits on the Datum model. The semi-automatic dig limits on the Simulation model are shown on the top right, and the semi-automatic dig limits on the Datum model are shown on the bottom right.

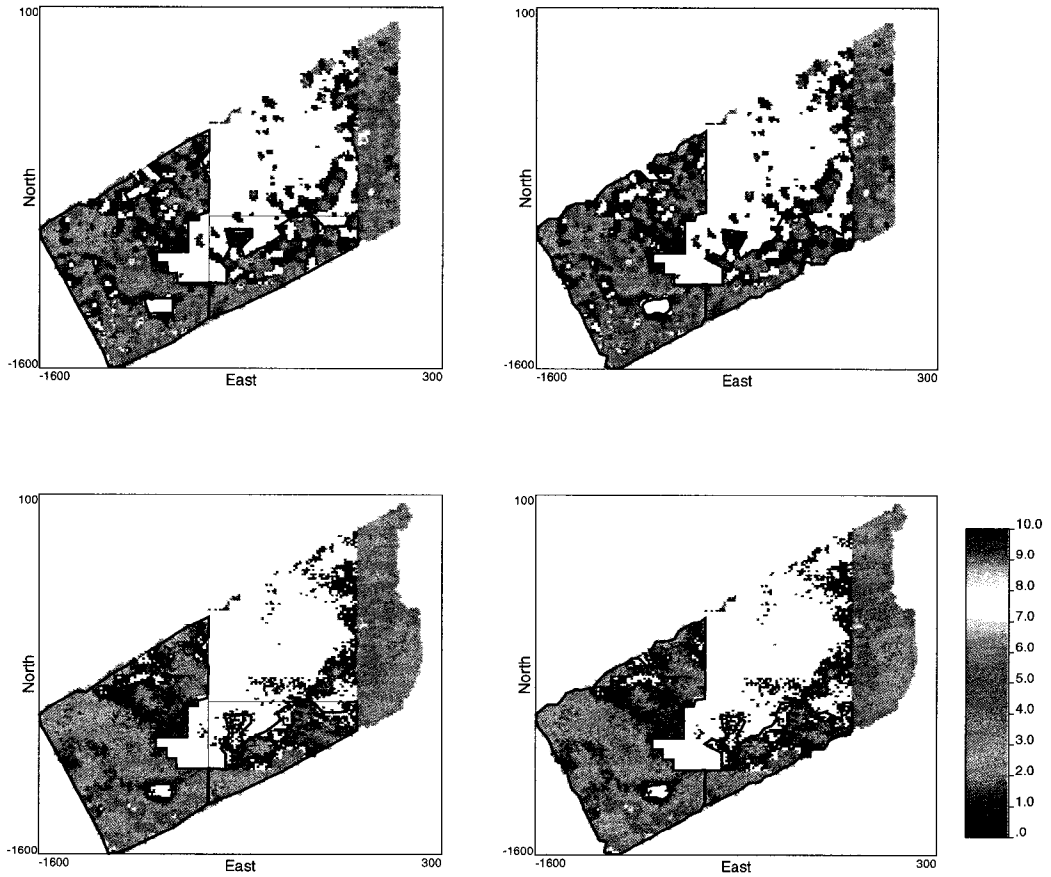


Figure 4.32: These maps are for the 4940 region results. The hand drawn dig limits are superimposed on the ID3 model are shown in the top left map. On the bottom left are the hand drawn dig limits on the Datum model. The semi-automatic dig limits on the Simulation model are shown on the top right, and the semi-automatic dig limits on the Datum model are shown on the bottom right.

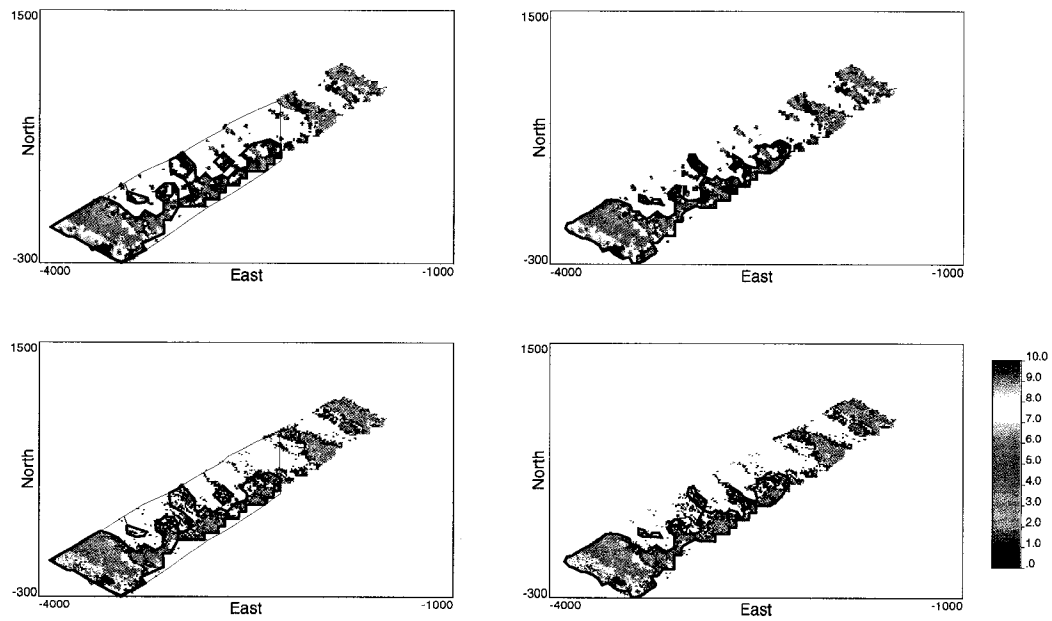


Figure 4.33: These maps are for the 5640 region results. The hand drawn dig limits are superimposed on the ID3 model are shown in the top left map. On the bottom left are the hand drawn dig limits on the Datum model. The semi-automatic dig limits on the Simulation model are shown on the top right, and the semi-automatic dig limits on the Datum model are shown on the bottom right.

Chapter 5

Semi-Automatic Well Location Selection Theory

This chapter presents the theoretical and implementation details for semi-automatic selection of well locations. The semi-automatic well location selection algorithm selects well locations that are jointly optimal with respect to the reservoir model of uncertainty.

5.1 Jointly Optimal Well Locations

In the petroleum industry well location selection rarely benefits from the consideration of the reservoir model of uncertainty because there are few tools to perform the task. Typical well location selection exercises performed by the Asset Team rely on a single realization of the reservoir model of uncertainty. The well plan is usually selected early in the project, and any subsequent uncertainty studies are constrained to observing the effect of the reservoir model of uncertainty on the previously selected well plan.

A challenging aspect of selecting well locations under uncertainty is the integration of multiple realizations into the decision making process. The well location selection algorithm considers all the realizations simultaneously and selects well locations that are jointly optimal with respect to the realizations of the reservoir model of uncertainty.

Consider a hypothetical well location selection problem for one well on a model of uncertainty consisting of five realizations. The goal of the selection problem is to maximize exposure to high quality reservoir. Figure 5.1 shows the five realizations on a 2D x - y plane. The shaded areas represent regions of high reservoir quality. The left and right sides show the same five realizations. The jointly optimal solution is presented on the left and an alternative approach is shown on the right.

The aim of the jointly optimal approach is to maximize exposure to high quality reservoir on all realizations at the same time. The alternative approach aims to select an optimal well location for each realization. This is the same approach referenced above where a single realization is used to select well locations.

The well locations are shown as black dots with derricks attached. Black dots without derricks attached are projections of other well locations from other realizations. The jointly optimal approach manages to intersect high quality reservoir on all 5 realizations. On inspection however, the jointly optimal location only barely manages to intersect high quality reservoir on some realizations. The alternative approach maximizes exposure to high quality reservoir on a realization-by-realization basis, and gives realization specific results only. The optimal well location for realization 1 has the best chance of intersecting high quality reservoir on all realizations because it intersects regions of high reservoir quality on realizations 1,2,3, and 4. The well location for realization 5 performs the worst because it only intersects a high quality reservoir on realization 5. The jointly optimal solution outperforms

the realization-by-realization approach because it intersects regions of high reservoir quality on all 5 realizations and has the best chance of meeting the goals of the Asset Team.

The other critical difference between the jointly optimal approach and the alternative approach is that there is no need to reconcile multiple solutions at the end of the location selection exercise. The maps at the bottom of the Figure show the jointly optimal solution on the left and the right side shows 5 realization specific optimal solutions. The jointly optimal solution can be used directly, whereas the realization specific approach has $L = 5$ realization optimal well locations. There is no straightforward technique for compiling the L well locations into a single optimal well plan.

Summary

The primary advantage to using the jointly optimal approach is that well locations are selected robustly with respect to the model of uncertainty and not constrained to a particular realization or interpretation. The joint optimal approach can be difficult to implement because the entire model of uncertainty must be read into memory and must be accessible at all times during optimization. The demands on computer memory or computing time increase with the variable space of the model of uncertainty (i.e., L realizations of structure, porosity, permeability, and water saturation). Despite these demands, the jointly optimal solution is considerably faster than performing flow simulation on a number of well locations on the L realizations.

5.2 Global Objective Function

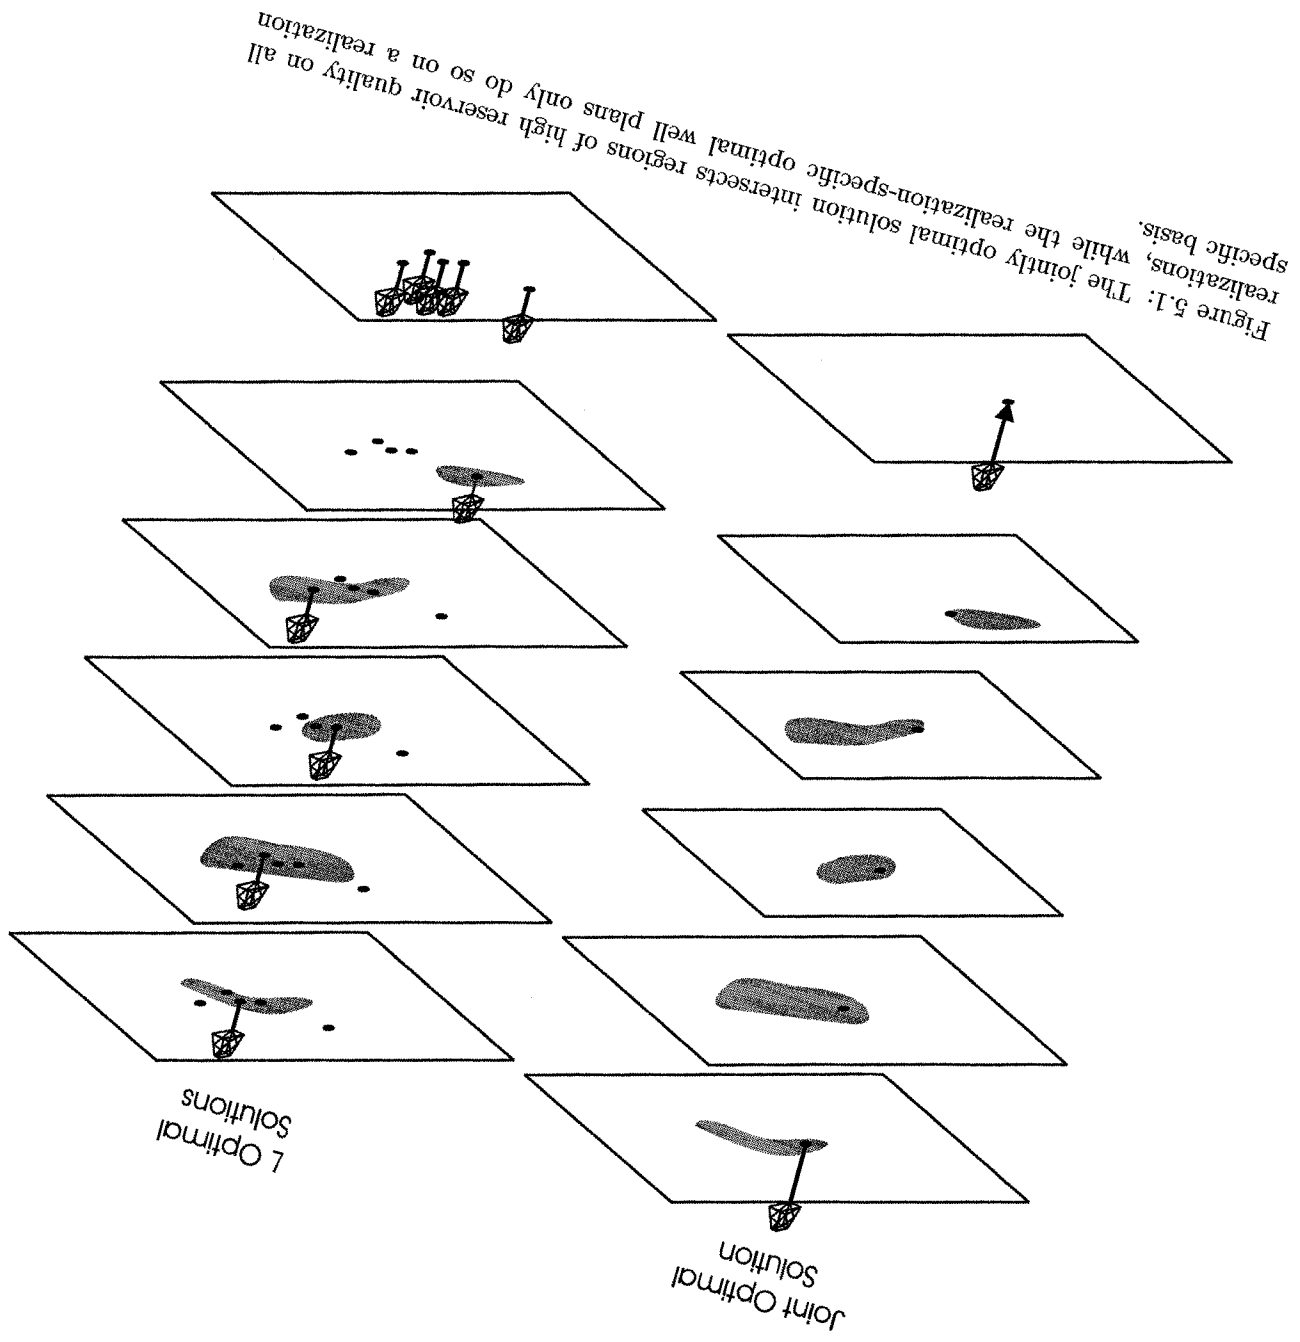
Water flooding is an important secondary hydrocarbon recovery technique, that is also becoming an increasingly important for primary hydrocarbon recovery. The number of discoveries of significant hydrocarbon fields is decreasing, and in general, the reservoir quality of new found reservoirs is decreasing. Water flooding early in reservoir development can lead to improved recovery efficiency over the life of the reservoir compared to a two stage approach that depletes the reservoir first under the existing drive mechanism and follows up with water flooding. This is particularly the case in deep water and ultra deep water reservoirs with low initial pressures. The selection of injector and producer locations is an important factor in maximizing recovery and shareholder value.

Some primary limiting economic constraints for a water flood are the cost of water injection, the cost of treatment of injected and produced water, and the cost obtaining or disposal of water. When water these costs equal the income of the water flood, it is uneconomic to continue operation.

The global objective function for the well location selection problem aims to quantify the economics of a water injection well plan and serve as a proxy to flow simulation. The global objective function is posed as a weighted linear combination of three component objectives that characterize key economic factors of a given well plan. The global objective function is written:

$$\begin{aligned} GOBJ &= \sum_{i=1}^I COBJ_i^{\lambda_i} \\ &= C_{prod} \cdot COBJ_{EOPT}^{\lambda_{EOPT}} - C_{treat} \cdot COBJ_{EWPT}^{\lambda_{EWPT}} - C_{inject} \cdot COBJ_{EWIT}^{\lambda_{EWIT}} \end{aligned}$$

where C_{prod} is the cost of producing oil in terms of barrels of oil equivalent, C_{treat} is the cost of producing water and separating it from the oil in terms of barrel oil equivalent, C_{inject} is the cost of pumping and pretreating water. The component objective functions are defined as follows:



- $COBJ_{EOPT}$ estimates the volume of produced oil for a given well plan.
- $COBJ_{EWPT}$ estimates the volume of water produced for a given well plan.
- $COBJ_{EWIT}$ estimates the volume of water required to produce oil for a given well plan.

The well location selection algorithm iteratively perturbs the well plan until it conforms to a well plan that minimizes the global objective function. Recall that in order to abide by the conventions of optimization, the global objective function is minimized. The global objective function written above is easily modified to satisfy this convention by multiplying it by -1. For ease of discussion, the term optimized will be used to minimize confusion about the minimization and maximization of the global objective function.

The global objective function is merely a proxy for the net present value of a well plan because the fluid volumes are estimates and the estimates are calculated without the flow equations. The weights, λ_{EOPT} , λ_{EWPT} , λ_{EWIT} , have great importance as they are used to calibrate the global objective function to flow simulation results and give a global objective function of improved ability to predict well plan performance.

Several factors affect hydrocarbon recovery under waterflood and strategic selection of injector and producer locations with these factors in mind can lead to maximized recovery with minimized cost. The following Section discusses the guiding principles for the selection of a water flood well plan. The construction of the component objective functions depends on these principles and is also discussed in the next Section. Following the discussion on the component objective functions, is a Section discussing the calibration of the global objective function. The calibrated global objective function is a good proxy for flow simulation and it is relied on to take the place of flow simulation in the well location selection technique.

5.2.1 Component Objective Functions

The following discussion highlights important factors for selecting well locations in a water flood project:

Reservoir Structure :

The structure and stratigraphy of the reservoir controls the location of the wells. Structure is the principal factor governing gravitational segregation. Producer wells down dip from injection wells realize shorter water breakthrough times than well plans that have the injector down dip from the producer.

Lithology :

Lithological factors that influence floodability are porosity, permeability, clay content, and mineral content. In some reservoirs, only a small portion of the total volume possesses sufficient permeability to be effective in water flooding. Clays can swell clog pores and subsequently affect permeability. Chemical reactions between injection fluids and resident minerals can change the wettability of the matrix and alter flooding performance. Lithology effects are usually captured by the facies model.

Reservoir Depth and Matrix Integrity :

Reservoir depth and matrix integrity govern injection pressure. For shallow reservoirs, high injection pressures can either exceed the fracture strength of the matrix, or open naturally existing fractures and reduce water injection effectiveness.

Porosity :

The total available hydrocarbon for recovery is controlled by porosity. In many reservoirs there is positive correlation between porosity and permeability.

Permeability (Magnitude and degree of variability) : The magnitude of the permeability controls rate of injection and production. Also, high permeabilities may form conduits for water flow. Thus, permeability and the spatial distribution of permeability is an important factor in considering the spacings and locations of wells, as well as determination of volumes required for injection. Homogenous permeabilities are preferred. High permeability regions that are laterally extensive will facilitate early water breakthrough; however, lack of spatial correlation of permeability among producers and injectors reduces the chance of early water breakthrough is reduced.

Continuity of Reservoir Rock Properties :

The lack of spatial correlation among the important petrophysical properties leads to reduced reservoir connectivity and reduced hydrocarbon recovery.

Fluid Saturations and Distributions :

High hydrocarbon saturation is preferred to low, not just because it translates to higher volumes of recoverable oil, but because the relative permeability to oil is higher when oil saturation is high. As a waterflood proceeds, saturation of water increases, and permeability to oil decreases.

Fluid Properties and Relative Permeability Relationships :

The fluid properties have important effect on water flood performance. Of major importance, is the viscosity of oil because it affects the mobility ratio. The mobility ratio is the ratio of the mobility of the displacing fluid (water) to the mobility of the displaced fluid (oil). The greater the viscosity of the oil, the larger the mobility ratio and the lower the recovery after water breakthrough.

The Well Plan :

The injector and producer well locations are essential factors in a water flood. If the injector and producer spacing is too short the producers will water out prematurely and reduce ultimate recovery. Producers that are spaced too close will compete for pressure and displacing fluid.

Generally, all the above factors must be considered jointly when selecting a well plan. Three commonly used and frequently cited techniques for predicting waterflood performance and for planning well locations are the Stiles [79], Dykstra-Parsons [27], and the Buckley-Leverett frontal advance equation [10]. These approaches have been used extensively in industry, but there are issues with implementation. None of the three techniques can easily cope with a model of uncertainty or handle the heterogeneity offered by a reservoir model. Many professionals apply flow simulation to the problem of selecting well plans because flow simulation does account for multiple realizations and heterogeneity; however, flow simulators are costly to use in terms of CPU and professional time and software licensing fees are high.

The well location selection technique aims to select optimal injector and producer locations under a reservoir model of uncertainty, and to minimize dependence on flow simulation and the flow equations. Satisfaction of this aim means that some of the factors important for consideration of a waterflood well plan cannot be incorporated into the well location selection technique. This is partly because some factors rely on knowledge of how things change with respect to producing the reservoir which ultimately depend on time. For instance, initial fluid saturation change as the reservoir is produced as do the fluid properties.

None of the aforementioned considerations that involve change are incorporated into the global objective function. For instance, initial fluid saturations are considered, but saturation changes over time are not. Fluid properties, such as changes in viscosity due to gas coming out of solution, are not considered because they change with time. The following considerations for selecting a well plan are incorporated into the global objective function: reservoir structure, permeability, initial fluid saturations and distributions, rock properties, and the well plan.

Although the well location selection algorithm does not use flow simulation or the flow equations, flow simulation results are required to prepare the component objective functions. There are two instances where flow simulation results are required. First, flow simulation results are used as dependent data, and observations from a subset of the full field model, called a mini model, are used as independent data to construct a numerical model of how an injector-producer pair perform under varying separation distances and over multiple realizations. Second, flow simulation is used to calibrate the global objective function against full field response to dynamic conditions.

The mini model minimizes dependence of the well location selection technique on flow simulation by requiring that only a few realizations be run, and by keeping the grid a small subset of the full field models. A small grid size mini model means that flow simulation results can be acquired with minimal computing investment. There is no additional investment for preparation of the flow simulation parameters, relative permeability tables for example because the mini model runs rely on the same parameter set as the full field flow simulation. As for the calibration of the global objective function, the minimum requirement is flow simulation results for a single realization on the full field model.

The flow simulation results from the mini model give total cumulative oil production, total cumulative water production, and total water injection volumes for a number of realizations and injector-producer separation distances. Inferences obtained from the mini model observations are used as input data for the construction of the three component objective functions, and subsequently the global objective function.

The three component objective functions, $COBJ_{EOPT}^{\lambda}$, $COBJ_{EWPT}^{\lambda}$, and $COBJ_{EWIT}^{\lambda}$, are constructed with the aid of multiple regression. Observations from the reservoir model of uncertainty given by the mini model are used as independent data and the flow simulation results are used as dependent data. For the dependent data, cumulative oil production data is used for the EOPT component objective function, cumulative water production for the EWPT component objective function, and cumulative water production for the EWIT component objective function. The independent information is collected from the mini model and constrained by a recoverable volume template and by the collection of connected cells that qualify as high reservoir quality. The collection of connected reservoir quality blocks is called a geobody. The recoverable volume template and geobodies are discussed in greater detail in Section 5.2.1, and Section 5.2.1, respectively. The individual component objective functions, that is, $COBJ_{EOPT}$, $COBJ_{EWPT}$, $COBJ_{EWIT}$, are discussed in Sections 5.2.1, 5.2.1, and 5.2.1 respectively.

There is no way to apriori identify the functional dependencies of the component objective functions on the independent observations. The approach used to cope with the lack of information is to collect a number of different data and use statistical tools to identify independent variables that could be considered significant contributors to a model. The particular data collected is covered in subsequent Sections. The collected information might be considered a little redundant, but, as mentioned previously, there is no way to be sure what data will contribute significantly to the model, so many data are collected.

The mini model data collection process consists of: (1) performing data analysis to obtain the parameters required to construct the static property reservoir model (the histograms, and semivariograms for example), (2) construction of the mini model reservoir model of uncertainty, (3) preparation of the parameters for the dynamic model, (4) processing of the multiple realizations through flow simulation, and (5) collection of the static data required for multiple regression.

The number of required mini model flow simulation runs can become large if there are a large number of realizations and separation distances; however, the flow simulation requirement is mitigated by using a small grid, or a subset of the reservoir model of uncertainty.

The effect of structure is accounted for outside of the mini model and multiple regression exercise. This approach is motivated by the ability vary the penalty due to structure independently of the reservoir properties. If the effect of structure were rolled into the multiple regression scheme, there would be no way to modify the importance of structure in

the optimization independently of the effect of the reservoir. An alternative approach that does roll the effect of structure into the multiple regression is presented in the Future Work Section, Section 7.3.2. The effect of structure is accounted for with a map of the locally varying direction of strike and dip for each block on the grid. The maps are discussed in greater detail in Section 5.2.1.

The application of the global objective function for well location selection requires the assumption that the observations taken from the mini model on a single injector-producer pair translate to multiple injector-producer pairs in the full field model. The assumption is easily verified by observing the correspondence between the global objective function and the flow simulation results given the Asset Team selected well plan prior to calibration. Low correlations may imply that the assumption is not valid.

Mini Model Definition

The mini model serves as a small scale version of the reservoir model from which inferences can be drawn about the dynamic behavior of the reservoir model and applied to the selection of well locations.

The mini model contains only a fraction of the blocks of the full field model, and has blocks of the same dimension as those in the full field model. The mini model has no structural features, and may, or may not, contain conditioning data. The mini model may contain any number of realizations. The grid size must be large enough to avoid data contamination by boundary effects. The mini model grid parameters can be inferred from professional expertise, or by flow simulation results. The pressure draw down radius is a good guide for selecting a drainage radius and can be observed directly from the flow simulation results that are used to calibrate the global objective function. Similarly, the injector influence radius can be inferred from professional expertise, or flow simulation results from the flow simulation run used to calibrate the global objective function. The width of the mini model grid should correspond to the drainage radius of the producer, and the length of the mini model should correspond to, at minimum, about half the full field model plus the injector influence radius and the producer drainage radius. The mini model should have the same number of vertical blocks as the full field model. The flow simulation model should contain the same number of time steps as the full field model.

Structure is not included in the mini model so that the number of flow simulation runs is kept to a minimum. Several additional mini model realizations would have to be considered under flow simulation in order to characterize the effect of structure. Structure information can obscure the inference of the relationship between the producer, injector, and the reservoir. Instead of accounting for structure within the component objective functions, it is deferred to calibration from the full field results and maps that account for structure. The mini model is built using the same population techniques and parameters, such as the semivariogram and histogram, used to populate the full field model.

Conditioning data may or may not be used to build the model because of the effect heterogeneity has on flow simulation. The decision to use conditioning data is ultimately up to the practitioner; however a cautionary note is in order. The heterogeneity issue of concern has nothing to do with spatial uncertainty as the semivariogram is reproduced and hence the spatial variability is reproduced by many of the techniques used to populate the reservoir model and the same semivariogram is used to build the full field model and the mini model. The issue of heterogeneity identified here has to do with the magnitude of the ergodic fluctuations in reservoir properties near conditioning data over multiple realizations. Near conditioning data, there is little variability in the reservoir rock properties over multiple realizations, and variability increases with distance from the conditioning data according to the semivariogram. Increasing heterogeneity may dramatically decrease cumulative production. Due to the effect of heterogeneity, there may be good reason to avoid using conditioning data in the construction of the mini model.

If the goal of the well location selection is to select locations for injectors only, then it




Figure 5.2: The chargeable volume is the sum of the blocks falling in the drainage radius around the producers and the line of sight volume that the injector(s) can motivate to the producers subject to HCPV from the model of uncertainty.

might be good practice to construct the mini model using production wells as conditioning data. Similarly for selecting producer locations, it may be good practice to use one or more injector wells for conditioning.

The component objective functions are the end destination for the dynamic information derived from the mini model. The cumulative fluid productions are used as dependent variable information and key static model parameters are used as independent variable information in a multiple regression constructed component objective functions.

Recoverable Volume Template

The recoverable volume template is a region delineated by the line-of-site-area from the injector to the drainage radius of the producer. The diagram in Figure 5.2 illustrates three recoverable volume templates. The shaded triangular regions are the line of site regions between the injectors and producer pairs. The circles around the producers represent the drainage radii. The large circle represents the radius of influence of the injector. Producers that fall outside the radius of injector influence are considered to be unaffected by the injector.

Producer-producer and injector-injector interaction are accounted for. Blocks that fall into a region of overlapping recoverable volume regions are assigned to the nearest well. The greater the overlap the fewer number of blocks are available to the wells and the recoverable volume region.

Neither the producer drainage radius or the radius of injector influence should be selected arbitrarily. Both the drainage radius and the injector influence radius are, at minimum, functions of permeability, permeability heterogeneity, the fluid saturations, the relative permeabilities, and fluid properties. Both radii are best inferred by observing flow simulation results. The approach used here is to obtain flow simulation results using the seed well plan, and take as the drainage radius the average pressure drawdown radius of the producers. Inference of an injector influence radius is more involved. Given enough time a water flood water front can travel a significant distance; however, many reservoir projects are contractually time constrained. To maximize value in the project, injector and producer separation distances are sometimes selected such that the water front falls just short of the separation distance to maximize sweep within the specified time frame. The technique used to infer the injector influence radius is to construct a mini model of the reservoir that consists of a number of realizations and observe the advance of the water front and the oil production over a number of different injector and producer distances.




Figure 5.3: The sectional view of the reservoir marked *a* has three geobody numbers because non-reservoir regions prevent block face connectivity across the section. The sectional view marked *b* has 2 geobody number because a break in the non-reservoir blocks permit block face connectivity across the centre and right hand side of the section.

Geobodies

A geobody is a collection of face connected blocks that share one or more common features. The paper by Deutsch et al. [24] describes the procedure for constructing a geobody model. The two sectional maps in Figure 5.3 demonstrate the notion of a geobody. In map *a*, the blocks in the three shaded regions have porosity greater than a specified threshold and are identified as reservoir. The blocks in the hatched areas have values that do not exceed the threshold and are not identified as net reservoir. Each of the three shaded areas represent a different geobody, and are shaded differently because there are non-net reservoir barriers that prevent block face connectivity. In map *b*, the non-reservoir barrier is not continuous. In map *b*, a single block is face connected to the region that in map *a* is identified as a different geobody, thus there are 2 geobodies in map *b*. Since the recoverable volume template changes when the well locations change, the geobody changes too and is re-evaluated on every well location change.

The following information is collected from the collection of cells constrained to the recoverable volume template and the geobody:

- Surface area to volume ratio. Fluid production is a function of how tortuous the path the fluids must take, and by how much surface area is available. If there is abundant surface area then the poor quality reservoir rock can more easily pass fluid and there will be more conduits to feed the producer. If the surface area is low, a more direct path exists with lower capacity to leach fluids from low quality blocks.
- Coefficient of variation. The coefficient of variation is the standard deviation of the permeability divided by the mean permeability. The coefficient of variation quantifies the magnitude of variability within the high reservoir quality rock. Increasing values

of the coefficient of variation indicates increasing variability, and high permeability variability affects fluid production.

- Size. Size measures the number of blocks in the constrained region.
- Hydrocarbon pore volume (HCPV). The hydrocarbon pore volume is constrained to the recoverable volume template and the geobody as well.

$$HCPV = \sum_{\mathbf{u} \in \mathcal{A}} \phi(\mathbf{u}) \cdot (1 - s_w(\mathbf{u})) \cdot dx(\mathbf{u}) \cdot dy(\mathbf{u}) \cdot dz(\mathbf{u}) \quad (5.1)$$

where the region \mathcal{A} is the region constrained by the recoverable volume template and the geobody, ϕ is the porosity of the blocks, s_w is the water saturation of the blocks, and dx , dy , dz are the block dimensions.

- Average permeability. The average permeability is a coarse measure of how well the geobody might flow fluid.

Calculating HCPV is a little redundant with the size calculation since HCPV and size increase as distance between the injector and producer increases. Average permeability is redundant with the coefficient of variation on permeability since the coefficient of variation requires average permeability for calculation. The geobodies are recalculated on each perturbation.

Estimated Oil Production COBJ

The estimated oil production COBJ is a proxy for measuring the hydrocarbon production for a given well plan. $COBJ_{EOPT}$ is formulated as follows:

$$COBJ_{EOPT} = EOPT + EOPT * \lambda_{STRUCT} \quad (5.2)$$

where $EOPT$ is the regressed function that estimates oil produced and λ_{STRUCT} is a penalty calculated according to the orientation of the injector-producer pair and the proximity of the well to the oil water contact. The penalty function λ_{STRUCT} is discussed in detail in Section 5.2.1. The general idea behind λ_{STRUCT} is to add or remove fluid volumes as a consequence of structure. Under favorable structural conditions, say the injector is maximally downdip from the producer, oil production is enhanced by λ_{STRUCT} . Under less favorable conditions, say the producer is down dip from the injector, oil production is penalized.

The dependent information that used in the multiple regression is the cumulative oil production taken from the flow simulation results of the mini model. The final EOPT component objective function is a proxy for hydrocarbon fluid production for any injector producer pair configuration. Producer-producer interaction is accounted for by the recoverable volume template, the affect of structure and gravity segregation is accounted for by the structure penalty.

Estimated Water Production COBJ

The estimated field water production COBJ is constructed in the same way as $COBJ_{EOPT}$, except that the structure penalty is subtracted from the COBJ to reflect the notion that under favorable structural conditions, a producer will see reduced water production. The formulation for $COBJ_{EWPT}$ is shown below:

$$COBJ_{EWPT} = EWPT - EWPT * \lambda_{struct} \quad (5.3)$$

where EWIT is the regressed function that estimates water produced and λ_{struct} is the previously mentioned structural penalty. Like the $COBJ_{EOPT}$, EWIT is constructed using

multiple regression on the separation distance of the injector and producer under consideration, the surface area to volume ratio of the geobody constrained to the recoverable volume template, and the coefficient of variation on permeability. The weight, λ_{struct} , is the same weight that is used in EOPT component objective function discussed above and it accounts for structure. If the producer sits below the injector, the anticipated effect is an increase in the water cut.

Like EOPT, EWPT uses observations from the mini model as input data for the construction of a regressed function to estimate the volume of produced water. EWIT does not use the EOPT function. The independent variables are the same, but the response variable is the FWPT taken from the flow simulation outputs over multiple injector and producer distances.

Estimated Water Injected COBJ

The estimated water injected COBJ is not prepared in the same fashion as $COBJ_{EOPT}$ and $COBJ_{FWPT}$ nor does it use the structural penalty. Structure has little effect on the injector wells. Of greater importance, is the material balance: the producers must make room for the water by pressure and fluid depletion. If there is not enough room for the injected water, the reservoir rock may fracture, or the reservoir may store water. In some instances it is desirable to over inject; however, it is an expensive procedure.

The satisfaction of the material balance constraint is achieved by making sure that the EWIT never exceeds the volumes of water produced plus the voidage replacement factor applied in the flow simulation. $COBJ_{EWIT}$ is written:

$$COBJ_{EWIT} = EWPT \cdot VR \quad (5.4)$$

where VR is the voidage replacement ratio used in the flow simulation.

Structural Penalty (λ_{STRUCT})

A couple of well known guidelines for selecting injector and producer locations are that producers should be up dip from the water injectors, and producers should be completed well above the oil water contact. Otherwise, gravity and pressure gradients motivate water directly to the producer and the producer waters out prematurely. If the producer is completed below the oil water contact then it will water out prematurely.

To abide by the above guidelines one cannot merely confirm that the completion elevations for the producer are always higher than that of the injector and the oil water contact. One reason is that there may be a crest between the injector and producer.

Structure is accounted for using three observations. The first is the azimuth of the dip of the plane defined by the block with the highest elevation within a specified radius centered on the block of interest. The penalty for azimuth is calculated as:

$$\lambda_{azi} = \begin{cases} \frac{360 - |\alpha_I - \alpha_P|}{180}, & \text{if } \alpha_I > 180 \\ \frac{|\alpha_I - \alpha_P|}{180}, & \text{otherwise} \end{cases} \quad (5.5)$$

with α_I being the azimuth of the injector and α_P being the azimuth of the producer. The second observation used to account for structure is the difference in elevation of the blocks intersected by the producer and the injector at the basal surface. The depth penalty is calculated as:

$$\lambda_{dep} = \frac{depth_P - depth_I}{LKE - HKE} \quad (5.6)$$

where $depth_I$ and $depth_P$ being the depth of the base surface for the inject and producer respectively, and LKE and HKE are the lowest known elevation and the highest known elevation of the base.

The final part of the penalty weight λ_{SWAT} accounts for the proximity of the well to the oil water contact. To some extent, ensuring that the producer is not completed in the water leg is accounted for in the estimated oil production COBJ where hydrocarbon pore volume is calculated. That calculation does not account for the possibility that the producer may draw water up from the water zone. A simple approach is used to account for water coning. A weight equal to the inverse of the distance between the basal elevation that the producer is penetrated and the elevation of the oil water contact is calculated:

$$\lambda_{owc} = \frac{1}{Z_{owc} - Z_{pen}} \quad (5.7)$$

where Z_{owc} is the elevation of the oil water contact, and Z_{pen} is the elevation of intersected block.

The penalty weight λ_{SWAT} is the weighted linear combination of the λ_{struct} and λ_{owc} penalties:

$$\lambda_{struct} = \lambda_{tune} \cdot (\lambda_{azi} + \lambda_{depth} + \lambda_{OWC}) \quad (5.8)$$

where λ_{tune} is a tuning parameters used to increase or decreases the relative importance of the structural penalty. The overall effect of λ_{SWAT} is that under favorable conditions, that is, the producer is up dip and well aligned with the injector and the producer is at a structural high, the weight increases: $COBJ_{EOPT}$ increases, $COBJ_{EWPT}$ decreases, and $COBJ_{EOPT}$ decreases because the producer is not producing as much water and there is a decreased need for voidage replacement.

5.2.2 Calibration Procedure

The general idea for the calibration procedure is to modify the weights, λ_{EOPT} , λ_{EWPT} , λ_{EWIT} , in the global objective function presented in Section 5.2, so that the correlation coefficient between global objective function values and the reservoir performance values from flow simulation is maximized. The calibrated global objective function is considered a good proxy, but not a replacement, to flow simulation and should give comparable results to flow simulation without having to incur the cost of having to actually perform flow simulation.

The flow simulation results for the well plan are combined in the same fashion as the global objective function:

$$BOEV = C_{prod} \cdot FOPT - C_{treat} \cdot FWPT - C_{inject} \cdot FWIT \quad (5.9)$$

where C_{prod} is the cost of producing oil, C_{treat} is the cost of producing water, and C_{inject} is the cost of injecting water. The costs are in terms of barrel oil equivalent. FOPT is the cumulative field oil produced. FWPT is the cumulative field water produced, and FWIT is the cumulative field water injected.

The relationship between global objective function and BOEV is not always linear, but the correlation coefficient is a measure of linear correspondence between two variables. The exponent weights help increase the correlation coefficient.

The algorithm for selecting the weights starts with seed values of 1 iteratively perturbs them using simulated annealing. The weights are perturbed simultaneously in the calibration algorithm. The calibration algorithm follows the typical simulated annealing steps:

1. Using the seed weights, calculate the initial correlation coefficient;
2. Randomly select a weight;
3. Perturb the selected weight;
4. Calculate the new correlation coefficient;

5. use the simulated annealing rules to accept or reject the new weight;
6. If the new weight is accepted then update the weights used to calculate the global objective function. If it is rejected, revert to the old weights and repeat the previous steps from step 2 on until the stopping criteria are met.

A maximum number of perturbations is set as the only stopping criteria. A large number of perturbations are used. The calibration algorithm is quite fast though, so a few hundred thousand perturbations take only a few minutes on a moderate PC.

The final weights are non-unique. A different random seed will give different weights. The final correlation coefficients over multiple random seeds will be nearly the same though. Non-unique weights should not be perceived as a problem. The end result is a global objective function that is calibrated to dynamic results and represents a good proxy to flow simulation results. The well location selection algorithm is not sensitive to how the quality of the well plan is prescribed. It is however sensitive to the differences in well plans though.

The well plan used to calibrate the global objective function may have an effect on the final optimal well locations, but none have been so far observed. Care should be taken to select a seed well plan that would give globally representative results - some wells should be placed in poor quality reservoir regions and some in high quality reservoir regions. The aim here is to give some palpable information on how poorly or how well the reservoir might respond to production.

5.3 Implementation Details

This Section discusses issues surrounding the implementation of the simulated annealing algorithm and the global objective function.

5.3.1 Well Topology

Well topology refers to the different well types used in optimization and how they are declared. Three well topologies are considered: vertical, deviated, and segmented, see Figure 5.4. Each well type is defined by one or more segments. Vertical and deviated wells have one segment, and segmented wells have more than one segment. Each segment is defined by 6 coordinates. The start of a segment is defined by the coordinates $(x_s, y_s, z_s)_{i,j}$, where $i = 1, \dots, I$ is the number of segments, $j = 1, \dots, J$ is the number of wells in the well plan. The end segment is defined by the coordinates $(x_e, y_e, z_e)_{i,j}$.

For a vertical well, $i = 1$ and $(x_s, y_s)_{1,j} = (x_e, y_e)_{1,j}$. The values z_s and z_e may or may not be fixed. This permits the algorithm the flexibility of selecting the optimal depth for a vertical well. A deviated well has one segment and none of the end coordinates are tied. A segmented well is defined by multiple segments. For any two adjacent segments, the coordinates $(x_e, y_e, z_e)_{i,j}$ and $(x_s, y_s, z_s)_{i+1,j}$ are equal.

5.3.2 Grid Topology

The grid system is based on the Euclidean coordinate system with the origin at the bottom left corner of the grid. The convention for principle directions has the y -axis associated with the north-south direction, and the x -axis associated with the east-west direction. The grid indices increase from 1 to nx in the positive x direction, and increase from 1 to ny in the positive y direction.

The grid topology uses corner point and centre point grids. These grid systems are described in Figure 5.5. For each grid system heavy dots indicate essential coordinates. An areal view of the corner point grid system, is shown in the top left corner of Figure 5.5. It shows a regular grid of blocks with coordinate locations at each corner of a block in the grid. The top right corner of Figure 5.5 shows the corner point grid in a sectional x - z view. The




Figure 5.4: Three different well types are considered: vertical, deviated, and segmented.

increments in the x and y directions are regular but the z coordinates may not be. Allowing the z coordinate to be irregular affords great flexibility and permits the specification of a grid that can closely follow geologic structure.

The bottom half of Figure 5.5 shows the centre point grid system. The centre point grid uses the same origin and increment architecture as the corner point grid in the x - y plane, but uses a regular z coordinate. Since the increments are regular in the x , y , and z directions only the x , y , and z coordinates for the centre of each regular block is required.

The corner point and centre point grids are used simultaneously to characterize the reservoir. The corner point grid specifies the reservoir structure and the centre point grid is used to populate the corner point grid with petrophysical properties.

5.3.3 Initial Well Plan

The initial well plan consists of any number of wells of the same topology. There may be any number of existing wells with any combination of well types. The initial well plan must have wells with segment coordinates located inside the grid. The initial well plan should not affect the final solution, and does affect the number of perturbations required to find a solution. If the initial well plan is, in terms of the solution space, far from the optimal solution more perturbations may be required to find the optimal solution.

5.3.4 Identification of Intersected Grid Blocks

The global objective function requires location specific information from the reservoir model for evaluation. Identification of the blocks intersected by the well trajectory is accomplished by discretizing the well path into a series of points as shown in Figure 5.6. In the Figure, the start of the well path is denoted by a large dot and the end is denoted with a small dot.

The x - and y -index for each well path point is easily solved using the x - and y -index that corresponds to the x and y coordinates. In the centre point grid topology, the z coordinate is also easy to find using the same technique. Finding the z -index in the corner point topology

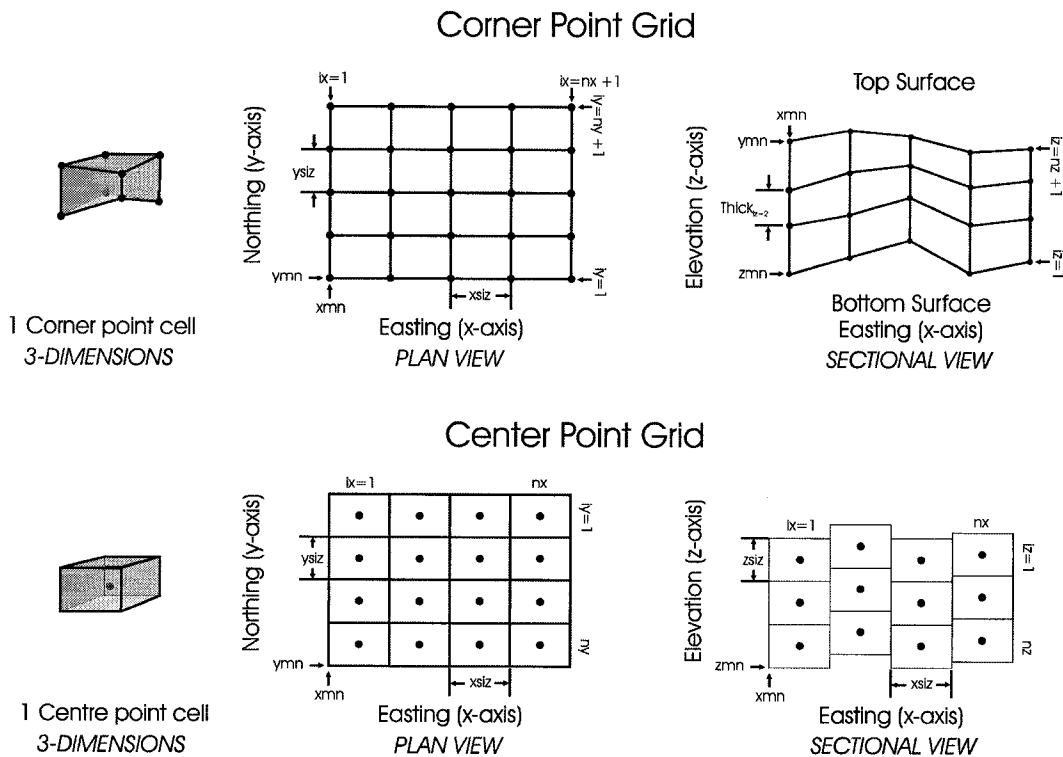


Figure 5.5: The corner point grid system is used to model reservoir structure. The centre point grid is used to populate the corner point grid with reservoir properties.

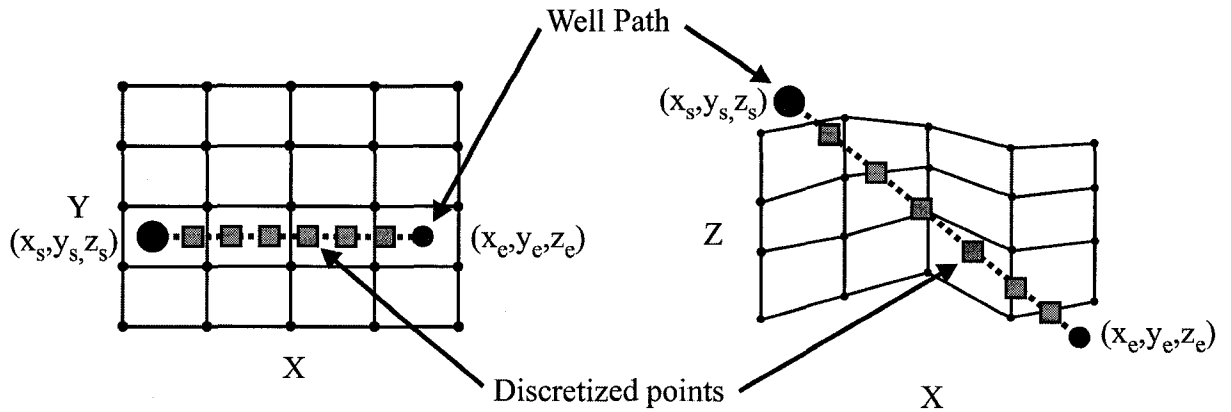


Figure 5.6: The well path is discretized into a number of points. The large dots indicate the start of the well path. The small dots indicate the end of the well path.

is more involved since the grid may not be regular in the z direction thus one cannot simply solve for the corresponding index.

The approach used to solve for the z -index consists of the following steps which are depicted in Figure 5.7. At the top of the Figure is an example well path and corresponding discretization points shown as squares.

1. Find the x and y index for the current discretization point. Find the block face that the discretizing point is closest to. The term block face refers to the $z-x$ or $z-y$ sectional projection of a block. Starting from the north most face and rotating clockwise, the faces are labeled 1, 2, 3, and 4. Starting from the top of the column of blocks corresponding to the x and y index of the discretization point, find the average elevation of the top of the face closest to the discretization point. This point is shown as a dot at the top of the column in Figure 5.7. A test is performed to see if the discretization point is below the average elevation for the top of the face.
2. If the discretization point is not above the average elevation for the face then increment a counter and find the average elevation for the top of the face one index below the current index. Test if the discretization point is below the average elevation for this face.
3. Repeat the step above until the discretization point is above the average elevation for the face. The value of the counter is saved as the z index if the test is satisfied.

The intersected blocks must be identified after each perturbation. The above steps are repeated after each perturbation because the well locations change. The global objective function needs this information to test the goodness of the perturbed well plan.

The algorithm keeps track of well intersected blocks for use by the global objective function with an array with the same dimensions as the reservoir model. The wells in the initial well plan are assigned identification integers when they are read into the algorithm. As the x, y, z indices for the discretization points are found the corresponding block in the intersection grid is assigned the well identification integer.

5.3.5 Perturbation Mechanism

The perturbation mechanism randomly selects one of the wells, then a segment (if the well is segmented), a segment end (if the well is not a vertical well), and finally a Δx , Δy , Δz to





Figure 5.7: The process for determining the z -index for the discretized point is a small loop that compares the average elevation for the top of each block to the elevation of the discretized point over the iz elevations. If the average top elevation is less than the elevation of the discretized point then z index is the index of the previous average elevation.

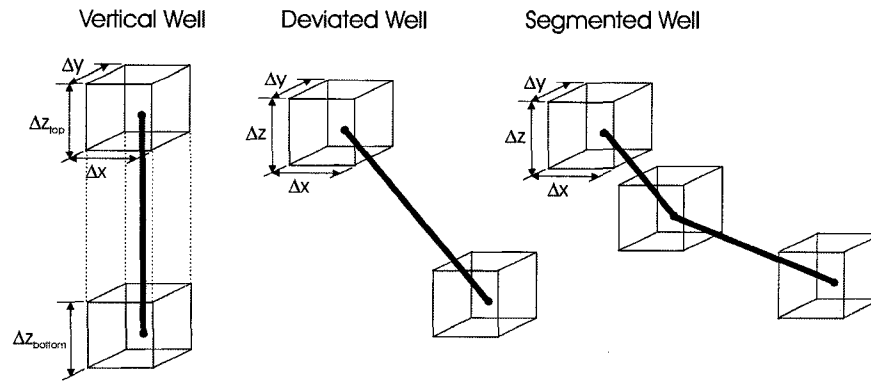


Figure 5.8: This Figure illustrates the perturbation mechanism on the three different well topologies.

move the candidate segment end to. Note that if the well to be perturbed is a segmented well, the selected segment end of the candidate segment and its adjoining segment are perturbed.

The perturbation mechanism is slightly different for each well type. Figure 5.8 is an illustration of the perturbation mechanism for each well type. The boxed regions indicate the range of possible new locations that can occur due to a perturbation. The dashed lines on the vertical well diagram indicate that both segment ends are moved simultaneously for a vertical well.

5.3.6 Well Selection Algorithm

This Section describes the algorithmic steps in the optimization. The steps are illustrated in the flowchart in Figure 5.9 and discussed below:

Read the Reservoir Model of Uncertainty and Initial Well Plan :

The algorithm starts with a reservoir model of uncertainty and an initial well plan. The algorithm requires a reservoir model of uncertainty that includes multiple realizations of the reservoir architecture and petrophysical parameters. Realizations of the top horizon and the thicknesses of each subsequent zone make up the reservoir architecture. Realizations of porosity, water saturation, permeability, and geo-object connectivity are required.

The initial well plan consists of the number, locations, and trajectories of the wells. The well selection algorithm can handle well plans with vertical, or deviated, or segmented wells.

Prepare the intersection grid :

The algorithm reads the well plan and determines which blocks are intersected by the well plan. The intersection information is stored in an array that has same dimensions as the reservoir model. Intersected blocks are encoded with the well identification integer that corresponds with the intersecting well.

The intersection information is used by the global objective function to evaluate how good the well plan is. The global objective function keeps track of the performance of each well using the well identification integer.

The intersection grid is rebuilt after each perturbation because the well plan changes after a perturbation.

Compute Initial Objective Function value: The global objective function quantifies the goodness of the well plan. In keeping with accepted optimization paradigm of minimizing the global objective function, the revenue component objective function is multiplied by -1 so that the component objective function decreases as revenue increases.

The value of the global objective function is calculated using the well plan on each of the L realizations of the reservoir model of uncertainty. The value of the global objective is recorded as the sum of L global objective function values. The initial value of the global objective function is saved as O_{old} . The initial global objective function value is used to calculate the probability of accepting a perturbed well plan.

Perform a Perturbation :

A well is randomly selected, then a segment end is selected, followed by new coordinates for the candidate well and segment.

Prepare perturbed well plan intersection grid :

The well plan has been modified by the perturbation mechanism so the intersection grid must be updated to reflect the changes.

Compute the new global objective function value :

The value of the global objective function is calculated using the perturbed well plan $Plan_{new}$ and the model of uncertainty. The new value for the global objective function is saved as O_{new} and is used to calculate the probability of accepting $Plan_{new}$ as the current well plan.

Accept or reject perturbed well plan? :

Calculate $\Delta O_g = O_{new} - O_{old}$. Use ΔO_g and the current temperature, T , from the annealing schedule, to calculate the probability of accepting the newly perturbed well plan $Plan_{new}$ according to:

$$P\{accept\} = \begin{cases} 1, & \text{if } \Delta O_g \leq 0 \\ e^{(-\frac{\Delta O_g}{T})}, & \text{otherwise} \end{cases} \quad (5.10)$$

A ΔO less than 0 indicates that the perturbation reduces the objective function and should be accepted: set P_{accept} to 1. If ΔO is greater than 0 draw a random number $\in [0, 1]$, and test it against P_{accept} . If the random number is less than or equal to P_{accept} then the perturbed well plan accepted. If the random number is greater than P_{accept} then the perturbed well plan is not accepted.

Update the initial global objective function and well plan: If the perturbation is accepted, save O_{new} as O_{old} and save $Plan_{new}$ as $Plan_{old}$. If the perturbation is rejected then restore the previous well plan and the previous value for the global objective function and go back for another change. If $Plan_{new}$ is rejected it will be replaced by the next accepted perturbed well plan. The algorithm only requires the results from the previous and perturbed well plan. The results from perturbations previous to the most recent perturbation are lost.

Stopping Criteria Satisfied? : If the stopping criteria specified in the annealing schedule, k_{max} or num , have been satisfied then stop the algorithm. If the change in the global objective function has been less than ϵ percent over $k_{attempt}$ then stop the algorithm. If the stopping criteria have not been satisfied then increment the counter for k_{max} and if the perturbation was not accepted and $k_{attempt}$ is satisfied then increment num .

Write Outputs If one of the stopping criteria are satisfied then write out the coordinates for the well plan and a grid specifying the cells intersected by the well plan.

Check Annealing Schedule : If k_{accept} or $k_{attempt}$ has been satisfied then apply the reduction factor λ to the temperature T and reset the counters for k_{accept} and $k_{attempt}$. If k_{accept} has not been satisfied and the perturbation was accepted then increment the k_{accept} counter. If $k_{attempt}$ has not been satisfied and the perturbation was rejected then increment the $k_{attempt}$ counter.

Loop Until the a Stopping Criteria is Satisfied :

The algorithm perturbs the well plan until one of the stopping criteria in the annealing schedule are satisfied.

5.3.7 Comments

In terms of CPU requirements, there is a potentially large memory requirement for the algorithm. All of the realizations for the attributes must be read and held in memory for evaluation of the global objective function after each perturbation. A reservoir model may consist of millions of blocks for each attribute.

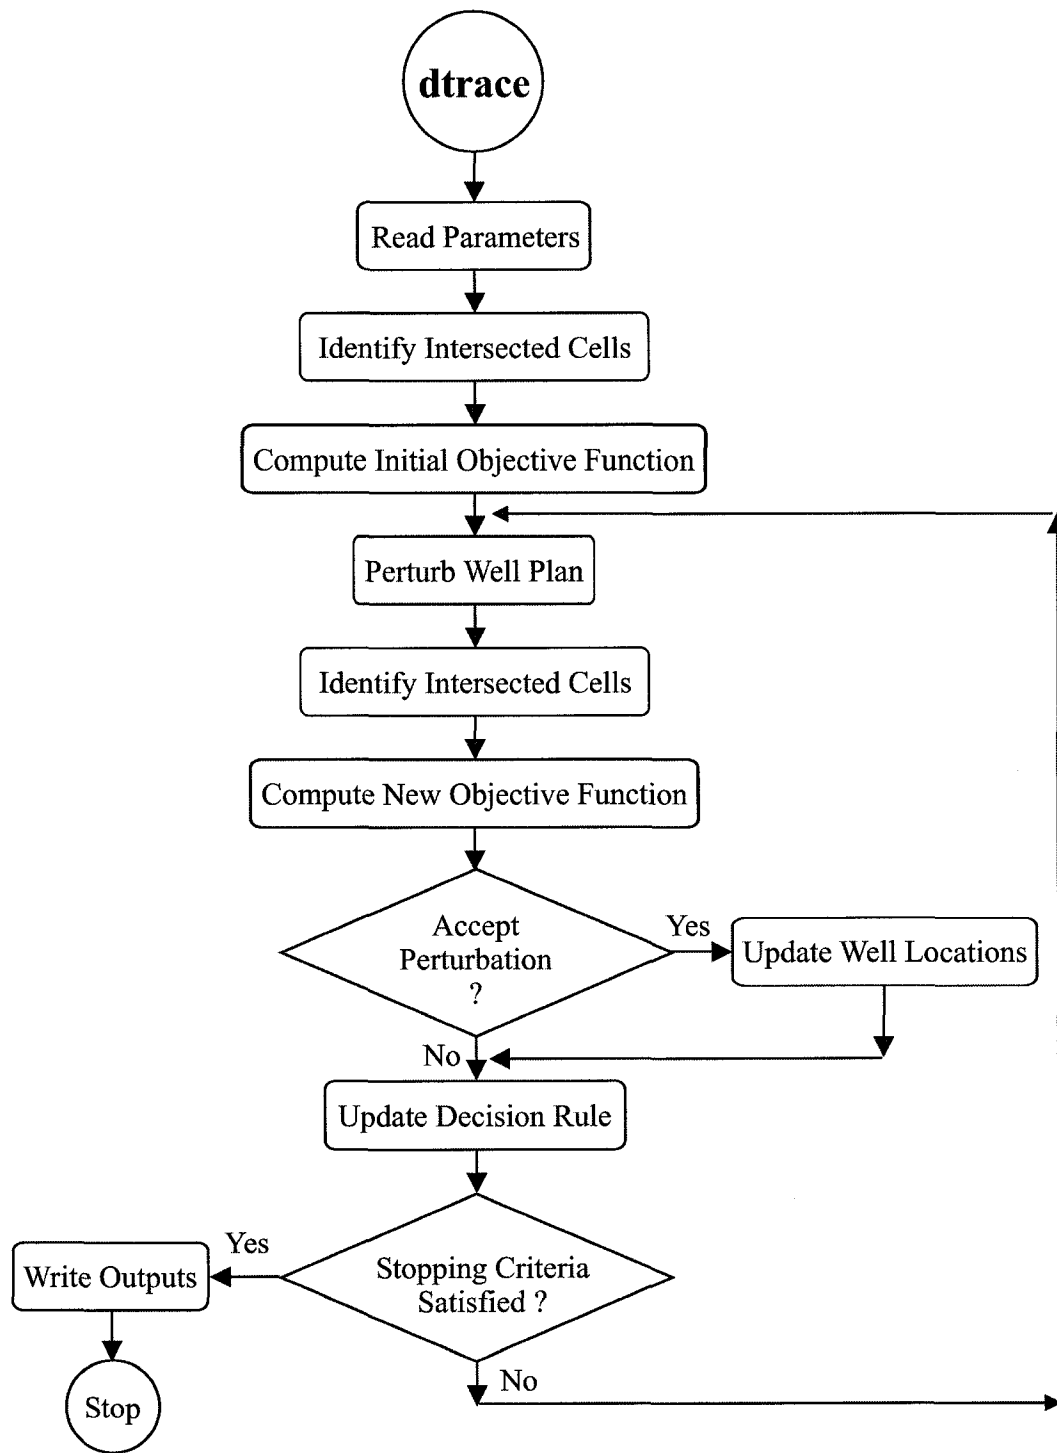


Figure 5.9: The flow chart for the well selection algorithm.

Chapter 6

Well Location Selection Case Studies

The objective of this Chapter is to present some well location selection case studies. The first Section reviews some important considerations for selecting a well plan. The last section presents a case study that selects a well locations for a waterflood on a heavy oil field.

6.1 Considerations for Selecting Well Locations

A critical step in geostatistical modeling is the selection of a modeling approach and required parameters such that the output model is consistent with all available information. The selection of a modeling approach includes decisions such as the facies population technique and the suitability of Gaussian techniques for continuous variables. Important parameters include the histogram of each variable, the spatial correlation of each variable, and the relationships between variables.

The selection of the modeling approach and the required parameters affects the decision making processes. If the model is in error, subsequent decisions may also be in error. Uncertainty is an important root cause for modeling errors. The best way to proceed with a modeling exercise is to use all of the data at hand to build the model and possibly revise the model as new information becomes available. The additional data can be used to incrementally update the model and aid in future decisions.

In this Section, two experiments are performed that examine the incremental modeling approach. The first experiment considers the notion of selecting all the well locations at the same time or selecting them sequentially. The second experiment is an extension of the first: is there an advantage to selecting well locations sequentially by updating the reservoir model as drilling information becomes available.

Both experiments rely on simplistic synthetic reservoir models. The experiments use a simplified version of the global objective function discussed in Chapter 5. The simplified global objective function is discussed within the experiment. The simplicity of the experiments and the models does not dilute the importance of the results though.

6.1.1 Sequential Well Planning

This case study shows that if an optimal well plan is appended with another well, selected optimally or otherwise, the appended well plan is a suboptimal solution compared to a well plan where all the well locations are selected simultaneously. For example, consider a well plan consisting of two wells. Appending the two well plan with a single well is suboptimal compared to a 3 well plan where all the wells were optimally selected simultaneously. The following experiment uses a synthetic reservoir to demonstrate this idea. The experiment




Figure 6.1: An unconditional realization is shown on the left and the corresponding geobody realization is on the right.

uses a simplified version of the global objective function than that described in Chapter 5. The global objective function aims to accumulate much hydrocarbon pore volume as possible within the drainage radius of the wells.

The Synthetic Model

The grid used in the experiment contains 50x50x10 blocks that measure 100mx100mx1m. The top and base of the reservoir is flat. Only porosity, water saturation, and geobody connectivity are considered. Porosity and water saturation are held constant at 25% and 10% respectively. The geobodies are derived from a single unconditional sequential Gaussian simulation realization of standard normal score (mean=0, variance=1) values. The realization was processed using the geobody connectivity algorithm described in Section 5.2.1. A threshold of 0 was used to give a grid where 50% of the values are reservoir quality. The first layer of the sequential Gaussian realization and resulting geobody realization are shown in Figure 6.1. The view is areal with the x used for the east and west directions and y axis is used for the north and south directions.

The first layer of the model described above is used as a single realization of a 2D reservoir model. The realization will be used to select horizontal well locations. Later, the model will be used in its entirety as a single realization of a 3D reservoir to select up to 3 deviated well paths.

Well Location Selection

The experiment uses a simplified global objective function that aims to maximize the accumulation of connected HCPV blocks. The global objective function is written below:

$$GOBJ = \sum_{l=1}^L HCPV$$

where HCPV is calculated using:

$$HCPV = \sum_{l=1}^L \sum_{\forall \mathbf{u} \in \mathcal{A}} \phi^l(\mathbf{u}) \cdot (1 - s_w^l(\mathbf{u})) \cdot dx(\mathbf{u}) \cdot dy(\mathbf{u}) \cdot dz^l(\mathbf{u}) \cdot i_r(\mathbf{u}) \cdot i_a^l(\mathbf{u}) \cdot i_g^l \quad (6.1)$$

with $s_w^l(\mathbf{u})$ and $\phi^l(\mathbf{u})$ being the water saturation and porosity respectively for the block, and $dx(\mathbf{u})$, $dy(\mathbf{u})$, $dz^l(\mathbf{u})$ are the dimensions of the block. Some well location studies consider

thickness as an uncertainty parameter. The superscript l indicates the dependence of the data on the realization $l, l = 1, \dots, L$. The categorical variable $i_r(\mathbf{u})$ is used to apply the drainage radius constraint. If the block falls within the drainage radius $i_r(\mathbf{u}) = 1$, and $i_r(\mathbf{u}) = 0$ otherwise. Well-to-well interaction is accounted for by the $i_a(\mathbf{u})$ variable. If a block has already been accounted for by another well, $i_a(\mathbf{u}) = 0$ and $i_a(\mathbf{u}) = 1$ otherwise. If $i_a(\mathbf{u}) = 1$, the well closest to the block claims the volume. Static connectivity is accounted with the use of the concept of a geobody and the categorical variable $i_g^l(\mathbf{u})$. In order for a block to be accumulated, it must have the same geobody number as at least one block intersected by the well path. The global objective function is not calibrated to flow simulation results, the study is supported using only static information.

Results

Figure 6.2 shows the progression of the well location selection for a single horizontal well on the first layer of the hypothetical reservoir model. The initial well plan is shown in map *a*. The well path and drainage radius is shown as a dark line with a shaded area surrounding it. Maps *b* through *h* show solutions after 1250, 2500, 3750, and so on, of accepted perturbations. The map on the bottom of the figure is the final optimal solution after 10000 perturbations.

By observation, the initial well plan is too short and in a suboptimal location: very little of the large geobody is accumulated in the drainage radius of the well. Subsequent perturbations improve the alignment of the well and accumulate more HCPV. Note that most of the improvements/changes occur in maps *b, c*, and *d* and that fine scale improvements occur in maps *e* through *h*.

The simulated annealing algorithm iteratively improves the solution as the optimization proceeds. Allowing a sufficient number of perturbations is essential for a globally optimal solution. Map *i* shows the final optimal solution. Observe that no other solution accumulates more geobody.

Figure 6.3 shows a map of the final solution for a well plan with 3 horizontal wells. Note that the three well solution is not merely the one well solution appended with two wells. The importance of this observation is that the optimal multi-well plan must be solved for by considering all proposed wells simultaneously; sequential well planning leads to suboptimal solutions. This contradicts the idea of selecting well plans sequentially with the goal of incrementally adding information to the model. Incremental information offers the opportunity to make improved subsequent decisions because decisions are made with incrementally more data. Section 6.2 shows that even when incremental information is added from new wells the increase in information content in the reservoir model may not lead to improved decision making.

Observations from the selection of horizontal well plan solutions can be applied to the selection of vertical wells. Figure 6.4 shows optimal well plans for up to 4 vertical wells. The maps are the same maps with the same perspective as the horizontal well location selections shown above. The well and associated drainage radius are depicted by a shaded disc with a black dot in the middle.

Note that the single vertical well problem shown in the top left corner of Figure 6.4 has more than one optimal solution. Figure 6.5 shows two optimal solutions side by side. Both maximize the global objective function. For any well location problem there may be more than one well plan that optimizes the global objective function.

6.2 Sequential Well Location Selection With Incremental Information

The idea for this experiment is to place the wells sequentially using information from previous wells to better select the next well location. Proceeding in this fashion should result in well

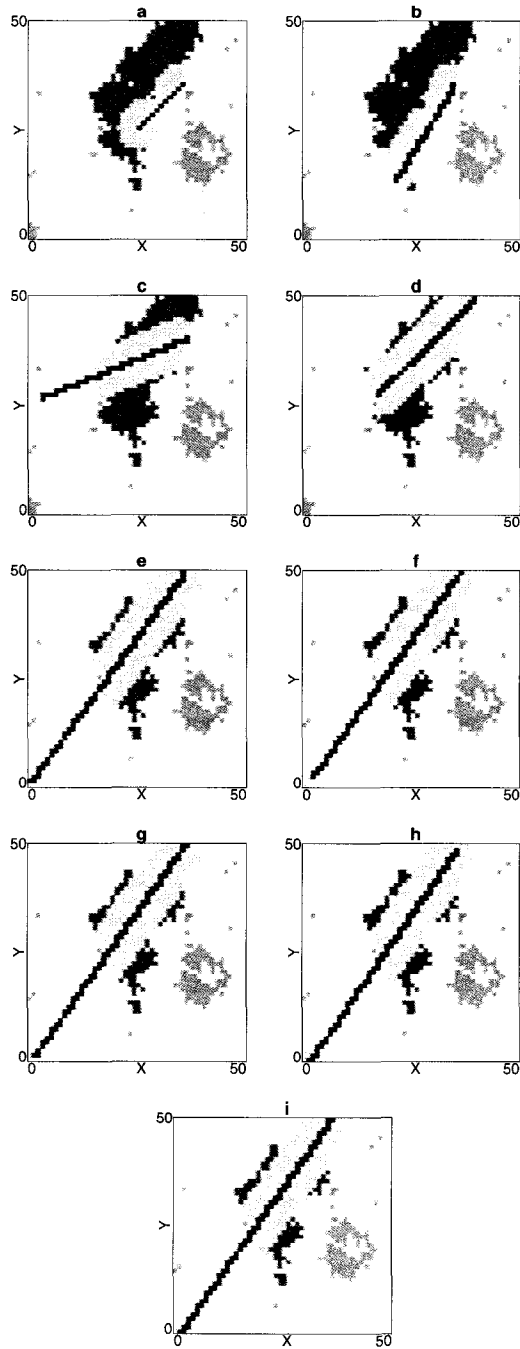


Figure 6.2: The maps shown depict the progress of the algorithm as it iteratively perturbs the well plan until it is optimal. The maps show an areal perspective and select an optimal horizontal well path. The top left map shows the seed well plan. Each subsequent map shows the well plan after a number of perturbations. The map at the bottom shows the final well plan.




Figure 6.3: The map shows the three horizontal well solution. The optimal well location for the one well solution is not aligned with any of the wells in the three well solution.




Figure 6.4: These maps show an areal perspective and show optimal solutions for, starting from the top left and moving to the bottom right, 1, 2, 3, 4 vertical wells




Figure 6.5: There is more than one optimal one well solution. The map on the left shows one, whole the map on the right shows another.

locations that improve recoverable reserves over well locations selected without the aid of the incremental information. Evaluating the value of added information requires a model for what that new data might be. A set of possible “truth” models is simulated and the incremental information is extracted from the set of truth models.

The experiment starts with an example that shows that simply adding incremental information to the model with no updating of the conceptual model or other parameters does not yield better decisions. Then, another example illustrates that a change in the conceptual model improves decision making.

6.2.1 Methodology 1

Consider the problem of selecting N wells in a petroleum reservoir development. Choosing N positions all at once with exploration data is not expected to be as good as using the well data incrementally. The idea for this experiment is to start with a model that represents our current state of uncertainty and proceed by adding information to update this uncertainty. At each step a decision is made and the value of the added information is assessed. Of course, the very best decisions are made with the inaccessible true reservoir description. The three cases, no incremental information, incremental information, and ultimate reservoir information, have been color-coded and shown histograms in Figure 6.6:

1. The *red* case (histogram on the left) is the decision using only exploration data for all wells.
2. The *blue* case (histogram in the middle) shows the distribution of oil in place where each well location is chosen using the exploration data and all previous development wells. This is the best approach in practice and should give better results than the red case.
3. The *green* histogram (on the right) shows the distribution of reserves if the true reservoir description were available. This will be better than the red or blue case.

The true reserves are unique for any particular reservoir and would appear as a “spike” on a histogram; however, the true reserves are not accessible in practice. A model of uncertainty is substituted as the “true” reservoir and each realization is represented as a plausible truth. The distribution of reserves is found by measuring the reserves obtained by considering the reservoir model of uncertainty. There is no other recourse, as a model of




Figure 6.6: This schematic illustration shows the histograms of uncertainty in recoverable reserves accounting for exploration data only (red on left side), updating the model and well locations with incremental information as it becomes available (blue in the middle). The distribution of reserves given “perfect” information is the green distribution on the right.

uncertainty is all that is available. If the model is representative, then the representation is valid.

The simulation approach to predicting the value of incremental information will be referred to as geostatistics while drilling or GWD. The GWD program is shown schematically in Figure 6.7, and explained in the following text:

1. The starting point is a set of truth reservoirs, R_l , $l = 1, \dots, L$. These truth reservoirs should be taken as the set of geostatistical realizations that represent the current state of uncertainty, that is, they reproduce all available exploration data. Note that each reservoir model R_l is a full specification of the reservoir structure and all internal heterogeneity.
2. Apply the best available algorithm to determine the optimal number of wells, N , and their locations over the set of true reservoir models $\{u_i, i = 1, \dots, N\}$. Note that there is only one optimal well plan. It is unreasonable to determine realization specific well plans because there is no practical way of reconciling the L different well plans afterwards.
3. Calculate the reserves on each of the L realizations and construct a distribution of reserves given the initial state of uncertainty. This represents the *red* histogram in Figure 6.6. Set $M = N$, where M will be the remaining wells to drill and N is the final number considering the well plan.
4. The following steps are to be considered for each of the $l = 1, \dots, L$ true reservoir models.
 - a. Choose one of the M wells to drill based on logistical drilling considerations or simply the one with the largest expected reserves. Extract the reservoir properties and structure at that location from the current (l) realization to be considered in building a new set of realizations with the added information. Reset M to $M - 1$ once one of the wells is chosen.
 - b. Construct another suite of geostatistical realizations conditional to all available information including the original data and the data from all previously drilled ($N - M$) simulated wells. Clearly, this procedure must be automated in a script or in some commercial software such as *Jacta*TM.
 - c. Run the optimal well placement to refine the position of the remaining M wells using the latest suite of geostatistical realizations. These new locations will be considered to select the next one. Loop back to *b* if $M > 1$; otherwise, proceed to *e* and then keep looping over L .

- d. Calculate the reserves using the updated well locations and the current simulated true reservoir R_t . This is one number to go into the blue histogram of Figure 6.6. The set of updated reserves make up the blue histogram.
5. Find the unique optimal well locations for the N well on the reservoir model of uncertainty. Use these locations to calculate the production for each realization. This set of L numbers make up the green histogram.

This simulation exercise provides a measure of how incremental information improves the ultimate reservoir decision-making or profitability. No new information is being considered in this exercise; in this sense it is similar to the bootstrap technique.

First Example

The setting of the first example is a model of uncertainty for a synthetic reservoir on a regular 50×50 grid. It is not likely that the synthetic reservoir would be observed in nature; however it serves the purpose of illustrating that adding incremental information does not always lead to better decision making. The model consists of porosity and porosity derived geobodies for incorporating a measure reservoir connectivity. Porosity is conditionally simulated using sequential Gaussian simulation (SGS) and the four conditioning data shown in Figure 6.8. Porosity is first simulated as a normally distributed variable and later transformed to porosity using the simple linear transform to avoid issues surrounding the Gaussian transform: $z = 5 \cdot y + 10$, where y is the simulated Gaussian value and z is the porosity value. A histogram of porosity over 100 realizations is shown in Figure 6.9. The geobodies represent a set of connected blocks having porosity greater than 7.5md. Figure 6.10 shows a simulated porosity map and the corresponding geobody map. The optimal well locations were selected using a global objective function that accumulates HCPV with the geobody constraint; the performance of the well locations is evaluated by calculating HCPV only from each block falling within the drainage radius and belonging to the same geobody intersected by the well path.

Figure 6.11 shows a geobody map with four well locations superimposed on the map. The well path is indicated by a dark pixel. The drainage radius is indicated by a shaded circular region around the well path. The well located in the bottom right of the map would have the highest quality because, as indicated in the corresponding porosity realization shown in Figure 6.10, it is in a location of high porosity and the drainage radius is almost completely occupied by a geobody. The well location of the lowest quality is shown in the top right.

100 realizations of the synthetic reservoir were submitted with the goal of sequentially placing four wells. Figure 6.12 shows the results for the steps 3 through 5 of the GWD algorithm for a single realization. The map at the top of the Figure corresponds to step 3: a realization of a geobody map with four optimally placed wells superimposed. The set of maps below the step 3 map are each one realization of the updated model of uncertainty with the updated optimal well locations corresponding to step 4c. Starting from the right, the maps show the first through the fourth fixed wells. A heavy circle is drawn around the fixed wells. The map on the far left is the map for step 4d: the updated well locations superimposed on the geobody realization in step 3. Below the set of step 4 maps is the map for step 5: the four well locations optimized to the realization.

Note that the realization used to represent the updated model does lead to better well locations. By observation the drainage radius of the updated well locations contain more geobody and are in locations of higher porosity than the globally optimal locations. This is a local observation, the final results are presented below.

The Red, Blue and Green Histograms for 100 realizations are shown in Figure 6.13. The mean of the Red Histogram is 7476.23, the Blue Histogram is 7302.46, and the Green Histogram is 10596.17. Over 100 realizations updating does not lead to better well locations. The results were verified and extensively tested.

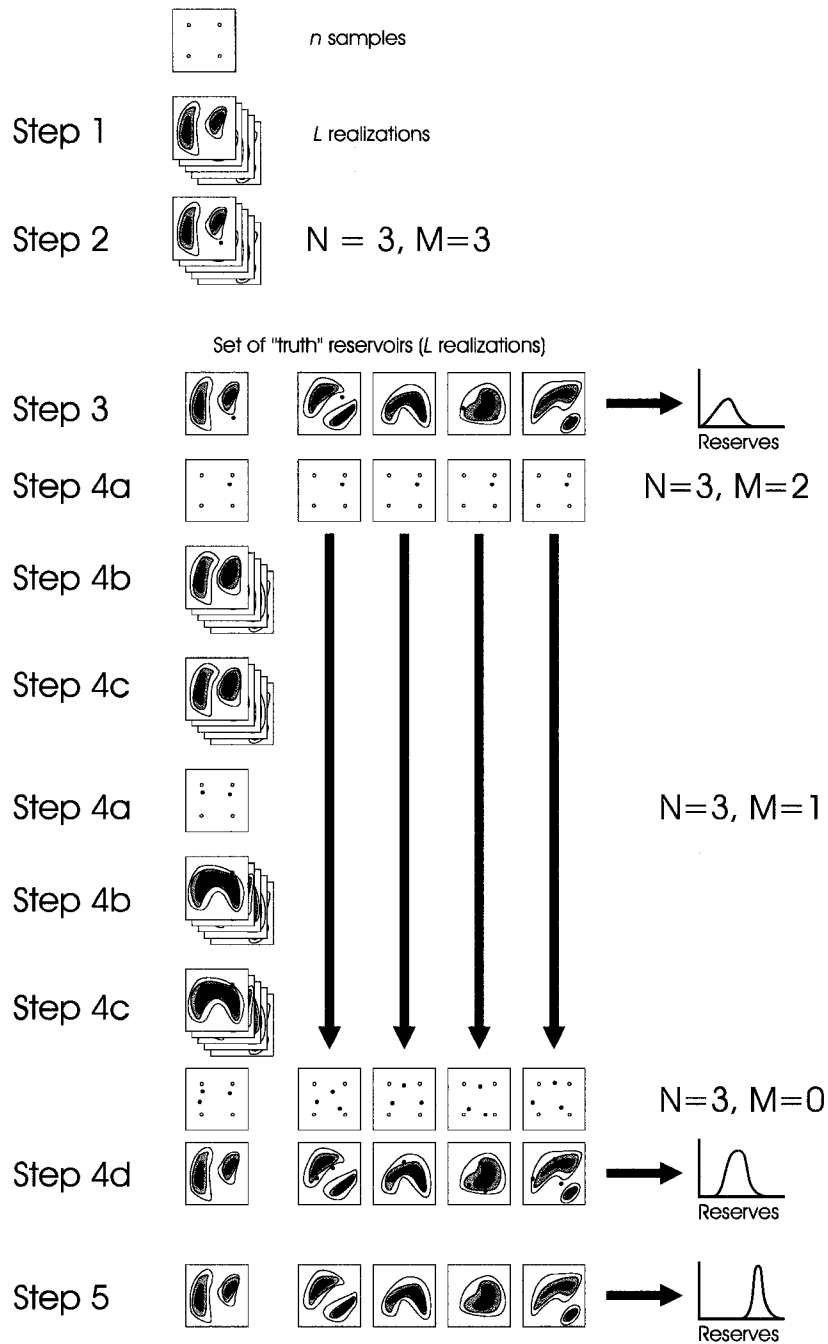


Figure 6.7: This schematic illustrates the procedure to assess the value of incremental data.




Figure 6.8: The map shows the positions of the four conditioning data and the conditioning values.




Figure 6.9: This histogram shows the distribution of transformed porosity for 100 realizations.




Figure 6.10: Shown is a single porosity realization and the corresponding geobody realization.




Figure 6.11: The map shows the geobodies and the four well locations. The well locations are denoted as a single dark pixel surrounded by a shaded circular drainage radius. Referring to the porosity map in Figure 6.10, the well location on the bottom right corner has the highest quality, the well on the top right has the lowest.




Figure 6.12: Together these maps represent one leg from Figure 6.7. The maps on the left are the geobody maps and the placed wells for one realization for step 3, step 4d, and step 5. The maps on right from the map showing step 4c show 1 to 4 fixed wells as indicated by the heavy circles.




Figure 6.13: The red histogram is on the left, the middle one is the blue histogram, and the right is the green histogram. Note that over 100 realizations updating does not lead to better decisions.




Figure 6.14: Here is a schematic illustration of the histograms of uncertainty in window quality accounting for exploration data (red on left side), updating the model and window locations with incremental information (blue second from the left), updating the model and window locations using the conceptual information and the incremental data (orange second from the right). The distribution of reserves given “perfect” information is the green distribution on the right.

Comments

The results of this show an apparent violation of common sense: incremental may not lead to improved decision making. The primary reason that incremental information may not help much is that the model is only refined within the range of correlation which is coincident with the drainage radius of the well and the range of the semivariogram. The updated portion of the model is never available for decision making because it is already occupied by a well. Another factor contributing to the failure of incremental information to lead to better decisions is sample bias. The well locations are preferentially located, by construction of the objective, in high valued locations. The model is successively conditioned by good data leading to bias in the updated models. The collected information used to construct the reservoir model must be debiased to account for biased sampling to avoid the apparition that the reservoir is increasing in quality.

6.2.2 Methodology 2

The results of the first experiment motivate a second more simplified experiment that examines the impact of changing the conceptual model and using incremental data versus using only the incremental data. The goal of this experiment is the sequential selection of two optimal locations for “windows” that maximize reservoir quality Q . The quality of a window location is quantified by taking the sum of all the values falling within the window. This is similar to optimally placing two wells of maximum quality except that in this experiment a square window is used instead of a drainage radius.

The color coding scheme outlined previously is retained but one more histogram is added: the histogram that represents the use of a revised conceptual model and incremental information for decision making. This distribution should have a mean greater than that of using only incremental information. Thus the distribution should fall between the blue distribution and the green distribution. We call this the orange histogram and it is shown in Figure 6.14.

The methodology for this experiment is a variation of the GWD algorithm that is called geostatistics with incremental data (GWID). The GWID algorithm starts by constructing a “true” conceptual model consisting of a map of secondary data (Y_{truth}), a map of primary data (Z_{truth}) and a randomly selected correlation coefficient (ρ^l) correlating the primary and secondary data. The Y_{truth} is intended to mimic the type of information that would be provided by seismic data. The pseudo-seismic information and the correlation coefficient are

used to construct a map of a primary reservoir parameter. The GWID algorithm proceeds as follows:

1. Draw a random correlation coefficient ρ^l either $+/-$, where l is the l^{th} trial for $l = 1, \dots, L$. Altering the correlation between the primary and the secondary information over a large number trials affords the opportunity to observe the effect of using an inappropriate conceptual model.
2. Simulate Y unconditionally and call this the true secondary information Y_{truth} .
3. Simulate a Z using y as collocated with ρ^l . Call this Z_{truth} .
4. From the corners of the Z_{truth} map collect four conditioning data. This data is considered exploration data. At this point the conceptual model consists of the exploration data and the same semivariogram as used in the construction of Z_{truth} .
5. Simulate 20 realizations of the Z variable using the four conditional data. At this stage we proceed with the construction a geostatistical model with all available data.
 - Find the two best locations, $i(ix, iy)$ and $j(jx, jy)$, for the windows that maximize Q^* , where Q^* is the window quality at locations i, j given one piece of information, the exploration data. The two best locations are found by searching the grid exhaustively. Overlapping windows are accounted for by assigning the blocks in the overlapping area to only the closest window.
 - Using the two optimal locations $i(ix, iy)$ and $j(jx, jy)$, calculate Q_{red}^l from the truth model. The distribution of the Q_{red}^l values is the red histogram.
6. Randomly select one of the two locations i or j . Go back to the truth model Z_{truth} and extract a new conditioning data.
7. Simulate 20 Z realizations using only the 5 conditioning data.
 - Find the second location that maximizes Q^{**} , $i(ix, iy)$, $j(jx, jy)$ respectively, where Q^{**} is the quality for the window locations using two pieces of information: the exploration data and the incremental information. This step represents decision making using only incremental data.
 - Calculate Q_{blue}^l from the truth model.
8. Cosimulate 20 Z realizations using the 5 data, Y_{truth} , and
 - Find the second location that maximizes Q^{***} , $i(ix, iy)$, $j(jx, jy)$ respectively, where the third piece of information is the conceptual model. This step represents decision making using both the conceptual model and the incremental data.
 - Calculate O_{orange}^l from the truth model. The distribution of the O_{orange}^l values
9. Using Z_{truth} find the two locations i, j that maximize O_{green} . The distribution of O_{green} values is the green histogram in Figure 6.14.
10. Repeat over L trials to observe the distributions of uncertainty.

Second Example

The experiment uses a 20x20 grid of data. A random correlation between the primary and the secondary information to be either 0.75 or -0.75 is drawn. The secondary map is constructed unconditionally using SGS and a Gaussian semivariogram having a long range and minimal nugget effect. The primary map is also constructed unconditionally but uses an exponential semivariogram with moderate nugget effect. The simulated values for both maps are normally distributed: no transform was applied. Figure 6.15 shows a map of Y_{truth} on the left and Z_{truth} on the right. The corresponding histograms are shown below. The observed correlation coefficient is 0.75.

The following mapped results are for one trial and for both of the optimal locations selected in the Red case. Collectively the following figures represent a single trial in the second experiment. For each of the figures, the histogram of Gaussian values over the 20 realizations is shown on the left. The map in the middle is the e-type estimate over 20 realizations, and the map on the right is a single realization. The optimal window locations are superimposed on the maps. The histograms illustrate the effect of bias sampling. The optimal window locations are preferentially located in high valued areas. Using the data extracted from Z_{truth} conditions the realizations and artificially increases the mean. The e-type map is shown because it illustrates what the optimization routine “sees” when searching for the two optimal locations.

Figure 6.16 shows the results for the Red case. Two optimal locations are in the bottom left and bottom right corners of the maps. The mean of the realizations is not the same as the mean of the Z_{truth} map (-0.236 vs. 0.027, respectively). In practice there is no way to know if the sample mean is equivalent to the true mean, but the sample mean represents one aspect of our conceptual model. Figure 6.18 shows the results after updating using the optimal window location on the bottom right in the Red case. The updated location is the same as in the Red case. The mean is 0.07. The mean is much higher in the Blue case because of the high values conditioning data. Figure 6.17 shows the results after using the conceptual model. The second window location is better than the selection made without the conceptual model. The mean is -0.009. Figures 6.19, and 6.20 are the same as above but use the incremental data from the bottom left window location in Figure 6.16. The respective means are -0.001, and 0.032. Note that all of the updated means differ significantly from the mean in the Red case. Despite the fact that they are closer to the mean of Z_{truth} and all represent differences in the conceptual model. Figure 6.21 shows the optimal location for the Green case: perfect information.

Comments

The resulting histograms for 100 trials are shown in Figure 6.22. The mean quality of the Red Histogram is 50.62, the Blue Histogram is 61.00, the Orange Histogram is 88.41, and the Green Histogram is 113.65. In this case using only the incremental information lead to better optimal location selection, but only marginally. Using both the conceptual model and the incremental information lead to significantly better location selection on average.

A summary of results is shown in Table 6.1. The table tabulates the number of times the Red case exceeded the Blue and Orange cases ($Red > Blue$, $Red > Orange$), the Red case was equivalent to the Blue and Orange case ($Red < Blue$, $Red < Orange$), the Blue and Orange cases were equivalent to the Green case ($Green = Blue$, $Green = Orange$). The tabulated results show that using just the incremental information does not clearly lead to better decision making even though the Blue Histogram showed higher average quality than the Red Histogram: only 41% of the time the Blue case exceeded the Red case. Comparing the Blue and Orange cases, the Orange case yielded better decisions 78% of the time.

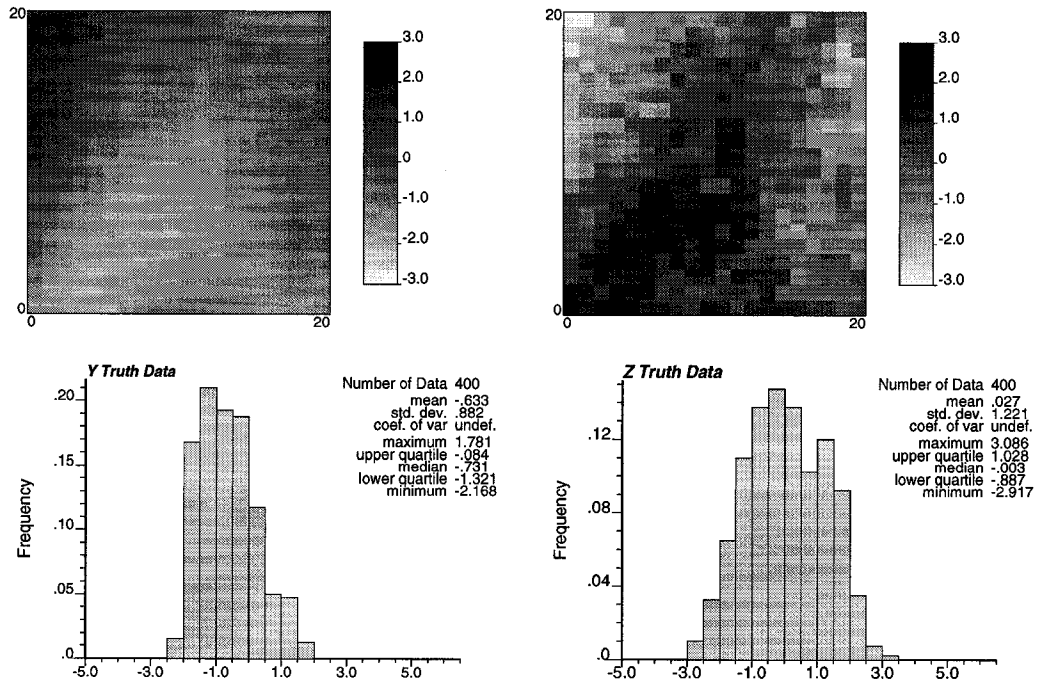


Figure 6.15: The map on the left is the secondary information, the Y information. The corresponding histogram is below. The map on the right is the Z map. The secondary information was used to construct this map. Its corresponding histogram is on the right. The correlation coefficient is -0.75 . The mean of Z_{truth} is 0.027




Figure 6.16: The histogram for the 20 realizations is shown on the left. The mean is -0.236 . The map in the middle is the e-type map conditional to four data taken from the corners of the Z_{truth} data shown in Figure 6.15. The selected locations are shown as a small circle with a square window surrounded it. The map on the right shows the locations superimposed onto the Z_{truth} map.




Figure 6.17: These results use the incremental data from the Red case. The location on the right was selected. On the right is the histogram of the realizations constructed with the 5 data. The mean is 0.07. The map in the middle is the e-type map over the 20 realizations and using the 5 conditioning data. The map on the left shows the locations superimposed on the Z_{truth} map.




Figure 6.18: These results are for the Orange case: the incremental data from the bottom right corner of the Red case and the conceptual model were used. The mean is -0.009.




Figure 6.19: These results use the incremental information provided by the window located on the bottom left corner of the Red case. The mean is -0.001. Note that the incremental information does not change the selected location.

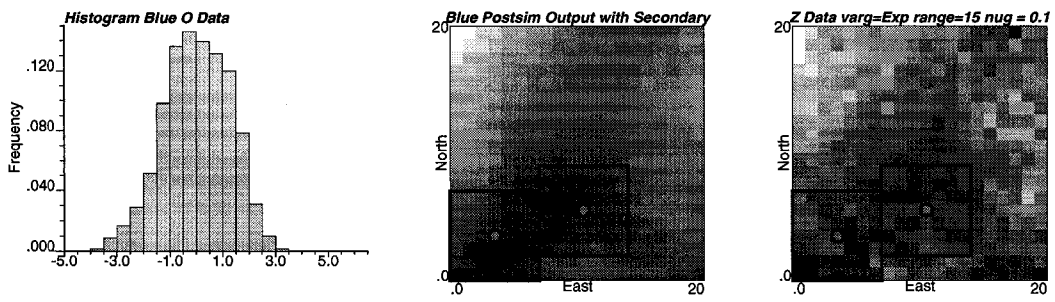


Figure 6.20: These results use the incremental information from the bottom left corner of the red case and the conceptual model. The mean is 0.032.

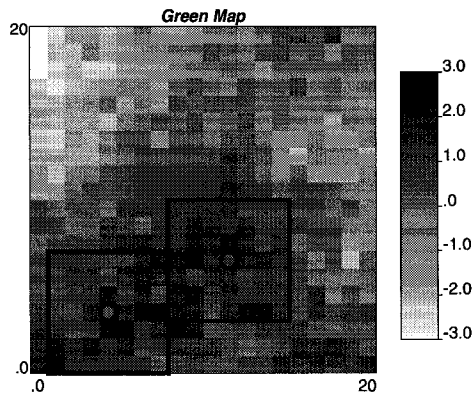


Figure 6.21: These are the final well locations for the green case.

	Blue (Percent)	Orange (Percent)
Red >	22	3
Red =	37	19
Red <	41	78
Green =	5	10

Table 6.1: This Table documents the results for 100 trials in the second experiment.

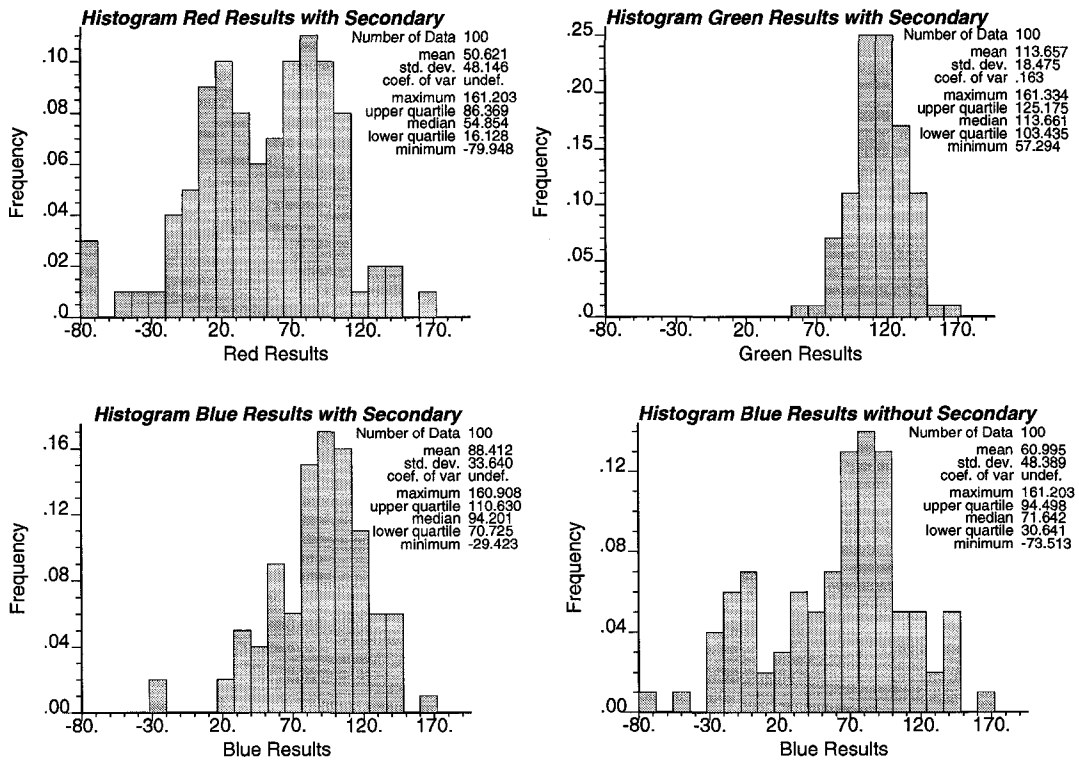


Figure 6.22: The top left histogram show the Red case results using the secondary information. The top left histogram shows the Green Case results using the secondary information. The histogram on the bottom left shows the Blue case results that consider the secondary information. The bottom right histogram shows the Blue case result without using secondary information.

6.2.3 Comments

The first experiment showed that appending well to an existing well plan can lead to sub-optimal well plans. It is better to select the wells all at once. The results of the second experiment showed that using incremental well data to update a reservoir for subsequent decision making does not always lead to better well locations. The results of the first experiment can be explained by: (1) due to the range of correlation the new data succeeded in updating the model near the existing well and provided little additional information for the rest of the model, and (2) the well locations are preferentially located in high valued areas and the extracted samples lead to bias in the simulations.

The results from the second experiment show that although the incremental information did lead to better decision regarding optimal window placement the improvement was not clear. Using conceptual information and incremental information clearly did lead to better window locations. The following summary statement states the findings of the experiments. Aim to select all well locations simultaneously if the conceptual model is well proven. If there is a chance that the conceptual model can be changed in light of new information, select well locations sequentially.

The effect of incremental information is surely affected by the size of the grid. For a larger grid the effect of informing a single node is proportionately smaller than the effect with would be if the grid were smaller: one node in the 50x50 grid represents 1/2500 of the total possible information whereas one node on 20x20 grid represents 1/400.

These findings may not apply to other decision making problems without further investigation. The problem of selection well locations is peculiar because the region of interest is not allowed to contribute to the next decision for a future well location.

One useful lesson that could be extended to almost any other decision making problems is that of revising the conceptual model as more information becomes available. The conceptual model is the foundation of the entire model; stochastic or deterministic. Revisions of the conceptual model have greater impact than that of simply adding data.

6.3 Well Selection on the Smiley Buffalo Field

This Section applies the well location selection algorithm to a waterflood program on a heavy oil field in Saskatchewan, Canada, called the Smiley Buffalo field. The Section begins with a definition of the experiment, moves on to a general geology review of the field, followed by a discussion of the static and dynamic models. The Section concludes with the well location selection. The geological, static, and dynamic model based on the paper by Mohr et al. [60].

6.3.1 Basis for Experimentation

The objective of the experiment is to demonstrate any economic advantage arising from applying the well location selection technique developed in Chapter 5 to the problem of selecting well locations on Section 15 of the Smiley Buffalo field.

Section 15 was selected for two reasons. First, Section 15 is compartmentalized from the rest of the reservoir as a consequence post depositional dissolution of the material just below the reservoir. The dissolution caused parts of the Smiley Buffalo field to collapse that resulted in compartmentalization of the field. Field observations and dye tests support the supposition that Section 15 is isolated from the rest of the reservoir. Using a small Section of the reservoir makes it easier to scrutinize the well location selection algorithm. Secondly, the timing of the development plan makes for a good case study.

Section 15 was produced for some time by a single well (8-15) under primary production. The plan for the reservoir was to increase productivity by turning Section 15 over to a water flood with well 8-15 as the injector. The field was under appraised, so 3 additional wells were drilled (1-15, 9-15, and 10-15). The map shown in Figure 6.23 shows the location of




Figure 6.23: The map shows the basal surface of the Smiley Buffalo Section 15 reservoir. The area of interest is delineated by a thick polygon. There are four conditioning wells indicated by black dots.

the wells and the topography of the base of the reservoir. The wells added new information and helped guide the selection of a new development plan. The new information did not change the conceptual geological model. After a year of drilling inactivity, the four more production wells were drilled (2-15, 4-14, 6-15, 7-15) and Section 15 was put on water flood.

The four new wells, 2-15, 4-14, 6-15, and 7-15 form the basis of comparison between 4 hand selected well locations and 4 semiautomatically selected well locations. Both well plans are selected using the same information, except that the semiautomatic well locations will be selected using a reservoir model of uncertainty and of course the semiautomatic location selection algorithm.

The two well plans will be compared by observing flow simulation results over the model of uncertainty given the Asset Team locations and the semiautomatic well locations. The semiautomatic well locations will be considered superior to the Asset Team locations if, on average, the semiautomatic well locations have greater value than the Asset Team locations. The value of the wells plan will be presented in terms of barrel oil equivalent value (BOEV). BOEV values are independent of price per barrel issues and are a convenient means for comparison. The calculation of BOEV is covered in Section 5.2.2. The experimental results are exhaustive, that is, flow simulation results are obtained for each realization for the Asset Team locations and the optimal locations. The time period for flow simulation is 20 years. This likely exceeds the economic limit for the reservoir; however, considering the reservoir over this time period does not dilute the validity of the results as long as the Asset Team

selected well locations and the optimal well locations are compared over the same time frame, and the value being compared is barrels of oil equivalent.

The reservoir model of uncertainty consists of twenty realizations and relies on the four initial well locations for conditioning data. The same 20 realizations are used for the two experiments. Flow simulation results are obtained from the well locations over these realizations. The reservoir model of uncertainty is discussed in greater detail below.

The remainder of the Section covers the following topics: The static model and the dynamic model; the flow simulation results given the Asset Team locations; the preparation of the mini model, the mini model results; the construction of the component objective functions, the calibration of the global objective function; the semiautomatic well location selection; and finally, a comparison of results.

6.3.2 Geology Review

Smiley Buffalo is located in the south mid-west part of Saskatchewan. The target zone for Smiley Buffalo is the Middle Bakken sandstone, which was deposited as an offshore sand ridge in late Devonian to early Mississippian time. It is one of a series of elongate tidal sand ridges trending NE-SW in west-central Saskatchewan.

The Middle Bakken formation is subdivided into the upper, middle, and lower zones. The Lower and Upper Bakken are composed of finely-laminated, highly organic rich black shale. The shales are good marker beds as they pervade the study area, are impermeable, and “sandwich” the zone of interest. The seismic picks of the Lower and Upper Bakken were used to build the structural model for Smiley Buffalo. The Middle Bakken reservoir zone will be referred to as Smiley Buffalo.

Smiley Buffalo was deposited as a series of ridges and swales that have been reworked and migrated by current flow from the north west. The reworking left the seaward side of the ridges with regions of sand of high reservoir quality. In general, reservoir rock quality degrades toward shore due to increasing occurrence of interbedded layers of fines.

The conceptual geological model describes four facies types:

Facies 1 Green Siltstone Facies. Non-flow unit capping the Middle Bakken sandstone sequence. This facies type is more an unconformity than a facies type and served to delineate the top of structure.

Facies 2 Ridge Sandstone. This facies type is the primary reservoir facies. It is generally of good to excellent reservoir quality (25-35% and multi-Darcy permeabilities). The Ridge Sand facies is a higher energy fine grained, moderate to well sorted, semi to unconsolidated quartz sandstone.

Facies 3 Inter-Ridge Sandstone 1. A sandstone having thick intervals (2-5 cm) of coarse sand intermingled with thick inter-bedded silts. This facies has generally low reservoir quality (20-25% porosity and 2-2000 md permeability).

Facies 4 Non-Reservoir Sandstone. This facies type is a merged facies type that contains precursor material, calcite cemented sandstone, and bioturbated sandstone. Thin sandstone intervals (0.5-1 cm) with thick inter-bedded silts. This facies is generally poor to very poor reservoir quality (17-23% porosity and highly anisotropic permeabilities of 0-750 md).

The Middle Bakken formation has been subjected to extensive post-depositional collapse. These post-depositional events are due to dissolution and collapse in the dolomitic limestones of the Torquay formation stratigraphically below Smiley Buffalo. Seismic studies reveal that the majority of dissolution took place after deposition of the Middle Bakken but prior to deposition of the Mannville formation above.

The collapsed areas are referred to as sinkholes. The sinkholes have complicated mapping and modeling the field. The sinkholes compartmentalize the reservoir, and in some




Figure 6.24: The map shows the isochore thickness of the reservoir zone and the conditioning well locations.

areas, permit intrusion of out-of-zone water. Fractures around the sinkholes have caused early breakthrough of injection water in producing wells and may also provide a direct path to thief zones for injected water. The sinkholes were modeled as a network of topographic lows and pinched out layers. Faults are not incorporated into the model.

The structural horizons were built using the well markers and the seismic information as secondary information under kriging with a trend. The structural model consists of 2 horizons, the bottom of the green siltstone and the top of the precursor or calcite cemented horizons (where pre-cursor had been completely eroded), thus the model has only a single zone.

6.3.3 Static Model

The structural model is based on the top and base surfaces isopached up and down from stratigraphically nearby seismic picks. An isochore map of the reservoir area of interest is shown in Figure 6.24. The map shows a clear elongate thickness trend with the long axis in the north-east and the short axis in the south-west direction.

The sinkholes have isolated Section 15 so that only Section 15 is populated with petrophysical properties. Four well data were used to condition the model. The map in Figure 6.24 shows the area of interest, an elevation map of basal surface, and the conditioning well locations.

None of the well locations showed a definitive oil/water contact, so a contact was inferred from a nearby analogue reservoir. None of the wells in Section 15 showed a gas/oil contact, but wells in Smiley Buffalo proper do show a contact. No gas/oil contact was applied to Section 15.

The fine scale grid for the whole of Section 15 consists of 130x115x18 blocks (269100




Figure 6.25: These histograms show the raw and upscaled facies proportions for Smiley Buffalo model

blocks total), that measure 20x20m and have, on average, 1m thickness. The facies model was derived from the complete data set for Smiley Buffalo proper, which consists of 134 wells. Of the complete data set, eighteen of the wells have core with a complete core analysis including total porosity, air permeability, grain density, and residual oil and water saturations. The facies model relies on a relationship between the lithological and petrophysical logs (gamma ray, neutron, density logs, and photo electric logs). A combination of statistical and deterministic tools were applied to gain acceptable calibration between the facies model and the core facies descriptions.

With the exception of the facies model, all of the parameters required for population of reservoir model of uncertainty were derived solely from the 4 conditioning wells.

The facies log data was upscaled to the grid using the most probable classification. The input and block averaged proportions Facies 2, 3, 4 are shown in Figure 6.25. About 45% of Section 15 is Facies 2, 30% Facies 3, 25% Facies 4.

Biased core sampling in Facies 3 and 4 was a concern. Facies 3 and 4 both have alternating beds of fine grained quartz sandstone and silt. Samples were extracted mostly from the sandy intervals in these facies and silts were ignored. This sampling practice gave optimistically biased distributions of permeability and porosity. The bias was treated in the calibration of the logs petrophysical attributes to the core data.

Log permeabilities were calculated using a relationship derived from multiple regression of core K_{max} values against wet clay volumes. Porosity and fluid saturations were determined using probabilistic log analysis. The input logs included neutron, density, deep resistivity, gamma ray, and photo electric logs in the newer wells.

There is no true wet zone observed in the log data at Smiley Buffalo which made it difficult to do any complex fine tuning of water saturation models. Median log-derived water saturations are: 36% for Facies 2, 46% for Facies 3, and 72% for Facies 4.

The upscaling of petrophysical attributes to the grid was arithmetic for porosity and water saturation. Geometric upscaling was used for permeability because it did the best job of reproducing the shape of the distribution. Upscaling was facies biased. If a block was upscaled from log data as facies 2 then only facies 2 petrophysical properties were used to populated the cell with data.

The distribution of raw and upscaled porosities for each facies is shown in Figure 6.26.




Figure 6.26: The raw and upscaled porosity distributions for Smiley Buffalo

Facies	Type	Nugget	Maj. Dir.	Maj. Range	Min Range	Vert. Range
2	Sph.	0.047	N41E	535	375	5
3	Sph.	0.046	N21W	750	390	2
4	Sph.	0.13	N135E	370	270	5

Table 6.2: The facies semivariograms

The global average porosity is 22.5% and the upscaled is 23.84%. The raw and upscaled distributions of permeability are shown in Figure 6.27. The global average permeability is 884mD, and the upscaled 827mD. It is acknowledged that there are more robust techniques for upscaling permeabilities, but the simplistic technique used here was satisfactory for this project. The global and upscaled distributions of water saturation are shown in Figure 6.28. The global average raw water saturation is 47% and the upscaled is 48%.

The modeling approach was kept simple. Sequential indicator simulation was used to construct the facies realizations, sequential Gaussian simulation was used to construct the porosity realizations, and colocated cosimulation was used to construct the immovable water and permeability realizations.

No clear areal trends were observed in the four well data, or over Smiley Buffalo proper, so none were used. The sinkholes might be masking a trend. A vertical proportion curve was used for facies modeling. There is poor correlation between porosity and depth ($\rho = 0.35$), and depth and permeability ($\rho = 0.25$). Part of the weak porosity and permeability trend is expected to be captured by the vertical proportion curves.

Tables 6.2, 6.3, 6.4, 6.5 show the semivariograms used to populate the facies, porosity, permeability, and water saturation models.

The column labeled Type specifies semivariograms numerical model. The column Nugget specifies the nugget effect for the semivariogram. The column Maj., Min., Vert. specify the ranges for the major, minor, and vertical directions.

These modeled semivariograms have little nugget effect. The selection of such small




Figure 6.27: The raw and upscaled permeability distributions for Smiley Buffalo




Figure 6.28: The raw and upscaled water saturations for Smiley Buffalo

Facies	Type	Nugget	Maj. Dir.	Maj. Range	Min Range	Vert. Range
2	Exp.	0.01	N47E	1365	615	8
3	Exp.	0.01	N60E	1065	610	10
4	Exp.	0.01	N140E	690	500	7

Table 6.3: The porosity semivariograms

Facies	Type	Nugget	Maj. Dir.	Maj. Range	Min Range	Vert. Range
2	Exp.	0.01	N35E	700	580	9
3	Exp.	0.01	N35E	640	425	8
4	Exp.	0.01	N120E	580	230	7

Table 6.4: The permeability semivariograms

Facies	Type	Nugget	Maj. Dir.	Maj. Range	Min Range	Vert. Range
2	Exp.	0.01	N40E	430	310	10
3	Exp.	0.01	N10E	640	475	9
4	Exp.	0.01	N50E	260	200	5

Table 6.5: The water saturation semivariograms

nugget effects may understate reservoir heterogeneity. Except for water saturation, the semivariogram parameters echo the geologic conceptual model; the direction of maximal continuity is along the long axis of the reservoir.

The correlation between water saturation and porosity is relatively low, $\rho = 0.3$, and the relationship between porosity and permeability is nearly functional, see Figure 6.29. The permeability values are colored by facies type. The Figure shows moderate overlapping of permeability between facies types. Facies 4 has consistently low permeabilities, and due to the high correlation between permeability and porosity, consistently low porosity values. Facies 2 and 3 have more overlap, with most of the Facies 2 points falling in high permeability regions. A functional relationship could have used a to model permeabilities but colocated cosimulation was used instead. Porosity is used as the secondary variable. For Facies 2, 3, and 4, the following correlation coefficients were used: $\rho_2 = 0.84$, $\rho_3 = 0.87$, and $\rho_4 = 0.85$, respectively. A variance reduction of 90% percent was applied because this value reproduced the permeability histograms best. Directional permeability, except for Kv/Kh ratio, was not modeled, thus $K_x=K_y$. In flow simulation, a global Kv/Kh ratio of 0.6 is used.

Water saturation is highly variable. Figure 6.30 shows a scatterplot of water saturation versus porosity for Facies 2, 3, and 4. There are no easily identified relationships on any facies type. The global correlation coefficient is $\rho = 0.35$. The global correlation coefficient was used for Facies 2 and 3 under colocated cosimulation with porosity as a secondary variable. Water saturation for Facies 4 was modeled as a constant 100%.

The reservoir boundary was determined by using a threshold on the kriging variance and the model was clipped according to this boundary. Uncertainty in the reservoir extent was not considered in the model of uncertainty.

20 realizations of the reservoir model were generated. The posterior distributions and relationships were checked to ensure that they were reproduced. Figures 6.31, Figure 6.32, Figure 6.33, show realizations of the facies, porosity and permeability models.

6.3.4 Dynamic Model

This Section discusses the features of the dynamic model. The fine scale model contains 269100 blocks. The model is not unmanageably large, but was upscaled areally to reduce flow simulation run times. The grid was upscaled by a factor of 2 in the both the x and y directions. The upscaled grid has 67860 blocks and 35958 active blocks. Arithmetic averaging was used for all properties over geometric or harmonic upscaling because it better reproduced the histograms and semivariograms, and because it better reproduced well test results. Figure 6.34 compares the fine and upscaled permeabilities for well 8-15. Geometric and harmonic upscaling procedures underestimate the well flow capacity. The fine model




Figure 6.29: The permeability/porosity relationship for Smiley Buffalo. The data points are colored by facies type. The darkest colors are facies 4, and the lightest are facies 1.

flow capacity is 14520 mD-m. The corresponding flow capacities for arithmetic and geometrical upscaled models at the same location are 14376 and 7810 mD-m, respectively. After arithmetical upscaling, the average porosity weighted permeability is 1.3 D. The porosity weighted average permeability for Facies 2 is 1.8 D.

Vertical permeability was not modeled in the static model. Instead a global vertical permeability ratio K_{ratio} of 0.6 ratio was applied for Facies 2 and Facies 3. Core information from a nearby field was used to guide this selection. A rock compressibility of 4.5×10^{-6} kPa⁻¹ was determined using data from an analog field.

Initial Conditions

Initial reservoir pressure and temperature are $P_i = 7205.8 \text{ kPa}$ at -87 m tvdss (true vertical depth sub-sea) and $T_i = 27.2^\circ$, as determined from DST (drill stem testing) test results. DSTs recorded in nearby Sections gave static pressures between 6500 and 7000 kPa, so 6750 kPa was used. Fluid properties measured were oil and gas gravities at 13.2° API and 0.59° API , respectively. Other reservoir fluid characteristics are: initial solution gas-oil ratio 20.03 sm³/sm³, initial oil volume formation factor 1.041 rm³/sm³; and oil viscosity at initial conditions of 743 mPa.s. Figure 6.35 shows the oil PVT properties.

Fluid Contacts

An initial gas-oil contact at -87 m tvdss was determined from well logs on Smiley Buffalo proper. This is higher than the highest elevation in the top surface of the model. Although water zones could not be detected in the petrophysical study, water production occurred in several wells before the start of water injection and evidence from nearby wells showed a water contact between -106.6 and -116.5 m. Also, water zones are located in the structurally low areas caused by sinkhole dissolution and collapse of underlying formations.




Figure 6.30: The permeability/porosity relationship for Smiley Buffalo. The data points are colored by facies type. The darkest colors are facies 4, and the lightest are facies 1.




Figure 6.31: The map shows a facies realization from the Smiley Buffalo reservoir model.

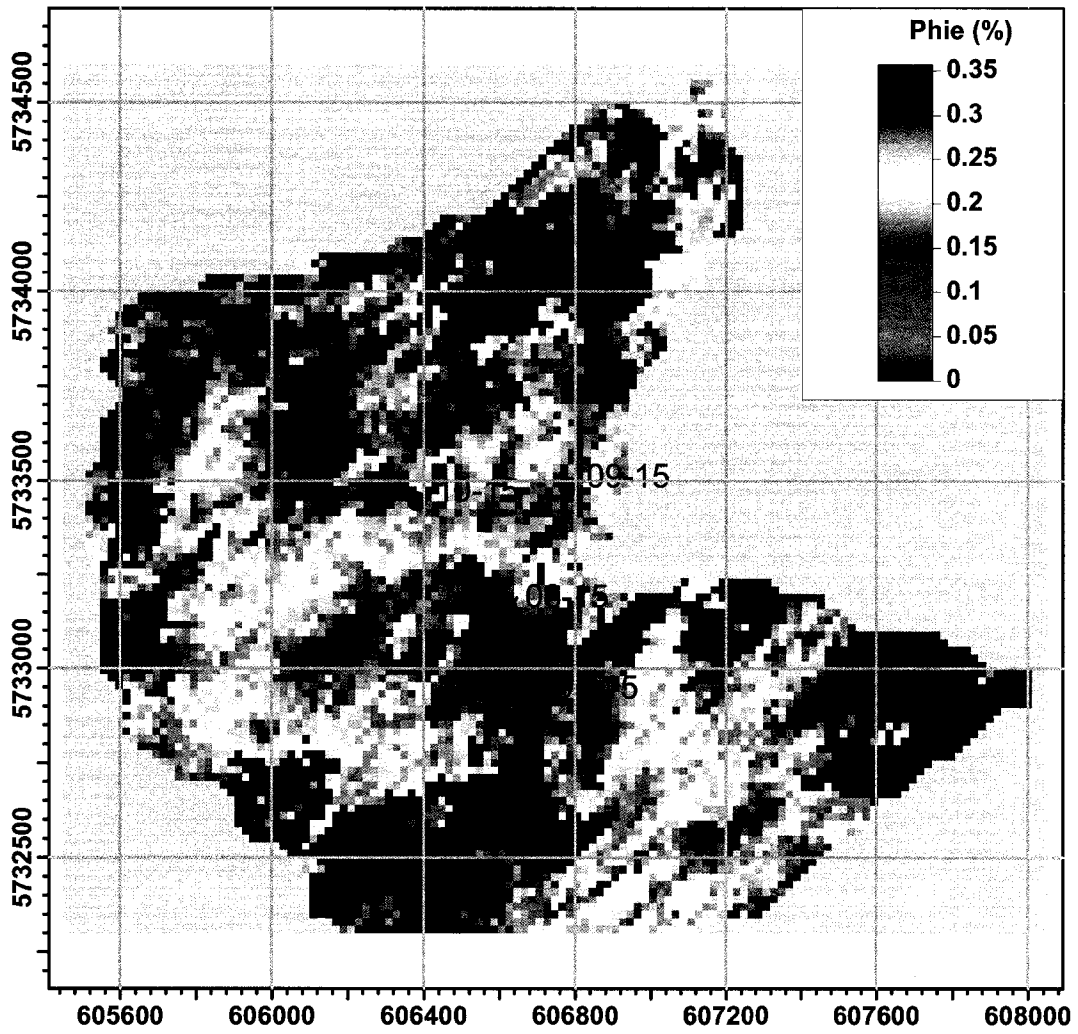


Figure 6.32: The map shows a porosity realization from the Smiley Buffalo reservoir model.

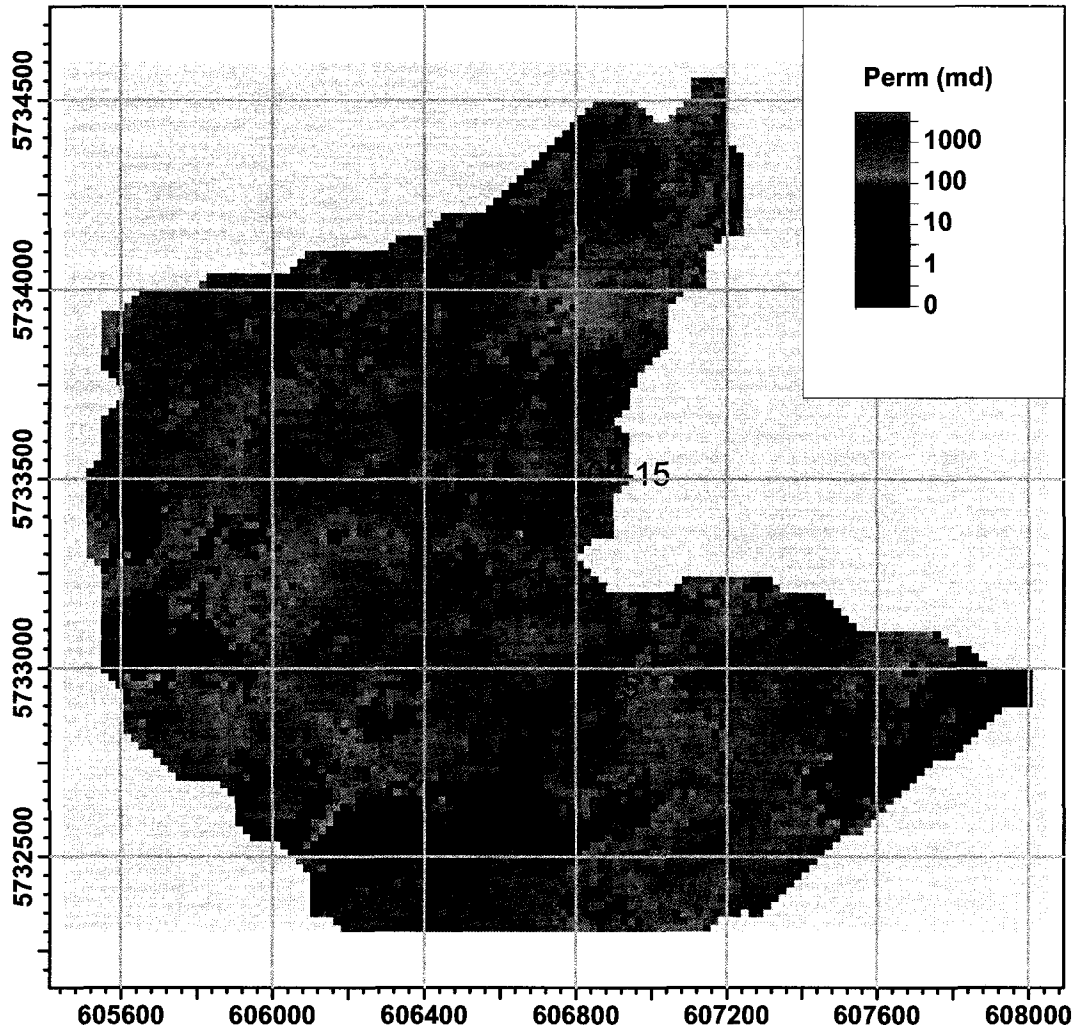


Figure 6.33: The map shows a permeability realization from the Smiley Buffalo reservoir model.




Figure 6.34: The chart compares the average permeability by layer using harmonic, geometric, and arithmetic upscaling on well 8-15. Source: Mohr (2005) [60].




Figure 6.35: The PVT properties for Smiley Buffalo. Analogue information from the nearby Court file was used to develop these curves. Source: Mohr (2005) [60].




Figure 6.36: The relative permeability curves used for reservoir simulation. Analogue information from a nearby reservoir was used to develop these curves. Source: Mohr (2005) [60].

Special Core Analysis

Two-phase, oil-gas and oil-water, unsteady state relative permeability curves were measured with Bakken reservoir fluid and rock samples in 1980. Unfortunately, those samples were damaged during the test procedures. The results were encouraging for Section 15, but water production behavior in other regions could not be satisfactorily explained.

Figure 6.36 shows the base oil-water relative permeability curves for the good and poor rock qualities used in the reservoir simulation.

The irreducible water saturations of the oil-water relative permeability curves were shifted to match the averages determined by well log analysis for the two rock types. $S_{wi} = 0.24$ was used for Facies 2 and $S_{wi} = 0.34$ for the rest of the reservoir facies. The capillary pressure functions were rescaled to match a 2 m transition zone in a nearby analog reservoir.

6.3.5 Asset Team Flow Simulation Results

This Section presents the flow simulation results for the well locations selected by the Asset Team. These results form the basis of comparison for the optimally selected well plan.

Figure 6.37 shows the field oil production versus time for the Asset Team selected wells. The thick dark line shows the average cumulative result, and the 20 realizations are shown as thin lines. There is quite a bit of variability among the runs (minimum = 1E6 BBL, maximum = 1.6E6 BBL). Clearly, heterogeneity has significant affect on recoverable oil. The average total field oil production is 1562E3 BBL. Figure 6.38 shows the field cumulative water production. The water production was quite high by the end of the life of the reservoir (24212E3 BBL), and it may not be economic to carry on the waterflood to this date because of the high water cut. The average cumulative injected water is 24919E3 BBL and the time step plot is shown in Figure 6.39.

Figure 6.40 shows the barrel oil equivalent value of the Asset Team well plan over 20 years. The formula for calculating BOEV is shown below:

$$BOEV = C_{prod} \cdot FOPT - C_{treat} \cdot FWPT - C_{inject} \cdot FWIT \quad (6.2)$$




Figure 6.37: This chart shows the field oil production versus time for the Asset Team Selected wells. The x -axis shows time and the y -axis shows cumulative barrels of oil.

The cost of producing oil, treating and pumping water are $C_{prod} = 0.95$, $C_{treat} = 0.006$, and $C_{inject} = 0.006$. These values are derived from field observations and will be used to calculate the value of the global objective function. The BOEV curves are expectedly similar to the cumulative oil production curves. As with cumulative oil, there is quite a lot of variability in BOEV (minimum = 752E3 BBL, maximum = 1308E3 BBL) of the field after 20 years is 1189E3 BBL.

The numerical values used to calculate BOEV are shown in Table 6.6, and cumulative distribution functions of each variable are shown in Figure 6.41. The values are in 1000s of barrels. At the bottom of the Table of results are the average values and the corresponding standard deviations.

The goal of the experiment is to select well locations using the semiautomatic well location selection algorithm that have, on average, greater barrel oil equivalent value than the well locations that the Asset Team selected. The Asset Team selected well plan is further examined in the Experimental Comparison Section.

6.3.6 Well Location Selection

In this Section, the semi-automatic well location selection algorithm is used to semiautomatically select well locations. The well location selection is called semiautomatic because initial well locations are required. The Asset Team locations are used as seed locations. The first step in the algorithm is calculate the value of the global objective function value for the seed well locations. The seed well locations are then iteratively perturbed until the well plan conforms a well plan that optimizes the global objective function.

The well location selection is presented in the following order. The drainage radius and injector radius of influence are important inputs to the mini model and the global objective function, so these results are presented first. The mini model is next. The mini




Figure 6.38: This chart shows the field water production versus time for the Asset Team Selected wells. The *x*-axis shows time and the *y*-axis shows cumulative barrels of water produced.




Figure 6.39: This chart shows the field water injected versus time for the Asset Team Selected wells. The *x*-axis shows time and the *y*-axis shows cumulative barrels of water injected.




Figure 6.40: This chart shows the barrel oil equivalent value for the field under the Asset Team well plan. The x -axis shows time and the y -axis shows cumulative value.

	FWIT (MBSL)	FOPT (MBSL)	FWPT (MBSL)	BOEV (MBSL)
	18152	1374	17084	1093
	22722	1378	21804	1042
	21857	1665	21189	1323
	21548	1339	20741	1018
	18310	1017	17538	751
	18003	1227	16975	956
	20642	1631	19638	1308
	20003	1434	19168	1127
	21176	1584	19726	1260
	21656	1613	20508	1279
	22472	1640	21287	1296
	20121	1347	19795	1040
	16245	1241	15626	988
	17877	1260	17683	983
	16202	1072	15580	827
	27098	1517	26042	1122
	17849	1512	16915	1228
	20040	1310	19310	1008
	22811	1519	21618	1176
Average	24919.42	1562.07	24212.94	1189.17
St. Dev.	2735.00	183.10	2644.08	141.67

Table 6.6: The tabulated flow simulation results for the Asset Team well plan. These used to calculate BOEV.

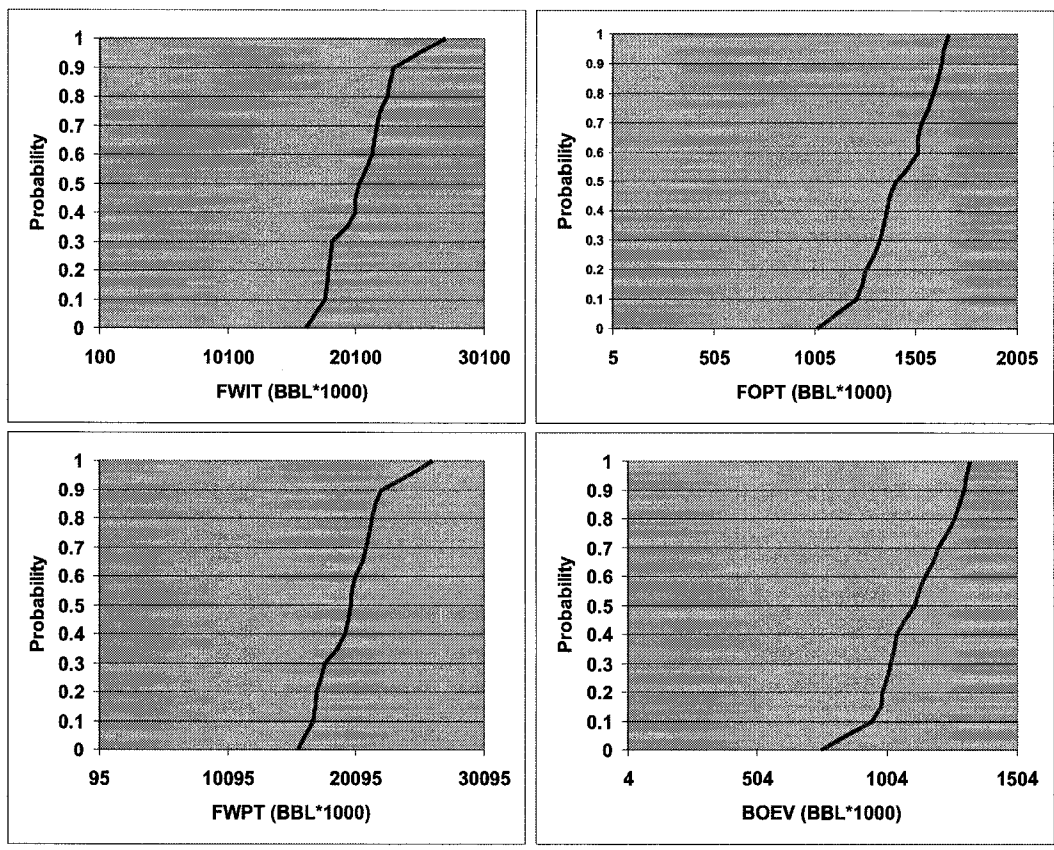


Figure 6.41: These charts are the cumulative distribution functions for the Asset Team well plan.

model results are used to build the functions that are used to estimate the produced fluid volumes and the injected water volume. The final ingredient for the component objective functions is the structural penalty. The newly constructed component objective function are compiled and the global objective function is calibrated to the flow simulation results using the Asset Teams well locations flow simulation results. The experimental results and concluding comments are presented last.

6.3.7 Drainage Radius and Injection Radius Selection

The drainage radius for the producer wells was obtained by observing the pressure draw down radius from flow simulation results of the Asset Team selected well locations. The drainage radius is used in the recoverable volume template and used to specify the width dimension of the mini model. The Asset Team well location flow simulation results are also used to calibrate the global objective function. Recycled use of the results minimizes the application of flow simulation programs.

A challenging aspect of selecting a drainage radius is that it is difficult to discern the range of pressure drop influence due to influences by structure, the reservoir, adjacent wells, injection parameters, and time. The selection of a drainage radius relies on professional expertise and judgment. There are likely more scientific or rigorous ways to select a drainage radius than the one used here, nevertheless, the selected approach worked well and is described below.

The drainage radius was selected using information from the pressure maps at the time step just after the peak oil production rate. The guiding philosophy employed is that at the time the peak oil production rate is realized, the system defined by the reservoir, the drawdown rate of the producer, and the injection rate has reached a state of equilibrium. Any changes to either the production or injection rate will expand or collapse the radius of pressure influence. It might be folly to use oil rate as a guide as it may be indicative of sweep. Observation of the water saturation maps over several timesteps showed that the water front was still quite some distance from the producer at the selected time step.

The chart shown in Figure 6.42 shows hydrocarbon production rates versus time for all 25 realizations with the average production rate indicated by a heavy line. The selected time step is 475 days (Time step 15).

For the selected time step, Figures 6.43, 6.44, 6.45 show vertically averaged pressures for realization 3 (the highest cumulative production), realization 5 (the lowest cumulative production), and the vertically averaged pressure over all realizations. Using these three maps, a drainage radius of about 100m was selected. Referring to realization 3, after 475 days, a pressure drawdown radius of about 100m can be observed. These estimates are supported by comparison with the pressure map for realization 5, and for the average pressure map over all realizations.

The injector influence radius was selected by observing the flow simulation results in the lower portions of the reservoir. Figure 6.46 shows the water saturation on the lowest layer of the reservoir for realization 5 at the end of the water flood. The injector had little influence beyond 400m so an injector influence radius of 500m was used.

6.3.8 Mini Model Results

This Section covers the construction and flow simulation results for the mini model. The mini model results provide the input information for the regression of the component objective functions. The parameters required to populate the model are derived from the full field model; however, no conditioning data are used. The dimensions of the grid blocks are the same as the full field upscaled model (40m x 40m x 1m). The number of vertical blocks is the same as the full field model.

The number of x and y blocks are guided by observations on the flow simulation results of the Asset Team well plan. The width of the grid is selected by observing the pressure




Figure 6.42: The chart shows the hydrocarbon production rate over time for all realizations and the average hydrocarbon production rate as a heavy line. The time used for observing the drainage radius is 475 days.

draw down maps discussed above. A radius of approximately 100m is observed giving a diameter of 200m. A grid dimension of 11 (twice the diameter of 200m plus one block so that the producer can be put in the middle of the grid) blocks in the x direction is selected to avoid boundary effects. The length of the mini model grid is important because observations are collected that characterize the injector and producer pair flow relationship over distance. If the grid is too short insufficient information is collected to ascertain the optimal injector/producer distance. If the grid is too long unnecessary flow simulations are run.

The length of the grid for the mini model could be as small as a single cell or as large as the greatest dimension of the field. A practical approach is used here, and a length of approximately the field length is used to give a grid length of 54 blocks. The final mini model grid measures 440m wide and 2160m long.

The injector well is located centrally at one end of the grid, $ix = 6$, $iy = 6$. The minimum separation distance for the injector producer pair is 0m (both wells in the same cell). Flow simulation results are obtained at separation distances of 80m to a maximum of 1680m ($iy = 48$) so that there are 6 blocks on three sides of the producer at the end of the exercise.

The chart in Figure 6.47 shows the average cumulative produced oil over 20 years versus separation distance. The most productive separation distances start at around 300m and end at around 550m. If the life of the field were longer the distances in the range of optimal separation distances would increase.

The chart in Figure 6.48 shows the average cumulative produced water over 20 years versus separation distance. Unlike cumulative produced oil, there is no ideal spacing: the further away the injector is from the producer the less water is produced. The only way to select the best well location is to employ an optimization algorithm and search for the best compromise.

The chart in Figure 6.49 shows the average cumulative injected water over 20 years versus separation distance. There is high correspondence between the cumulative water injected and produced charts. Figure 6.50 shows both injected and produced water on the




Figure 6.43: The map is the vertically averaged pressure for realization 3 after 475 days.




Figure 6.44: The map is the vertically averaged pressure for realization 5 after 475 days.




Figure 6.45: The map is the vertically averaged pressure averaged over all realizations after 475 days.

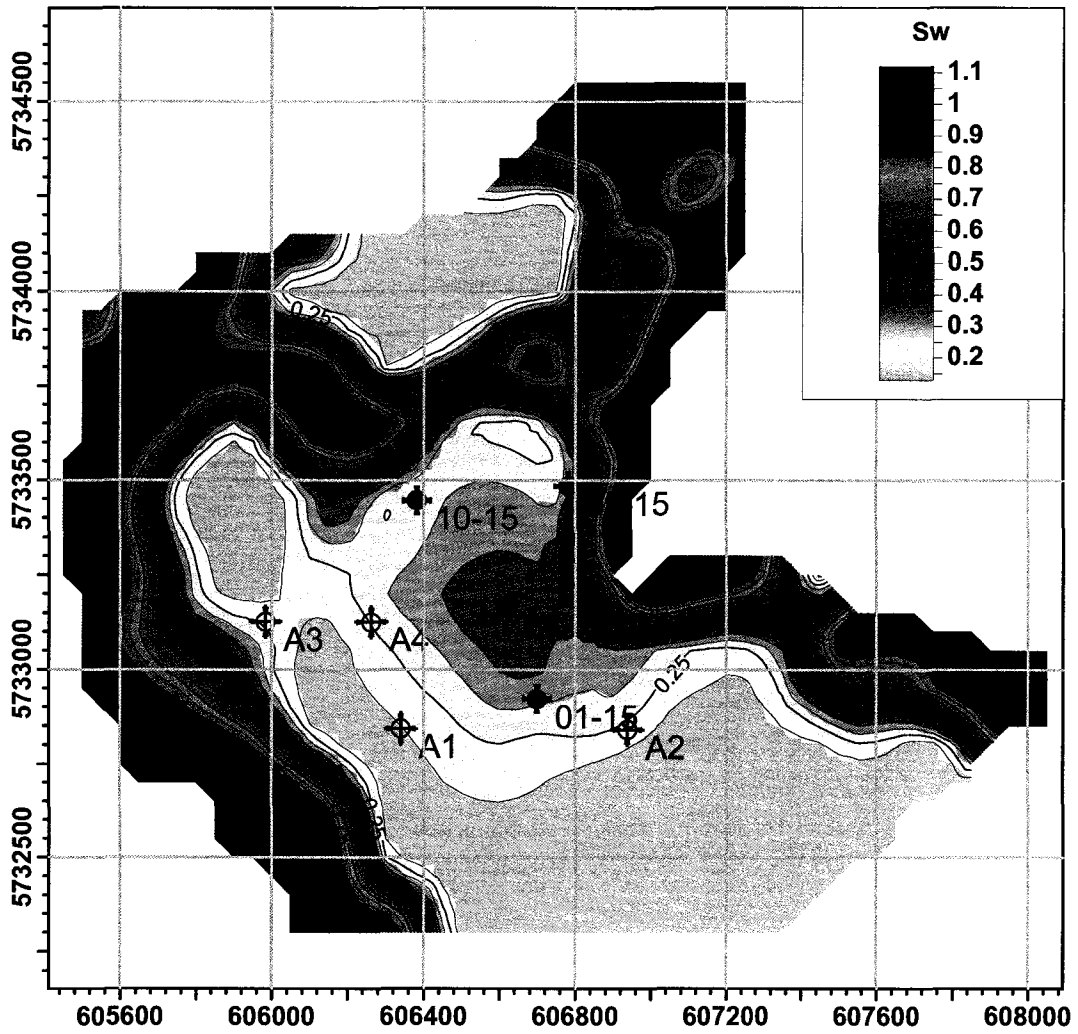


Figure 6.46: The map shows the water saturation on layer 18 (the bottom of the reservoir) after 20 years for realization 5.




Figure 6.47: The x -axis shows separation distance, and the y -axis shows the cumulative oil production averaged over all 20 realizations. The best separation distance is in the range of 300-550m.




Figure 6.48: The x -axis shows separation distance, and the y -axis shows the cumulative oil production averaged over all 20 realizations. The best separation distance is in the range of 300-550m.




Figure 6.49: The x -axis shows separation distance, and the y -axis shows the cumulative water injected averaged over all 20 realizations.

same chart but using a log scale. Beyond the optimal separation distance the injected water is no longer efficiently pushing oil to the producer: the injection volume remains the same, but the produced oil and water decreases, and the volume of injected water decreases. The unaccounted for water is sweeping regions of the reservoir that do not contribute to recovery efficiency.

Figure 6.51 shows cross sections of the water saturation model (realization number 1) at time steps 3 (at the top), 72 (in the middle), and 252 (at the bottom). The separation distance is at the maximum distance. At time step 72, gravity takes over and the water flows along the bottom of the reservoir until in time step 252 it begins coning up the producer. This is valuable information for selecting well locations as it demonstrates the effective distance of the water front. This is not to say that once water begins segregating to the bottom of the reservoir the water flood should stop; the injector may still be providing pressure support to the producer.

For each injector and producer pair, separation distance, and realization, a geobody model is calculated. The collection of cells used to calculate the geobodies start at the origin (6 cells behind injector) to 6 blocks beyond the location of the producer. Using the distance specific geobodies, the following parameters are calculated: surface area to volume ratio, the coefficient of variation on permeability, HCPV, average permeability, and the size (number of blocks) of the geobody. All of these parameters are collected because there is no way to know a priori what combination of variables will lead to the best multiple regression proxy for EOPT, EWPT, EWIT. There is some built in redundancy in the collected data. For example, the size of the geobody and HCPV are a little redundant because as the size of the geobody increases, HCPV increases. Similarly, as the size of the geobody increases the distance must increase. Nevertheless, the data was collected to maximize the potential of constructing regressed proxies of the highest quality.

Figure 6.52 shows how the coefficient of variation on permeability is affected by separation distance. As the separation distance increases the size of the geobody increases, the more data is incorporated into the geobody, and the distribution of permeability becomes better behaved.

Figure 6.53 shows how the ratio of the surface area to the volume changes with respect to increasing separation distance. At a distance of around 500m, the separation distance that yields the maximum volume of recovered oil, there is an inflection in the distance versus




Figure 6.50: The x -axis shows separation distance, and the y -axis shows log scale cumulative water produced (diamonds) and injected (triangles) averaged over all 20 realizations. The departure at large distance is due to water injected and not being produced at the well.

surface area to volume chart, and near a separation distance of about 1100m, the relationship between separation distance and surface area to volume ratio ceases to be meaningful.

Figure 6.55 and Figure 6.54 show a clear linear relationship between separation distance and HCPV and geobody size. These two data are redundant, but will be included in the examination anyway. Figure 6.56 shows average permeability versus separation distance. Despite the apparent trend in the plot, there is really very little variability in permeability - at permeabilities of greater than 1200md, 100md has little effect on the flow performance of the rock. Nevertheless, the data will be considered in the construction of the component objective functions.

6.3.9 Structural Penalty

The well location selection algorithm accounts for the topography of the basal surface. The map shown in Figure 6.57 shows the map of the direction of strike for each block in the grid. The azimuth is calculated by finding the direction of the block with the highest elevation within a 100m radius around the well and using the centre of the block under consideration as the origin. The Asset team wells are shown for reference. In the Asset Team well plan, wells 1-15 and 9-15 watered out the fastest. The elevation of the base at the injector (well 8-15) is -108.68. Well 1-15 is at elevation -111.0m, and well 9-15 is at -114.83. Also, according to the direction of strike map, all three wells are on the same slope. The well plan has the producers down dip from the injector. Table 6.7 shows the penalty values assigned to well in the Asset Team well plan.

The map of azimuths is shown in Figure 6.57. The injector is near the top edge of a bowl that services a number of producers. The map of dips in the direction of strike is shown in Figure 6.58. The wells 8-15, 1-15, and 9-15, are on similar slopes, and there are no valleys or crests between the injector and down dip producers. Well 4-14 is the best performing well. The elevation at well 4-14 is -108.35, nearly the same elevation as the injector. The dip map shows that a crest lay between the injector and producer. This situation is preferred over the well locations of 1-15 and 9-15.




Figure 6.51: The x -axis shows separation distance, and the y -axis shows the cumulative oil production averaged over all 20 realizations. The best separation distance is in the range of 300-550m.




Figure 6.52: The x -axis shows separation distance, and the y -axis shows the coefficient of variation on permeability for the geobodies constrained by the distance between the injector and producer.




Figure 6.53: The x -axis shows separation distance, and the y -axis shows the surface area to volume ratio for the geobodies constrained by the distance between the injector and producer.




Figure 6.54: The x -axis shows separation distance, and the y -axis shows the size for the geobodies constrained by the distance between the injector and producer.




Figure 6.55: The x -axis shows separation distance, and the y -axis shows the HCPV for the geobodies constrained by the distance between the injector and producer.

Well	Azi.	Basal	Azi. Pen.	Deep Pen.	OWC Pen.	Final Pen.
8-15	270	-99.7	NA	NA	NA	NA
1-15	270	-102.53	0.5	-0.060	-0.080	0.359
9-15	270	-106.96	0	-0.154	-0.124	-0.278
10-15	180	-103.41	1	-0.078	-0.086	0.834
2-15	0	-96.97	1	0.058	-0.055	1.0026
4-14	180	-102.82	1.5	-0.066	-0.082	1.35
6-15	45	-102.5	2	-0.059	-0.08	1.86
7-15	90	-96.17	0.5	0.0751	-0.0537	0.522

Table 6.7: This Table lists the individual penalties that comprise the structural penalty and the final penalty.




Figure 6.56: The x -axis shows separation distance, and the y -axis shows the permeability constrained by the distance between the injector and producer.




Figure 6.57: A map of the azimuth of strike for the Smiley Buffalo reservoir with the Asset Team selected wells.




Figure 6.58: This map shows the dip at each location on the basal surface.

6.3.10 COBJ Construction

This Section discusses the multiple regression input data, the fitment results, and the regressed function for the EOPT and EWPT component objective functions. There is no way to know apriori what combination of data and interactions that will construct a COBJ function that minimizes estimation error. For this reason many data, sometimes redundant data, are collected from the mini model. Since the functions are merely proxies for flow simulation, there will always be error in the estimate; however, it is the contention of this thesis that the estimation errors would be less than the potential loss due to selection of well locations by hand.

EOPT

The EOPT function estimates the cumulative volume of oil that an injector producer pair might produce given certain observations from the mini model. The data are collected from the mini model with the aim of uncovering relationships between the injector and producer. The relationships are modeled using multiple regression, and estimates provided by the regressed function are used in lieu of the flow equations. The data are presented in chart form and compare the static data observations to the flow simulation results.

Figure 6.59 shows the relationships between FOPT and the static information taken from the mini models. The posted results are averaged at each separation distance and over all realizations for clarity. The y axis on all the charts are for FOPT because these values are not truly estimates of production, they are actual production volumes from the mini model. The chart on the top left shows distance versus FOPT. An observation that can be made is that there is high correspondence between distance and average FOPT. According to the chart, on average, the optimal separation distance is around 500m. The chart on the top right shows the surface area to volume ratio. There is high correspondence here too. For small surface area to volume ratio values, the volume of the good reservoir blocks is small and the surface area is large giving more conduits available to flow oil to the producer. This is evidenced by the increase in FOPT for small values of surface area to volume ratios. According to the chart, well locations should be selected such that the ratio of surface area to volume is near 1.95. The chart on the left and in the middle row shows the relationship between the coefficient of variation on permeability and FOPT. The COV is a measure of the heterogeneity on permeability within good reservoir rock. There is a clear relationship between the independent and dependent variables. Internal heterogeneity clearly pays a big role in recovered hydrocarbon. If there is too little, or too much heterogeneity, FOPT decreases. The chart in the middle row, and on the right shows the average permeability and FOPT relationship. If there is too much permeability hydrocarbon production falls because these regions are swept early and become water conduits to the producer. Regions of low, to moderate permeability are preferred. The two charts on the bottom row are for HCPV and Size versus FOPT. These two charts offer nearly the same information. HCPV will be used preferentially to construct the component objective function EOPT.

For the regression, a full quadratic fitment was explored and allowed the consideration of the interactions between the independent variables as well as relationships the squared terms might have with the dependent variable. The regression was performed using the Essential Regression software [78]. The regressed function is shown below:

$$EOPT = -30.78 + 0.37 * COV * HCPV - 0.00084 * Dist * SZ + 0.00028 * Dist^2 + 0.19 * HCPV \quad (6.3)$$

where COV is the coefficient of variation, Dist is the distance between the injector and producer, SZ is the size of the geobody, HCPV is the coefficient of variation. The function is simple in that it does not rely on too many factors. The factors were selected first by allowing the program to automatically select a function. Then the fitment of the function was quality checked and modified (factors added or removed) until a model that offered minimal

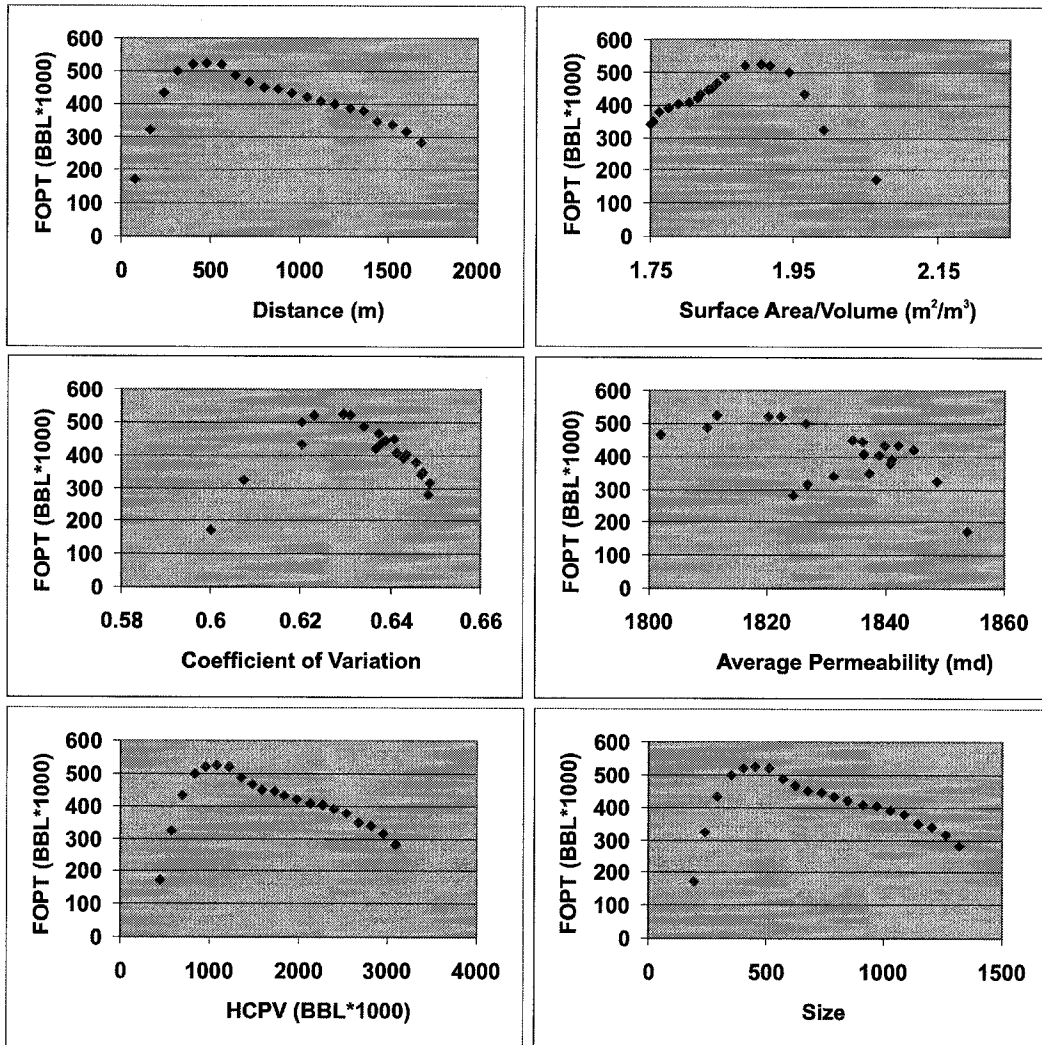


Figure 6.59: The 6 charts show how field oil production changes with respect to separation distance, surface area to volume ratio, the coefficient of variation on permeability, HCPV, and the size of the geobody.

Multiple Regression Summary	
R^2	0.792
R^2 for Prediction	0.765
Standard Error	1970.1

Table 6.8: This Table shows some fitment statistic for the regressed EOPT component objective function. The R^2 and R^2 for prediction values could be made higher by attempting to over fit.

Multiple Regression Summary						
Coef.	Coef. Val.	P value	Std Error	-95%	95%	t Stat
b0	-30.78	0.576	54.89	-139.75	78.19	-0.561
$COV \cdot HCPV$	0.374	4.45E-05	0.08732	0.200	0.547	4.280
$Dist \cdot SZ$	0.000825	4.22E-06	0.000169	-0.00116	-0.000490	10.76
$DIST^2$	0.000278	0.00205	8.77E05	0.00104	0.00452	5.329
$HCPV$	0.188	0.01527	0.07621	0.0370	0.340	-3.257

Table 6.9: This Table tabulates the fitment statistics for the coefficients of the EOPT component objective function.

heteroscedasticity between the estimate and the residuals, maximized R^2 for prediction, and minimized standard error was arrived at. The fitment parameters are shown in Table 6.8.

Common metrics for function performance are R^2 , R^2 for prediction, and the standard error. R^2 for prediction is calculated by leaving one independent variable out of the equation and using the remaining independent data and dependent data to calculate the missing value. The R and R^2 values are quite good, and the standard error is low, indicating a good fit by the EOPT function. The value of R^2 for prediction is the average value over all the independent variables. An indicator of good fit is the similarity of the values of R^2 and R^2 for prediction. The F factor is large (44.2) and the F significance is small (6.57E-21) indicating that the model is statistically acceptable.

None of the coefficients contributes significantly to the standard error and significance of each coefficient is quite high. The regressed function for EOPT is considered a good model for estimating produced oil.

Figure 6.60 shows a plot of EOPT versus the predicted EOPT. There is some error in prediction, but the magnitude of the errors are small.

Figure 6.61 shows a plot of the predicted EOPT versus the residuals. The thing to look for here is the increasing deviation of the residuals when the predicted values increase. The model shows good scatter of the residuals around 0 at all predicted values of EOPT.

$COBJ_{EWPT}$

The EWPT function is used to estimate the cumulative volume of water produced by the producer. This function on its own accounts only for the relationship between the injector and producer. The producer may produce water other than water injected by the injector as a consequence of structure of proximity to the oil water contact. Production of water other than injected water is accounted for by the structural penalty λ_{STRUCT} . Like the EOPT component objective function, EWPT is modelled using a function built with multiple regression and aim to serve as proxies for the flow equations.

The charts showing the relationships between FWPT and distance, surface area to volume ratio, coefficient of variation, average permeability, HCPV, and size are shown in Figure 6.62. The top left Chart shows the distance relationship is clearly non-linear and some-




Figure 6.60: The x -axis shows the EOPT values, and the y -axis shows the predicted EOPT values. There is good correspondence between the actual and predicted values.




Figure 6.61: The x -axis shows the predicted EOPT values, and the y -axis shows the standardized residuals. The residuals do not exhibit heteroscedasticity.

Multiple Regression Summary	
R^2	0.897
R^2 for Prediction	0.882
Standard Error	915.46

Table 6.10: This Table shows some fitment statistic for the regressed EWPT component objective function.

Multiple Regression Summary						
Coef.	Coef. Val.	P value	Std Error	-95%	95%	t Stat
b0	41446.1	2.41E-10	5849.2	29834.1	53058.2	7.086
$COV \cdot HCPV$	-1.987	7.45E-26	0.137	-2.258	-1.715	-14.52
$Dist \cdot SZ$	0.0420	9.14E-09	0.007	0.028	0.055	6.302
$DIST^2$	-87.26	3.69E-10	12.47	-112.02	-62.50	6.995
$HCPV$	3.314	1.66E-17	0.317	2.685	3.942	10.47

Table 6.11: This Table tabulates the fitment statistics for the coefficients of the EWPT component objective function.

what different than that of the FOPT and distance. There is no optimal distance, as distance increases water production decreases. There is also a strong relationship between FWPT and the surface area to volume ratio. Big conduits, that is geobodies that have small surface areas but large volumes, seem to produce less water than geobodies that have more surface area. This might have something to do with permeability heterogeneity. The chart in middle left shows the coefficient of variation versus FWPT. When the standard deviation is small, low heterogeneity, less water is produced. The chart in the middle right shows the average permeability of the geobody. The chart shows some strange behavior in that there are high average permeabilities where there is little water produced. This is a consequence of distance: as distance increases less water is produced, but the average permeability does not change. The two bottom charts show the relationships between HCPV and Size versus FWPT. The two charts are redundant because they merely show that as HCPV or Size increases (distance increases) less water is produced.

The procedure for modeling EWPT is the same as that used to model EOPT. The regressed function is shown below:

$$EWPT = 414456.1 - 1.99 * Dist * SAV + 0.042 * AVGP^2 - 87.3 * AVGP + 3.31 * SAV * AVGP \quad (6.4)$$

where SAV is the surface area to volume ratio, AVGP is the average permeability for the geobody. The function is simple in that it does not rely on too many factors. The function relies heavily on the average permeability of the geobody and the separation distance between the injector and the producer. The fitment parameters are shown below in Table 6.10.

The R^2 , R^2 are for prediction, and the standard error. The R and R^2 values are quite good, and the standard error is low, indicating a good fit by the EWPT function. The F factor is large (207.33) and the F significance is small (5.07E-46) indicating that the model is statistically acceptable.

As with the EOPT function, none of the coefficients in the function for EWPT contributes significantly to the standard error and significance of each coefficient is quite high. The regressed function for EWPT is considered a good model for estimating produced water.

Figure 6.63 shows a plot of EWPT versus the predicted EWPT. There is some error in prediction, but the magnitude of the errors are small.

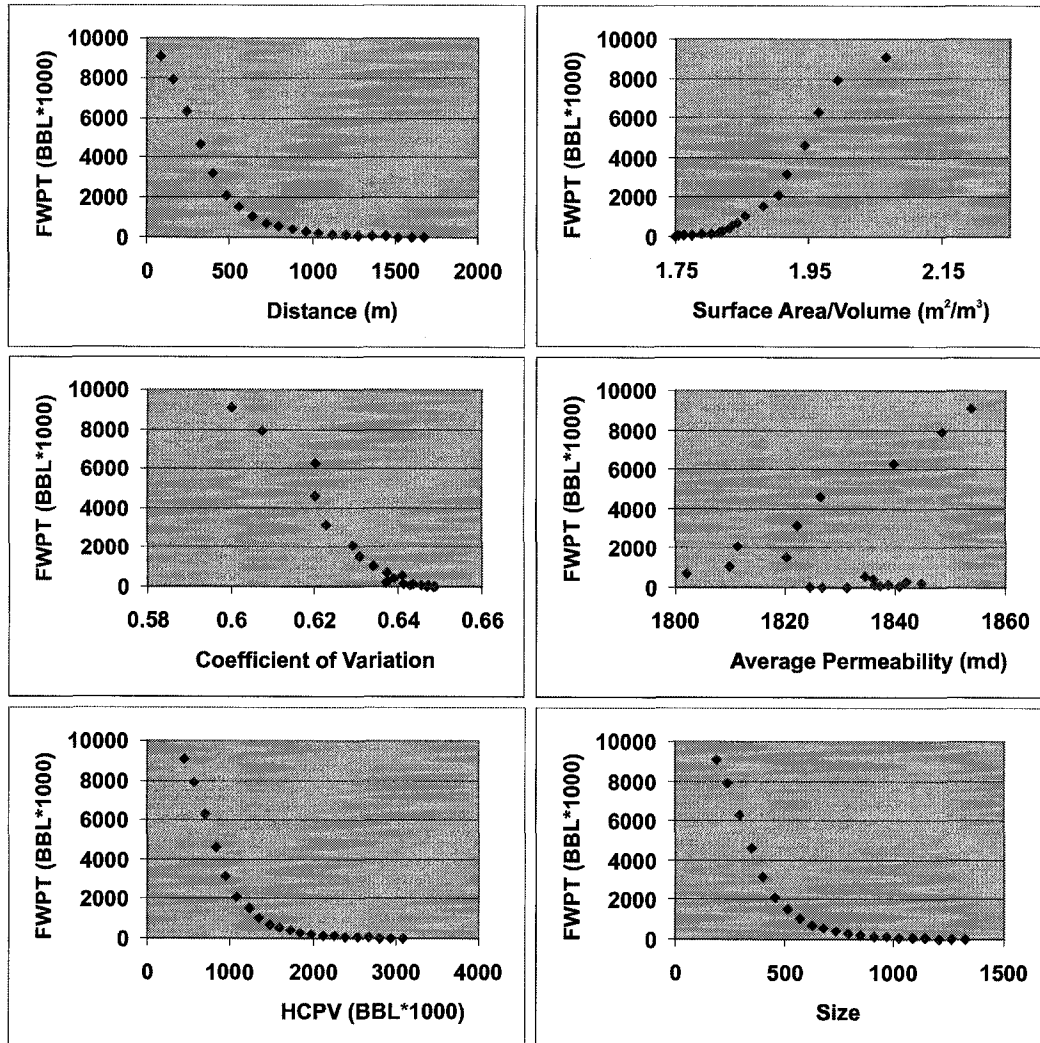


Figure 6.62: The 6 charts show how field water production changes with respect to separation distance, surface area to volume ratio, the coefficient of variation on permeability, HCPV, and the size of the geobody.




Figure 6.63: The x -axis shows the EWPT values, and the y -axis shows the predicted EWPT values. There is good correspondence between the actual and predicted values.




Figure 6.64: The x -axis shows the predicted EWPT values, and the y -axis shows the standardized residuals. The residuals do not exhibit heteroscedasticity.

Figure 6.64 shows a plot of the predicted EWPT versus the residuals. There is minimal heteroscedastic behavior and the model shows good scatter of the residuals around 0 at all predicted values of EWPT.

COBJ_{EWIT}

Modeling of the expected total injected water relies heavily on the estimate of water produced. The relationships between EWIT and distance, Surface area to volume ratio, Coefficient of variation on permeability, HCPV, and Size, are shown in Figure 6.65 and demonstrate the validity of this assumption; the charts are nearly the same as those for EWPT. For ease of comparison the EWPT values are shown in the chart as light colored triangles. The EWPT are all slightly less than the EWIT values due to a voidage replacement factor of 1.1 for the water injection.

COBJ Concluding Remarks

The models for EOPT, EWPT, and EWIT are reasonable proxies for flow simulation. The approach of using input data and multiple regression to construct a proxy for the action of performing flow simulation is not new, it has seen application in numerous fields under very

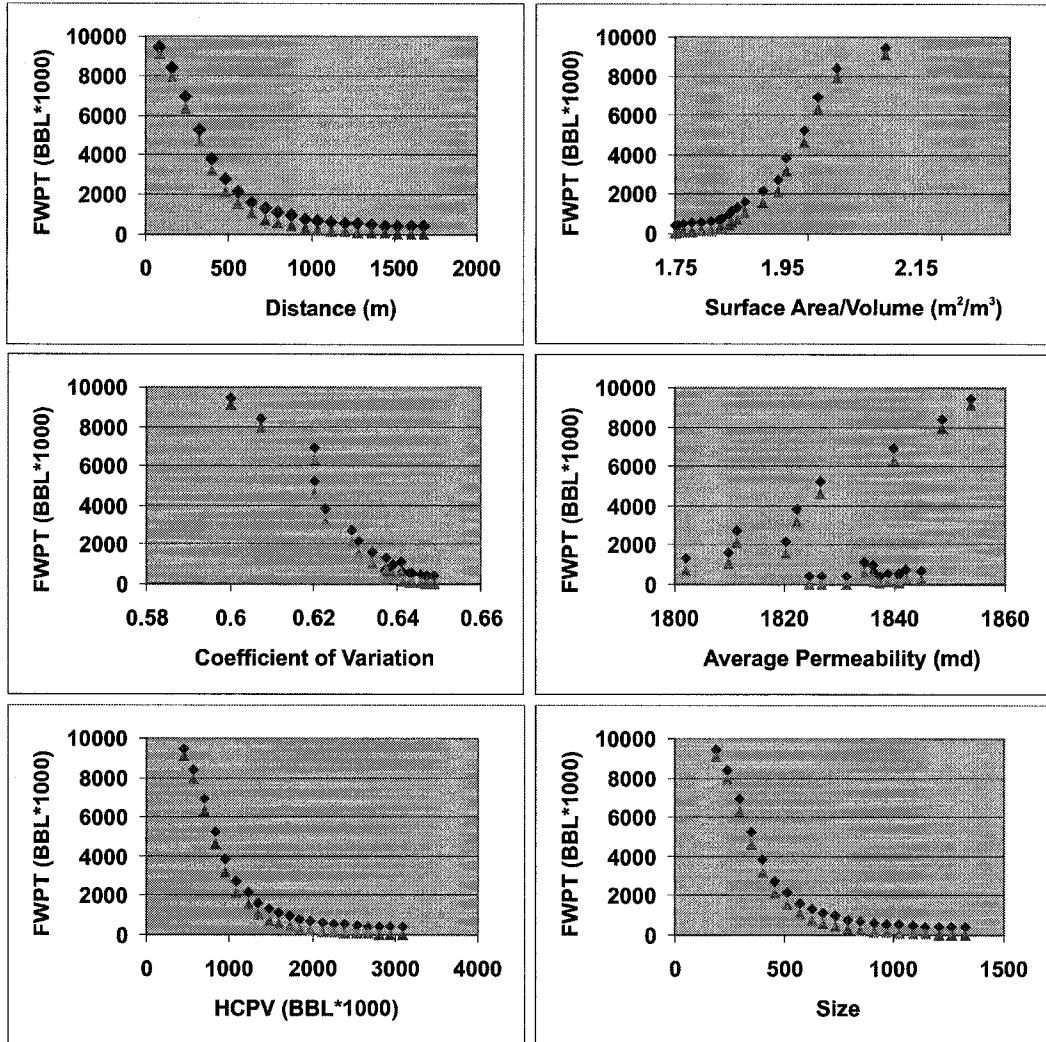


Figure 6.65: The 6 charts show how field water injected and produced changes with respect to separation distance, surface area to volume ratio, the coefficient of variation on permeability, HCPV, and the size of the geobody.

EOPT (MBBL)	FOPT (MBBL)	EWPT (MBBL)	FWPT (MBBL)	EWIT (MBBL)	FWPT (MBBL)	GOBJ (MBBL)	BOEV (MBBL)
1726	1374	18152	17740	19514	17084	1416	1093
1920	1378	22722	21063	23170	21804	1559	1042
2081	1665	21857	27198	29918	21189	1634	1323
1473	1339	21548	21249	23373	20741	7430	1018
1444	1017	18310	15728	17301	17538	9545	751
1459	1227	18003	15743	17317	16975	9440	956
1961	1631	20642	20742	22816	19638	1601	1308
1775	1434	20003	21681	23849	19168	1413	1127
2110	1584	21176	18418	20260	19726	1773	1260
2020	1613	21656	21699	23869	20508	1645	1279
1899	1640	22472	20506	22556	21287	1546	1296
1752	1347	20121	23580	25938	19795	1367	1040
1462	1241	16245	14372	15809	15626	1208	988
1780	1260	17877	20113	22124	17683	1438	983
1343	1072	16202	16252	17877	15580	1071	827
1970	1517	27098	26595	29254	26042	1536	1122
1839	1512	17849	20611	22673	16915	1487	1228
1882	1310	20040	19162	21078	19310	1547	1008
2079	1519	22811	19971	21968	21618	1723	1176
1943	1562	24919	22870	25157	24212	1558	1189

Table 6.12: This Table of values compares the component objective function values with the corresponding field flow simulation values as well as the global objective function values (before calibration) with the BOEV values from flow simulation

similar conditions. The next step is to calibrate the global objective function to the flow simulation results. If the component objective functions are good proxies the calibration weights will be small.

6.3.11 Calibration to Flow Simulation Results

In this Section, the Asset Team well plan results presented above are used to calibrate the global objective function. Table 6.12 shows the values for the component objective function values and the corresponding values from the Asset Team well plan. The values are quite close, and credit to this results can be assigned to the good fit of the regressed EOPT and EWPT functions. Recall that EWIT is actually a function of EWPT and the voidage replacement factor used in the simulation runs. The two last columns present the global objective function values and the BOEV values from the flow simulation. The pre-calibration correlation coefficient between the global objective function and BOEV is quite high at 0.80. The calibration exercise will improve the correlation, but probably not by much.

The values of the global objective function pre- and post-calibration as well as the BOEV values are shown in Table 6.13. The selected weights are $\lambda_{EOPT} = 1.44$, $\lambda_{EWPT} = 0.946$, and $\lambda_{EWIT} = 1.01$. The correlation coefficient increased to 0.83 using the weights.

The calibrated global objective function values are somewhat different from the pre-calibration global objective function values and the BOEV values. The well locations selection algorithm is not affected by the difference as it aims to find the best solution no matter the numerical values enumerated by the global objective function. A drawback of using the calibrated values is that they do not give a direct estimate of the BOEV; however, one can be backed out given the weights.

Pre. GOBJ (MBBL)	Post. GOBJ	BOEV (MBBL)
1416	17203	1093
1559	19966	1042
1634	22924	1323
1131	13971	1018
1174	13581	751
1187	14545	956
1601	20383	1308
1413	17868	1127
1773	22589	1260
1645	21525	1279
1546	19705	1296
1367	17771	1040
1208	13777	988
1438	17848	983
1071	12148	827
1536	21336	1122
1487	18831	1228
1547	19252	1008
1723	22135	1176
1558	20472	1189

Table 6.13: The Table of values compares the pre- and post-calibration global objective function values with the BOEV values calculated directly from flow simulation results. There is a large difference in the magnitude of the values.

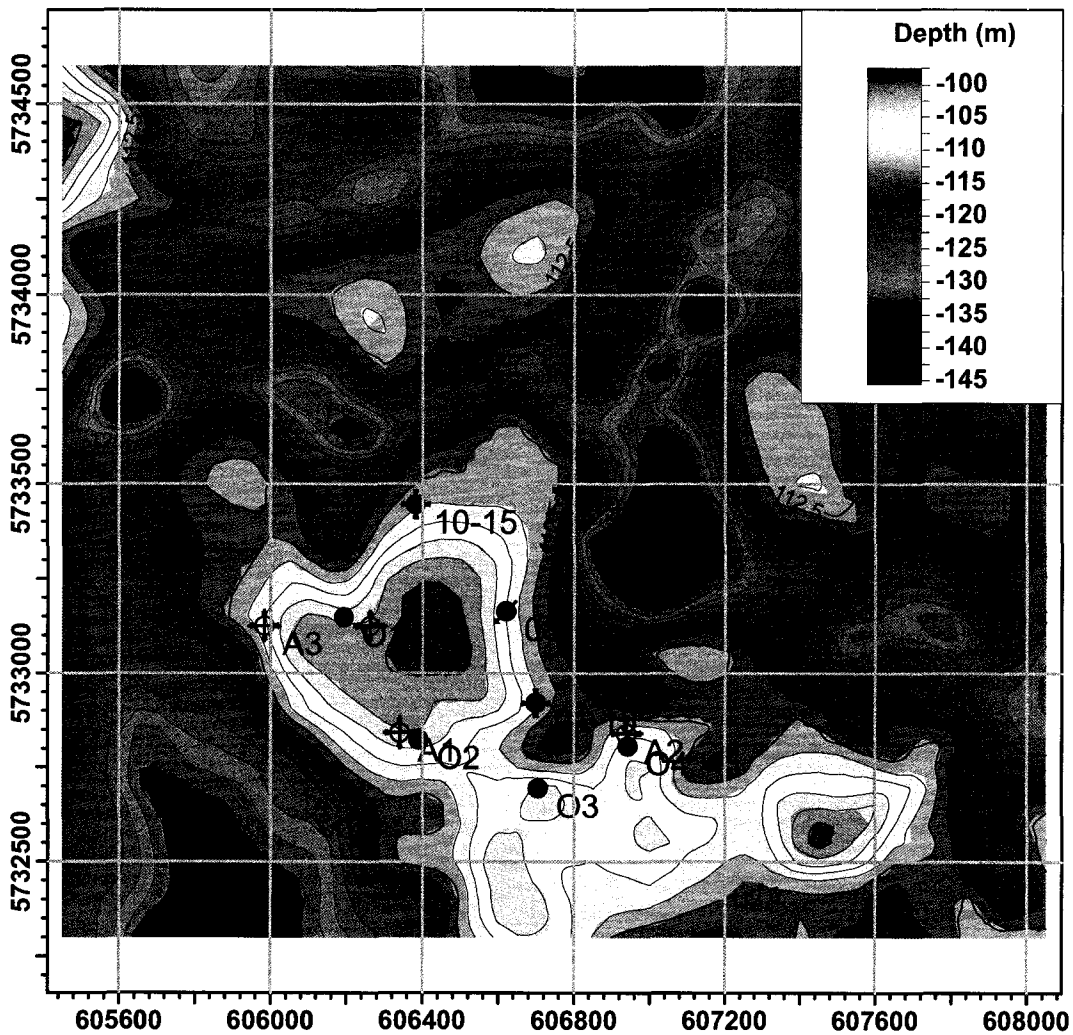


Figure 6.66: This map shows the optimal well locations (solid dots), the Asset Team locations (circles with crosshairs), the conditioning data (dots with crosshairs), and the injector (dot with arrow).

6.3.12 Experimental Results and Comparison

The map shown in Figure 6.66 shows the basal surface, the injector (dot with arrow), the conditioning wells (dots with crosshairs), the optimal wells (dots), and the Asset Team wells (circles with crosshairs). For additional clarity the wells are AX to indicate an Asset Team location and OX to indicate an optimal well location. The optimal wells are placed not much differently than the Asset Team locations. The most obvious difference is the locations of wells A3 and O3 and the spacing of the wells.

Figure 6.67 shows a map of the recoverable volumes for the Asset Team (top) and optimal well plans (bottom). The most obvious difference is the location of the most western well in the Asset Team well plan.

Figure 6.68 shows a map of the sweep for each well plan. All of the wells are serviced by the injector in the optimal well plan compared to the Asset Team well plan. In the Asset Team well plan, the most western well (well named A1) is not reached by the injector. The topology of the base for the path to well A1 increases in elevation then decreases by the





Figure 6.67: This map shows the recoverable volume template for the Asset Team well plan (top) and the optimal well locations (bottom)

	FWIT (MBBL)	FOPT (MBBL)	FWPT (MBBL)	BOEV (MBBL)
	21814	1734	20327	1394
	24042	1452	22817	1098
	26145	2098	24770	1687
	26084	1505	24976	1123
	22146	1289	21104	965
	23942	1533	22861	1175
	25212	1940	23868	1549
	22803	1700	21320	1350
	25088	1840	23295	1458
	26287	1986	24681	1581
	29553	2022	28120	1575
	23699	1712	22741	1348
	20454	1466	19422	1153
	19860	1423	19182	1117
	20068	1316	19082	1016
	29957	1666	28299	1233
	20544	1658	19222	1336
	25162	1564	24205	1190
	29367	1868	27880	1431
	28535	1895	27208	1466
Average	24538	1683	23269	1312
St. Dev.	3046	226	2904	179

Table 6.14: This Table enumerates the flow simulation results for the optimal well plan. The average BOEV exceeds the BOEV for the Asset Team well plan.

time A1 (*A* is intended to represent Asset Team) is reached, and Well A4 screens well A1 and produces much of the water that might have reached well the producer A1. A similar situation exists in the optimal well plan. Well O3, is screened by well 1-15. The difference with this situation is that well 1-15 is down dip of O3. Well 1-15 sees fast water breakthrough but due to the volume of water injected, and two adjacent producers pulling the water front, well 1-15 is bypassed by the water front and well O3 is serviced by the injector. Also, well A3 is 638m away from the injector and well O3 is 473m away. According to Figure 6.47, cumulative oil drops significantly beyond about 500m. It could be that well A3 was not solely intended as a producer and may have had the partial intent of being a step out well.

The cumulative oil production, water production, and water injected are shown in Figures 6.69, 6.70, 6.71, respectively. The cumulative for each of the fluids are higher in the optimal well locations versus the Asset Team selected locations. Figure 6.72 shows an increase at each time step for BOEV. This means that the optimal well plan outperforms the Asset Team selected locations.

Table 6.14 shows the numerical values for the cumulative fluids produced and injected for the last time step. Each parameter shows increases. Water injection at this field is relatively cheap, so, although there is an increase in water handling costs, there is added recovered hydrocarbon that might offset the extra costs.

Table 6.15 compares the average cumulative volumes for the Asset Team and optimal well locations. The last two rows of the Table show the increases in terms of percent and relative to the Asset Team well plan. There is an increase of more than 19% cumulative oil. Despite the increase in water handling, BOEV increased by more than 19%.

Figure 6.73 shows the cumulative distributions for FWPT, FOPT, FWIT, and BOEV. The functions are all shifted and reflect the increased performance from the optimal well

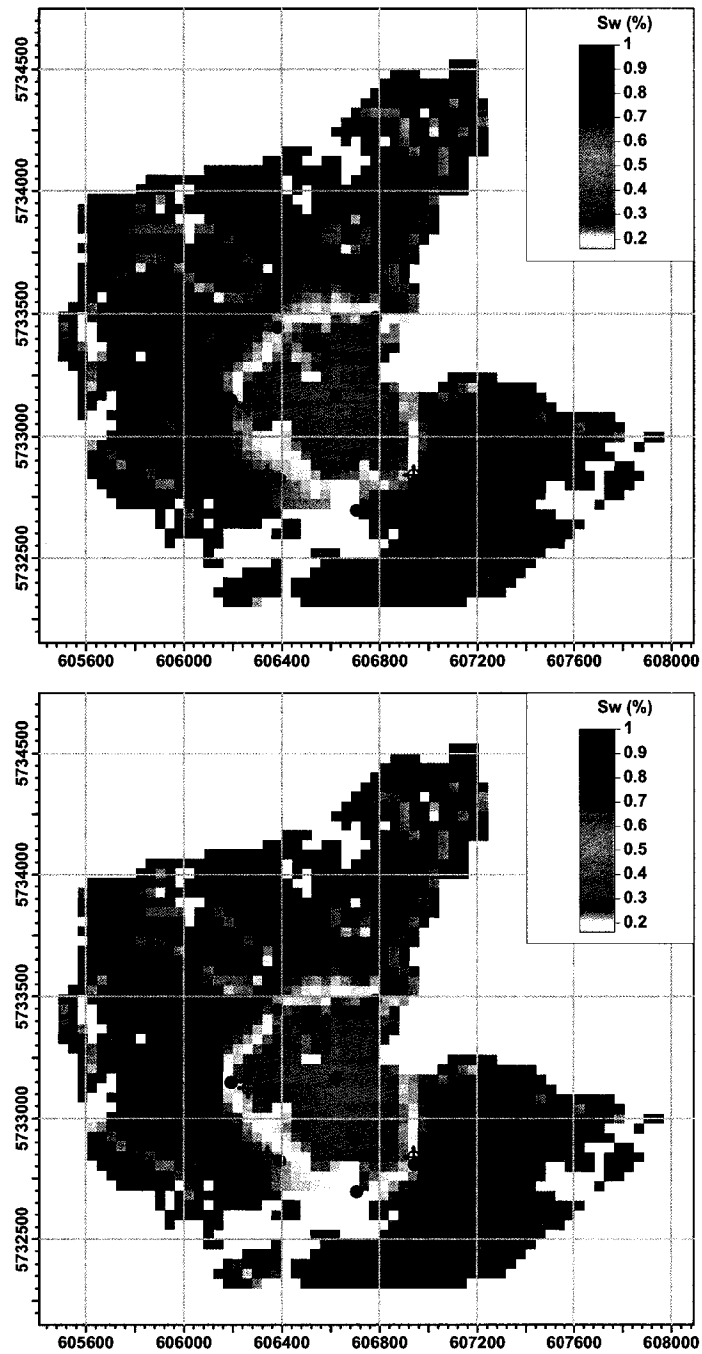


Figure 6.68: This map shows the sweep pattern for the Asset Team (top map) selected well plan and the optimal well locations. The Asset Team and optimal well locations are also shown.




Figure 6.69: This chart shows the field oil production versus time for the Asset Team Selected wells. The *x*-axis shows time and the *y*-axis shows cumulative barrels of oil.




Figure 6.70: This chart shows the field water production versus time for the Asset Team Selected wells. The *x*-axis shows time and the *y*-axis shows cumulative barrels of water produced.




Figure 6.71: This chart shows the field water injected versus time for the Asset Team Selected wells. The *x*-axis shows time and the *y*-axis shows cumulative barrels of water injected.




Figure 6.72: This chart shows the barrel oil equivalent value for the field under the Asset Team well plan. The *x*-axis shows time and the *y*-axis shows cumulative value.

	FWIT	FOPT	FWPT	BOEV
	(MBBL)	(MBBL)	(MBBL)	(MBBL)
Optimal Average	24538	1683	23269	1312
Optimal St. Dev.	3046	226	2904	179
Asset Average	20486	1413	19622	1101
Asset St. Dev.	2735	183	2644	142
Change Average	19.78	19.17	18.58	19.25
Change St. Dev.	11.41	24.04	9.87	19.19

Table 6.15: This Table compares the summary statistics for the Asset Team well plan and the optimal well plan. Despite an increase in water handling costs, there is almost 20% improvement in BOEV.

Asset Well	WOPT (MBBL)	WOPT (MBBL)	Optimal Well
01-15	260	238	01-15
09-15	115	136	09-15
10-15	228	275	10-15
A1	214	272	O1
A2	238	257	O2
A3	136	259	O3
A4	218	244	O4

Table 6.16: This Table compares the cumulative produced oil volumes for each producer for the Asset Team well plan and the optimal well plan

plan.

Source of Improved BOEV

The optimal well plan shows improved performance over the Asset Team well plan, but it would be interesting to know if the improvements are solely due to the change in location of wells A3 and O3. This Section discusses the cumulative volumes for all the producer wells.

Table 6.16 shows the cumulative oil volumes for all the producers. There is little change among the existing producers - just a little reorganization of oil among the wells. Most of the improvement does come from the new well locations and in particular A3 and O3 where the cumulative oil production is almost doubled.

Table 6.17 shows the cumulative water volumes for all the producers. For the most part all the wells in the optimal plan saw increased water production. The volume of injected water increased from 21016 BBL*1000 to 24711 BBL*1000, yet the total water produced increased from 19622 BBL*1000 to 23269 BBL*1000. This means that the Asset Team well plan produced 93.4% of the water injected and the optimal well plan produced 94.2% of the injected water so there is no change in the efficiency of the injection.

6.3.13 Comments

The optimal well plan outperforms the Asset Team well plan by about 19% in terms of barrel oil equivalent value. The water flood sweep was more efficient by observation in that all of the wells were in communication with the injector. There were only moderate changes in well location. This was mostly due to the limited available area for the selection of new well locations and that was a consequence of dissolution and collapse of material stratigraphically below the reservoir.

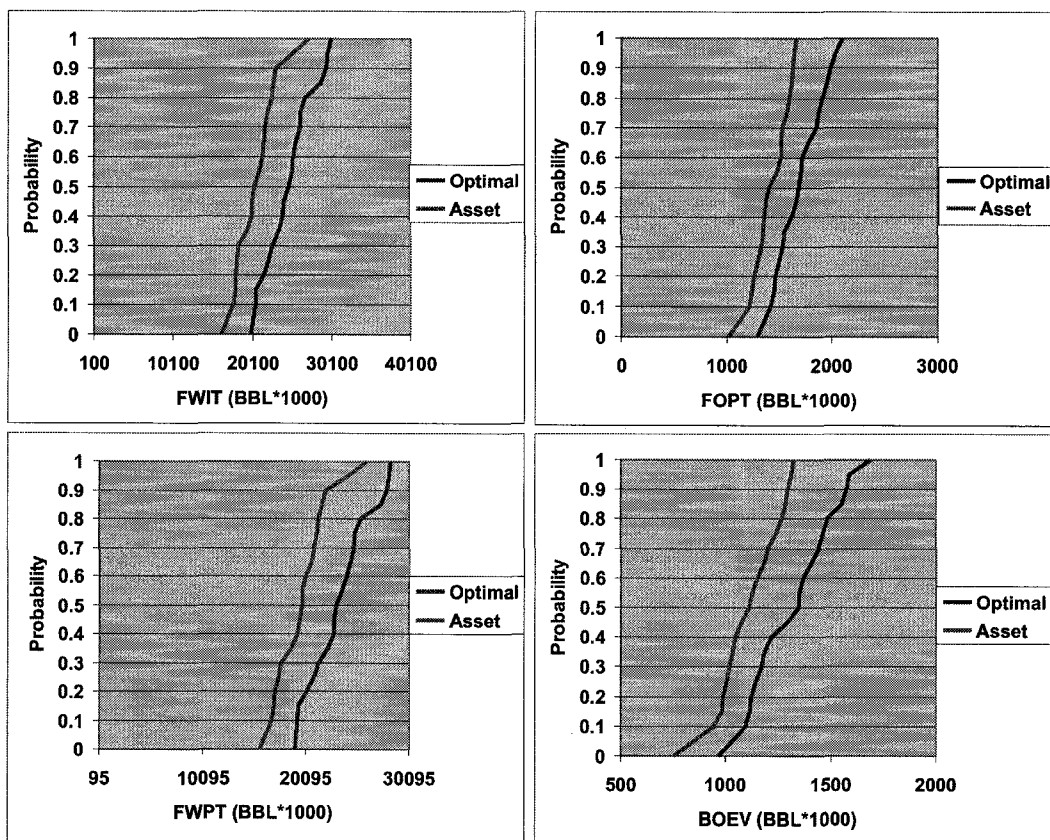


Figure 6.73: These charts are the cumulative distribution functions for the optimal well plan (dark line), and the Asset Team well plan (light line).

Asset Well	WWPT (MBBL)	WWPT (MBBL)	Optimal Well
01-15	7561	8365	01-15
09-15	2163	3183	09-15
10-15	2049	3209	10-15
A1	1541	2352	O1
A2	2194	2662	O2
A3	1153	1160	O3
A4	2957	2335	O4

Table 6.17: This Table compares the cumulative produced water volumes for each producer for the Asset Team well plan and the optimal well plan

Chapter 7

Concluding Comments

In the mining and petroleum industries, engineers make decision with the goal of maximizing value for the stakeholders. A common factor in selecting dig limits and well locations is subsurface uncertainty. Geostatistics can be used to construct models of uncertainty, but there are few tools for selecting dig limits or well locations given a model of uncertainty. Decision making under uncertainty is a challenging problem: the global objective function may not be amenable to many well known optimization techniques, the relationships between component objectives may be competitive and non-linear, and the solution space is large. The semi-automatic dig limit and well locations selection techniques developed in this thesis are effective in making optimal decisions in light of uncertainty. This Chapter provides summaries of the dig limit and well location selection algorithms and discusses limitations of the techniques.

7.1 Dig Limit Selection

The dig limit selection algorithm accounts for uncertainty and the limitations of mining equipment in the selection of dig limits. As input information, the algorithm requires a mineral model of uncertainty, a quantification of mill response to mineral grade and contaminant content, and an understanding of how the mining equipment responds to contours in the dig limits, or the digability of the dig limits. The procedure for applying the dig limit selection technique is as follows:

1. Construct models of uncertainty on each important mineral including contaminants. Any mine should have these models readily available.
2. Convert the mineral models of uncertainty to an expected profit map using the expected profit transform.
3. Select an initial dig limit.
4. Use the dig limit selection algorithm and a set of digability factors to build a digability catalog. Select the dig limits from the digability catalog that represent the selectivity of the mining equipment and survey these to the bench.

The dig limit selection algorithm is a semi-automatic technique because it requires an initial dig limit solution to start. The algorithm starts by measuring the performance of the initial dig limit, then the algorithm perturbs the dig limit until it conforms to a dig limit that strikes the best compromise between ore content and ease of mining for the mining equipment. A single perturbation consists of randomly drawing a vertex on the dig limit, randomly drawing a change in location for the vertex, measure the performance of the new dig limit, and accepting or rejecting the new dig limit according to the simulated annealing decision rule.

Experiments were conducted that examined the features of the dig limit selection algorithm. Three case studies were performed. The first case study selected dig limits on a real data set for a copper mine in Chile. The second compared semi-automatically selected dig limits to dig limits drawn by hand. The first comparison experiment involved a synthetic data set. Mining professionals were given a map of expected profit and the specifications of available mining equipment and asked to select the dig limits that would maximize profit. The third case study involved a real data set on a copper mine in the United States and compared dig limits drawn by the mining engineer to semi-automatically selected dig limits. In both comparison experiments, the semi-automatic dig limits improved the profit in the order of 1-1.5

7.1.1 Dig Limit Selection Application Comments and Limitations

The dig limit selection technique is a new approach for selecting dig limits. It requires collecting and processing of information that is not commonly practiced. For this reason, some comments on using the technique are provided below:

- Precise information is required from the mill on how it responds to changes in grade and contaminant concentration. This information is usually known by the mill operators and engineers. Integrating this information into the expected profit transform may be difficult. Aim to identify significant contributing factors and eliminate factors that have poorly understood or minimal effect on recovery.
- Characterizing the digability penalty function may be difficult. There is little or no documentation on how mining equipment performs under particular mining conditions. In this thesis it was assumed that mining efficiency varies smoothly from poor efficiency to high efficiency. This may not be the case. There may be a critical threshold where mining efficiency collapses to 0 rather than the assumed asymptotic approach to low efficiency. The only option is to observe and tabulate the performance of the mining equipment and construct mine specific digability penalty functions.
- Selecting a digability factor is challenging. The digability catalogue is a good solution, but it relies on a subjective understanding of the selectivity of the mining equipment. An iterative approach must be used where changes and differences in head grade and estimated head grade are documented and used to guide the selection of a good digability factor.

The dig limit selection algorithm is not recommended for setting long term production targets or long term planning. The equipment assumptions, i.e., the penalty function and selected digability factor used to select the dig limits may not be valid at the time of mining. Information content changes as mining proceeds and the new information may change the dig limits.

Semi-automatic dig limit selection is not recommended for recoverable reserves estimation. Dig limit selection requires assumptions about the mining equipment and the data that may not be valid by the time the mining process actually excavates the volume.

7.2 Well Location Selection

The well location selection algorithm integrates static and dynamic reservoir information to select optimal well locations. The selected locations are jointly optimal, that is, the locations consider all realizations in the model of uncertainty simultaneously in the search for the optimal well plan. As input information the well location selection algorithm requires, a reservoir model of uncertainty, a minimal set of flow simulation results, and a seed well plan. The procedural steps for applying the algorithm are as follows:

1. Construct the full field model of uncertainty and obtain flow simulation results for at least one realization.
2. Construct the mini model using the parameters used to construct the full field model.
3. Collect the static and dynamic information required for multiple regression of the component objective functions.
4. Build the map of azimuth of dips for calculating the structural penalty.
5. Use the full field model flow simulation results to calibrate the global objective function.
6. Use the seed well plan and the calibrated global objective function to select an optimal well plan.

The essential steps of the well location selection algorithm can be described as follows: 1) start with an initial guess at the well locations, 2) measure how this reference well plan performs using a global objective function, 3) propose a random change to the well plan and measure how the modified well plan performs using the global objective function, 4) use the simulated annealing decision rules to accept or reject the new well plan, 5) if the new plan is accepted, update the old plan with the well locations and go the first step, or if the new plan is rejected go to the first step. Each loop over the 5 steps is a single perturbation. The algorithm evaluates perturbs the well plans until a large number of perturbations yield no improvement, or some other stopping criteria are invoked.

The case studies provided interesting insight into the process of selecting well locations. The merging of optimal well plans does not yield an optimal well plan. In other words, the best approach is to select all well locations at the same time because incrementally adding wells to an existing well plan gives suboptimal results. The added information supplied by adding well locations incrementally to the well plan may not lead to superior well locations. The added well information only serves to update the near well bore region, a region unavailable for future locations due to well to well interactions. In summary the best approach is to select all well locations simultaneously and review the data as it arrives. If the new data changes the conceptual model or indicates a change is required in the modeling practice, update the model and select new locations in light of the incremental data.

7.2.1 Well Location Selection Application Comments and Limitations

The well location selection algorithm integrates static and dynamic information. The selection algorithm relies on a proxy to flow simulation for the selection of locations. Due to the complexities of fluid flow and of reservoir management, there are limitations to the semiautomatic well location selection algorithm:

- The global objective function is calibrated to the flow simulation results and is a proxy for flow simulation for quick evaluation of potential well plan. The final well plan may require slight modification to optimal results under flow simulation.
- The algorithm depends strongly on flow simulation results to obtain a global objective function. The algorithm could be adapted to work without flow simulation results, but there is no replacement for the flow equations and without the flow simulation results the global objective function is a significantly compromised proxy for flow simulation.
- The construction of the component objective functions depends on multiple regression and it may be that the collected static information will not facilitate the construction of a regressed function with acceptable fitment statistics. In this case one could use one of the distance and fluid production relationships directly instead of an unsatisfactory regression model.

- The well location selection algorithm is most easily applied to reservoirs that have seen little or no production. The need to incorporate multiple realizations that history match is the primary impediment here. The saturation functions change with respect to time and the only way to construct plausible realizations is to account for the physics of flow using a flow simulator.
- There is no facility for incorporating completion intervals. As it is currently set up the algorithm assumes that the whole interval is completed and this constraint was propagated to the flow simulation parameters. This may not be realistic and will certainly bias the calibration procedure because completion intervals are often selected specifically to enhance performance. Furthermore, different completion intervals can have significant effect on the behavior of the reservoir.
- The technique does not consider tubing effects or near well bore issues such as skin or formation damage.
- In the flow simulation part of the technique assumptions must be made about dynamic reservoir characteristics, bottom hole pressures for example. Thus the results can not be used to predict future well performance or for booking reserves.
- One must be cautious about data honoring in the reservoir model of uncertainty when selecting well locations. Data locations have the least uncertainty in the reservoir model of uncertainty. The ergodic fluctuations may cause regions without data control to look less appealing than regions with data control to global objective function. Thus the technique may have reduced effectiveness in reservoirs with many wells.
- The well location selection technique is labor intensive and for some projects it may present more work than is worthwhile. For example, it may not be worthwhile to apply the technique to a field where drilling is cheap, production is low, or there are already many wells in place.
- The algorithm only works for a specified number of wells of the same type (vertical, deviated, or segmented) of wells. The problem that might be raised is changing well types is that it would add a significant number of perturbations if a well is changed from a vertical to a segmented well, or from a segmented well to a vertical well. Also, changing well types might add significant noise to the global objective function and consequently increase the number of perturbations required for finding a global optimal solution.
- The algorithm is limited to a specified number of wells. The anticipated problem with allowing the algorithm to optimize the number of wells as well as the locations of the wells is noise or drastic changes in the global objective function. Simulated annealing is somewhat resistant to problems such as this, but it would add to the number of required perturbations and add to the complexity of the problem: should the algorithm add a vertical, deviated, or segmented well, where should the seed location be, should a well be converted from its current well type to another rather than adding a well? It can be readily seen that the problem could quickly assume huge proportions.

7.3 Future Work

7.3.1 Dig Limit Selection

This Section outlines in point form possible avenues for future research in dig limit selection:

- The semi-automatic dig limit selection algorithm would benefit from studies on equipment performance on dig limit tortuosity. Such studies would provide the information

required to avoid relying professional understandings of how equipment performs and eliminate the need to construct a digability catalog.

- Extensive experimentation was undertaken on selecting the initial dig limit. None of the algorithms developed in the course of the research were very good at it, particularly when the initial dig limit had to be tortuous. A technique for selecting good initial dig limits would eliminate the need for human involvement and make the algorithm fully automatic.
- While there are software packages that simulate the mining process, these rely on mining out the deposit on a block-by-block approach. Mining does not proceed block-by-block. Mining proceeds shovelful-by-shovelful. A mining simulator based on the shovelful-by-shovelful approach integrating fully automatic dig limit selection would be useful for planning.
- Increased selectivity may be made possible by integrating the grade control procedure with a global positioning system (GPS). For example, GPS sensors on the shovel boom communicating with a computer would be able to precisely locate where the shovel is digging and inform the operator about the optimal classification of the material. GPS connectivity with the haul trucks would specify the destination of its load. The key to success here is the accuracy of the dig limits and hence the information used to select them.
- Incorporating GPS into the dig limit selection algorithm may open up the option of in situ, or in pit, blending by keeping track of the grades excavated by a shovel with a GPS beacon on it and the estimated grade loaded on the haul truck.
- The initial and optimally selected dig limits represent a volume of material equivalent to the area inside the dig limit polygon multiplied by the height of the bench to be mined. The mining face will not be perpendicular because blasted material is unconsolidated and sloughing will occur. Some dilution occurs at the dig limits. The extent of dilution is subject to how it is mined. Because there is a human element to the dilution there really is no way to account for it in selecting dig limits. The algorithm for selecting dig limits does not explicitly account for the additional room that equipment may require to sloughing of material. The user is expected to account for sloughing by selecting appropriate parameters.

7.3.2 Well Location Selection

This Section outlines areas for possible future research involving well location selection.

- The well location selection algorithm is most easily applied to reservoirs that have never seen production. This is because the saturation functions, and possibly the petrophysical properties, change with respect to time. Much research has been undertaken to construct realizations that automatically history match. Application of the well location selection algorithm to a set of history matched realizations would be an important extension to this research.
- Finite element analysis flow simulators are computationally intensive. Streamline flow simulators are far less intensive. It might be possible to integrate a streamline flow simulator into the perturbation/global objective evaluation steps of the well location selection algorithm.
- The work included in this thesis covers applications in the hydrocarbon reservoirs. The same principles might be applicable to other types of reservoirs.

- Considered in this research was an application to a water flood. This is one of many different enhanced oil recovery (EOR) techniques. An important research avenue might be other EOR techniques.
- Considered in this research, was a water drive with water injection reservoir. This is but one type of drive mechanism. Some deep water reservoir are being produced under water, gas and compaction drive with water and gas injection. Application of the technique to such reservoir may provide important changes in development planning that increase productivity.
- This research did not include uncertainty in the dynamic parameters in the selection of the optimal well plan. A possible avenue for research might be the construction of a set of calibration weights that reflect the span of dynamic property uncertainties.

Bibliography

- [1] E. Aarts and J. Korst. *Simulated Annealing and Boltzmann Machines*. John Wiley & Sons, New York, 1989.
- [2] F. G. Alabert. The practice of fast conditional simulations through the LU decomposition of the covariance matrix. *Mathematical Geology*, 19(5):369–386, 1987.
- [3] F. G. Alabert and G. J. Massonnat. Heterogeneity in a complex turbiditic reservoir: Stochastic modelling of facies and petrophysical variability. In *65th Annual Technical Conference and Exhibition*, pages 775–790. Society of Petroleum Engineers, September 1990. SPE Paper Number 20604.
- [4] M. Armstrong. Improving the estimation and modeling of the variogram. In G. Verly et al., editors, *Geostatistics for natural resources characterization*, pages 1–20. Reidel, Dordrecht, Holland, 1984.
- [5] K. Badiozamani. Computer methods. In H.L. Hartman, editor, *SME Mining Handbook*, volume 1, pages 598–625. SME, Littleton, Colorado, 1992.
- [6] M. Barnes. *Computer Assisted Mineral Appraisal and Feasibility*. SME, New York, 1980.
- [7] R. J. Barnes. Teachers Aide: The variogram sill and sample variance. *Mathematical Geology*, 23(4):673–678, September. 1991.
- [8] A. Basu and L. N. Frazer. Rapid determination of the critical temperature in simulated annealing inversion. *Science*, 249:1409–1412, September 1990.
- [9] B. L. Beckner and X. Song. Field development planning using simulated annealing - optimal economic well scheduling and placement. In *1995 SPE Annual Technical Conference and Exhibition*, Dallas, TX, October 1995. Society of Petroleum Engineers. SPE Paper Number 30650.
- [10] S. Buckley and M. Leverett. Mechanism of fluid displacement in sands. *Transactions of The AIME*, 146:107, 1962.
- [11] J. P. Chilès and P. Delfiner. *Geostatistics Modeling Spatial Uncertainty*. John Wiley & Sons, Inc., New York, 1999.
- [12] P. da Cruz, R. Horne, and C. Deutsch. The quality map: A tool for reservoir uncertainty quantification and decision making. In *SPE Annual Technical Conference and Exhibition*, Houston, TX, October 1999. Society of Petroleum Engineers. SPE Paper Number 56578.
- [13] L. P. Dake. *The Practice of Reservoir Engineering*, volume 1. Elsevier, Amsterdam, ND, 4th edition, 2004.
- [14] M. David. *Geostatistical Ore Reserve Estimation*. Elsevier, Amsterdam, 1977.

- [15] C. Deutsch. Geostatistical methods for modeling earth sciences data. Unpublished MIN E 612 Course Notes, University of Alberta, 1999.
- [16] C. V. Deutsch. A Fortran 77 subroutine for determining the fractional area of rectangular grid blocks within a polygon. *Computers & Geosciences*, 16(3):379–384, 1990.
- [17] C. V. Deutsch. *Annealing Techniques Applied to Reservoir Modeling and the Integration of Geological and Engineering (Well Test) Data*. PhD thesis, Stanford University, Stanford, CA, 1992.
- [18] C. V. Deutsch. *Geostatistical Reservoir Modeling*. Oxford University Press, New York, 2002.
- [19] C. V. Deutsch, T. Faechner, and S. Zanon. Self-healing sequential gaussian simulation for integration of secondary data. In C. V. Deutsch, editor, *Third Annual Report of the Centre of Computational Geostatistics*. University of Alberta, 2001.
- [20] C. V. Deutsch and A. G. Journel. The application of simulated annealing to stochastic reservoir modeling. In *Report 4, Stanford Center for Reservoir Forecasting*, Stanford, CA, May 1991.
- [21] C. V. Deutsch and A. G. Journel. *GSLIB: Geostatistical Software Library and User's Guide*. Oxford University Press, New York, 1992.
- [22] C. V. Deutsch and A. G. Journel. *GSLIB: Geostatistical Software Library and User's Guide*. Oxford University Press, New York, 2nd edition, 1998.
- [23] C. V. Deutsch, E. Magri, and K. P. Norrena. Optimal grade control using geostatistics and economics: Methodology and examples. *SME Transactions*, 308, February 2000.
- [24] C. V. Deutsch and S. Srinivasan. Improved reservoir management through ranking stochastic reservoir models. In *SPE/DOE Tenth Symposium on Improved Oil Recovery, Tulsa, OK*, pages 105–113, Washington, DC, April 1996. Society of Petroleum Engineers, Society of Petroleum Engineers. SPE Paper Number 35411.
- [25] C. V. Deutsch and T. T. Tran. Fluvsim: A program for object-based stochastic modeling of fluvial depositional systems. *Computers & Geosciences*, 28(3):525–535, May 2002.
- [26] C. V. Deutsch and L. Wang. Hierarchical object-based stochastic modeling of fluvial reservoirs. *Math Geology*, 28(7):857–880, 1996.
- [27] H. Dykstra and R. Parsons. The prediction of oil recovery by water flood. *Secondary Recovery of Oil in the United States, 2nd ed.*, pages 160–174, February 1950.
- [28] C. Farmer. The generation of stochastic fields of reservoir parameters with specified geostatistical distributions. In S. Edwards and P. R. King, editors, *Mathematics in Oil Production*, pages 235–252. Clarendon Press, Oxford, 1988.
- [29] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-6(6):721–741, November 1984.
- [30] I. M. Glacken. Change of support by direct conditional block simulation. In *Fifth International Geostatistics Congress*, Wollongong, September 1996.
- [31] D. J. Goggin, J. Gidman, and S. E. Ross. Optimizing horizontal well locations using 2-d scaled-up geostatistical reservoir models. In *SPE Annual Conference and Exhibition, Dallas*. Society of Petroleum Engineers, October 1995. SPE Paper Number 30570.

- [32] P. Goovaerts. *Geostatistics for Natural Resources Evaluation*. Oxford University Press, New York, 1997.
- [33] E. Gringarten and C. V. Deutsch. Variogram interpretation and modeling. *Mathematical Geology*, 33(4):507–534, May 2001.
- [34] B. Gyagler. *Optimization of Well Placement and Assessment of Uncertainty*. PhD thesis, Stanford University, Stanford, California, 2002.
- [35] C. Hansen and J. Fanchi. Producer/injector ratio: The key to understanding pattern flow performance and optimizing waterflood design. *Society of Petroleum Engineering*, pages 317–327, 2003. SPE Paper Number 86574.
- [36] Richard F. Hartl. A global convergence proof for a class of genetic algorithms. Unpublished manuscript, Institute of Econometrics, Operations Research and Systems Theory, Vienna University of Technology, Argentinierstraße 8, A-1040 Vienna, Austria, December 1989. Presented at the ÖGOR annual conference, Dec. 6th, 1989.
- [37] J. H. Holland. *Adaptation in Natural and Artificial Systems*, volume 1. University of Michigan, University of Michigan, Michigan, 1975.
- [38] M. G. Ierapetritou, C. A. Floudas, S. Vasantharajan, and A. S. Cullick. A decomposition based approach for optimal location of vertical wells. *AIChE Journal*, 45:844–859, February 1999.
- [39] L. Ingber. Very fast simulated re-annealing. *Mathl. Comput. Modelling*, 12(8):967–973, 1989.
- [40] L. Ingber. Simulated annealing: practice versus theory. *Mathematical Computer Modelling*, 18(11):25–57, March 1993.
- [41] L. Ingber and B. Rosen. Genetic algorithms and very fast simulated reannealing: A comparison. *Oper. Res. Management Sci.*, 33(5):523, 1993.
- [42] E. H. Isaaks. *The Application of Monte Carlo Methods to the Analysis of Spatially Correlated Data*. PhD thesis, Stanford University, Stanford, CA, 1990.
- [43] E. H. Isaaks and R. M. Srivastava. *An Introduction to Applied Geostatistics*. Oxford University Press, New York, 1989.
- [44] A. G. Journel. Geostatistics for conditional simulation of orebodies. *Economic Geology*, 69:673–680, 1974.
- [45] A. G. Journel. The lognormal approach to predicting local distributions of selective mining unit grades. *Mathematical Geology*, 12(4):285–303, 1980.
- [46] A. G. Journel. *Fundamentals of Geostatistics in Five Lessons*. Volume 8 Short Course in Geology. American Geophysical Union, Washington, D. C., 1989.
- [47] A. G. Journel and F. G. Alabert. Focusing on spatial connectivity of extreme valued attributes: stochastic indicator models of reservoir heterogeneities. SPE Paper Number 18324, 1988.
- [48] A. G. Journel and F. G. Alabert. Non-Gaussian data expansion in the earth sciences. *Terra Nova*, 1:123–134, 1989.
- [49] A. G. Journel and F. G. Alabert. New method for reservoir mapping. *J. of Pet. Technology*, pages 212–218, February 1990.
- [50] A. G. Journel and Ch. J. Huijbregts. *Mining Geostatistics*. Academic Press, New York, 1978.

- [51] A. G. Journel and E. H. Isaaks. Conditional indicator simulation: Application to a Saskatchewan uranium deposit. *Mathematical Geology*, 16(7):685–718, 1984.
- [52] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. RC 9355, IBM Research Report, 1982.
- [53] S. Kirkpatrick, C. D. Gelatt Jr., and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.
- [54] L. Kornze, T. B. Faddies, J. C. Goodwin, and M. A. Bryant. Geology and geostatistics applied to grade control at the mercur gold mine, mercur utah. *Applied Mining Geology: Problems of Sampling and Grade Control*, pages 45–56, 1985.
- [55] O. Leuangthong and C. V. Deutsch. Stepwise conditional transform for removing spatial correlation. *Mathematical Geology*, 35(2):155–173, February 2003.
- [56] G. Matheron. *Traité de géostatistique appliquée*. Vol. 1 (1962), Vol. 2 (1963), ed. Technip, Paris, 1962.
- [57] G. Matheron. The intrinsic random functions and their applications. *Advances in Applied Probability*, 5:439–468, 1973.
- [58] N. Metropolis, A. Rosenbluth, M. Rosenbluth, A. Teller, and E. Teller. Equation of state calculations by fast computing machines. *J. Chem. Phys.*, 21(6):1087–1092, June 1953.
- [59] N. Metropolis and S. Ulam. The monte carlo method. *Journal of The American Statistical Association*, 44(247):335–341, September 1949.
- [60] R. Mohr, C. Estracda, and K. Norrena. 3D Geomodelling and Flow Simulation of the Late Devonian Bakken Formation in South-Central Saskatchewan - Case Study: Smiley Buffalo Heavy Oil Waterflood. In *SPE International Thermal Operations and Heavy Oil Symposium*, Calgary, Canada, November 2005. Society of Petroleum Engineers. SPE Paper Number 97733.
- [61] C. Murray. Indicator simulation of petrophysical rock types. In A. Soares, editor, *Geostatistics Tróia '92*, volume 1, pages 399–413. Kluwer, 1993.
- [62] A. Noble. *Ore Reserve/Resource Estimation*. Society for Mining, Metallurgy, and Exploration, Inc., Littleton, Colorado, 1992. SME Mining Handbook.
- [63] K. Norrena and C. V. Deutsch. Using the critical temperature to improve the speed of geostatistical applications of simulated annealing. In *Seventh International Geostatistics Congress*, Cape Town, South Africa, April 2000.
- [64] H. Omre. The variogram and its estimation. In G. Verly et al., editors, *Geostatistics for natural resources characterization*, volume 1, pages 107–125. Reidel, Dordrecht, Holland, 1984.
- [65] A. Ouenes, S. Bhagavan, P. H. Bunge, and B. J. Travis. Application of simulated annealing and other global optimization methods to reservoir description: Myths and realities. In *SPE 69th Annual Conference and Exhibition, New Orleans, LA*, pages 547–561, Washington, DC, September 1994. Society of Petroleum Engineers. SPE Paper Number 28415.
- [66] Y. Pan. Application of least squares and kriging in multivariate optimizations of field development scheduling and well placement design. *Reservoir Development and Design Optimization*, 11, 1995. SPE Paper Number 38895.

- [67] V. Pinto, W. Pottenger, and W. Thompkins. A survey of optimization techniques being used in the field. In *Third International Meeting for Research in Logistics*, Trois Riviers, 2000.
- [68] M. Prats, P. Aime, E. Hazebroek, and E. ALLEN. Effect of off-pattern wells on the performance of a five-spot water flood. *Society of Petroleum Engineering*, pages 173–178, 1962. SPE Paper Number 170.
- [69] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes*. Cambridge University Press, New York, 1986.
- [70] A. Richmond. Applying four different risk models in local ore selection. *Natural Resources Research*, 11(4):299–314, 2002.
- [71] A. Richmond. Financially efficient dig-line delineation incorporating equipment constraints and grade uncertainty. *Journal of Surface Mining, Reclamation and Environment*, 18(2):99–121, 2004.
- [72] G. W. Rosenwald and D. W. Green. Geology and geostatistics applied to grade control at the mercur gold mine, mercur utah. *Applied mining Geology: Problems of Sampling and Grade Control*, 1974.
- [73] O. Leuangthong S. Zanon, L. Cunha and C. V. Deutsch. A short note on the quantification of the effect of input variable uncertainty in oil reservoir performance prediction. Annual Report 6, University of Alberta, Centre of Computational Geostatistics, Edmonton Alberta, Canada, 2004.
- [74] N. Schofield, P. Hellman, N. Johnson, and R. Hague. An introduction to the optimization of grade control practices. Technical report, Hellman and Schofield Pty Ltd, 2001.
- [75] D. Seifert, J. J. M. Lewis, C. Y. Hern, and N. C. T. Steel. Well placement optimisation and risking using 3-d stochastic reservoir modelling techniques. In *SPE/NPF European Conference*, Stavanger, Norway, April 1995. Society of Petroleum Engineers. SPE Paper Number 35520.
- [76] R. M. Srivastava. Matheronian geostatistics: where is it going? In *Fifth International Geostatistics Congress*, Wollongong, September 1996.
- [77] R. M. Srivastava, D. R. Hartzell, and B. M. Davis. Enhanced metal recovery through improved grade control. In *23 Application of Computers and Operations Research in the Mineral Industry*, pages 243–249. SME, Tuscon, AZ, April 1992.
- [78] D. Steppan, W. Werner, and W. Yeater. *Essential Regeression and Experimental Design for Chemists and Engineers*, October.
- [79] W. Stiles. Use of permeability distribution in waterflood calculations. *Transactions of The AIME*, 186:9, 1949.
- [80] H. Szu and R. Hartley. Fast simulated annealing. *Physics Letters A*, 122(157):157–162, August 1987.
- [81] S. Vasantharajan and A. S. Cullick. Well site selection using integer programming. In V. Pawlowsky-Glahn, editor, *Proceedings of IAMG'97*, volume 1, pages 421–426. CIMNE, 1993.
- [82] V. Černý. Thermodynamical approach to the travelling salesman problem: an efficient simulation algorithm. *Journal of Optimization Theory and Applications*, 45:41–51, 1985.

- [83] Hans Wackernagel. *Multivariate Geostatistics*. Springer-Verlag, Berlin, 1995.
- [84] T. Wagenhofer and D. G. Hatzignatiou. Optimization of horizontal well placement. *SPEJ*, 1996. SPE Paper Number 35714.
- [85] B. Wang. Well site selection algorithm considering geological, economical and engineering constraints. Master's thesis, University of Alberta, Edmonton, Alberta, 2000.
- [86] Tong Xian-zhang. A comparative study of the characteristics and susceptibility of pattern-type water-injection well networks from the viewpoint of balanced waterfloods. *Society of Petroleum Engineering*, pages 893–900, 1983. SPE Paper Number 10567.