

**Improving Development Planning in a Natural Gas Storage Field
Through Simulation-Optimization:**

Uncertainty Analyses

February 7, 2001

Virginia M. Johnson
Lawrence Livermore National Laboratory

DISCLAIMER

Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract number W-7405-ENG-48.

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government thereof, and shall not be used for advertising or product endorsement purposes.

Table of Contents

Introduction	1
Modifications to the FORGAS/IMEX Model	1
Alternative Hypotheses for the Reservoir's Permeability	2
Expected vs Pessimistic Remediation Outcomes	4
Recreating the Knowledge Base of Simulations	5
Modifications to the Objective Function	5
Retraining the ANNs	5
Search Procedures	7
Optimized Solutions for Alternative Permeability Models	7
Optimized Solutions for Expected vs Pessimistic Remediation Outcomes	9
“Handicapping” New Well Locations	10
Other Issues	12
Conclusions	13
References	15

Tables and Figures

Table 1—FORGAS parameters controlling the operation of a fictitious production plant compressor	2
Table 2—Percentage improvement over baseline in contract shortfall produced by each remediation candidate under two outcomes	4
Table 3—Optimized development plans under three permeability models	8
Table 4—Optimized development plans under expected vs. pessimistic remediation outcomes	9
Table 5—Each new well’s mean distance from existing neighbors and individual impact on contract shortfall	10
Table 6—Top 15 development plans from two cost/risk/distance weighting schemes under two permeability models	11
Table 7—Optimized development plans obtained from ANN- vs. simulator-reliant searches under various conditions	13
Fig. 1—Natural log scaling of permeabilities in the gas storage field	3
Fig. 2—Relative locations of wells and prospective well locations	3
Fig. 3—Scatter plots of predicted vs. actual contract shortfall	6

Introduction

This is the second of two papers describing the application of simulation-optimization methods to a gas storage field development planning problem. The first paper¹ began by giving a detailed description of the field and earlier efforts to model the effects of selected field development options on the field's productivity. It then outlined the basic steps required to apply a combination of artificial neural networks (ANNs) and the genetic algorithm (GA) to explore a much larger universe of field development planning options. Familiarity with the contents of the first paper is a prerequisite for understanding the material presented in this second paper.

The optimized solutions to the planning problem presented in the first paper were based on a deterministic, "best guess" view of the field's reservoir properties. However, practical field development planning dictates that at least some of the uncertainties associated with these properties be taken into account. This second paper describes procedures and presents results showing how the ANN-GA approach to optimization can be extended to accommodate three sources of uncertainty pertinent to the field being studied:

1. Alternative hypotheses regarding the permeabilities in a key region of the field
2. Uncertainty regarding the likely success of remediating existing wells
3. Risks associated with siting new wells in relatively unknown regions of the field

The first two sources involve physical properties (permeabilities and skin factors, respectively) that are embedded in the simulation of the reservoir response and, therefore, require substantial changes to the knowledge base of simulations. The third source of uncertainty is examined simply by making changes to the objective functions driving the optimization.

To streamline the presentation of results, only one of the two injection/withdrawal scenarios described in the first paper is included here, namely the 30-day peak service at the 2.5 Bcf baseload.

Modifications to the FORGAS/IMEX Model

Prior to undertaking the uncertainty analyses, two alterations to the modeling of operations in the field were made. These changes were unrelated to the uncertainty issues and were intended to improve the realism of the simulations.

The first of these changes concerned the manner in which the impact of low wellhead pressures was accounted for. The first version of the model contained no simulation of compression facilities. Any flowing wellhead pressures below 70 psia occurring anytime during the period of simulation were simply recorded and subsequently included in the objective functions as a cost of doing business. This was unsatisfactory because converting low pressure conditions to dollar costs was an arbitrary process and because it failed to account for real engineering constraints such as minimum tolerable suction pressure. In the absence of plant data on which to base a history match, a fictitious compressor was modeled within FORGAS to govern withdrawal operations. The parameters selected to simulate the compressor were judged to be reasonable approximations of operating conditions at the actual plant:

Table 1—FORGAS Parameters Controlling the Operation of a Fictitious Production Plant Compressor

Maximum capacity	40,000.00 bhp
Delivery pressure	130.00 bhp
Initial power installed	0.00 bhp
Maximum added compression	40,000.00 bhp
Compressor efficiency	0.85
Power per unit	1.00 bhp
Fuel per unit	0.26 Mcfd
Number of compression stages	1
Minimum suction pressure	50.00 bhp

In the withdrawal scenario used for this study, compression is only required during the peak service period itself. For the baseline no-action case, fuel usage, i.e. the volume of withdrawn gas which must be diverted to power the compressor, jumped as high as 436.5 Mcfd before the compressor reached capacity. Because of this overhead, adding compression facilities to the model increased the size of the cumulative contract shortfall on the baseline case from 0.2551 Bcf to 0.3579 Bcf, clearly making the field development planning problem more challenging than before.

A second alteration also made the management problem more difficult to solve. In the earlier model, the effect of drilling a new vertical well or remediating an existing well was simulated by changing the skin factor of the affected well to 2.0. This value is reasonable for new wells but was judged to be too optimistic in the case of well remediations. Consequently, a skin factor of 5.0 was employed as a more realistic expected outcome of remediation. By reducing the effectiveness of well remediation, it became possible for the more costly but also more effective strategy of drilling new wells to appear in the optimal solutions.

Alternative Hypotheses for the Reservoir’s Permeability

Figure 1 illustrates the permeabilities associated with each grid block of the IMEX model of the field, overlaid with white squares to indicate the location of existing wells. The field is divided into a main section in the north, west, and south and a second, smaller section well to the east. The low-permeability (0.1 md) region separating the sections is indicated by three connected rectangles. If flow in this middle region were, in fact, higher than is currently thought, the 30-day peak service scenario might be achieved at considerably less cost. To estimate if it is worth going to the expense of gathering further data in this region, two alternatives to this most conservative permeability model were created. These alternatives simply consisted of setting the permeabilities in the three-rectangle region to 1.0 md, creating the “medium-permeability” model, and 100.0 md, to create the “high-permeability” model. An idea of how the increase in flow over this region affects the management problem can be seen from the improvement in contact shortfalls for the baseline no-action case: 0.3579 Bcf for the original low-permeability model, 0.2378 Bcf for medium permeability, and 0.1048 Bcf for high permeability.

Selecting two alternative points from the permeability continuum and assigning them, as a block, to the entire region is obviously not representative of the geostatistical universe of possibilities for that region. However, the goal of this analysis is to pick reasonable alternative views of the reservoir, re-conduct the

optimization based on those alternatives, and decide if the improvements in the cost of field development justify the expense of drilling observation wells to test the validity of the alternatives.

Figure 2 presents the relative locations of the existing wells and the various planning options (vertical well remediations, drilling new vertical wells, and adding injection capability to the existing horizontal well), that were selected for study and which are described in detail in the first paper. Increasing the permeability of the questionable region opens up the possibility of siting new vertical wells in that region. However, any solutions that included wells in the questionable region would only apply to the higher-permeability models of the field. A more practical approach is to attempt to identify optimized solutions that hold across all three permeability models.




Fig. 1—Natural log scaling of permeabilities (in md) in the gas storage field. White squares indicate the locations of existing wells.

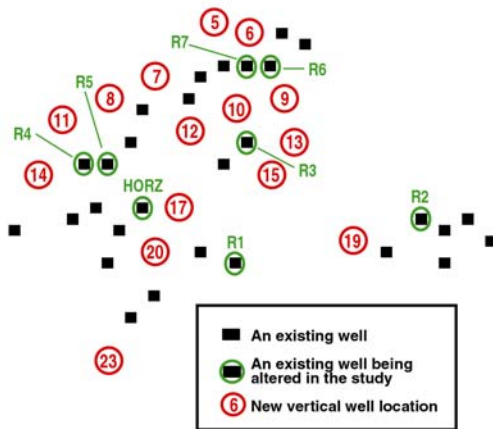


Fig. 2—Relative locations of wells and prospective well locations being manipulated in the optimization study.

Expected vs Pessimistic Remediation Outcomes

As was mentioned above, the effectiveness of remediation was implemented by setting the skin factor of the remediated well to 5.0. This is referred to as the “expected-outcome” model of remediation because the value 5.0 seemed a reasonable expectation for remediation. Since there exists some unknown likelihood that the remediation will not produce such a favorable reduction in skin factor, some procedure was required to account for the impact of unfavorable remediation outcomes on the optimized solutions.

The first plan to incorporate degrees of remediation success involved multiple simulations of the impact of each candidate well based on a Monte Carlo sampling over the distribution of likely skin factors for that well. In the revised problem formulation, a remediation candidate well would not simply be included or not included in the field development plan but would be included with some expectation of success. If the reservoir’s response to each well were linear, the construction and evaluation of such a response function would be straightforward. However, in a system where each well’s impact is affected by the presence or absence of other wells in the development plan, the amount of simulation required to build the response function becomes prohibitive.

Instead, the same reasoning that guided the selection of alternative hypotheses for the permeabilities problem was applied to the remediation outcomes problem. For decision-making purposes, it is probably sufficient to pick one or more contrasting views of the outcomes to determine 1) how sensitive the optimized solutions are to variation in the remediation outcomes, and 2) how much additional cost is required to accommodate a more pessimistic set of outcomes. To compare with the “expected-outcome” model described earlier, a “pessimistic-outcome” model was created. In this model, when a remediation candidate was included in the field development plan, its skin factor was set to 50% of its unremediated skin factor. For example, R1’s unremediated skin factor was 40.0. When selected for inclusion in a development plan, its skin factor was set to 20.0. Since the same rule applied to any and all remediation candidates, this model truly represents the case in which all attempted remediations have relatively unsuccessful outcomes.

An idea of the impact this change in skin factors has on reservoir response can be seen in the table 2.

Table 2—Percentage Improvement Over Baseline in Contract Shortfall
Produced by Each Candidate Under Two Remediation Outcomes
(Given the low-permeability model)

Well	Original Skin	Improvement in Contract Shortfall	
		Expected Outcome	Pessimistic Outcome
R1	40.0	14.0%	4.4%
R2	20.0	4.1%	2.5%
R3	20.0	11.2%	5.7%
R4	40.0	14.6%	5.0%
R5	31.0	14.4%	6.9%
R6	37.6	9.2%	1.8%
R7	31.0	12.5%	4.5%

Recreating the Knowledge Base of Simulations

Because of the changes made to incorporate compression into the model of the field, the entire knowledge base was recreated from scratch. This involved generating a set of 524 field development plans according to the same set of rules described in the first paper. The set consisted of:

1. The no-action case
2. 23 single-action plans (one for each of the 23 possible development options)
3. 500 multiple-action plans

Each plan underwent four separate simulations, all involving the 30-peak service at 2.5 Bcf baseload scenario but varying the permeabilities and remediation-outcome conditions, as follows:

1. Expected-remediation outcomes with the original low-permeability reservoir model
2. Expected-remediation outcomes with the medium-permeability reservoir model
3. Expected-remediation outcomes with the high-permeability reservoir model
4. Pessimistic-remediation outcomes with the low-permeability reservoir model

Since only one injection/withdrawal scenario was being considered, the entire knowledge base was recreated in about 100 hours of CPU time on a dedicated Pentium III processor.

Modifications to the Objective Functions

The analyses involving alternative permeabilities and remediation outcomes do not, in themselves, require any changes to the objective function used to evaluate the effectiveness of a given field development plan. However, the addition of compression facilities to the simulation made the old method of assigning a penalty for low pressure conditions at the wellhead obsolete. With compression included in the model, any contract shortfall displayed by a development plan already includes most of the economic penalty (i.e. withdrawn gas having to be diverted to power the compressor rather than being available for delivery to the customer) associated with low pressures. There may also be a certain amount of wear on equipment which could be added as an additional penalty. But since compression is only required for 30 days out of the withdrawal season and the period over which the optimization exercise extends is only three years, this cost of doing business was considered too trivial to include in the objective function. The general objective function continues to be defined as follows:

$$\text{Total} = w_1\$Cost + w_2\text{Risk}\% \quad (1)$$

The only change is that $\$Cost$ no longer contains the low pressure conditions penalty described in the first paper. To simplify the presentation of results, only one set of weights ($w_1 = 0.33$, $w_2 = 0.67$) was applied. A major advantage of emphasizing risk-reduction is that the objective function is less sensitive to variations in the market price of natural gas which the company may have to purchase to make up shortfalls.

Retraining the ANNs

Under the modified objective function, only one attribute – contract shortfall – needed to be estimated. However, the attribute could be expected to behave differently under the four different reservoir models. So, four ANNs were trained and tested from the recreated knowledge base of simulations, using the

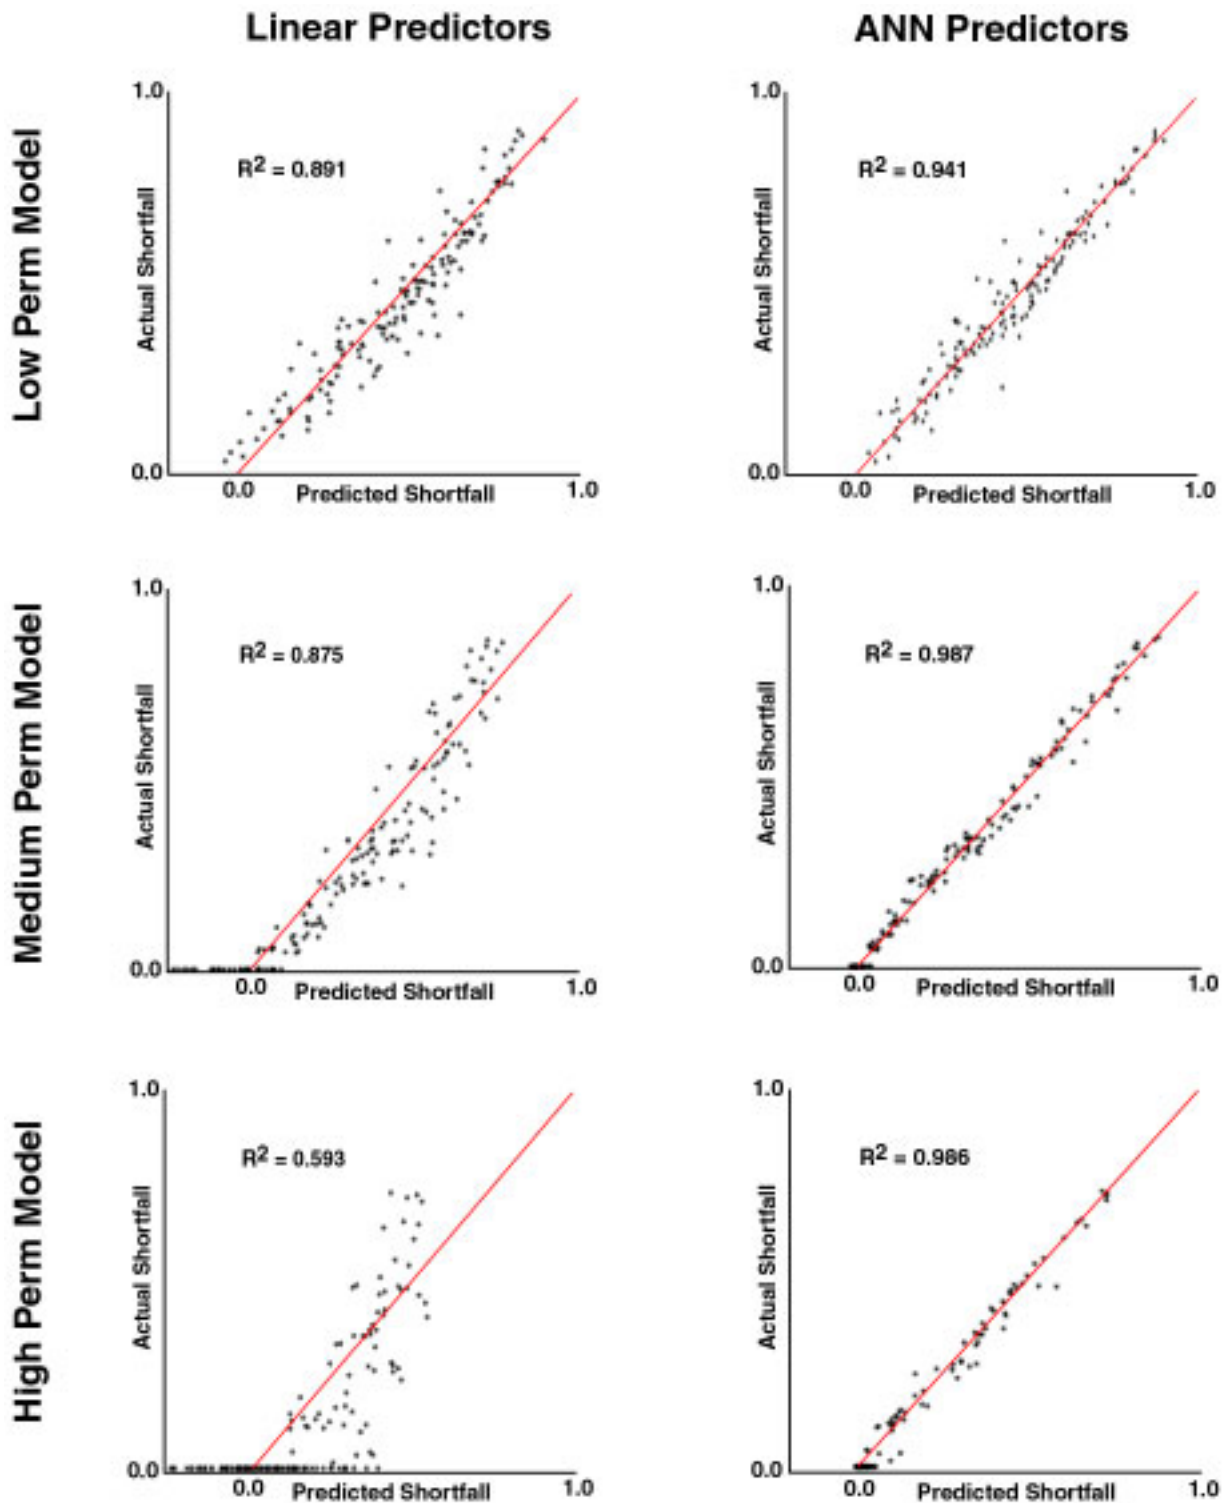


Fig. 3—Scatter plots of predicted vs. actual contract shortfall comparing linear and ANN predictors over the three permeability models of the reservoir; shortfalls are scaled from 0.0 – 1.0.

same procedures described in the first paper. Of the 524 cases in the knowledge base, 150 of the multiple-action cases were randomly selected for the test set, leaving a training set size of 374. Given the nonlinearities potentially introduced by adding compression facilities to the model, it was expected that these training and testing set sizes would be insufficient to capture the underlying relationships. However, the test results indicated otherwise. The scatter plots in Fig. 3 are intended to illustrate the apparent ease with which the ANNs are estimating contract shortfall with compression under the three different permeability models. Scatterplots for the corresponding linear predictors are also included.

This first item of interest in the figure concerns the accuracies for the low permeability model. Despite any complexities introduced by the addition of compression, both the linear predictor and the ANN yield high predictive accuracies (R^2 's of 0.891 and 0.941, respectively) on the 150-case test set. Visual examination of the plots also suggests that accuracy is good at the low end of the scale, where the search for optimal field development plans will be concentrated, as elsewhere.

Consistent with the notion that increasing flow in the disputed section makes the withdrawal contract easier to meet (and, consequently, easier to predict), the ANNs' predictive accuracies on the medium- and high-permeability models increase to 0.987 and 0.986, respectively. Note, however, that this is not the case for the linear predictors, whose R^2 's drop to 0.875 and 0.593. Inspection of the plots clearly shows the problem encountered with the linear predictors. In the low permeability model, no plans in the test set are able to meet the withdrawal contract three years in a row, thereby reducing contract shortfall to 0.0. In the higher permeability models, however, there are several cases in the test sets that reduce contract shortfall to 0.0. The ANNs correctly capture this bottoming-out phenomenon but the linear predictors generate negative estimates of contract shortfall.

For space reasons, the scatter plots for the linear and ANN predictors of contract shortfall under the pessimistic remediation outcome were excluded from the figure. The R^2 's associated with these predictors are 0.950 for the linear predictor and 0.972 for the ANN.

Search Procedures

The same genetic algorithm (GA) implementation details described in the first paper were also employed for the uncertainty analyses. It will be important to bear in mind that when the results of optimized searches based on ANN predictions of contract shortfall are presented, the tables display the scores calculated from the simulators' verification of the optimal solutions, not the original scores calculated from the ANN predictions.

Optimized Solutions for Alternative Permeability Models

Table 3 presents the three lowest-scoring field development plans obtained by two search techniques (simply querying the cases in the knowledge base and GA search employing ANNs as predictors) under the three permeability models. The scores associated with the baseline no-action case under each model are also included for reference. Where remediation candidates appear in the solutions, it can be assumed that scores reflect the expected remediation outcome (a skin factor of 5.0). Some notable features of these results are as follows:

1. Although the expected outcome of remediation is less optimistic than in the first paper, the cheapness of remediation relative to drilling new wells combined with the fact that all the

remediation candidates are sited in high-permeability areas means that remediation options continue to dominate the solutions.

2. Even so, we do observe that the minimization of risk necessitates the drilling of two new wells, at least under the original, low-permeability model.
3. The effect of increasing permeability in the disputed section is, as anticipated, to make the management problem cheaper to solve: the need for expensive new wells disappears.
4. As the management problem becomes easier, there is also less difference between the lowest-scoring cases in the knowledge base and the optimized solutions. In fact, under the high permeability model, no optimization was conducted because there were already low-cost cases from the knowledge base capable of meeting the withdrawal contract.

Table 3. Optimized Development Plans Under Three Permeability Models
(Risk-emphasized Objective Function)

Permeability Model	Search Method	Total Score	Risk%	\$Cost	Development Plan
Low	--	0.7972	100.0	715,800	No-action case
	Knowledge Base Query	0.2205	6.0	1,016,600	R1, R3, R4, R6, R7, NV06, NV13, NV14
		0.2388	0.8	1,315,400	R1 - R7, NV05, NV06, NV14, NV19
		0.2402	1.2	1,305,800	R1, R3 - R7, NV13, NV15, NV19, NV20, IHOR
	Optimization/ ANNs	0.1296	0.7	703,000	R1 - R7, NV06, NV20
		0.1403	2.4	700,000	R1, R3 - R7, NV06, NV13
		0.1774	6.6	747,600	R1 - R7, NV13, NV20, IHOR
Medium	--	0.7673	100.0	475,600	No-action case
	Knowledge Base Query	0.1758	21.4	157,000	R1, R2, R5, R7
		0.1992	0.0	974,000	R1, R3, R4, R6, R7, NV06, NV13, NV14
		0.2020	1.1	950,400	R1, R6, NV13, NV17, NV20, IHOR
	Optimization/ ANNs	0.0176	0.0	86,000	R1 - R7
		0.0322	1.9	96,800	R1 - R7, IHOR
		0.0744	0.0	364,000	R1, R3 - R6, NV13, IHOR
High	--	0.7207	100.0	209,600	No-action case
	Knowledge Base Query	0.0097	0.0	40,000	R1, R5, R7
		0.0133	0.0	55,000	R1, R2, R5, R7
		0.0789	0.0	326,000	R1, R4, NV13

From the standpoint of practical decision-making, a plan of action to be constructed from Table 3 might look as follows:

1. Remediate wells R1, R5, and R7 under the high permeability assumption. It does not matter if this permeability model is too optimistic because the best solutions found under the lower permeability models include the remediation of these wells.

2. If step 1 fails to produce the desired performance results, add the rest of the remediation candidates to the plan, as indicated by the optimal solution under the medium permeability model. Again, nothing is lost if the medium permeability model is also too optimistic.
3. Only if the field still does not perform as desired would drilling of two new wells (NV06 for certain, paired with either NV20 or perhaps NV13) be indicated.

Optimized Solutions for Expected vs Pessimistic Remediation Outcomes

Table 4 contrasts the solutions obtained when the outcomes of remediation are expected (resulting in skin factors of 5.0) vs. when they are poor (achieving skin factors which represent only a 50% improvement). In both cases, the low permeability model of the reservoir is in force.

Table 4. Optimized Development Plans Under Expected vs. Pessimistic Remediation Outcomes
(Risk-emphasized Objective Function)

Remediation Outcome	Search Method	Total Score	Risk%	\$Cost	Development Plan
Expected ^a	Knowledge Base Query	0.2205	6.0	1,016,600	R1, R3, R4, R6, R7, NV06, NV13, NV14
		0.2388	0.8	1,315,400	R1 - R7, NV05, NV06, NV14, NV19
		0.2402	1.2	1,305,800	R1, R3 - R7, NV13, NV15, NV19, NV20, IHOR
	Optimization/ ANNs	0.1296	0.7	703,000	R1 - R7, NV06, NV20
		0.1403	2.4	700,000	R1, R3 - R7, NV06, NV13
		0.1774	6.6	747,600	R1 - R7, NV13, NV20, IHOR
Pessimistic	Knowledge Base Query	0.4232	24.4	1,462,000	R1, R3, R5 - R7, NV13, NV17, NV20, NV23, IHOR
		0.4255	24.5	1,472,400	R1, R3 - R7, NV13, NV15, NV19, NV20, IHOR
		0.4279	24.8	1,474,600	R1 - R6, NV12, NV13, NV17, NV20, IHOR
	Optimization/ ANNs	0.3735	24.5	1,179,400	R1 - R7, NV13, NV14, NV20
		0.3784	25.5	1,171,200	R1 - R5, R7, NV13, NV14, NV20
		0.3806	25.4	1,185,800	R1 - R7, NV13, NV17, NV20

^a These are the same results as presented in the Low Permeability section of Table 3

The pessimistic model is a fairly extreme one since it assumes that all the attempted remediations will be relatively ineffective. However, it is still possible to apply reasoning similar to that outlined for the three permeability models. In fact, the optimized solutions for the pessimistic outcome assumption can be incorporated into the plan of action as follows:

4. If the results of step 3 fail to produce the desired field performance, drill a third new vertical well at location NV14.

A flaw in this plan is that the solutions under the pessimistic assumption do not include NV06, which seems to be the preferred location under the expected outcome assumption. A compromise might consist of drilling NV13 and NV20 at step 3, rather than coupling NV06 with either NV13 or NV20.

A further concern is raised by the magnitude of the scores under the pessimistic outcome assumption. Note that it appears impossible to reduce the Risk% values beyond 24.4% and that optimization

produces only a small improvement over the cases found in the knowledge base. This may indicate that the management problem is not really solvable under these most pessimistic assumptions. If the decision-makers have reason to believe that the pessimistic assumption is realistic, these results might indicate that no attempt should be made to upgrade the field to meet 30-day peak service demands at the 2.5 Bcf baseload.

“Handicapping” New Well Locations

Table 5. Each New Well’s Mean Distance from Existing Neighbors and Individual Impact on Contract Shortfall Under Two Permeability Models			
Location	Mean Distance (in ft.)	Improvement Over Baseline	
		Low Permeability	Medium Permeability
NV06	1601	9.9%	10.3%
NV07	1795	1.3%	1.3%
NV10	1901	4.2%	6.5%
NV20	1934	21.5%	27.4%
NV12	1952	2.7%	3.4%
NV09	2067	3.7%	7.4%
NV17	2143	14.8%	15.6%
NV08	2272	1.1%	1.2%
NV05	2370	2.7%	3.8%
NV14	2418	15.9%	15.4%
NV15	2712	16.9%	20.9%
NV11	2742	5.0%	4.8%
NV19	2925	7.0%	8.9%
NV13	2934	31.9%	37.0%
NV23	3954	5.9%	4.8%

The third source of uncertainty being addressed in this paper concerns decreasing confidence in the permeabilities assigned to grid blocks as their distance from known wells increases. Rather than taking the approach applied for disputed permeabilities in the mid-section, which involved creating alternate knowledge bases to reflect different views of reality, this form of uncertainty was handled by simply adding a distance penalty term to the objective function:

$$\text{Total} = w_1\$Cost + w_2Risk\% + w_3Distance \quad (2)$$

where Distance was merely the sum, over each new well location in the field development plan, of the mean distances in table 5. This sum was then normalized between 0.0 and 1.0, with the upper bound determined by the largest such sum in the cases in the knowledge base. Since the remediation candidates and the injection facility for the horizontal well would always be contributing a distance of 0.0 to this term, this method by itself would greatly favor development plans consisting only of those

particular options. To counteract this bias, the scaled Distance term of any plan with a distance of 0.0 was set to the scaled value of its Risk% term. This was effectively the same as applying the following function to these cases:

$$\text{Total} = w_1\$Cost + (w_2+w_3)Risk\% \quad (3)$$

The mean distances in Table 5 were calculated by identifying the three nearest neighbors of each target well and averaging their distance from the target well. Table 5 also includes the individual impact on contract shortfall of each new well location under the lower two permeability models. This information will be needed to interpret the optimized solutions.

Table 6. Top 15 Development Plans Obtained From Two Cost/Risk/Distance Weighting Schemes Under Two Permeability Models

Low Permeability Model

W_{Cost}	: 0.1	W_{Cost}	: 0.1
W_{Risk}	: 0.9	W_{Risk}	: 0.6
$W_{Distance}$: 0.0	$W_{Distance}$: 0.3
Existing Wells	New Locations	Existing Wells	New Locations
R1, R2, R4 – R7	NV06, NV13, NV20	R1 – R7	NV06
R1 – R7	NV09, NV13, NV20	R1, R2, R4 – R7	NV06
R1 – R7	NV06, NV13, NV20	R1 – R7	NV06, NV20
R1 – R7	NV06, NV13, NV17	R1, R3 – R7	NV06, NV20
R1 – R7	NV06, NV13, NV19, NV20	R1 – R7	NV06, NV17
R1 – R7	NV06, NV13, NV14, NV17	R1, R2, R4 – R7	NV06, NV20
R1 – R7	NV06, NV10, NV13, NV20	R1 – R4, R6, R7	NV06
R1 – R7	NV06, NV13, NV15, NV20	R1 – R4, R6, R7	NV06, NV20
R1 – R7	NV06, NV09, NV13, NV17	R1, R3 – R7, IHOR	NV20
R1, R3 - R7	NV06, NV13, NV19, NV20	R1 – R4, R6, R7, IHOR	NV06, NV20
R1 – R7	NV06, NV13, NV15	R1 – R4, R6, R7	NV01, NV13
R1, R2, R4 – R7	NV06, NV13, NV15, NV20	R1 – R7	NV20
R1 – R7	NV06, NV15, NV17, NV20	R1 – R7, IHOR	NV13
R1 – R7	NV06, NV12, NV13, NV20	R1 – R7	NV09, NV20
R1 – R7	NV10, NV13, NV20	R1 – R4, R6, R7	NV20

Medium Permeability Model

W_{Cost}	: 0.1	W_{Cost}	: 0.1
W_{Risk}	: 0.9	W_{Risk}	: 0.6
$W_{Distance}$: 0.0	$W_{Distance}$: 0.3
Existing Wells	New Locations	Existing Wells	New Locations
R1 – R7, IHOR	--	R1 – R5, R7	NV06
R1, R3 – R5, R7, IHOR	NV13	R1, R2, R4 – R7, IHOR	NV06
R1, R2, R4, R5, R7, IHOR	NV13	R1 – R7	NV06
R1, R4 – R7, IHOR	NV13	R1 – R7, IHOR	NV06
R1 – R5, R7	NV13	R1 – R5, R7, IHOR	NV06
R1 – R5, R7, IHOR	NV13	R1 – R7	NV07
R1, R3 – R7, IHOR	NV13	R1 – R7, IHOR	NV07
R1 – R6, IHOR	NV13	R1, R3 – R5, R7, IHOR	NV20
R1 – R4, R6, R7, IHOR	NV13	R1 – R3, R5, R7, IHOR	NV20
R1 – R5, R7, IHOR	NV20	R1, R3, R5 – R7, IHOR	NV20
R1, R3 – R7, IHOR	NV20	R1, R4 – R7	NV20
R1 – R6, IHOR	NV20	R1, R4 – R7, IHOR	NV20
R1 – R4, R6, R7, IHOR	NV20	R1 – R7	NV10
R1 – R3, R5 – R7, IHOR	NV13	R1, R3 – R7	NV20
R1 – R3, R5 – R7, IHOR	NV20	R1 – R4, R6, R7	NV20

Table 6 presents sets of optimized solutions obtained by applying two different sets of weights to the three terms in the objective function. To simplify the table, specific scores associated with each plan have been suppressed. Instead, the components of each plan are divided into existing well and new well locations. Optimized solutions obtained when the 0.1/0.9/0.0 set of weights is applied to the terms in the objective function are given on the left side of the table. These results represent the case where distance is not a factor for any plan. The extreme emphasis on Risk% over \$Cost is needed to force the

appearance of new wells in the solutions. Any significant weight associated with \$Cost produces the usual slate of solutions in which the remediation candidates are dominant. \$Cost requires at least a small positive weight, however, to avoid the appearance of ludicrously expensive plans that would never be seriously considered. Optimized solutions obtained when some weight (0.3) is given to the distance term appear on the right side of the table. The most useful way to interpret the table is to determine if there is a change in the frequency of occurrence of a particular new well location between the left- and right-side solutions.

Beginning with the solutions obtained under the low permeability model, NV06, NV13, and NV20 are dominant in the solutions in which distance is not a factor. When distance contributes to the objective function, NV13's frequency of occurrence is greatly reduced. This is consistent with the mean distances associated with each of these wells in Table 5. Although it is clearly a very influential location – single-handedly producing a 31.9% improvement in contract shortfall – NV13 is the second to most distant well in the set of new locations. In contrast, NV06 and NV20 remain in the solutions because they are much closer to existing wells.

The effect of including a distance penalty is even clearer in the medium permeability model. When distance is not factored in, only NV13 and NV20 appear in the optimized solutions. Apparently, increasing flow in the mid-section of the field makes it possible to meet the withdrawal contract without assistance from new wells in the northern portion of the field, which is where NV06 is sited. When distance becomes an issue, however, NV13 drops out of the solution set, NV20 is retained, and NV06 makes an appearance.

Unlike the solutions obtained when alternative remediation outcomes and mid-section permeabilities are contrasted with each other, the results in Table 6 do not represent different realities that must be somehow reconciled with each other to develop a plan of action. Rather, they simply illustrate the trade-offs decision-makers may have to confront, given their tolerance for unknowns. Drilling at NV20 requires no trade-off because it is both an influential well and relatively centrally located. Drilling at NV06 and NV13, however, require that the decision-maker choose between the greater confidence in the permeability assigned to NV06's vicinity and NV13's apparently greater influence on the reservoir.

Other Issues

Although it is not directly related to uncertainty issues, the introduction of compression facilities into the basic model of the field necessitated a re-examination of the ANN accuracy question. That is, are the solutions obtained from searches relying on ANN predictions as good as those obtained by using the original simulators to generate predictions during the course of search?

This question was revisited with the revised model employing the same procedures as described in the first paper. The consumption of computer resources by the two methods also remained the same. Each search employing the FORMGAS/IMEX simulator required about six days (150 hours) to complete on a dedicated Pentium III processor. Searches using the ANNs to generate predictions usually required 5 minutes or less on the same equipment and about 5 hours to conduct verification runs on the top 100 solutions. In fact, in the 150 CPU hours required to complete one simulator-reliant search, it was possible to create the entire knowledge base, train/test ANNs, conduct searches, and verify optimal solutions for all the analyses in Tables 3, 4, 5, and 6. With that caution in mind, examination of the results in Table 7 does show that the top solutions obtained from the simulator-reliant searches were slightly superior to the top solutions from the ANN-reliant searches under two of the three conditions. Under the low permeability model, the simulator-reliant search found a subset of the best solution from the ANN-reliant search which, while raising Risk% slightly, was a clearly more economical development plan. Under the pessimistic remediation outcomes condition, the advantage shown by

the simulator-reliant search's best solution was very slight. Under the medium permeability model, the two forms of search generated almost identical best solutions.

Table 7. Optimized Development Plans Obtained from ANN- vs Simulator-reliant Searches Under Various Conditions

Condition	Prediction Method	Total Score	Risk%	\$Cost	Development Plan
Low Permeability Model	ANNs	0.1296	0.7	703,000	R1 - R7, NV06, NV20
		0.1403	2.4	700,000	R1, R3 - R7, NV06, NV13
		0.1774	6.6	747,600	R1 - R7, NV13, NV20, IHOR
	FORGAS/IMEX	0.1139	5.6	432,000	R1 - R7, NV06
		0.1296	0.7	703,000	R1 - R7, NV06, NV20
		0.1325	1.2	685,600	R1, R2, R4, R6, R7, NV06, NV13
Medium Permeability Model	ANNs	0.0176	0.0	86,000	R1 - R7
		0.0322	1.8	96,800	R1 - R7, IHOR
		0.0744	0.0	364,000	R1, R3 - R6, NV13, IHOR
	FORGAS/IMEX	0.0176	0.0	86,000	R1 - R7
		0.0322	1.8	96,800	R1 - R7, IHOR
		0.0755	0.0	369,000	R1, R2, R3, R6, R7, NV13, IHOR
Pessimistic Remediation Outcome	ANNs	0.3735	24.5	1,179,400	R1 - R7, NV13, NV14, NV20
		0.3784	25.5	1,171,200	R1 - R5, R7, NV13, NV14, NV20
		0.3806	25.4	1,185,800	R1 - R7, NV13, NV17, NV20
	FORGAS/IMEX	0.3652	23.4	867,400	R1 - R7, NV13, NV17, NV20, IHOR
		0.3735	24.8	1,168,600	R1 - R5, R7, NV13, NV14, NV20, IHOR
		0.3735	24.5	1,179,400	R1 - R7, NV13, NV14, NV20

Conclusions

Basic Model Changes. The primary purpose of this paper was to demonstrate ways in which the ANN-GA approach to simulation-optimization could address some of the uncertainties that face petroleum engineers when making decisions regarding field development planning. Before doing so, however, it was important to correct two weaknesses in the basic model of the field that had been implemented in the first paper.

First, a fictitious production plant compressor was added to the model to reflect more realistic engineering constraints on low flowing wellhead pressures and a better estimate of the economic impact of operating the compressor to counteract low pressures. The effect of this facility on the baseline no-action case was, as expected, to increase the amount by which the field fell short of contract requirements under the 30-day peak service at 2.5 Bcf baseload. This is a reasonable outcome since a portion of withdrawn gas was now being diverted to operate the compressor.

The second “fix” to the base model was a change in the skin factors assigned to remediated wells from the optimistic value of 2.0 to a more realistic value of 5.0. It was expected that this change would reduce the dominance of the remediation wells in the sets of optimal field development plans. This did happen to a certain extent; but the low cost of well remediation compared to drilling new wells still meant that the remediation field development options retained their popularity in the optimal solutions.

Uncertainty Issues. The fundamental problem in addressing uncertainty within the context of simulation-optimization is to find a way to incorporate the uncertainty without overwhelming the process with a massive number of alternatives. One way to do this is to look at the matter from the decision-maker's point of view, which entails breaking the uncertainty down into a small number of manageable scenarios and looking for solutions which can be reconciled across them.

This approach was illustrated in the handling of two forms of uncertainty in the test field. The first of these involved disputed permeabilities in an important mid-section of the field. Rather than viewing the problem geostatistically, which might have involved hundreds of alternative realizations of the permeabilities, two reasonable alternative values were selected and assigned to the section as a block. The knowledge base was augmented to include evaluation of each field development plan under all three permeability models and the entire process of ANN training, GA optimization, and solution verification was carried out three times. Results were then examined to see if an overall plan of action could be constructed which would hold across all three permeability models. In this particular case, such a plan was possible because the optimal solutions to the higher permeability models were subsets of the solutions to the original, low permeability model.

The same approach was applied to the problem of anticipating the relative success of remediation. All plans in the knowledge base involving remediation wells were evaluated a fourth time under the assumption that remediation would only produce a 50% improvement in skin factor. Optimal solutions obtained under this very pessimistic set of assumptions were compared with those obtained under the expected remediation outcomes and a plan reconciling the two views of the world was devised. It should be noted, however, that the best solutions under the pessimistic view were pretty far from meeting the injection/withdrawal contract. If the pessimistic view were to be taken seriously, the answer to the management problem would probably be to *not try* to provide 30-peak service at the 2.5 Bcf baseload with the existing array of development options.

Not all approaches to the management of uncertainty require extending the knowledge base and repeating the ANN-GA process under different sets of assumptions. To illustrate a much simpler approach, the locations of new wells in the field were "handicapped" according to their distance from existing wells where the permeabilities are better known. This could be accommodated by simply adding a distance penalty to the objective function and applying different weights to the terms in the objective function to observe the effects. Once again, it proved possible to locate at least one new well location that contributed to optimal solutions under different weighting schemes. Results from only two weighting schemes were presented in the tables. But, given the speed with which these analyses can be performed, any number of trial values could and were applied.

Flexibility of the ANN-GA Approach. The analyses in this paper illustrate both the weakness and strength of the ANN-GA approach to simulation-optimization. Most of the simulation runs are conducted during the knowledge base creation stage, which essentially fixes the underlying model and the field development options. When assumptions change, as occurred, for example, when compression facilities were added, it was necessary to rerun the cases. This is the main reason why it is better to keep the total number of cases in the knowledge base rather small (around 500) and possibly live with a certain degree of error in the ANN predictions. It is also preferable to postpone uncertainty analyses until the base problem and array of field development options have stabilized.

On the other hand, nowhere is the power of the method more apparent than in the handling of the distance penalty for new wells. All within a few hours, it was possible to construct a new objective function and run many different analyses varying the weights applied to individual terms, whereas completing a single

analysis relying on the full FORGAS/IMEX model to generate predictions would have required about six days. By enabling planners to experiment freely, the ANN-GA methodology greatly increases the value of reservoir simulators as decision-making tools.

References

1. Johnson, V.M., Ammer, J.R., and Trick, M.D.: “Improving Gas Storage Development Planning Through Simulation-optimization”, SPE 65639, Society of Petroleum Engineers, Tulsa (2000).