EFFECTIVENESS OF ANTITHETIC SAMPLING AND STRATIFIED SAMPLING IN MONTE CARLO CHRONOLOGICAL PRODUCTION COST MODELING

by

Chris Marnay
Utility Planning and Policy
Lawrence Berkeley Laboratory

Todd Strauss
Industrial Engineering and Operations Research
University of California, Berkeley

BERKELEY, CA 94720

ABSTRACT

Sampling error detracts from the usefulness of estimates of mean hourly marginal cost produced by Monte Carlo chronological production cost models. Variance reduction techniques commonly used in other Monte Carlo simulation applications can significantly improve the precision of estimates. Two such techniques, antithetic sampling and stratified sampling, are tested for a fictitious system. The number of iterations needed to reach a precision target falls significantly. The estimated savings in total computing time could exceed 50 percent for a full one-year forecast. Both techniques are easily implemented and should be used in Monte Carlo production costing efforts to estimate hourly marginal cost.

KEYWORDS

convergence criteria, least-cost planning, marginal cost, mean and variance of production cost, production costing, system planning, variable operating cost, variance reduction

I INTRODUCTION

A. Modeling Problem

At any point in time, each resource on an electric utility system has a capacity level that is available for dispatch as necessary. The total available capacity and the roster of generating resources available, then, represent just one of many possible *states*, or combinations of resources, of the generation system. The operator has to cope only with the short-run changes in this state that arise as individual resources become available or go on outage. The objective in production cost modeling is to simulate the cost-minimizing operation of the system through time, and to estimate system values of interest to the modeler, without attempting to forecast the actual state of the system that will obtain at every point in time. That is, the modeler cannot reasonably attempt to forecast the available capacities of every resource on the system through the study period, while, on the other hand, he or she needs to take proper account of the randomness of generator failures so that the simulation produces valid estimates of the system parameters of interest.

90 SM 446-5 PWRS A paper recommended and approved by the IEEE Power System Engineering Committee of the IEEE Power Engineering Society for presentation at the IEEE/PES 1990 Summer Meeting, Minneapolis, Minnesota, July 15-19, 1990. Manuscript submitted August 31, 1989; made available for printing April 24, 1990.

B. Monte Carlo Imprecision

Monte Carlo chronological (MCC) models represent the system state as a vector of resource availabilities.

$$[a_1, a_2, \dots, a_I]$$

where, a_i is the available capacity of unit i

If unit i has a *forced outage rate*, that is, if it is not perfectly reliable, a_i is a random variable. A simulation for some time period for one sample system state is called an *iteration*. The system state vector is established for each iteration by randomly drawing an availability from the *outage distribution* of each unreliable resource. This outage distribution typically has just two spikes, the probability of the unit being available at the unit's capacity, and the probability of it being on outage at zero capacity. After establishing a sample system state, MCC models attempt to accurately simulate the operation of the system through time as if this state obtained, applying whatever rules of unit commitment and dispatch the modeler chooses to impose. The use of sampling introduces *imprecision*, or *sampling variance*, into all the key outputs from a MCC model. Variance reduction techniques improve the precision of all such outputs, although the improvement will vary depending on the influence of rare events in the estimation. Because outputs are derived by averaging the outcomes of several simulations, those parameters that are heavily influenced by rare events, such as reliability indices, will be the most imprecise and will benefit the most from variance reduction. Using a control variable variance reduction technique, Oliveira, Pereira, and Cunha demonstrate a dramatic improvement in the precision of estimates of loss of load probability and expected power not served.[13] System hourly marginal cost, the variable of interest here, is influenced by rare events, although less so than reliability indices. Therefore, variance reduction will be less effective in marginal cost estimation than in reliability index estimation.

C. Antithetic and Stratified Sampling

Antithetic sampling and stratified sampling are two common methods that improve the precision of Monte Carlo estimates. This paper reports the results of a test of the effectiveness of these two techniques for *variance reduction* in the estimation of system hourly marginal cost using MCC production costing of a fictitious system.[10]

II MARGINAL COST

A. Role of Marginal Cost Estimation

The marginal cost concept has considerable economic importance. The literature on marginal cost pricing of electricity is well summarized by Crew and Kleindorfer.[2,4]

Considerable interest in improving marginal cost estimation has resulted from its increasing role in regulatory hearings.[8,9,14,15,20,21] This paper furthers the goal of providing better computing tools for utilities and other participants in rate proceedings.

B. Marginal Cost Distributions

Forecasting models traditionally treat hours as discrete steps, and input and output values are considered fixed for the hour. True system marginal cost in any hour depends on the true system state and must be a random variable, here called M. The true system marginal cost in any hour has a discrete distribution with a spike of probability at the variable operating cost of each resource that could be marginal in the hour. These distributions tend to be severely skewed. The MCC production costing literature is moving towards greater efforts to estimate these entire distributions, rather than just their means.[6,11,16,19]

C. LDC Models

Mazumdar and Yin, and Feng, Sutanto, and Manhire show that an estimate of the true distribution of hourly marginal cost falls directly from the load duration curve (LDC) approach.[7,11] Estimates from LDC models, however, are compromised by the existence in the real system of any operating constraints that are not respected by the model, and by the aggregation of loads.[1,5]

D. Monte Carlo Chronological Models

In current MCC models, unplanned outages are assumed to be random in a probabilistic sense. Since the availability of a unit is a random variable, so is total system available capacity. If complex real-time operating constraints of a system are respected, marginal cost does not depend on the system state through any functional relationship that could be tractable analytically. Hourly marginal cost estimates, then, can be derived only by detailed simulations in which operating constraints, such as ramp rates, are embodied in heuristic unit commitment rules. Since the MCC production costing framework can better incorporate real-time operating constraints, it promises more accurate estimates of the true system hourly marginal cost distribution. This work addresses the statistical precision of estimates and does not directly address the accuracy of MCC models; nonetheless, by improving precision, the potential benefits of MCC models can be captured at lower computing cost.

MCC simulation of a random sampling of system states leads to an estimate of the true distribution of hourly marginal cost, in the form of a sample distribution. But the sample distribution is, in reality, an estimate of a *population* distribution formed by the marginal costs corresponding to all system states possible in the model's representation of the system. The empirical distribution obtained from multiple iterations eventually converges to the population marginal cost distribution that would emerge from the systematic simulation of all system states.

To summarize, the three levels of hourly marginal cost distribution means and variances in the MCC framework are:

1. the *true* system parameters, here $\boldsymbol{\mu}$ and $\boldsymbol{\sigma}^2$, which modeling is attempting, ultimately, to estimate;
2. the *population* parameters, μ and σ^2, which are actually calculable, but not at reasonable computing cost; and
3. the *sample* parameters, m and s^2, which are actually taken from the simulations of a random sample of model system states.

MCC models output the mean of their sample hourly marginal cost distribution, m, and this is taken by the modeler as a point estimate of the true system mean marginal cost for the hour, $\boldsymbol{\mu}$. What the models are actually reporting consists of a sample estimate, m, of the population mean, μ. How good μ is as an estimate of $\boldsymbol{\mu}$ is the *accuracy* problem, and how good m is as an estimate of μ is the *precision* problem.

Because sampling is used, results have a sampling variance; that is, m is also a random variable and has a distribution. The variance of m, here called VAR[m] to clearly distinguish it from the variance of the marginal cost distribution, measures the imprecision of m. If the system states are randomly picked, statistical theory holds that m forms an *unbiased estimator* of μ, and VAR[m] is simply the ratio of the population variance, σ^2, and the sample size, n.[17]

$$E[m] = \mu \qquad \text{eq.1}$$

$$VAR[m] = \sigma^2/n \qquad \text{eq.2}$$

The imprecision of the estimator m, measured by VAR[m], compromises the usefulness of m as an estimator of μ, and, ultimately, as an estimator of $\boldsymbol{\mu}$. Applying variance reduction techniques can improve the statistical precision of hourly mean marginal cost estimates. The improvement results from replacing the simple mean estimator, m, with alternative estimators that have sample variances lower than VAR[m] for the same number of iterations, n.

III CONVERGENCE CRITERION

A minimum acceptable level of precision can be specified before the simulation begins, the model being instructed to continue iterating until the desired target is reached. The unbiased sample mean marginal cost, m, must be converging towards the population mean, μ, so such specifications are usually called *convergence criteria.* The Central Limit Theorem holds that m must be normally distributed, asymptotically, i.e. as $n \to \infty$.[17] Therefore, convergence criteria can be expressed as confidence intervals based on the normal distribution. A convergence criterion might, for example, be that the width of the 95 percent confidence interval for μ must be less than a constant.

$$1.96 \cdot \frac{s}{\sqrt{n}} < \frac{w}{2} \qquad \text{eq.3}$$

where, s is the sample standard deviation
n is the number of iterations completed
w is a constant specified by the modeler.

Alternative convergence criteria that reflect the modeler's specific precision goals can easily be devised. While confidence intervals are convenient, they should be interpreted with care because the marginal costs of the various hours of a simulation period are not strictly statistically independent.

Since no established standard for setting w exists, any criterion adopted for this study must be arbitrary. The rule applied is that 90 percent of the hours in a weekly simulation period must have *standard errors* less than 0.1 cents/kWh, which is equivalent to a rule, as above, with w = 0.392. To implement this criterion, the hours of the week are sorted in ascending order of standard error. The target is met when, for the 152nd ordered hour, $s/\sqrt{n} < 0.1$. If the large sample properties of the Central Limit Theorem hold, there is a 95 percent chance m is within ±0.2 cents/kWh of μ.

IV ANTITHETIC SAMPLING

The technique of antithetic sampling exploits the basic probabilistic result that the variance of the sum of two random variables, X and Y, depends upon their covariance.

$$VAR[X + Y] = VAR[X] + VAR[Y] + 2COV[X,Y] \quad \text{eq.4}$$

If the two random variables are negatively correlated, then the last term on the right hand side is negative. Thus, the variance of the sum will be less than the sum of the variances, that is, less than the variance under independence. This result suggests finding two negatively correlated estimators, $m_{at,1}$ and $m_{at,2}$, thus ensuring a negative covariance term. The average of the two can provide a new estimator.

$$m_{at} = \frac{m_{at,1} + m_{at,2}}{2} \quad \text{eq.5}$$

In production cost modeling, antithetic sampling can be implemented as follows. The system state vector for the first iteration is established using a vector of random numbers from a random number generator. This system state vector forms an input to the MCC simulation which yields an hourly marginal cost estimate m_{11} as output. Then the antithesis of the random number vector is used to generate the second system state. Since random numbers are, by convention, between zero and one, the antithesis of a random number u_i is $1-u_i$. This antithetic system state is used in the second iteration, which yields marginal cost estimate m_{21} as output. This process is repeated until a total of n system states have been simulated, where n is an even number. The n/2 hourly marginal cost estimates, $m_{11}, m_{12}, \ldots, m_{1,n/2}$, are used to calculate the mean marginal cost estimate $m_{at,1}$, while the antithetically sampled states, $m_{21}, m_{22}, \ldots, m_{2,n/2}$, yield the antithetic marginal cost estimate $m_{at,2}$.

The marginal cost random variable M_1 and its estimator $m_{at,1}$ correspond to the randomly sampled states, while the marginal cost random variable M_2 and its estimator $m_{at,2}$ correspond to the antithetically sampled states. The random variables M_1 and M_2 are identically distributed but not independent. The antithetic estimator m_{at} is unbiased, being a linear combination of the unbiased estimators $m_{at,1}$ and $m_{at,2}$. Because m_{at} is algebraically identical to the sum of n/2 independent pairs of observations, the Central Limit Theorem still applies and m_{at} is asymptotically normally distributed. The expectation and variance follow.

$$E[m_{at}] = \mu \quad \text{eq.6}$$

$$VAR[m_{at}] = \frac{\sigma^2 + COV[M_1, M_2]}{n} \quad \text{eq.7}$$

The expression for $VAR[m_{at}]$ in eq. 7 differs from that for VAR[m] in eq.2 only by the appearance of the covariance term.

For antithetic sampling to work, the negative correlation between the antithetic number $1-u_i$ and the random number u_i must be reflected in the output marginal costs. That is, the outputs must be weakly monotone with respect to the inputs.[18] The system state inputs are clearly weakly monotone functions of the random numbers. However, the marginal costs must also be weakly monotone functions of the system states to guarantee that the random variables M_1 and M_2 are negatively correlated. Intuitively, the monotonicity between system states and marginal costs should hold because less capacity available should lead to higher costs. Indeed, under unconstrained conditions in which resources are dispatched in strict economic order, which immediately ensures global cost minimization, the monotonicity holds. In this circumstance, marginal cost is weakly monotone with respect to every element in the system state vector, which is a sufficient condition for no variance increase to result from negatively correlated inputs. In actual MCC simulations, however, real-time unit commitment constraints often preclude strict economic dispatch. Although intuitively unlikely, therefore, it is not inconceivable that a forced outage could lead to lower marginal cost than would exist if this resource were available. Consider the example of two almost identical iterations. They differ only in that, in the second iteration, the unavailability of a resource necessitates the continuous operation of a more expensive, but less rampable, unit. If minimum load conditions arise because of the diminished load-following ability of the system, although a unit has been lost in the second iteration, most models would report a zero marginal cost, breaking the monotonicity assumption. In other words, in marginal cost estimation, antithetic sampling cannot ensure a reduction in variance under all conditions.

Antithetic sampling has three highly desirable properties. First, it is easily implemented. Second, it may be used in conjunction with other variance reduction techniques because it changes only the random drawing procedure, not the actual estimators $m_{at,1}$ and $m_{at,2}$. These can further be replaced by other estimators. In this study, for example, stratified sampling estimators replace them. Third, and most importantly, antithetic sampling requires no additional prior knowledge of the output random variables beyond monotonicity. For these reasons, antithetic sampling shows great promise as a variance reduction technique for MCC modeling, and, indeed, it has already been both proposed and tested.[3]

V STRATIFIED SAMPLING

Stratified sampling is a familiar technique in many statistical applications. Each observation is classified according to a disjoint subset of the sample space, called a *stratum*, from which the sample was drawn. The probability of each stratum occurring, p_k, must form a discrete random variable taking on a limited number of values, so the sample space must be readily and clearly divisible into a manageable number of strata. A subset of observations is taken from each stratum.

$$n_1 + n_2 + \ldots + n_K = n \qquad \text{eq.8}$$

where, n_k are the sample sizes from K strata
n is the total sample size.

A conditional population distribution resides in each stratum, and its mean, μ_k, and variance, σ^2_k, can be computed.[18] The *stratified estimator*, m_{ss}, is the weighted average of the sample means, $m_{ss,k}$, of all the strata.

$$m_{ss} = \sum_k p_k \cdot m_{ss,k} \qquad \text{eq.9}$$

Being the weighted sum of independent asymptotically normally distributed random variables, m_{ss} is asymptotically normal, and being a linear combination of unbiased estimators, m_{ss} is unbiased. Its variance appears below.

$$E[m_{ss}] = \mu \qquad \text{eq.10}$$

$$VAR[m_{ss}] = \sum_k p_k^2 \frac{\sigma_k^2}{n_k} \qquad \text{eq.11}$$

The stratified estimator tends to have smaller variance than the raw estimator because the observations within a stratum are more similar to each other than to observations in other strata. In other words, the within strata variance is smaller than the between strata variance.[18]

In practice, the choice of n_k can be tricky and significantly affects the improvement in precision. Assuming the overall sample size, n, is fixed, two values, the stratum weight, p_k, and the stratum variance, σ^2_k, determine the most effective choice of n_k. In MCC applications, the obvious strata are the outage states of large unreliable resources because they substantially influence marginal costs. A full set of p_k is readily available, therefore, in the form of the forced outage rates. As in most applications, however, little or nothing is known *ex ante* about σ^2_k. In this study, three methods of choosing n_k are considered.

A. Optimal Stratified Sampling

This first case, ss1, forms an unattainable upper bound on the variance reduction because it is assumed that the stratum variance, σ^2_k, is known. In this case, the value of n_k that yields the greatest reduction in variance is as follows.[18]

$$n_k = n \cdot \frac{p_k \sigma_k}{\sum_j p_j \sigma_j}$$

$$j = 1, 2, \ldots K \qquad \text{eq.12}$$

B. Proportional Stratified Sampling

The second case, ss2, uses proportional sampling. The n_k are selected as follows.

$$n_k = n \cdot p_k \qquad \text{eq.13}$$

Applying this common approach in production costing is particularly straightforward because values of all likely p_k are known.

C. Equal Stratified Sampling

In the equal sampling approach, ss3, no information about either the p_k or the σ^2_k is required. The sample is simply divided equally among the strata as follows.

$$n_k = \frac{n}{K}, \ \forall k \qquad \text{eq.14}$$

VI TEST PROCEDURE

A. CalECo

A fictitious California utility, the California Electric Company, or CalECo, has been developed.[10] About a third of CalECo's energy comes from its own hydro capacity. This is an energy limited resource primarily available in the spring. In addition, it runs a two-unit 2.0 GW nuclear station, which is the resource used for stratification. Below the nuclear unit in the merit order are a 1.0 GW out-of-state coal station, and an array of gas, oil, and combustion turbine units. CalECo also purchases 1.0 GW of power under fixed-price contracts from qualifying facilities under the terms of the Public Utilities Regulatory Policies Act. This is a *must run* resource as far as the CalECo dispatch is concerned. CalECo also purchases power under two economy purchase agreements. In all, CalECo has 20 resources, of which 5 have no forced outage rate. That is, CalECo has a population of system states with 2^{15} members, each of which has an associated marginal cost for each hour. The only operating constraints imposed are ramp rates, and the spinning reserve requirement is a fixed 5 percent of load. Hypothetical CalECo loads for the 1990 test year are based on actual historic loads for a California utility scaled to an arbitrary peak of 12.5 GW. Because the number of iterations required is so large as to preclude the possibility of simulating an entire year, 4 example weeks, one from each quarter, are selected.

B. Test Steps

The analysis has three distinct steps.

1. Only two kinds of minor code changes are needed in the MCC model. First, the random draws that determine the availability of each unit are all saved, so that the state of any unit on any iteration is known with certainty. Second, key outputs are recorded at an iteration-by-iteration level. The MCC model serves only as a source of data sets. The model used is POWRSYM PLUS (P+), developed by Energy and Control Consultants (ECC) of San Jose, CA. Running CalECo simulations with P+ on an engineering workstation rated at 10 MIPS and 1 Mflop requires approximately 10 cpu seconds per iteration-week.
2. FORTRAN programs read the large data base files and generate summary statistics.
3. The summary statistics are analyzed in a standard sta-

tistical package, which leads to estimates of the variance reduction potential of each technique. Unlike the above steps, this involves only modest computing power and is done on a Macintosh II.

C. Antithetic Sampling (at)

On the first iteration, a vector of random numbers is used, whereas on the second, the vector of antitheses is used. The output consists of pairs of iterations; the first one is identical to the base case while the second one corresponds to the antithetic vector.

D. Stratified Sampling (ss1-3)

Approximating the results of a fully implemented stratification poses a more difficult challenge than does antithetic sampling. The strata used consist of the outage states of the two nuclear units. Separate data sets containing simulations with both units available, with one available and one not, and with both on outage, are needed. Analyzing these separate output files replicates the effect of implementing stratification because it is identical to subdividing the sample space within a simulation.

E. Joint (as1-3)

Estimating the precision improvement possible using both techniques simultaneously requires repeating the stratified sampling procedure, but with a separate data set compiled using antithetic sampling.

F. Data Base

These are the actual P+ runs conducted:

1. a year-long 100-iteration pilot run that serves as the basis for the selection of example weeks
2. a 1000-iteration base case for the four example weeks
3. a 1000-iteration antithetic base case that differs from 2. only in that antithetic sampling is used
4. a 1000-iteration stratified sampling run with both nuclear units available, using a random number stream independent of all other cases except 5.
5. a run identical to 4. but using antithetic sampling
6. a 1000-iteration stratified sampling case with one nuclear unit available using a random number stream independent of all other cases except 7.
7. a run identical to 6. but using antithetic sampling
8. a 1000-iteration stratified sampling case with neither nuclear unit available, using a random number stream independent of all other cases except 9.
9. a run identical to 8. but using antithetic sampling

The data sets created by these nine runs consist of large random samples of estimated hourly marginal costs, weekly production costs, random number streams, and the outage states of all CalECo resources.

VII RESULTS

The figure shows the marginal cost results from the base simulation run. Week 22 has the lowest marginal costs, but they are erratic, while week 8 has fairly flat costs. Weeks 36 and 47, during which resources are more typical of thermal based systems, exhibited the clear daily cycle that one might expect. The rows of the table below show the results for each of the four example weeks. All the variance reduction cases should be compared to the base case that appears in section A of the table. The bold number, 26 percent in the antithetic only case, shows the percent reduction achieved in the total number of iterations. This number provides the best guide here of what computational savings might be achieved in actual practice.

CalECo Hourly Marginal Cost

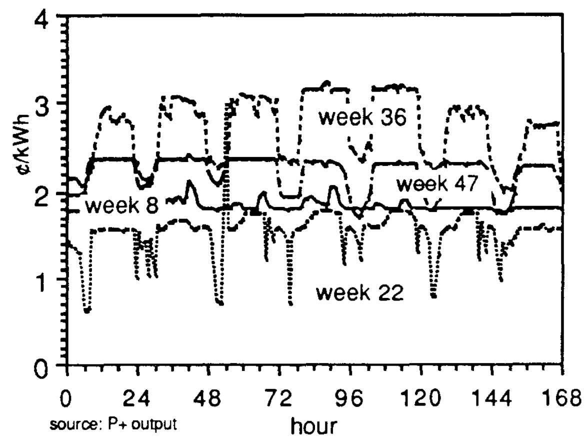

A. Antithetic Only

week	base d	base its	at d_{at}	at its	%
8	0.24	6	0.22	6	8
22	0.71	51	0.63	40	11
36	1.04	109	0.84	72	19
47	0.57	33	0.52	28	9
wt		203		150	**26**

B. Stratified Only

ss1 d_{ss1}	its	%	ss2 d_{ss2}	its	%	ss3 d_{ss3}	its	%
0.19	4	21	0.22	5	8	0.19	4	21
0.37	14	48	0.42	18	41	0.54	30	24
0.80	65	23	0.81	67	22	0.93	88	11
0.38	15	33	0.47	22	17	0.44	20	23
wt	104	**49**		117	**42**		148	**27**

C. Joint

as1 d_{as1}	its	%	as2 d_{as2}	its	%	as3 d_{as3}	its	%
0.18	4	25	0.21	6	12	0.18	4	25
0.36	14	49	0.42	18	41	0.54	30	24
0.67	46	36	0.67	46	36	0.77	60	26
0.34	12	40	0.40	18	30	0.39	16	32
wt	82	**60**		92	**55**		116	**43**

The d values are measurements, akin to the standard deviation, which permit useful comparisons of the variance reduction achieved in the various test cases. The d values provide estimates of the imprecision of the estimators at any sample size according to the following formula.

$$\widehat{VAR}\,[m_{xxx}] = \frac{d^2_{xxx}}{n} \qquad \text{eq.15}$$

The hat signifies that this is an estimate. The similarity of eq.15 to eq.2 should be obvious and, in the base case, d = s. Derivation of the formulas for the d value for the other cases can be found elsewhere.[10] The % column contains the percentage reduction in the d value that the variance reduction technique achieves. The estimated number of iterations required to reach the convergence target is also shown. The wt row sums the number of iterations for the 4 weeks except that numbers of iterations of less than 10 are counted as 10. The motivation for this sum is that in practice a minimum number of iterations should always be conducted to avoid small sample problems. Since the four weeks are carefully chosen representatives of each quarter, 13 times the sum is a rough estimate of the total number of iteration-weeks needed for a full-year simulation.

A. Antithetic Only

The results clearly show the surprising effectiveness of antithetic sampling. In this test, the results alleviate concerns that monotonicity might not hold. Two aspects of the results particularly impress. First, a substantial computational savings results from an easily implemented procedure. Second, the benefits of antithetic sampling are focused in the weeks of highest variance, exactly what is needed to deliver high overall computational savings.

B. Stratification Only

The results rank as expected, with ss1 delivering the largest results. However, proportional sampling, ss2, follows closely. Given the convenience of this technique in the production costing context, this result should encourage the modeler. Equal sampling, ss3, proves less effective than the other two approaches, and by a significant margin. Overall, it is actually barely more effective than antithetic sampling, and requires more prior knowledge of the system. An interesting result, however, is that ss3 delivers more variance reduction in the lower variance weeks, particularly week 8. Clearly, too much of the power to reduce variance is spent in times when the net effect is minimal, or zero, as far as the weighted sum is concerned.

C. Joint

The joint results are better than either technique used independently but a diminishing returns effect exists. The as1 result, a 60 percent reduction in total iterations, serves as the best achievable result for the test case. Remarkably, however, the feasible technique as2 delivers a reduction of 55 percent.

VIII CONCLUSION

The major caveat to the results concerns their specificity. CalECo is not intended to be a typical utility. It represents a contrived test case that mirrors California conditions, while being of a manageable size. Other utilities have quite different marginal cost distributions, and the results of variance reduction in the simulation of their systems would be quite different. Similarly, the choice of production cost model affects the variance of marginal cost because of inaccuracy. A different unit commitment logic may deliver a different marginal cost distribution, that is, a different μ and σ. The effectiveness of variance reduction, then, will differ for each system and each model. Notice, in particular, that P+ follows in the tradition of models that simulate weeks individually, as independent periods. Within this context, the measure of computing requirements as iteration-weeks makes sense. But the total computational savings possible depends heavily on the length of the simulation period. If, for example, a model simulated the whole year in an indivisible manner, then the number of iterations must ensure that the desired level of precision is achieved for all hours. Finally, note again the arbitrariness of the convergence criterion, which determines all the results regarding estimated numbers of iterations.

While many approaches to variance reduction in chronological production cost models are possible, antithetic sampling and, to a lesser extent, stratified sampling are two relatively easy techniques to implement. Furthermore, they can be used simultaneously, and they effectively improve the precision of hourly marginal cost estimates for the CalECo test utility. While their effectiveness cannot be guaranteed under all circumstances, considerable computational savings can be captured with little effort. Further, the danger that the use of these techniques in simulation will actually be disadvantageous is small.

The joint procedure that incorporates antithetic sampling and proportional stratification emerges as the method of choice. For the CalECo test case, this approach not only cuts estimated total iterations in half, it performs almost as well as the optimal joint estimator.

IX ACKNOWLEDGEMENTS

This paper is based on a report prepared by Chris Marnay, Electricity Economist, and Todd Strauss, Statistician, pursuant to a contract to the California Public Utilities Commission (CPUC). The study was financed by the CPUC in preparation for the Fourth Report to the California State Legislature on Implementation of Assembly Bill 475 (1985). The statements, findings, conclusions, and recommendations presented in this paper and the original report, available from the authors, do not necessarily reflect the views of the staff nor the commissioners of the CPUC. This work, including results and data, is entirely in the public domain and are available to other researchers; however, the models used are proprietary products available only from their respective vendors. This project has involved many helpful people, but the authors would like to express their particular thanks to three anonymous reviewers and the following: Chris Danforth of the CPUC; Ron Davis of Henwood Energy Services; Jerome K. Delson of Applied Analytics Associates; Joseph H. Eto and Ed Kahn of Lawrence Berkeley Laboratory; Don Imamura of the Northern California Power Agency; Prof. R. Taber Jenkins of the University of Tennessee at Chattanooga; Prof. Shmuel Oren of the IEOR Dept. at U.C. Berkeley; Chok Pang of ECC; and Nyla Marnay of home.

X BIBLIOGRAPHY

[1] Bell, A.M. Comment to Mazumdar, [11].

[2] Boiteux, Marcel. "Marginal Cost Pricing." Ch. 3 of Nelson, James R. *Marginal Cost Pricing in Action.* Englewood Cliffs, NJ: Prentice-Hall, 1980.

[3] Carney, T. "Methodology for Specifying Number of Monte Carlo Trials Required for BENCHMARK Simulation and Variance Reduction Recommendations," Appendix D of *BENCHMARK - A Computer Program for Simulation of Hourly Generation.* Draft Final Report, EPRI Research Project 1808-4, January 1985.

[4] Crew, Michael A., and Paul R. Kleindorfer. *The Economics of Public Utility Regulation.* Cambridge, MA: MIT Press, 1986.

[5] Delson, Jerome K., ed. *Proceedings of the 1985 Chattanooga Conference on Production Simulation of Electric Power Production.* EPRI Monograph on Simulation of Electric Power Production, Book 4, EPRI Research Project 1808-7, 1987.

[6] Delson, Jerome K., Xiaoming Feng, and William C. Smith. "A Validation Process for Probabilistic Production Costing Programs." Paper delivered at the 16th IEEE Power Industry Computer Application Conference, Seattle, WA, May 1989.

[7] Feng, X., D. Sutano, and B. Manhire. "Evaluation of Variances in Production Cost of Electric Power Systems." Department of Electrical and Computer Engineering, Ohio University, ECE report #046, April 7, 1989.

[8] Kirsch, Laurence D., Robert L. Sullivan, Theresa Flaim, Joseph J. Hipius, and Margaret G. Krantz. "Developing Marginal Costs for Real-Time Pricing." *IEEE Transactions on Power Systems,* vol. 3(3), 1988.

[9] Marnay, Chris, and Todd Strauss. *Chronological Model Comparison.* Report to the California Public Utilities Commission, February 1989.

[10] Marnay, Chris, and Todd Strauss. *Variance Reduction in Chronological Production Cost Modeling.* Report to the California Public Utilities Commission, February 1990.

[11] Mazumdar, M., and C.K. Yin. "Variance of Power Generating System Production Costs." *IEEE Transactions on Power Systems,* vol. 4(2), May 1989.

[12] Morgan, B. J. T. *Elements of Simulation.* New York: Chapman & Hall, 1984.

[13] Oliveira, C., M.V.F. Pereira, and S.H.F. Cunha. "A Method for Reducing Computational Effort in Monte-Carlo Based Composite Reliability Evaluation." *IEEE Transactions on Power Systems,* vol.4(4), Oct. 1989.

[14] Oyama, Tatsuo. "Theoretical Aspects of an Electricity Marginal Cost Model." IEEE Transactions on Power Apparatus and Systems, vol. PAS-102 (5), 1983.

[15] Oyama, Tatsuo. "Applying Mathematical Programming to Measure Electricity Marginal Costs." IEEE Transactions on Power Systems, vol. PWRS-2 (2), 1987.

[16] Rau, Narayan S., and Yousef Hegazy. "The Variance of Production Cost." Paper presented at the IEEE/PES 1989 Summer Meeting, Long Beach, CA, July 9-14, 1989. 89 SM 706-3 PWRS.

[17] Ross, Sheldon M. *Introduction to Probability Models.* Third Edition, Orlando, Florida: Academic Press, 1985.

[18] Rubinstein, Reuven Y. *Simulation and the Monte Carlo Method.* New York: John Wiley & Sons, 1981.

[19] Ryan, S.M., and M. Mazumdar. "Effect of Frequency and Duration of Generating Unit Outages on Distribution of System Production Costs." *IEEE Transactions on Power Systems,* 89 SM 708-9 PWRS.

[20] Stremel, John, and William Stillinger. "How an Electric Utility Production Cost Model Can be Validated." *Public Utilities Fortnightly,* December 8, 1988.

[21] Scherer, Charles R.. "Estimating Peak and Off-Peak Marginal Costs for an Electric Power System: An *Ex Ante* Approach." *Bell Journal of Economics,* Autumn 1976, vol 7(2).

Chris Marnay was born Chris Pignone 3 May 1950 at Stourport-on-Severn, U.K. He received a B.A. in Development Studies and an M.S. in Agricultural and Resource Economics from the University of California, Berkeley, and is a doctoral candidate in Energy and Resources at Berkeley. He is a consultant to the California Public Utilities Commission and a Research Associate at the Lawrence Berkeley Laboratory. He was formerly a participant in the Resource Systems Institute at the East-West Center, University of Hawaii, and has acted as a consultant to CESEN of the Ansaldo Group, Genova, Italy. His current research interests include non-traditional pricing schemes for electric power and the incorporation of environmental operating constraints into production cost models. He is a member of the American Association for the Advancement of Science and the International Association of Energy Economists.

Todd Strauss was born 2 September 1962 in New York City. He received the S.B. in mathematics and the S.M. in statistics from the Massachusetts Institute of Technology. He is currently a doctoral student in Industrial Engineering and Operations Research at the University of California, Berkeley. Mr. Strauss has been a student intern at the California Public Utilities Commission since 1988. His current research interests include modeling the pricing and dispatching of electric power, particularly under chronological operating constraints associated with purchased power, hydro resources, and demand-side management strategies. Mr. Strauss is a member of the Operations Research Society of America and the American Statistical Association. He has been a statistical consultant since 1984 and has worked on projects with the International Nutrition Foundation, University Hospital (Boston), Chemical Bank, and Statistical Innovations Inc.