Title
Least-Cost Planning for Pacific Gas and Electric

Authors
Kahn, E.
Pignone, C.
Comnes, G.A.

Publication Date
1987-02-01

LBL-22702
UC-350

Lawrence Berkeley Laboratory

UNIVERSITY OF CALIFORNIA

APPLIED SCIENCE DIVISION

Least-Cost Planning for Pacific Gas & Electric

STAGE 1: LMSTM Calibration

E. Kahn, C. Pignone, and G.A. Comnes

February 1987

APPLIED SCIENCE DIVISION

Prepared for the U.S. Department of Energy under Contract Number DE-AC03-76SF00098.

LBL-22702

Least-Cost Planning for Pacific Gas & Electric

STAGE 1: LMSTM Calibration

Edward Kahn, Chris Pignone, and G. Alan Comnes

Energy Analysis Program
Lawrence Berkeley Laboratory
University of California
Berkeley, CA 94720

February 15, 1987

This work was supported by the Pacific Gas & Electric Company and the Assistant Secretary for Conservation and Renewable Energy, Office of Building Systems, Building Energy Research Division of the U.S. Department of Energy under contract DE-AC03-76SF00098

TABLE OF CONTENTS

List of Tables

EXECUTIVE SUMMARY

The goal of this project is to improve planners' ability to evaluate demand-side programs in the electric utility resource planning process. The project is designed as a case study of the Pacific Gas and Electric Company (PG&E). The analysis is based on the use of a mainframe integrated utility planning model called the Load Management Strategy Testing Model (LMSTM). This model accepts a disaggregated chronological representation of electric utility load shapes and a representation of its demand, rates, and other significant characteristics. It carries out an integrated production simulation that meets the demands forecast on the basis of the company's salient features. Because the inputs to integrated models must approximate the true company parameters, there is some concern that the simplifications will be excessive and important information will be lost. This concern is greatest for the production simulation element. PG&E has a particularly complex power system whose operation is subject to many constraints and which must integrate a wide variety of generation resources. Thus, the first task of the project involves the calibration of LMSTM's production simulation to the more detailed models used by PG&E for electric resource planning purposes, notably the company's own GRASS model. Following the completion of the calibration, we will conduct the least-cost resource planning evaluation based on large-scale implementation of a demand-side program.

PG&E supplied LBL with the company's LMSTM input and output files. The primary task of the calibration exercise is to aggregate and respecify this data in a form compatible with the output of a recent GRASS run. Aggregation is necessary because LMSTM simulates seasons compared to the monthly simulation periods used in the more complex models, and the dispatch of LMSTM is based on a limited number of resource groups. The choice of season definitions is a user option in LMSTM. For the PG&E system, the principal constraint of this choice involves the substantial annual fluctuations of the hydroelectric resources. LBL chose a specification that captured the main qualitative variations as well as possible, but this choice required a re-characterization of the load data which must be specified in LMSTM as four typical day-types per season. The model exhibits intense sensitivity in its reported marginal costs with respect to the load inputs, especially the depth of the off-peak trough and hence, choice of the correct base period for the load shapes is a critical decision.

Detailed calibration tests were conducted for three representative test years over the simulation period, 1989, 1996, and 2003. Annual energy production for each resource type was compared as well as the annual value for marginal energy cost calculated by LMSTM and the detailed model. Production by fuel type was within 2% for base load resources. For marginal resources the deviations were large, averaging 15%. The results for annual energy marginal cost, however, were better. The average deviation was 5%. These results are sufficiently close for use in strategic planning studies of the type envisioned. They are substantial improvements over previous efforts of this kind, due largely to more accurate specification of system constraints.

In the second phase of this study, LBL will use the calibrated model to evaluate future expansion plans based on large-scale implementation of a typical demand-side program. The program will emphasize either thermal energy storage applications or

improved efficiency of residential appliances. The value of such a program will be estimated by calculating the value of the supply-side resources that it displaces.

I. INTRODUCTION

This report describes the activities undertaken by LBL to calibrate the Load Management Strategy Testing Model (LMSTM) for use in a least-cost planning exercise with the Pacific Gas and Electric Company (PG&E) system.[1] There is great general interest in analytical procedures that facilitate the integration of demand-side activities into utility resource planning. It is widely recognized that this is a complex problem which requires sophisticated modeling. A reasonable question to ask at the outset is what kind of model should be used for this purpose. This is a difficult problem for which there is little guidance and many competing claims and constraints. LBL has chosen to use the LMSTM model (for a general description of LMSTM see DFI, 1982). The motivation for that choice and the procedures necessary to achieve consistency with the results of PG&E's detailed production simulations are the subject of this report. Understanding these issues requires a survey of the production simulation field, paying special attention to the special problems of representing demand-side resources in the same framework traditionally used to evaluate supply side resources.

We begin our analysis in section II with an overview of the production cost modeling process and its role in least-cost planning. This overview surveys the literature on standards and practices for this field and defines the concept of calibration used in the remainder of the report. This discussion explains the methodological background within which the choice of model dilemma arises.

Following this general methodological review, we focus in section III on the generic issues involved with calibrating LMSTM to the detailed models used by PG&E, notably GRASS. The two principal issues of concern are the definition of system boundary and seasonal aggregation. The system boundary issue arises because PG&E dispatches all electric resources in Northern California including those owned by municipalities. The economic consequences of demand-side programs, however, fall only on PG&E ratepayers. Accounting for these differences is the system boundary problem. The seasonal aggregation issue arises because LMSTM uses four seasons, whereas the detailed models simulate twelve months per year. The resolution of these two generic issues determines how individual resources are represented in LMSTM.

The next task is the characterization of the generic resource types available to supply energy to PG&E. One of the special characteristics of the PG&E system is the wide variety of resource types available. In section IV, we define the problems associated with each resource type and characterize the approach we have adopted for representing it. These discussions are placed in the context of modeling conventions used by PG&E itself and those used by regulators and other producers who sell power to PG&E. The resource types discussed are: (1) Hydro, (2) Northwest Imports, (3) Oil and Gas Units, (4) Geothermal, (5) Nuclear, (6) Qualifying Facilities under PURPA, and (7) Generic Resources.

Finally, in section V we report detailed results of the calibration process for selected representative years. These results include not only the commonly cited figures on annual production by fuel type, but also details of the marginal costs. The importance of marginal cost cannot be over-emphasized; demand side analysis requires an adequate representation of time-varying marginal costs. The value of demand-side resources can be measured only if the

[1] All references to LMSTM in this paper refer to version 2.0, which is the one exclusively used for this project. The version previously used at PG&E is a hybrid between version 1.4 and version 2.0. Further, as of January, 1987, the primary version in use will be 3.0, which contains considerable enhancements over version 2.0.

marginal value of different load shape impacts is calculated reasonably. We report some results on how marginal cost varies with different input assumptions.

II. MODELING STANDARDS AND PRACTICES

It is remarkable how little attention has been given to the development of standards for production simulation. Despite the importance of this kind of modeling to electric utilities, there have been relatively few detailed studies comparing various techniques and simulation features. In this section we provide a technical review of such studies aimed at the professional with some experience. Attention is paid to important features of simulation models. This review motivates a definition of the calibration process used for LMSTM. Finally, we discuss those aspects of production simulation that are most relevant to least-cost planning issues.

1. Review of Features and Their Evaluation

Production simulation is a field dominated by traditions of practice rather than empirical testing. Models cannot be easily "tested against reality" because the probabilistic nature of the simulation is not strictly comparable to particular outcomes. The models calculate an expected value of an underlying random process. The real world results of power system operations represent an individual outcome of that process. It is a common predilection in such circumstances to increase the complexity of the model simulation to approximate reality more closely. Under these circumstances, the most complex model becomes *de facto* the most correct model. The only limitation on this process is computational resources. Our discussion in this section describes the evolution of production simulation techniques and the evaluation of features by pragmatic or systematic procedures.

a) Monte Carlo Chronological Models vs. Equivalent Load

The earliest tradition of probabilistic simulation was the Monte Carlo technique. This approach is very computation intensive. It involves a large number of deterministic simulations which reproduce the utility system dispatch under assumed conditions on the outage status of generating units. To account for all possible outages and produce an expected value, a large number of "draws" is required. The computational burden of Monte Carlo makes it difficult to vary the resource parameters of the simulation. This lack of flexibility inspired the introduction of the equivalent load approach.

The equivalent load technique is based upon a non-chronological load representation. By rearranging hourly loads for a given period in a monotonically declining order, the demand for power can be interpreted probabilistically. This places it on the same conceptual basis as the randomly available supply, and allows a computationally efficient procedure for dispatching generation against load. The equivalent load is the load requirement facing a given unit after considering both the dispatch of all previous units against load requirements and the "outage loads" created by the random failure of those previous units. This technique was introduced by Balerieux and associates in Belgium and Booth in Australia in the late 1960's and early 1970's (Balerieux, *et al.*, 1967 and Booth, 1971). It quickly won acceptance among utility practitioners. The great advantage was that a multiplicity of resource scenarios could be evaluated.

The equivalent load approach has dominated production simulation until quite recently. Constant improvements in computational efficiency and power, however, have eroded the pragmatic justification for the equivalent load procedure. Concerns about the chronological constraints on power system operation have brought attention back to the hourly framework. Among the models which adopt a chronological framework are EPRI's BENCHMARK, TVA's POWERSYM and the models based on it, and LMSTM.

One of the unique aspects of LMSTM is its "smart" approach to Monte Carlo simulation. Instead of relying on a large number of random draws to produce the outage distribution, it selects a small number of cases based on an intelligent selection from the outage distribution. The outage distribution is represented analytically, in a manner discussed in section II.1.*b)*, and this representation allows accurate characterization of the outage states with only a small number of direct simulations (Stremel and Bubb, 1986).

b) Analytic Approximation Techniques

A major part of the probabilistic aspect of production simulation involves characterization of the outage distribution. The most burdensome procedure involves the explicit enumeration of all possible outage states and their probabilities. The equivalent load procedure makes this enumeration implicitly and iteratively, accounting for a good part of its efficiency. An alternative approach can be found by appealing to the more sophisticated elements of probability theory. Because the outage distribution is the sum of random variables, in the limit, it approaches a normal distribution. In most practical cases, however, the naive normal approximation is inadequate because the systems are too small and the individual outage distributions are too skewed. There are several techniques that correct for this.

The Edgeworth and Gram-Charlier approximations are orthogonal series based on the sums of moments of the individual outage distributions. They are reasonably well behaved, but do exhibit pathologies (Levy and Kahn, 1982). When combined with such a representation of the load duration curve, these approximations are less pathological and can be used for production costing in the equivalent load context. Comparisons with numerical versions of the equivalent load method published for a single case are reasonably favorable (Caramanis, Stremel, Fleck, and Daniel, 1983). This approach is used in EPRI's Electrical Generation Expansion Analysis System (EGEAS) model, among others.

Another approach to correcting the naive normal approximation is the large deviations technique. This involves a transformation of the outage distribution away from the "tail," or the low probability region, to the central region of the normal curve. Here the normal approximation works better. The outage distribution can be accurately approximated and the results "shifted back" to where they belong. The large deviations method has been applied to production costing in the equivalent load context, using a numerical representation of the load duration curve (Mazumdar and Wang, 1985). Test results are consistent with numerical tests reported above for the Edgeworth approximation, but there is no commercial model available that uses this approach.

A third analytic approach is called the mixture of normals approximation, or MONA (Pacific Gas and Electric, 1985). The motivation of this approach was to produce a better fit of an analytic function to the load duration curve than the Edgeworth or related series. The solution technique is to divide the load duration curve into a small number of regions and fit a cumulative normal curve to each region. These are then spliced together for a representation of the entire curve. The outage distribution is represented by a mixture of normals. This method weights outages of n units by the probability of n units being out, these probabilities being approximated by a normal distribution. Highly detailed numerical studies of this method show favorable results compared to numerical or Edgeworth type approximations for the equivalent load context. The UPLAN production cost model uses the MONA method.

Any of these approximations to the outage distribution could be applied to the "smart" Monte Carlo approach as well as to the equivalent load approach to production simulation. In fact, there are any number of ways to approximate the outages analytically so that an intelligent selection of representative outage cases can be selected. LMSTM uses what amounts to an Edgeworth type characterization of the outage distribution.

c) Marginal Cost Definitions and Reports

Production simulations contain enough information internally to report on marginal costs as well as total production cost, but, remarkably the models have not typically made such calculations explicitly until quite recently. Indeed, literature on the definition of marginal cost concepts in production simulation is sparse. In the last few years, more attention has been focused on this problem.

Within the equivalent load context, there are two reasonably explicit definitions of instantaneous marginal cost.

1. The first definition focuses on marginal cost at a particular load level. The probabilistic nature of the simulation implies that no one unit is the marginal producer at a given load with certainty. Units only have a probability of being marginal at a given load level. These probabilities can be computed from the dispatch order. For a given unit at a specified place in the dispatch order, the probability of being marginal is the difference between the loss of load probability (LOLP) without that unit and the LOLP with that unit (see Bloom, 1984).[2] A unit is marginal to the degree that it reduces this probability. The marginal cost then is just the operating cost of marginal units weighted by the probability that the unit is marginal. Several models use this definition of marginal cost. These include Pacific Gas and Electric's GRASS, UPLAN, and PROMOD.
2. An equivalent definition is based on the expected fraction of marginal cost attributable to each unit. This is just the average probability of a unit's being marginal when all loads in a period are considered. This definition is just a notational variant of the previous one, but it is computationally less demanding and conceptually more valuable to the model user. The average probability of a unit being marginal in a given period can be read from projections onto the probability axis of the equivalent load duration curves before and after the dispatch of that unit. This involves fewer computations than the averaging of unit marginal probabilities over all loads.

More important than the computational differences are the conceptual ones. It is useful to decompose the marginal cost in a given period into contributions from different resources. This is especially helpful if the resource mix is varied with corresponding variation in cost. Aggregating the unit average marginal probabilities by fuel type provides a convenient representation of the dispatch and the cost structure. This procedure is implemented in the Environmental Defense Fund's Electric-Financial (ELFIN) model and is available in the PRINT 10 report of PG&E's GRASS model. LMSTM also provides output which can be interpreted in this manner, but the output report is not designed with this in mind, and is, therefore, not easily usable in this manner. Conceptual differences between marginal cost in the equivalent load framework and the chronological framework are discussed below.

There are numerous subtleties involved in the definition of marginal cost in the equivalent load context. One of the more prominent involves the role of storage hydro resources. Energy of this type does not change in response to load changes. Nonetheless, it is dispatched at the

[2] LOLP is the probability that load cannot be met with a fixed randomly available supply.

''high'' end of the load duration curve, i.e. against peak loads. The methods outlined above assign marginal probabilities to storage hydro. It is common to allocate this probability to the thermal block immediately ''below'' storage hydro in the dispatch. Such an allocation is necessary so that a marginal cost may be associated with these probabilities since there is essentially no operating cost for storage hydro. The convention of assigning storage hydro marginal probability to the next lowest thermal unit is based on various opportunity cost arguments (EDF, 1986). It is also worth noting that hydro pumped storage output is also treated the same way for marginal cost purposes.

In principle, the concept of marginal cost is more straight-forward in the chronological domain. In a deterministic simulation, the marginal unit at a given load is the last unit serving load. When the probabilistic aspect is added, however, then marginal cost again becomes a probability weighted average of various contributions. Within the chronological approach, it is not necessary to use Monte Carlo simulation. An alternative is to adapt the Balerieux-Booth approach to a chronological load shape. This is the method used by POWERSYM. In this case, the LOLP-type definition of marginal cost formalized by Bloom for the equivalent load domain is also used. It is typical in all chronological models to ''shave'' loads with hydro resources so that only thermal resources and their outages are considered in the probabilistic dispatch and marginal cost calculations.

d) Unit Commitment and Non-Economic Constraints

We have not discussed the dispatch order in production simulation as yet, but this is a source of wide variation among models. Although all simulations purport to perform an ''economic'' dispatch of resources, there are many engineering and contractual constraints that limit the use of resources. We discuss two of the more important ones, which are related to each other. These are unit commitment requirements for ''slow-start'' units and ''must-run'' constraints.

The unit commitment problem is associated with the need to meet peak load requirements and the inability to turn units off and on quickly. Consequently, if a unit is required to meet peak loads, it may not be possible to turn it off during low load periods. This means it may have to run at some minimum level of operation even at times when less expensive energy is available. The standard representation of such slow-start units is the definition of a minimum capacity block with a high operating cost, and one or more larger capacity blocks with lower operating costs. A production simulation model which simulates the commitment problem identifies the number of such slow-start units required to meet peak loads in a given period, and dispatches all the minimum blocks of these units in the ''base load.'' The remaining capacity blocks of these units are dispatched in economic order.

Apart from slow-start units there are other circumstances in which capacity must be dispatched out of economic order. One prominent example is the set of Qualifying Facilities (QF's) producing power under PURPA. Because utilities are required to purchase QF power, their output must be taken ahead of other resources. If it were not, then ordinary economic dispatch would result in the curtailment of QF's, violating the obligation to purchase under PURPA. To model this requirement, production simulations designate QF's as must-run units. This means they are dispatched in the base load ahead of all those resources which are dispatched economically.

It should be noticed that both the minimum blocks of committed units and the entire capacity of must-run units are treated identically. This can result in some semantic confusion since it is possible to designate the minimum blocks of committed units as must-run blocks. Indeed, in production simulation models which have the must-run feature, but not the commitment feature, there is very little distinction between the two phenomena. Among the models in this category are ELFIN Version 1.30 and EGEAS Version 3.0. In models of this type, the user can approximate commitment by an appropriate choice of must-run blocks from among the slow-start units (Kahn, 1985a and Kahn 1986).

It is possible that power systems are so constrained that the set of units that must be dispatched uneconomically is large relative to load. One reason for this may be the contractual terms of purchase contracts. PG&E, for example, purchases geothermal steam for its Geysers units under a contract that limits curtailment to those situations in when nuclear units have already been curtailed maximally. We discuss this contract in section IV below. In situations such as this, it becomes necessary for the production simulation model to employ devices that order the dispatch sequence of the must-run units. One common device for this purpose is the use of "dispatch penalty factors." These act like shadow prices in the sense that they tell the model which resources are more or less valuable within the must-run class.

The importance of these details can be seen in their impact on marginal cost. The presence of must-run constraints has the effect of lowering marginal cost compared to the case where there are no such constraints. By forcing large amounts of energy into the base load, these constraints also force lower cost resources up the loading order and "onto the margin." In the PG&E case, for example, the Geysers must-run constraint forces nuclear units, which have lower variable costs than geothermal, to be marginal producers more of the time than geothermal. This lowers marginal cost. Detailed examples can be found elsewhere (Kahn, 1985a and Kahn, 1986).

e) Endogenous Pricing

Not all resources have exogenously determined prices which can be input to a production simulation. Examples of those with endogenous prices are QF's, split-savings purchases, and PG&E geothermal units. In each case, the pricing formula differs. What they have in common is that the production simulation model must perform some calculation either before or after the simulation to determine the price of the resource. We discuss these cases briefly.

The simplest case of endogenous pricing involves pre-processing before a particular simulation. Split-savings purchases based on an average of known fuel costs is one such case. The PG&E geothermal price is essentially also of this kind with the added complexity that the price in a given year depends on the dispatch results from the previous year. This calculation could be made by successive model runs and an exogenous calculation, or it could be done internally by the model as long as the simulation period started one year ahead of the period of interest. The only model which currently calculates the price in this latter fashion is ELFIN Version 1.30.

More complex endogenous pricing requires the results of a given simulation to be fed back into the same simulation. A simple example of this kind is marginal cost pricing of QF's. This case is simple because all that must be done is the marginal cost for a given period be assigned to a non-marginal resource. There is nothing about the simulation that is affected by the pricing. The more difficult cases involve multiple model runs to determine price. This is necessary if QF's are priced under an "avoided cost" rule that assigns price on the basis of cost differences with and without QF's in the supply mix. In this case, it is necessary to perform two simulations, but at least the outcome does not affect dispatch results. Examples of this calculation are given elsewhere (Weisenmiller and Yardas, 1986 and Kaiser Engineers, 1986). An alternative

QF pricing formulation based on the "zero-intercept" method also requires multiple simulations (see Kahn, 1985b).

Probably the most complex endogenous pricing is associated with fully integrated power pools. Such pools commonly use a split saving formula for pricing, which requires simulation of both the pool as a whole and the individual company on an isolated basis. These simulations can get quite complex when different rules are used to dispatch resources in each context. Details of these issues including the necessity of using a chronological approach for this problem are given in Bloom, 1984b.

2. Calibration of Simpler to More Complex Models

The general rule cited above that more complex simulation models are preferable to less complex models runs into pragmatic limits. Complex models can be difficult to debug, require substantial computing time and have too many free parameters to yield uniform results. For strategic analysis or the study of special problems, it may turn out to be more efficient to use a simpler rather than a more complicated model. Least-cost planning, in which supply side resources are traded off against demand-side activities, may well be such a situation. A principal concern associated with the use of simple models, however, is the extent to which their results are consistent with the more complex models. To address this concern, we describe a process for the calibration of models.

As any experienced modeler knows, there are always procedures available to make the results of one model appear consistent with those of another model. The task of calibration consists of defining the extent of legitimate procedures and measuring the deviations in results. In the domain of production simulation, it is useful to separate procedures into the specification of inputs and the specification of features. Measurement of results should include both the outcomes on production by fuel type and marginal cost.

a) Specification of Inputs

Model calibration often requires changes in input specification to account for differences in the level of detail between models. In passing from the more complex to the simpler representation it is necessary to aggregate some aspects of the input data. The principles on which the appropriate level of aggregation are chosen may be either formal limitations of the simpler model or the pragmatics of scenario design. Block representation of units provides an example of formal limitations. Complex models may allow many capacity blocks for a given thermal unit; the simple model may allow only one or two. To model or approximate the unit commitment problem discussed in section II.1.*d)*, it is necessary to have at least two capacity blocks per unit, one being the minimum block, the second aggregating all others together. Other LMSTM aggregations worthy of mention include the seasonal aggregation of hydro resources, the aggregation of hydro categories, the multiple unit representation of Northwest and geothermal power, and the aggregation of oil and gas units into classes.

The pragmatics of scenario design generally work in the opposite direction, limiting the amount of aggregation. We separate Northwest purchases made by PG&E from those made by municipalities for pricing purposes; i.e. municipal purchases are not charged to PG&E customers. Distinctions among municipal purchasers are unimportant. The same logic applies to geothermal units. For hydro units, it is important to treat QF separately both because they are paid differently than the rest of hydro, and because the bulk of the hydro expansion occurs in this category over the relevant simulation period. Even if there were no pricing difference, it might be convenient to separate the expandable resources from those which are near their development limits.

There is a subtle distinction between input specification determined by aggregation decisions or scenario pragmatics, and those associated with the approximation of features from one model to another. Because models differ in the features they embody, it is common practice to adjust input parameters to approximate the operation of features that the simple model lacks and the complex model embodies. Because these parameter specifications are more associated with model features, we discuss them separately.

b) Feature Approximation

The real differences between models involve algorithms which represent various features of power system operation. Simple models tend to suppress constraints which limit the flexibility of real systems in optimizing operations. Complex models purport to represent more of these constraints. There are two distinct approaches that can be taken to reconciling differences in algorithmic representation of system features. We call one of these approaches "output calibration" and the other "qualitative calibration." By the term output calibration, we refer to the process of specifying as input to the simple model every possible detail of output ascertainable from the complex model. An exercise of this kind is pedestrian in the sense that a complete implementation of it would effectively remove the role of simulation from the simple model, and reduce the concept of calibration to arithmetic verification. Although intellectually uninspiring, output calibration has its defenders (Gerber, 1985). The ideal model for such purposes would have no internal algorithms to represent system features, but only a completely flexible capability to accommodate any pre-specified set of inputs.

The term qualitative calibration refers to a process in which the modeler defines the main simulation properties of a complex model and selects approximations to them appropriate for the simple model. The distinction is best illustrated by a few examples. We focus on two important production simulation phenomena for this purpose: maintenance scheduling and unit commitment. We begin by discussing the maintenance scheduling algorithms used in simple and complex models.

The most detailed approach to the maintenance problem seeks to construct a unit by unit schedule of maintenance. The problem is constrained by manpower limits, geographical constraints, economic and reliability considerations, among others. The problem has commonly been formulated in the dynamic programming framework where the objective is to levelize and minimize system risk as measured by LOLP (Wu and Gross, 1977). At the other extreme is the simple procedure known as "distributed maintenance" (Stremel and Jenkins, 1981). In this approach all units are derated in such a way as to levelize reserve margins from period to period.

Simple models generally use the distributed maintenance approach as the default, but allow the user the option of specifying maintenance in a more particular manner as an input choice. The qualitatively most important units to schedule for maintenance are the nuclear plants. Their capacity is large and their variable costs are low, so that their presence or absence from the supply mix has a large impact on cost. A reasonable approach to qualitative calibration in this case therefore would be to specify nuclear maintenance only and let the simple model use the default algorithm for everything else. It might be argued, however, that large oil and gas units should also have specific unit maintenance schedules. The question then becomes one of defining the boundary between large and small. As this process continues in the direction of greater and greater detail, we begin to converge upon output calibration. This is not desirable.

One guideline appropriate for the "how far to go" problem in qualitative calibration is the matching of levels of detail. For the maintenance scheduling problem in a model such as LMSTM, there is a natural boundary at exogenous scheduling of the nuclear units. The reason for this is the seasonal aggregation feature of LMSTM. Detailed models typically report simulation results at the monthly level. This corresponds approximately to the maintenance requirement of oil and gas units. LMSTM uses four "seasons" of a length determined by the user. If the average season were three months, this would correspond to the nuclear unit maintenance requirement. Thus in LMSTM it would be difficult to represent maintenance for oil and gas units because the simulation period is too long for a meaningful disaggregation. Features cannot be used at a level of detail greater than the "least common denominator" of a model's representation. Thus, for the maintenance problem in LMSTM we have a natural guideline for qualitative calibration. It is necessary to define such natural guidelines for other features.

Next we consider the unit commitment problem. It is instructive to compare the approach of a detailed model such as GRASS with simpler models such as ELFIN and LMSTM. In GRASS, the commitment problem is solved each month of the simulation, with the assumption that units are committed for the entire month. At the opposite extreme, LMSTM solves the problem for each daytype. In Version 1.0 of LMSTM, it was assumed that units could be shutdown for periods on the order of several hours and then restarted later in the day if necessary and economic (DFI, 1984b). The amount of operational flexibility implied by this representation exceeds what is characteristic of the PG&E system. Version 2.0 of LMSTM allows a must-run treatment of committed units which is much closer to the reality modeled by GRASS (Stremel, 1985). Thus a feature enhancement to LMSTM effectively eliminates the algorithm gap between GRASS and Version 1.0. The user of LMSTM must still know which feature to use; i.e. the SHUTDOWN or MUST RUN characterization.

ELFIN lacks compatibility with GRASS in the area of unit commitment. There is no internal logic which chooses committed units. There is only exogenous, user-specified must-run capability (EDF, 1986). Even this is limited to a single list of must-run blocks for committed units that is invariant over the simulation period. In practice the user must specify (or guess) the average level, of committed units. This is the standard procedure (Weisenmiller and Yardas, 1986; Kahn, 1985a, 1985b, 1986). Because of seasonal variation over the year, and year to year changes in the average level this exogenous choice of an average commitment level is sometimes too high and too low at other times. There is no *a priori* way of knowing what the net effect of these deviations is. *Ex post*, of course, one can always check the details in the manner of output calibration, if the appropriate results are available.

Where feature approximation relies exclusively on exogenous choice of input specification, i.e. where the model lacks internal logic to approximate a more detailed feature, then calibration over a range of variables is uncertain. To continue with the unit commitment example, suppose we want to examine a significant change in the resource mix from a calibrated case to one which is substantially different. With a different resource mix, a different pattern of committed units emerges. How different will this pattern be? How should the approximating choice made in the calibrating case be changed, if at all? Does it matter if no change is made? Unfortunately, there are no compelling answers to these questions.

c) Decision Criteria

Given the calibration difficulties outlined above, how should one proceed in choosing a simple model for the desired tasks? Common sense suggests that perfect calibration and certainty is unattainable. The following tests provide reasonable guidance.

Suppose the simple model has been selected, and that a "base case" simulation is available for the detailed model. This is the generic case. Then the input specifications are made on pragmatic grounds. Feature approximation is constrained by differences in the level of aggregation between the models and the degree to which features differ. To the degree possible, output calibration is to be avoided. This is not always possible, but it is equally true that crucial intermediate results from the detailed model are not always available. Some exogenous calibration choices have to be made. These should be noted explicitly and sensitivity tests performed to explore the choices. The decision criteria is a best fit to annual production by fuel type and marginal cost. Let us discuss the use of these criteria.

All production simulation models report annual energy production aggregated to some level or other. It is unreasonable and unnecessary to expect that models replicate results at the level of individual units or even small groups of units. The crucial results for costing purposes are production at the level of fuel types. Models cannot be said to be calibrated if results are wildly divergent at this level. Broadly speaking, there are two kinds of fuel types, those which are base loaded and those which are marginal. While this distinction is not iron-clad, it is usually identifiable. Base loaded resources ought to be replicable to within 1-2% at the level of annual energy production. Marginal resources are harder to replicate because feature differences come into play. Deviations of more than 10% in annual production by fuel type for marginal resources mean that calibration is not particularly good.

Most models calculate marginal cost, or can be used to approximate marginal cost. We do not discuss the case where approximation is necessary (see Kahn, 1985b). For calibration purposes, we are interested in both the level of marginal cost and its structure. A reasonable target for replication of marginal cost averaged over a simulated year is 5-10% deviation from the detailed model. It is important to investigate seasonal and diurnal variation in marginal cost because this provides insight into what causes deviations between the simple and the detailed models. A good calibration captures the broad pattern of these variations. The closer the match on the seasonal and diurnal level, the more confidence one can have in the robustness of results with respect to variations in the resource mix from the base case. Because we usually cannot explore resource variations we have only the marginal cost results to provide indications of what can be expected as the supply/demand balance changes.

A useful calibration examines the details outlined above for a number of years in the simulation period. The years selected should differ as much as possible with respect to the supply/demand balance. Because a number of "tunable parameters" is available for which exogenous choices must be made, it is prudent to examine the sensitivity of results to these choices. These tests also provide guidance in defining the nature and extent of mismatches between the models.

The generic calibration exercise which we have just defined does not address the most difficult question, the choice of the "optimal" simple model. This problem is in some basic sense indeterminate. A good simplification is good only relative to a particular question. In practice there are always real world constraints on the choice of a simple model, which are outside the domain of calibration criteria. These constraints include cost, compatibility with other users or models, ease of use, and computer resources. Real world constraints are often the determining factors in model choice. Calibration is usually an exercise in evaluating how bad a choice was forced upon the user. Only after the calibration exercise do we know what we would have liked to know before we began. Even then, we probably don't know enough about alternative models to make a choice that is guaranteed to be better. It is with this reality in mind that we turn explicitly to the choice of model question for least-cost planning.

3. Application to Least-Cost Planning

In this section we describe the properties of a production simulation model that are useful for least-cost planning and compare LMSTM to other models for this purpose. The three issues we focus on are: (1) chronological characterization of marginal costs, (2) valuation methods for demand-side programs, and (3) integration of production simulation with the fixed costs of utility investment. It is useful at the start to say something about what we mean by least-cost planning.

The term least-cost planning has entered the energy lexicon recently in a specific context, namely, the interest in trading off demand-side activities against new supply-side investments. The original thesis advocated by the proponents of least-cost planning was that the full social cost of demand-side initiatives in the electricity market is less than the cost of new supply (Sant, 1979 and Lovins, 1981). There have been many subsequent twists and turns in this discussion, including the very different economic conditions resulting from the excess supply of electricity which developed in the 1980's, and the uncertainties of accounting for consumer response to demand-side programs. The excess supply issue is a transient phenomenon that disappears sometime in the 1990's. The consumer behavior issues are more fundamental and obscure. They make it difficult to perform the social cost analysis, which is at the heart of the least-cost planning problem. It is not necessary to resolve these, however, to perform an interesting and useful analysis of demand-side programs. Much can be learned by examining the value of demand-side impacts and leaving the questions associated with achieving those impacts, i.e. the consumer behavior issues, for a separate analysis. This factorization of the least-cost planning problem does not represent a solution to the conceptual difficulties, but it does define a tractable research agenda. It is in this spirit that we investigate least-cost planning by systematically studying the value of demand-side programs.

a) Chronological Characterization of Marginal Costs

Demand side programs have load impacts which vary by time of day and season of the year. Further, the value of a unit load impact varies more or less with the chronological variation in marginal costs. Production simulation in the chronological domain, therefore, is the procedure best adapted to analysis of demand-side resources. It is the chronological feature of LMSTM that is cited by utilities that use this model for the analysis of demand-side programs (Nordell, 1986 and Niagara Mohawk, 1985).

It is instructive to examine load duration curve models. Many of these models have procedures which allow for chronological representation of resources. Such resources are commonly called "time-dependent." The standard procedure in the load duration curve domain is to preprocess time dependent resources before the simulation. After these resources are subtracted from the chronological base case load model, a load duration curve is constructed and simulation begins. Models which use this procedure include GRASS, ELFIN, UPLAN, and EGEAS. There is no possibility of assigning marginal value to the various elements of the time dependent load shape. This is really a consequence of the equivalent load formulation of the problem. By the time the thermal resources and their outages have been convolved into the load duration curve, there is no natural mapping from the outputs back to the chronological loads. The marginal costs defined in section II.1.*c)* cannot really be attached to particular hours except by modeling convention. Although these translations are often claimed for load duration curve models, when any serious need for time-differentiated marginal costs arises, planners use a chronological model. Transactions representation is the main area in which this arises (Bloom, 1984b).

There is an intermediate position occupied by models which segment loads into three load duration curves per month. Each such curve represents a period in which cost conditions are more homogeneous than the single load duration curve model. This is the approach taken by PROMOD. EGEAS also has this capability. While some movement in the direction of chronological representation is desirable, it is difficult to say *a priori* how much is enough.

The LMSTM load representation, while chronological in spirit, is not without aggregation problems of its own. The simulated year is decomposed up to 16 typical daytypes, four for each season. The user is free to specify the decomposition for his or her own system, but this involves some work and no small amount of judgment. The result is an aggregation of an average 23 actual days per daytype (22.81=365/16), and typically a lot more for some of the "average" days in a season. As we discuss in section III, there are many issues involved in the aggregation from the supply side perspective. Combined with the demand-side issues, it becomes plain that LMSTM is not a chronological model in the literal way that POWERSYM or BENCHMARK are. The LMSTM vendor is developing a new version of this model which embodies the typical week for each month approach that is the load model basis of many equivalent load production simulations. This reduces aggregation problems now inherent in LMSTM.

b) Valuation Methods for Demand-Side Programs

Least-cost planning must be distinguished from cost-effectiveness analysis of demand-side programs. The latter is based on static evaluation methods which assume that demand-side activities are sufficiently small that they do not affect the utility's overall cost structure. Large scale demand-side programs can be thought of as analogous to Qualifying Facilities under PURPA. QF's in large supply reduce the marginal costs of a utility compared to the case without them. The avoided cost pricing rule under PURPA defines the value of QF's to be the production cost savings due to their presence in the supply mix. This amounts to the difference between production costs without the QF's and production costs with the QF's. This "QF In/Out" value is generally greater than marginal cost with QF's in the supply mix, and increasingly greater as the relative contribution of QF output to total production increases (Jabbour, 1986; Kahn, 1985a, 1985b; Kaiser Engineers, 1986).

Cost-effectiveness analysis for demand-side activities is based on marginal costs after the load impacts are already accounted for. This is equivalent to pricing QF's on the basis of marginal cost with QF's in the resource mix, rather than on an avoided cost basis. If the load impacts are small, there would not be much difference. In the case of large load impacts, the cost-effectiveness method undervalues the demand-side resources.

From a modeling perspective, an important benefit of demand-side programs is a reduction in unit commitment requirements compared to the case with supply side resources. To capture this benefit, it would be convenient for the production simulation model to have the commitment feature. LMSTM has this capability, so we do not have to resort to exogenous manipulation to capture changes in commitment due to demand-side programs.

c) Integration with the Fixed Costs of Investment

Production simulation is only half the resource planning problem. The other half involves the fixed costs of investment. The traditional capacity expansion analysis procedure involves comparing the fixed investment costs of supply alternatives against their production cost benefits. To perform this kind of analysis it is useful to have a model in which the fixed costs of electricity production are taken into account.

Detailed production simulation models usually lack the ability to account for fixed costs. Models such as GRASS or PROMOD, are designed to be used in a decentralized environment where separate departments perform each of the analysis tasks involved in expansion planning. Highly detailed financial models custom tailored to the regulatory rules of individual companies are used for fixed cost analysis. The philosophy of this approach is based upon the desirability of substantial detail in analysis. The cost of this decentralization is a lack of flexibility. Strategic analysis conducted over a wide range of input variables is not possible with separate detailed models. It was largely to accommodate the need for analysis flexibility that the simpler production simulation models we have been discussing were created. These simpler models generally come equipped with a fixed cost component which allows for an integrated expansion analysis.

The critical choice involved in representing financial variables is whether the model is used for optimization. It is tempting to conceive of the expansion planning process as an optimization problem. The objective of minimizing utility system costs naturally suggests this framework. Implementing an optimization approach, however, involves many simplifications. It is not our intention to make a complete survey of these problems. A useful discussion can be found in the description of the EGEAS model, which is designed with optimization as a primary aim (Caramanis, Schweppe, and Tabors, 1982). We focus primarily on the financial representation, using EGEAS as a convenient example.

EGEAS uses the fixed charge rate concept to represent the fixed costs associated with investment. Fixed charge rates are a simplification of the capitalization, rate-basing and depreciation flow of funds that are the actual process through which investment is financed and costs are recovered. The fixed charge rate is a one-parameter representation of this process which facilitates the comparison of engineering alternatives (Gulbrand and Leung, 1975; Leung and Durning, 1978). The computational advantages gained by this process come at the price of suppressing the financial dynamics. This price is high when financial constraints limit the investment choices of firms.

In practice, users of EGEAS find that capital intensive plans are selected which minimize total costs, but are viewed as infeasible by financial executives of the utility. There are many ad hoc procedures that the user can employ to escape this outcome. They all amount to imposing some exogenous constraint on what the model can choose. None of these can really be based on financial criteria because EGEAS does not represent the financial detail with enough explicitness. Thus we are left with an optimization that is either abstracted from real world constraints, or fails to represent them internally.

The alternative to optimization is a financial representation that is more complete, but does not allow internal decision logic on investments to be incorporated into the modeling. In this framework, the model represents at some level the three basic financial reports of the firm, the income statement, the balance sheet, and the sources and uses of funds statement. There is substantial latitude possible in the level of accounting detail. The basic purpose of this approach is to allow for sensitivity study of various alternatives in a way which provides year by year financial information, as well as present value costs. Among the models which take this approach are ELFIN, LMSTM, and UPLAN. UPLAN also has optimization capability. LMSTM provides rate class data calculated by allocating total costs to different classes of customers.

III. REPRESENTING PG&E IN LMSTM

The previous discussion provides background for our choice of LMSTM as a model to use for least-cost planning of the PG&E system. As should be clear, there are no definitive arguments determining the "best" model for these purposes. The *a priori* benefits of LMSTM are its chronological dispatch, the broad scope of its simulation capability (including the unit commitment and marginal cost features), and the integration with financial data. A pragmatic benefit is the existence of LMSTM users in the Economics and Forecasting Department at PG&E. LMSTM was installed at PG&E in 1984 (DFI, 1984) and has been used since then. In this section, we turn to the generic questions involved in modeling PG&E in LMSTM. Section IV is devoted to the detailed issues of representation for different resource types.

We concentrate on three issues in this section. First, we discuss the definition of the system boundary. This is the problem of accounting for non-PG&E resources that are dispatched by PG&E in Northern California. Second, we address the definition of seasons. The issue is defining rules which map the twelve months per year simulated by detailed models into four LMSTM seasons. Third, we discuss how to use load data to construct the generic daytypes represented in LMSTM.

1. System Boundary

PG&E either dispatches centrally or coordinates all the interconnected generation resources in Northern California. These include federal projects that deliver power to municipalities and the resources owned by municipalities. This dispatch optimizes the operation of the power system within the physical and contractual constraints, of which there are many. This close integration means that it is not possible to separate ownership in the dispatch so as to distinguish the marginal costs due to PG&E resources only from those of the area. It is only possible to speak of marginal cost on an area basis.

Total operating cost is, of course, a different matter. Here ownership and control of resources determines who is responsible for costs. This means that we must dispatch PG&E with all Northern California loads and resources in the mix, but only carry the production costs of PG&E resources over to the financial module of LMSTM. The technical means of achieving this joint objective is the use of LMSTM's indirect cost capability.

LMSTM allows generating resources to be characterized by both direct and indirect operating costs. The original motivation of the indirect cost concept was to model the social costs of air pollution. It can be adapted to our purpose because its formal properties are essentially what we require. The indirect costs are used, in addition to the direct costs, to determine the economic place to dispatch a resource. They are not reported to the financial module. In our case, we simply assign a zero direct cost to municipally owned resources and an indirect cost equal to their cost to the agency involved. This yields a correct dispatch and cost accounting.

2. Seasonal Definition

It is difficult to reduce 12 seasons into four LMSTM seasons because the PG&E system exhibits such large monthly variations in resource availability, load, and cost. This is one of the most difficult choices involved in using LMSTM to model PG&E. Table 1 shows the final distribution of seasons and daytypes chosen. Appendix A details the process by which this choice was made. In this section, however, we discuss the major issues to be considered in defining a seasonal representation. A particular choice weighs these factors differently depending on the application of the model. Because the application we have in mind is of a long-term planning nature, we place more weight on these factors than on short-term factors.

Table 1
Seasonal Breakdown For LMSTM Representation of PG&E

ID	season name	months	daytype name	ID	no. days	no. hours
1	WINTER	jan. - apr.			120	2880
			WEEKEND	1	37	888
			NORMAL	2	39	936
			PEAK	3	38	912
			EXTREME	4	6	144
2	SUMMER	july - sep.			92	2208
			WEEKEND	1	29	696
			NORMAL	2	45	1080
			PEAK	3	14	336
			EXTREME	4	4	96
3	FALL	oct. - dec.			92	2208
			WEEKEND	1	30	720
			NORMAL	2	22	528
			PEAK	3	34	816
			EXTREME	4	6	144
4	SPRING	may - june			61	1464
			WEEKEND	1	20	480
			NORMAL	2	35	840
			PEAK	3	3	72
			EXTREME	4	3	72

The first consideration is resource variability. The main source of such variability is hydrological fluctuation. Water resources vary substantially over the year. Late spring is the most abundant period due to the mountain snowmelt. Correspondingly, late summer and early fall exhibit the lowest hydro conditions. This pattern involves not only California resources, but also the availability of energy imports from the Pacific Northwest, hydro energy being the main source of seasonal variability in the Northwest as well.

The other principal source of resource variation over the annual cycle is the maintenance schedule of nuclear plants. The maintenance scheduled for nuclear plants occurs predominantly in the period from January to April and October to December. Exogenous constraints, due to the nuclear fuel cycle are presumably responsible for the details of forecast schedules.

Another approach to defining seasons rests on marginal cost variations. The monthly fluctuations in marginal cost are the net outcome of the supply and demand balance. The resource variations discussed above, and others of lesser magnitude, are paralleled by load variations. A segmentation of the months into four seasons could be designed for a particular resource plan by attempting to minimize the variance of marginal cost within a season and maximizing it across seasons. Among the difficulties with this approach for demand-side planning is that marginal costs change as a result of the programs analyzed. Thus, an optimal segmentation before these programs may not necessarily be a very good one after them.

The real difficulty is that LMSTM does not offer enough choices. It would probably be best to have six "seasons" to capture the marginal cost variation. Suggestive data in this regard were assembled in the PG&E compliance filing in response to CPUC Decision 86-07-004 in OIR-2 (PG&E, 1986). Here PG&E developed a resource plan and attendant marginal costs that had the structure given in Table 2 below. This table shows the variations of marginal cost across months in column 1, normalized to the annual average. In columns 2 and 3, the variation within the month is shown, this time normalized to the monthly average. The peak period represents about the top quarter of the hours and the off-peak represents about the bottom quarter of the hours.

Table 2
Representative Marginal Cost Structure

Month	(Monthly Average)/ (Annual Average)	(On-Peak)/ (Month Average)	(Off-Peak)/ (Month Average)
January	1.00	1.33	0.53
February	0.90	1.37	0.44
March	0.80	1.40	0.35
April	0.80	1.34	0.27
May	0.70	1.52	0.17
June	0.70	1.74	0.16
July	1.00	1.52	0.30
August	1.10	1.36	0.54
September	1.30	1.28	0.68
October	1.30	1.22	0.65
November	1.20	1.18	0.67
December	1.20	1.19	0.64

Table 2 shows that for the representative conditions depicted, it would be desirable to average every two adjacent months starting at the beginning of the year. This would minimize the variation in monthly marginal cost. The data on the structure of marginal cost within the month confirm this. The ratio of on-peak to off-peak costs goes from about a factor of two in September to December all the way to a factor of about ten in May and June. Bi-monthly grouping would tend to minimize the variation in the on-peak to off-peak ratio. Given that we must choose a four season decomposition, Table 2 is a useful guide.

The last remaining factor to be considered is the role of retail rates. It is useful to estimate the changes in revenue collected from rates when demand-side programs are analyzed. Since

PG&E rates are seasonally differentiated, it is important to keep seasonal rate boundaries consistent with seasonal definitions motivated by resource or cost variation. The rate constraint makes the season definition problem more difficult. Table 2, for example, suggests that September might naturally be grouped with October-December. But from the rate perspective, it belongs with May-August. Although rates are subject to revision on time cycles that are short compared to planning periods, we nonetheless treat the rate constraint as binding.[3]

On the basis of these considerations, we can define a seasonal decomposition for PG&E. The rate constraint requires that the period from May through September be undivided across our definition of seasons. Marginal cost considerations suggest that May and June be treated together. We call these two months spring. July, August and September are aggregated into summer. These two seasons together have the summer rates. This leaves January through April as winter, and October through December as fall. These two seasons have the winter rates.

3. Daytype Definition

LMSTM represents each season by up to four typical daytypes. The user must specify how many days of each season belong to each of the four daytypes.[4] This choice is a critical step in the modeling process because the best possible specification of system loads should be sought and the definitions, once made, will most likely remain fixed for the duration of the study. Choosing the relative number of days to assign to each daytype also raises some weighty issues. While the primary goal of the study is to derive meaningful financial results for the period, and this favors even numbers of days in each daytype, a load control program exhibits more interesting results on some daytypes than others, and this favors focussing on these daytypes which will usually be the extreme ones. The LMSTM manual recommends that one daytype be assigned to weekends, and that the three remaining ones be defined to represent NORMAL, PEAK, and EXTREME weekday conditions (DFI, 1982). The 5% highest output days in the period are suggested for the EXTREME daytype. The next 10% are suggested for the PEAK daytype. The remaining 85% subdivide into WEEKEND and NORMAL.

Not satisfied with this formulaic approach, we attempted to derive a more natural division of each season into appropriate daytypes. LBL examined historic PG&E hourly loads from the years 1979-1984 to derive daily load shapes that would be consistent with our seasonal definitions. We believe that the days should be divisible by use of a suitable clustering technique. However, using the clustering feature of a standard statistical package yielded unsatisfactory results because the splits delivered differ wildly from one year to the next. We resorted to a graphical approach wherein we plotted the distribution of daily outputs for each season and then visually tried to subdivide the distributions into natural-looking groupings.

[3] As this report was going to press, we learned that that CPUC decisions sent down in December 1986 will change PG&E's rate seasons for all rates classes to two six-month periods, summer now ending on October 31. Thus, our seasonal definition is already obsolete which highlights the folly and frustration of making 15-year forecasts on assumptions such as fixed rating periods.

[4] This issue is taken up in considerable detail in Appendix A.

IV. RESOURCE PLAN CHARACTERIZATION

In this section, we describe in detail our characterization of the PG&E resource plan used in this study. This plan corresponds to the 1985 Long Term Plan (PG&E, 1985b). This is the latest publicly available long range plan for PG&E. It reflects expectations of mid-1985. Since that time there has been a major change in the world oil market which would cause some of the features of that plan to be revised. We comment briefly on these features when they arise in our discussion.

The principal burden of our discussion is on issues that affect the dispatch order of resources. We contrast our characterization with the data file created for DFI's 1984 calibration exercise. Our conclusion in general is that the 1984 calibration paid insufficient attention to non-economic constraints. The effect of such constraints in lowering marginal cost was only beginning to be recognized at that time. The 1984 calibration appears to have been based on the 1983 Long Term Plan in which these constraints were less important.

We review the resource plan by discussing the main features of each resource type. We take the representation of the 1985 Long Term Plan as given in the GRASS simulation as the detailed model to which we are calibrating LMSTM. The resource types can be divided into the following categories: (1) Hydro, (2) Northwest Imports, (3) Oil and Gas units, (4) Geothermal, (5) Nuclear, (6) Qualifying Facilities under PURPA, and (7) Generic Resources. The issues discussed include pricing and ownership, aggregation, and dispatch.

1. Hydro

There are two kinds of hydro resources: uncontrolled hydro that flows when available and is therefore dispatched in the base load, and controlled hydro that can be dispatched at times of high value. A convenient terminology for these categories adopted in ELFIN is ''run-of-river'' for the former, and ''pondage'' for the latter. This distinction in operational features cuts across ownership categories. Therefore, it is reasonable to divide all hydro resources into these two different types. We have built our hydro representation around this distinction. For pricing purposes, however, it is necessary to separate QF hydro and Northwest firm power, which is modeled by GRASS as pondage, into separate categories. We discuss how we adapt the GRASS representation into LMSTM format beginning with those hydro resources other than QF's and Northwest firm imports. For convenience we refer to this as the basic hydro system.

The GRASS simulation uses approximately fifteen categories to represent Northern California hydro resources owned by PG&E or public agencies. Each category is parameterized by a monthly maximum and minimum capacity and a monthly energy output. The monthly minimum capacity times the number of hours in the month gives the monthly run-of-river energy. All other energy in that month is pondage. We aggregate the three parameters for each category to produce monthly parameters for the Basic Hydro System. We then average the capacities over the months in each of our LMSTM seasons, and add the energies for each season. This characterizes the Basic Hydro System as we want LMSTM to represent it. The next problem is to put this parameterization into the data format of LMSTM, where the elementary unit is called a technology type.

The main difficulty in translating the three parameter representation of each LMSTM season is handling the varying capacities and energies. LMSTM does not easily accommodate resources which have characteristics that vary over the year. We could represent separately each season as a technology type. But for run-of-river and pondage, this would mean eight types for the Basic Hydro System, eight more for the QF's and four more for the Northwest firm imports (which are pondage only). Instead, we take advantage of the seasonal maintenance allocation factors to shape the resource definition to account for the annual variations. This parameter is enough to account for the run-of-river hydro. For pondage, we need to use the energy limit parameter as well. LMSTM allows energy limited resources to be characterized by a user designated MWh output per MW of capacity. To capture seasonal variations in pondage, we must also use the maintenance allocations. Details showing how the three parameter representation is translated into LMSTM format are given in Appendix B for the Basic Hydro System, QF Hydro and Northwest firm hydro.

LMSTM correctly handles the dispatch of run-of-river in the base load. Pondage is treated in the "Peak-shaving" mode. This means it is taken off the top load hours. The model first shaves the EXTREME daytypes, then the PEAK daytypes, then NORMAL weekdays and finally WEEKEND's. This sequencing can be verified by noting in the outputs that the least amount of pondage energy is dispatched on the weekends in any season, and increasing amounts for normal, peak and extreme weekdays. It should be noted that peak shaving pondage contributes only indirectly to marginal cost by reducing the load which remains to be served by thermal resources. This is roughly equivalent to other dispatch algorithms for storage resources.

The complex representation of hydro we have adopted better reflects the relatively large amount of off-peak hydro than the simpler model adopted in the 1984 calibration. In that version, hydro was separated into two types, PG&E and other. Each type was represented as an energy limited resource with a seasonal distribution of energy that corresponds to the aggregate hydro distribution. This makes all the hydro into pondage, and as a result over-optimizes the dispatch. The net effect of such a representation is to raise off-peak marginal costs relative to our representation.

2. Northwest Imports

There are two kinds of Northwest Imports, economy energy and firm energy. In some instances, this may be a contractual distinction; in the 1985 Long Term Plan it represents a qualitative distinction. By this we mean that while there is no contract which specifies a level of firm imports, the supply/demand balance in the Northwest indicates to PG&E planners that a certain amount of energy can be expected with a high degree of confidence to be available for on-peak dispatch. These firm quantities should be modeled differently than economy energy. We begin our discussion with economy energy since it is simpler.

Economy energy is treated in the 1985 Long Term Plan as if it were any other kind of thermal generating unit. It has a cost and a rated capacity, and is dispatched without constraint in its economic order. The thermal capacity rating changes from month to month reflecting variations in the availability of the resource. These availability variations are essentially driven by hydrologic fluctuations. Therefore, there is greater capacity in the first part of the year than the last part of the year.

The 1985 Long Term Plan also distinguishes the price of economy energy in the first six months of the year from the last six months. This seasonal price variation is incorporated into LMSTM. The latter price is about 20% higher than the former. The higher priced economy energy is available in substantially smaller quantities than the lower priced energy. Economy energy in the latter half of the year is less likely to be marginal since the supply/demand balance is tighter in the second half of the year.

Firm imports present more choices for representation. In the 1985 GRASS simulation, these resources are modeled as pondage. Because they are firm, they are not curtailed and only contribute indirectly to marginal cost. The LMSTM data file used by Economics and Forecasting includes Northwest Pumped Storage units. This representation provides a way to get Northwest power serving peak loads. Incorporating the pumping representation is consistent with contractual relations that allow for the return of energy to the Northwest as payment for energy received. In practice, California utilities do not generate energy to return to the Northwest, but simply pay for what is received. For this reason, we prefer not to use a pumping representation. Another possibility for firm purchases used on peak would be a load modification representation.

We opt for a pondage representation of Northwest firm purchases. This is most consistent with the GRASS representation. There is some question about how to price this resource due to more complex ownership issues. For economy energy we use the direct and indirect cost approach. This means we have one set of technology types for PG&E economy energy that have direct costs that vary with the seasonal price. Municipal Northwest economy energy is modeled in the same way except that its costs are indirect. This puts it in the correct spot in the dispatch order, but not in PG&E ratepayer costs. To price firm energy we approximate this outcome by pricing it all at the weighted average of PG&E and municipal quantities of firm energy. This gives the correct total cost to PG&E ratepayers. The price is not directly relevant to dispatch or marginal cost due to the pondage representation. Appendix B describes, in more detail, the characterization of Northwest economy and firm energy.

3. Oil and Gas Units

The main issue with oil and gas resources is modeling the unit commitment feature appropriately. In the 1984 calibration the minimum blocks of these units were correctly specified, based on data available at that time. As the problem of excessive must-run capacity became increasingly apparent, however, reductions were made in the minimum blocks of Moss Landing 6 and 7 from 200 MW each down to 50 MW. LBL made these changes in the LMSTM file. In addition, there was a typographical error in the specification of the minimum block size for the 122 MW class of oil units. Instead of 80% of capacity, this should be 8% of capacity. This error was corrected.

The more serious concern involves the question of minimum downtimes for oil and gas units. Version 1.0 of LMSTM did not include the must-run feature. Therefore, units which were committed to meet peak loads of any daytype could be shutdown during the off-peak hours. GRASS uses a commitment logic which essentially keeps units running for the entire month if they are necessary to meet the monthly peak load. This mismatch in features may well have contributed to problems in the 1984 calibration. Version 2.0 of LMSTM, released in 1985, does incorporate the must-run feature. This keeps the minimum block of units operating over all hours of a daytype if the unit is committed to meet the peak load of that daytype. We use the must-run specification, called either Technology Group Class 7 or 8, for oil and gas units. This specification is also be used for other resource types. Its unavailability for the 1984 calibration, and the implications of that, are discussed in section V.

4. Geothermal

There are two problems involved in the representation of geothermal resources, getting the correct price and dispatching them in proper order relative to nuclear units. The dispatch problem is essentially a non-economic constraint due to the contract between PG&E and the Geysers steam suppliers. This contract says that geothermal units cannot be curtailed until nuclear units have been curtailed first (Kahn, 1986). PG&E models this by using dispatch penalty factors applied to the geothermal price. Dispatch penalty factors are like shadow prices. They characterize the actual relations among resources as if they were economic instead of being constrained by exogenous factors. In particular, although the variable cost of geothermal is greater than the variable cost of nuclear, we want the model to act as if it were less. LMSTM has this feature, which is called a "dispatch weight factor." We use this feature to achieve the appropriate dispatch. By assigning a dispatch weight factor of 0.1 to geothermal (where the default value for all other resources is 1.0), the dispatch algorithm treats the geothermal variable cost as if it were 10% of whatever value has been assigned to it. This is enough to have geothermal dispatched ahead of the nuclear (except for the minimum nuclear block; see section IV.5. below). In addition, we designate the geothermal resources as must-run units to assure that they are curtailed only after all other dispatchable resources are curtailed. The 1984 calibration does not appear to have captured the appropriate dispatch order for geothermal.

The appropriate price for geothermal resources is determined endogenously by a contractual formula. It is not be necessary to perform a preliminary model run to determine that price for our base case calibration assuming that we end up with a dispatch close to the GRASS simulation. In general, however, such a preliminary analysis is necessary, unless the production simulation has the pricing formula endogenously built in. Currently only ELFIN has that feature.

In analyses where the base case simulation is perturbed, it is necessary to recalculate the geothermal price. This is a simple computation that is easily handled on a spreadsheet. An illustrative example of the magnitude of the price change effect shows that 15-25% changes in the geothermal price can result from 15-20% changes in the availability of base load energy (Kaiser Engineers, 1986). This shows a relatively high price elasticity with respect to supply changes.

5. Nuclear

The principal issue associated with the representation of nuclear units is the dispatch order relative to geothermal. We have discussed above the requirement that nuclear be curtailed before geothermal. PG&E has interpreted this requirement to mean that Rancho Seco (owned by the Sacramento Municipal Utility District, but dispatched by PG&E) is included in this constraint. Curtailment, however, is not interpreted to mean complete shutdown. The nuclear units are allowed to operate at a minimum level, equal to 25% of capacity, at all times. By designating the nuclear units as must-run, this minimum block is dispatched in the base load. This representation together with the dispatch weight factor of 0.1 for geothermal gives the appropriate dispatch order.

It is also important to be careful about the scheduling of maintenance for nuclear units. Examination of the 1985 GRASS simulation shows that 85% of all nuclear maintenance is scheduled during our winter (January-April) and fall (October-December) seasons. Due to the importance of nuclear units in the overall supply/demand balance, it is important represent their maintenance. We specify for each unit the long-term average maintenance as indicated in the

GRASS simulation for the class of nuclear units. This results in 44.6% allocated to WINTER, 6.6% allocated to SPRING, 8.4% allocated to SUMMER, and 40.6% allocated to FALL.

6. Qualifying Facilities

There are three kinds of Qualifying Facilities (QF's) from the dispatch point of view. The largest category is base load must-run units such as cogeneration or biomass, which are essentially just thermal resources. Wind and solar units have a time-dependent pattern of output that is ideally represented as a load modification. This is the procedure adopted in the 1985 GRASS simulation. Load modification is the best procedure to capture the time varying availability of energy from these resources. We do not adopt this procedure in LMSTM for lack of a convenient data representation of time dependent units. Such a representation would have to be compatible with our daytype and season definitions. Instead, we use the simple approximation of treating these resources as base-loaded in each season, but with seasonal variation in energy. We take the seasonal variation data from CPUC ELFIN files (House, 1986). This is achieved by finding that amount of capacity and those maintenance allocators which would produce the expected annual energy. Finally, the 1985 Long Term Plan includes dispatchable cogeneration QF's. GRASS models these as thermal resources with characteristics resembling oil and gas units (except there is no minimum block). We use the same approach in LMSTM.

It is also important to differentiate QF's by pricing terms. It is common to distinguish those QF's which are paid under long-term price offers from those which are paid under short-term price arrangements (Weisenmiller and Yardas, 1986; House, 1986). Among the long-term contracts, there are fixed price contracts, usually for hydro and wind projects, and contracts indexed to oil and gas prices, usually for cogenerators. The long-term contracts are due to expire between 1995 and 2000 and the producers on them switched to short-term pricing.

It is difficult to represent the short-term pricing of QF's in LMSTM. The short-term price for QF's is an endogenous price (see section II.1.*e)* above). LMSTM is not capable of calculating endogenous prices on the basis of a given simulation. QF short-term pricing based on QF In/Out methods (see section II.3.*b)*) requires two simulations. Where necessary, we perform these calculations.

7. Generic Resources

The LMSTM data file used by Economics and Forecasting has two generic resources represented. One is a generic base load resource that is essentially an unidentified coal plant. The other is a peaking turbine using distillate oil. The 1985 Long Term Plan shows a substantial amount of such resources starting in the late 1990s and playing a substantial role in the period 2001-2005. In this latter period, 2500 MW of generic base load is added and 1000 MW of generic peaking. We replicate this supply scenario in our file. This involves slight changes from the scenario specified in the file currently used by PG&E.

V. RESULTS

In this section, we describe the results of our calibration exercise. The discussion is divided into two parts. First, we address in section V.1. qualitative issues in the context of a simplified LMSTM database. The purpose of this discussion is to demonstrate the effect of various features described in section IV upon the marginal cost results and the total production by fuel type. These results do not represent the definitive calibration, but only a representative illustration. Details of the complete data set are given in section V.2. for three representative years in the simulation period 1989, 1996 and 2003.

1. Qualitative Features

The data set with which we began used only LMSTM Version 1.0 features (i.e. no must-run), had a different season and daytype definition than that suggested by LBL, and was based on PG&E loads and resources (i.e. excluding municipalities loads and resources). To test out the resource type descriptions given in section IV, we performed a series of test runs implementing these features step by step. For simplicity, we kept the initial season and daytype definitions and examined only 1989 conditions. To facilitate marginal cost comparisons, we wrote software which aggregates the hourly marginal costs reported by LMSTM, averages them over the daytypes, and weights them over the seasons. This makes it possible to compare results with GRASS.

For convenience, we express marginal energy cost as an Incremental Energy Rate (IER). The IER is used for pricing QF's; it is just the marginal energy cost divided by the price of oil or gas. IER's have dimensions of Btu/kWh and express the mix of contributions to marginal energy cost from oil and gas resources and other resources. Because marginal cost is a weighted average of contributions from various resources, IER will be lower to the degree that non-oil and gas resources contribute to the total. This follows from the relative cost of resources and the definition of IER, because we divide by the oil/gas price. If marginal cost consists only of contributions from oil and gas resources, the IER is just a standard engineering heat rate. Since the most efficient oil and gas units have heat rates of around 8500 Btu/kWh, if IER has a lower value, we know that non-oil and gas resources are contributing to marginal cost. It is convenient to use IER because it abstracts from the particular level of oil and gas prices and focuses attention on the supply/demand balance.

To structure the discussion, we focus attention on one particularly unsatisfying result from the 1984 calibration that carries over to the data set as used by Economics and Forecasting. This is the relatively limited operation of the Helms Pumped Storage plant. In the 1984 calibration, the LMSTM results for 1991 showed only 14 GWh of pumping output compared to 1159 GWh from GRASS. We claim that discrepancies of this magnitude are due to the lack of the must-run feature in Version 1.0 for representing non-economic constraints. This results in marginal cost being higher than it would be with such constraints included, particularly in the low load off-peak hours.

Table 3 summarizes selected results from some of these test runs. We indicate the significant changes associated with each run and a few statistics relevant to marginal cost and pumping loads. To provide some degree of comparability with Table 2, we report the ratio of maximum and minimum marginal cost for the WINTER season NORMAL weekday period from these runs. In this data set, the WINTER season is defined as the seven month period spanning

January through April and October through December. The NORMAL weekdays total 118 days or 32% of the year.

Table 3
Evaluating the Effect of Constraints
Selected LMSTM Test Runs and Corresponding GRASS Run (1989)

Run Number	Feature	Pumping Output (GWh)	IER (BTU/kWh)	Max MC/ Min MC
1.	Area Basis	289	9417	1.04
13.	Run of River Hydro QF Must Run	506	8535	2.17
19.	Oil/Gas Must Run Geothermal Dispatch	1464	6947	3.61
26.	Seasonal Pondage	1139	7236	5.76
35.	Increase QF's Reduce Oil/Gas Must Run	1188	6160	5.87
36.	Nuclear Maintenance	1187	6445	8.16
	GRASS '85 LTP	1664	6713	5.2

Run #1 represents the data set received by LBL augmented only by the additional loads and resources associated with the municipalities of Northern California not served by PG&E. On balance, these changes lower marginal costs because we are adding all low cost resources such as hydro, nuclear and geothermal. As we can see, however, the IER shows that oil and gas is marginal all of the time. The ratio of maximum to minimum marginal cost is low, and so is pumping output. Indeed, 90% of the pumping in this run is by Northwest pumped storage units, which were subsequently removed (see section IV.2).

Run #13 represents the first significant change in the dispatch logic. QF resources in the supply plan are designated must-run, and the hydro resources which have run-of-river characteristics are dispatched in the base load according to their seasonal availability. These changes lower annual average marginal cost; the IER drops by 900 Btu/kWh. The ratio of maximum to minimum marginal cost doubles, and pumping output goes up to 506 GWh. All of the pumping comes from Helms, since the Northwest pumping units have been removed.

Run #19 includes both the oil and gas must-run constraint and the geothermal dispatch requirements discussed above. Geothermal resources are designated must-run, and dispatch weight factors are used to have them dispatched after the minimum block of nuclear units, but before the remaining nuclear blocks. This additional set of constraints has a major effect on the cost structure. The annual IER drops about 1600 Btu/kWh compared to the previous case. Pumping output nearly triples and the ratio of maximum to minimum marginal cost goes from 2.17 to 3.61. This last result shows that the economic advantages of pumping depend on a substantial difference between the input price and the output value. Only when non-economic constraints lower off-peak costs does it become worthwhile pumping.

At this point, we have gotten more or less within the range of outcomes of the GRASS simulation for 1989, at least as far as these highly aggregated and selective parameters are concerned. Yet many details remain inconsistent and need to be adjusted systematically. For comparison purposes, it is useful to note that the 1985 Long Term Plan GRASS simulation of 1989 shows annual pumping output of 1664 GWh, an annual average IER of 6713, and a ratio of maximum to minimum marginal cost over the same seven month period of 5.20. The LMSTM results are still not well matched to the GRASS simulation with regard to production by fuel type. We discuss the detailed results of calibration with regard to these variables in the next section. To develop some intuition for the magnitude and direction of changes, we discuss Runs #26, 35 and 36, which show "fine-tuning" cases. All these cases involve details of resource representation that are less dramatic in their impact on marginal cost than the dispatch features discussed above.

Run #26 involves the representation of pondage hydro on a seasonal basis. Previous representation just allocated hydro over the year on total energy basis. But the pattern of run-of-river and pondage varies substantially. pondage is "saved" for the SUMMER PEAK period to the degree possible. This means there is less of it in WINTER than run of river and than the average. By making this adjustment, we are optimizing the system, which translates into an increase in marginal cost. Notice that annual IER increases in this case. The increase in the ratio of maximum to minimum marginal cost in this case is due to an increase in the maximum. We have removed pondage energy from the NORMAL weekday period, thereby raising marginal cost in the high load period. The decrease in pumping in this run is another indication of the move toward optimality from the seasonal perspective.

Run #35 includes two effects which work in opposite directions. First, we correct some input errors with respect to the minimum block size of oil and gas units. The minimum block of the Moss Landing 6 and 7 units is reduced by 150 MW each. The minimum block on the 122MW class of units is reduced to 8% of capacity from the incorrectly designated 80%. These changes tend to increase marginal cost because they reduce the magnitude of the must-run constraint. Working in the opposite direction are the increases in QF production. We have added 228 MW of thermal QF production, 250 MW of QF hydro that is run-of-river, and 150 MW of QF pondage Hydro. All of this tends to lower marginal cost.

Table 3 shows that the net of these two effects is to reduce annual average marginal cost. The IER drops almost 1100 Btu/kWh. Some of this change occurs within the WINTER NORMAL period. Although there is a small change in the ratio of maximum to minimum marginal cost in the NORMAL daytype period, this is due to relatively equal reductions in both terms. There is also a general reduction of marginal cost in the SPRING and WINTER WEEKEND periods.

Run #36 incorporates the qualitative features of the maintenance schedule for nuclear units. This involves allocating most of the maintenance to the WINTER season. Since it is optimal to have the nuclear units available for the SUMMER season, this change raises marginal cost. The IER increases by about 300 Btu/kWh. Since the on-peak costs increase in the WINTER period due to reduced nuclear output, the ratio of maximum to minimum marginal cost increases in this run to 8.16. This is greater than the corresponding ratio in GRASS. The annual IER remains about 250 Btu/kWh lower than GRASS and the pumped storage output is only about 70% of the GRASS level. These results should not be taken literally since resource representations in these tests are not intended to match closely in detail.

The purpose of these runs was to test the capability of LMSTM to capture broad general features of the GRASS simulations and to indicate the sensitivity of results to various input changes. The detailed calibration is reported below.

2. Results of Detailed Calibration

The following describes the work performed in the final calibration and how the results of the calibrated LMSTM model compare to GRASS.

a) Calibration

The complete calibration effort included a re-specification of the 16 daily load shapes and improvements in the specification of the supply side. The daily load shapes are based on historical PG&E hourly area loads. Several methods were tried in the effort to translate these historical data into LMSTM typical daytype load shapes. The methods used are described in detail in Appendix A. It was found that using hourly area loads from 1983 gave us the most sensible load shapes. For each season in 1983, the loads from all WEEKENDS were averaged to make the first daytype load shapes. The remaining three daytypes for each season were made by finding reasonable groupings of the days according to daily sales (GWh). Hourly data were averaged across the days in each of the resulting 16 groups (4 daytypes for 4 seasons) to produce the load shapes required by LMSTM. The result of this grouping process is summarized in Figure 1 of Appendix A. Figure 1 shows, for example, that the WINTER season (Jan-Apr) was grouped such that 37 days were WEEKENDS, 39 days were NORMAL days, 38 days were PEAK days, and 6 days were EXTREME. Note that the terms *NORMAL, PEAK,* and *EXTREME* are used primarily due to convention and that the magnitudes of the loads for particular daytypes (e.g., PEAK) are considerably different across seasons.

On the supply side, a comprehensive effort was made to get an accurate input and feature correspondence between the GRASS 1985 LTP and LMSTM. The important characteristics of the GRASS supply inputs are described in section IV. The major effort required in calibrating LMSTM's SUPPLY.IN file was in the hydro specification. This calibration effort is described in detail in Appendix B.

An important issue brought up during the calibration was the persistent failure of LMSTM's dispatch algorithm to meet load in later years (approximately 1997 on). During many hours of PEAK and EXTREME SUMMER days, LMSTM would leave 3 to 5 MW of unserved load even though sufficient capacity existed. While this dispatch result does not significantly impact generation results, it greatly changes marginal costs because the cost of unserved energy is very high. It is not clear why LMSTM misses peak loads by such a small amount, but a practical way around this problem was found by adding to the supply file a large amount of extra peaking capacity in years where unserved energy was a problem. This "phantom" capacity has characteristics similar to other distillate peaking capacity in the model except that its heat rate is marginally higher (so that it is dispatched last) and its fixed costs are zero. The addition of phantom capacity is an acceptable solution to this dispatch problem so long as this capacity is used only to get the model to serve this small amount of unserved energy and *not* to take the place of planned (and expensive) resources in the supply file.

b) Results

The results of the detailed calibration effort are shown in Tables 4-6. For each of the years 1989, 1996, and 2003, the tables show important generation and marginal cost outputs from

GRASS and from two LMSTM runs, #53 and #54. In the following discussion, we first compare the two LMSTM runs with each other. Second, we compare the LMSTM runs with GRASS.

Table 4
Comparison of Supply Model Outputs
(YEAR = 1989)

	GRASS	LMSTM #53	LMSTM #54	#54 REL. TO GRASS
	YEARLY PRODUCTION (GWh)			
RESOURCE				
NUCLEAR	17,555	17,415	17,504	-0.3%
OIL/GAS	8,466	10,532	9,612	13.5%
GEYSERS	9,634	9,543	9,543	-0.9%
OTHER GEO	2,581	2,374	2,374	-8.0%
QF-EOR	1,669	2,374	2,604	56.0%
GEN BASE	270	320	332	23.0%
QF-OTHER	15,389	15,576	15,576	1.2%
NW THERM	7,433	5,628	6,176	-16.9%
PUMP STG	1,664	853	696	-58.2%
HYDRO	27,794	26,996	26,996	-2.9%
	MARGINAL COST (mills/kWh)			
PERIOD				
WINTER	40.2	32.1	38.1	-5.2%
SPRING	40.6	23.3	25.2	-37.9%
SUMMER	50.7	53.2	55.2	8.9%
FALL	54.1	58.3	59.9	10.7%
ANNUAL	46.4	42.5	45.7	-1.5%
	INCREMENTAL ENERGY RATE (BTU/kWh)			
ANNUAL	6,720	6,155	6,618	-1.5%

Runs #53 and #54 are identical except that in run #54 the oil and gas units are allowed to shutdown their minimum blocks on WEEKENDS. This change in the WEEKEND must-run status is made by setting the CLASS number of the oil and gas units to "7" in run #53 and "8" in run #54. As this is the only difference between the two runs, the generation of base load resources (nuclear, geysers, other geothermal, QF-other, and hydro) is identical between runs #53 and #54. Because run #54 is allowed to shutdown minimum blocks on WEEKENDS, it is no surprise that this run shows less oil and gas generation and more generation of marginal resources (QF-EOR, generic base load, and Northwest thermal). Also, pumped storage in run #53 is utilized approximately 18% more than in run #54. This difference in pumping occurs because run #53 requires the minimum blocks on the oil and gas units to run on WEEKENDS which increases the variation of marginal costs on those days. Increased daily variation in marginal costs increases the opportunities for pumping.

Table 5
Comparison of Supply Model Outputs
(YEAR = 1996)

	GRASS	LMSTM #53	LMSTM #54	#54 REL. TO GRASS
	YEARLY PRODUCTION (GWh)			
RESOURCE				
NUCLEAR	16,542	17,537	17,579	6.3%
OIL/GAS	10,298	11,482	10,553	2.5%
GEYSERS	12,095	12,659	12,659	4.7%
OTHER GEO	2,588	2,374	2,374	-8.3%
QF-EOR	7,770	8,256	8,799	13.2%
GEN BASE	640	685	702	9.7%
QF-OTHER	20,557	20,546	20,546	-0.1%
NW THERM	6,771	5,307	5,606	-17.2%
PUMP STG	1,153	694	594	-48.5%
HYDRO	31,074	29,648	29,648	-4.6%
	MARGINAL COST (mills/kWh)			
PERIOD				
WINTER	88.4	77.6	80.4	-9.0%
SPRING	71.2	42.5	45.1	-36.7%
SUMMER	92.9	103.0	106.8	15.0%
FALL	107.4	107.9	108.8	1.3%
ANNUAL	91.4	85.8	88.3	-3.4%
	INCREMENTAL ENERGY RATE (BTU/kWh)			
ANNUAL	7,444	6,988	7,191	-3.4%

In terms of features, LMSTM run #53's inputs most closely match GRASS's inputs. That is, GRASS effectively runs the minimum blocks of its oil and gas units on the WEEKENDS like run #53. However, LMSTM run #54 generally shows generation results closer to GRASS's outputs. The exception to this general result is pumped storage generation which is closer to GRASS in run #53. Also, Run #54 generally does a better job than run #53 of matching GRASS's marginal costs.

From a theoretical view point, the most desirable run for future work would be run #53 as it has the best feature correspondence with GRASS. From a pragmatic viewpoint, however, it is wise to use run #54 for future work. This pragmatic viewpoint is based on the fact that Run #54 is the best match to GRASS with respect to resource generation and marginal costs and that a credible analysis of demand-side programs depends upon the use of credible marginal costs.

An evaluation of LMSTM run #54's performance relative to GRASS should be made by comparing both generation and marginal cost outputs of the two models. Using marginal resource generation as a measure of the goodness of the LMSTM calibration effort, LMSTM's performance appears somewhat disappointing. Averaged over the three test years, LMSTM

Table 6
Comparison of Supply Model Outputs
(YEAR = 2003)

	GRASS	LMSTM #53	LMSTM #54	#54 REL. TO GRASS
	YEARLY PRODUCTION (GWh)			
RESOURCE				
NUCLEAR	18,159	17,623	17,623	-3.0%
OIL/GAS	17,020	19,316	18,705	9.9%
GEYSERS	13,101	12,659	12,659	-3.4%
OTHER GEO	2,591	2,374	2,374	-8.4%
QF-EOR	9,847	10,523	10,855	10.2%
GEN BASE	8,932	9565	9617	7.7%
QF-OTHER	24,477	24,174	24,184	-1.2%
NW THERM	6,442	5,951	6,125	-4.9%
PUMP STG	1,289	731	648	-49.7%
HYDRO	31,193	29,711	29,711	-4.8%
	MARGINAL COST (mills/kWh)			
PERIOD				
WINTER	170.0	179.3	183.7	8.1%
SPRING	151.8	141.9	144.9	-4.5%
SUMMER	199.8	234.2	237.2	18.7%
FALL	200.7	215.2	216.9	8.1%
ANNUAL	182.1	195.9	199.1	9.3%
	INCREMENTAL ENERGY RATE (BTU/kWh)			
ANNUAL	7,966	8,569	8,709	9.3%

utilizes approximately 9% more oil and gas and 13% less Northwest thermal compared to GRASS. LMSTM better matches GRASS in its average annual marginal cost. Compared to GRASS, the annual marginal costs of LMSTM are 2% too low in 1989, 3% too low in 1996, and 9% too high in 2003. On a seasonal basis, however, LMSTM's marginal costs deviate significantly from GRASS. In general, LMSTM shows lower marginal costs in WINTER and SPRING and higher marginal costs in SUMMER and FALL. In the 12 seasonal marginal costs presented in Tables 4, 5, and 6, run #54 deviates from GRASS by more than 10% 5 times.

Some of the resource generation discrepancies between run #54 and GRASS can be explained in more detail. The match between run #54 and GRASS for nuclear generation varies. In 1989 LMSTM is quite close; in 1996 it is 6% too high; in 2003 it is 3% too low. This discrepancy is caused by the nuclear units having maintenance schedules in GRASS that cycle in periods greater than a year. The maintenance schedules for the nuclear units in LMSTM, however, are based on long-run averages exhibited in GRASS. Thus, a good year-to-year match should not be expected.

The hydro generation shown by LMSTM is approximately 4% too low over our three test years. This is due to the as-yet unexplained behavior of LMSTM to persistently under-utilize the Northwest pondage hydro.

Some discrepancy between GRASS and LMSTM is likely caused by differences in load shape input data. As noted in section V.2.*a)* and Appendix A, the load shape data in LMSTM is based on 1983 PG&E area hourly loads. Presumably GRASS's load shape data was based on a forecast made for the 1985 LTP. This forecast of load shapes may have been significantly different from the load shapes created by historical data. The load shape data used in GRASS was not available for comparison with LMSTM. Thus, it is not possible to quantify the effect of this possible input discrepancy.

Of most concern, however, is the discrepancy between the utilization of pumped storage in GRASS and LMSTM. LMSTM uses pumped storage less by approximately a factor of 2. An examination of hourly dispatch outputs of LMSTM shows that the model is not utilizing pumped storage even when adequate daily variations in marginal cost exist. The cause of this relative under-utilization of the pumped storage resource is unclear. We conjecture that something is wrong in LMSTM's "look-ahead" capability. A *possible* cause of the lack of pumping may be a result of LMSTM being fooled into thinking that PG&E's large amount of pondage hydro takes care of all peaking requirements when, in fact, the pondage hydro's energy limits cause oil and gas to be used for peaking. As a result, LMSTM does not pump at night because it thinks storage capacity is unnecessary.

VI. CONCLUSIONS

We have achieved a plausible simulation of the PG&E supply system with LMSTM using the GRASS 1985 LTP as a basis of comparison. Annual average marginal costs produced by LMSTM were within 10% of GRASS for all three test years. However, LMSTM exhibited greater discrepancies with GRASS in its seasonal average marginal costs and in the generation of the two major marginal resources (oil and gas and Northwest thermal).

While we would like to do better, the results are in an acceptable range for using our calibration of LMSTM for strategic analysis. The model as it now stands is a substantial improvement over the version used previously by PG&E.

These results were achieved by an accurate feature-to-feature correspondence between LMSTM and GRASS. While our final choice of assumptions deviates from GRASS with respect to must-run constraints on the minimum blocks of the oil and gas units and possibly with respect to load shapes, overall there is a good fit among the inputs and features of the two models. This calibration puts us in a good position to update some of the assumptions in the model which have changed since the 1985 LTP and to perform our demand-side planning case study.

These results only apply to version 2.0 of LMSTM, the model used exclusively for this project. The vendor has recently announced the release of version 3.0 which is significantly different to current versions. Notably, the new version will be based on 12 typical weeks per year, each corresponding to one month. This change will eliminate some of the aggregation issues discussed here and substantially increase computer run time. With the weekly representation, pumped storage operation should increase. The supply simulation is not expected to change otherwise. The extent to which results obtained here would change in the new version is open to speculation.

VII. BIBLIOGRAPHY

1. Baleriaux, H., E.Jamoulle, Fr. Lenard de Guertechin. "Simulation de l'exploitation de'un parc de machines thermiques de production d'electricite couple a des stations de pompage." *Review E (edition S.R.B.E.)*, pp.3-24, Vol. V, No.7, 1967.
2. Bloom, J. "Generation Cost Curves Including Energy Storage." *IEEE Transactions on Power Apparatus and Systems*, V.PAS-103, No.7, pp. 1725-1731, 1984a.
3. Bloom, J. "Modeling Economic Operations of Power Pools." Paper presented at TIMS/ORSA Meeting, 1984b.
4. Booth, R. "Power System Simulation Model Based on Probability Analysis." *IEEE Transactions on Power Apparatus and Systems*, V. PAS-91, No.1, pp.62-69, 1972.
5. Caramanis, M., F.Schweppe and R.Tabors. *Electric Generation Expansion Analysis System*, Vol. 1, EPRI EL-2561, 1982.
6. Caramanis, M., J.Stremel, W.Fleck and S.Daniel. "Probabilistic Production Costing: An Investigation of Alternative Algorithms," *International Journal of Electrical Power and Energy Systems*, V.5, No. 2, pp.75-86, 1983.
7. Decision Focus Inc. *Load Management Strategy Testing Model*, EPRI EA-2396, 1982.
8. Decision Focus Inc. *User's Guide to LMSTM Version 1.0*, EPRI EA-3653-CCM, 1984a.
9. Decision Focus Inc. *LMSTM: PGandE Base Case Supply Calibration with PGandE Models*, 1984b.
10. Environmental Defense Fund Inc. *ELFIN Overview and Summary of Simulation Methods*, 1986.
11. Gulbrand, K. and P.Leung. "Power System Economics: A Sensitivity Analysis of Annual Fixed Charges," *Transactions of ASME Journal of Engineering for Power*, V. 97, pp. 465-472, 1975.
12. Gerber, M. Testimony before the California Public Utilities Commission in Appl. No. 85-04-019, 1985.
13. House, L. ELFIN Files in CPUC Appl. No. 85-12-50, 1986.
14. Jabbour, S. "The Short-Run Value of Nonutility-Generated Electricity," *IEEE Transactions on Energy Conversion*, V.EC-1, No.1, pp.11-16, 1986.
15. Kahn, E. *Production Cost Modelling of the Pacific Gas and Electric System.* California Energy Commission P500-85-010, 1985a.
16. Kahn, E., D.Kirshner and D.Yardas. *Production Cost Modelling of the Southern California Edison System*, California Energy Commission P500-85-011, 1985b.
17. Kahn. E. *Comparison of ELFIN, GRASS and SAM*, Pacific Gas and Electric Company, 1986.
18. Kaiser Engineers Inc. *Oversupply and the Valuation Method for QFs*, 1986.
19. Leung, P. and R.Durning. "Power System Economics: On Selection of Engineering Alternatives," *Transactions of ASME Journal of Engineering for Power*, V.100, pp.333-346, 1978.

20. Levy, D. and E.Kahn. "Accuracy of the Edgeworth Approximation for LOLP Calculations in Small Power Systems," *IEEE Transactions on Power Apparatus and Systems*, V.PAS_101, No.4, pp.986-996, 1982.
21. Lovins, A., L.Lovins, F.Krause, and W.Bach. *Least-Cost Energy: Solving the CO2 Problem*, Brick House, 1981.
22. Mazumdar, M. and Y.Wang. "On the Application of Esscher's Approximation to Computation of Generating System Reliability and Production Costing Indexes," *IEEE Transactions on Power Apparatus and Systems*, V.PAS-104, No.11, pp.3029-3036, 1985.
23. Nordell, D. "An Approach to Conservation and Load Management Economic Analysis," *IEEE Winter Power Meeting*, 1986.
24. Pacific Gas and Electric. *The Mixture of Normals Approximation of Equivalent Load Duration Curves*, EPRI EA/EL-4266, 1985a.
25. Pacific Gas and Electric. Exhibit 3 - Long Term Planning, CPUC Appl. No. 85-12-50, 1985b.
26. Pacific Gas and Electric. Compliance Filing in OIR-2, CPUC Appl. No. 82-04-044, 1986.
27. Sant, R. *The Least Cost Energy Strategy: Minimizing Consumer Costs Through Competition*. Carnegie-Mellon University Press, 1979.
28. Stremel, J. *LMSTM Version 2.0 User's Guide Supplement*, 1985.
29. Stremel, J. and S.Bubb. *Comparing Direct Simulation and Probabilistic Production Costing Models*, 1985.
30. Stremel, J. and T. Jenkins. *Maintenance Scheduling Methods for Planning Models*, 1981.
31. Weisenmiller, R. and D.Yardas. Testimony before the California Public Utilities Commission in Appl. No. 86-04-012, 1986.
32. Wu, F. and G.Gross. "Probabilistic Simulation of Power System Operation for Production Cost and Reliability Evaluation," *IEEE International Symposium on Circuits and Systems*, 1977.

Appendix A

The Specification of Daytypes for LMSTM and the

Derivation of System Load Shapes

1. Introduction

As discussed above, LMSTM carries out its dispatch for up to 16 typical days of every year in the forecast period.[1] The year is split into up to four seasons and then 4 representative daytypes are defined for each season. The division of the year into seasons has been described above, and these seasons have to be further divided into 16 daytypes. The final division made is portrayed in Figure 1.[2] This section addresses the issues that arise in the definition of the daytypes for each season and the derivation of the critical LMSTM inputs for each daytpe, relative hourly load shapes and energy use weights that relate energy use each day to the other daytypes.

The choice of seasonal split and daytype definition is important for two fundamental reasons:

1. In the interests of accuracy, the definitions of seasons and daytypes must reflect as precisely as possible the true annual variations in operating conditions, and, therefore, costs.
2. The choice of seasons and daytypes is essentially fixed for the duration of the study at hand.[3]

In other words, the specification of seasons and daytypes is a critical step in the analysis and the researcher gets only one shot. His or her choice should be the best possible, not only for historic years, but also forecast years.

2. LMSTM jargon

It is necessary to digress a little here to explain some confusing LMSTM terminology. Most of the confusion arises from a significant divergence between the designers' intended use of LMSTM and its current use in practice. LMSTM is written as a *bottom up* model, on the demand-side, that *sums* up load curves across the various enduses to arrive at the system load shape, against which the supply-side dispatches its resources.[4] Changes in load patterns originate at the enduse level through changing technologies, numbers of customers, and policy initiatives and only appear at the system level as results from the summing up process. In this way, the forecast of hourly demands for each daytype are *endogenously* determined. In reality, however, adequate data on the sub-demands of several enduses and is rarely available and the common practice is for the researcher to define one, or a very few, and enduses. Falls in outputs and

---

[1] The failure of the model to provide the user with an easy way to restrict the computation and output of seasons and/or years of little interest is one of its most annoying weaknesses.

[2] Note that the four groupings of months in Figure 1 correspond to the definitions of the four seasons, WINTER, SPRING, SUMMER, and FALL, that appear in Table 1.

[3] This is not strictly true because LMSTM has a start-stop feature that permits the editing of input files midstream and, as always, a subsequent sensitivity analysis on the load inputs is possible. Both of these alternatives are, however, very clumsy in practice.

[4] We are using ''enduse'' to mean either the customer class level or the enduse level of LMSTM. In this case we are defining only one system shape so both are identical.

loads then have to be estimated *exogenously* and *subtracted* from the system shape through definition of a negative enduse.[5] This procedure effectively turns LMSTM on its head and makes it a *top-down* model.

Let us state explicitly that we take the latter tack. Our intention is to derive only 16 shapes that portray, as accurately as possible, the variations in PG&E's system loads. In the next phase of this study, a policy initiative will be represented as a decrement to these loads. Further, we have great difficulty considering the question of how loads may change over time without the policy change. We have concentrated on representing the loads reported to us by PG&E for the historic period 1979-84. We have so far made no efforts to assess likely changes in load shapes over the study period. Certainly, we have made no efforts in the direction of respecifying the daytype division in the midst of a forecast run. Our choice of seasons does take qualitatively into account likely changes in hydro availability in future years, although it remains fixed thereafter, but our specification of daytypes is based solely on historic data and is, likewise, fixed thereafter. While this approach is not in keeping with the conception of LMSTM, we believe we are in the mainstream of LMSTM users. Further, while taking more exhaustive advantage of LMSTM's forecasting prowess is a laudable goal in itself, since our mission was largely concerned with the the issue of calibration of this model's dispatch to that of other models, it was beyond the scope of this effort. By definition, we assume the forecasts of energy use, peak, etc. are exogenously determined. They are a given for our analysis, and our objective is merely to compare the the dispatch of LMSTM to that of other models, notably, GRASS.

A less substantive, but equally frustrating, disadvantage with LMSTM's terminology concerns the naming and numbering of seasons and daytypes. Seasons 1 and 2 *have* to be called WINTER and SUMMER; this much is clear from the manual. Seasons 3 and 4 can be named as the user chooses. For the purposes of this study, the same names as previously used by PG&E are maintained, i.e. season 3 is FALL and season 4 is SPRING. That is, remember that the numbering of seasons is: WINTER = season 1, SUMMER = season 2, FALL = season 3, and SPRING = season 4. This numbering is perpetually confusing but derives from the sensible assumption that if a user wanted to define only two seasons, 1 and 2, he or she would most likely call them WINTER and SUMMER. Less clear in the manual is the restriction on the numbering of daytypes, which we discovered by trial and error. Daytype 1 *must be* the lowest output daytype and 4 the highest, etc. Changes in this ordering result in pathological model behavior.

3. PG&E load data

In the body of the text, the highly constrained nature of the seasonal split is described, and our seasonal and daytype definitions appear as Table 1. One of the constraints imposed on the split requires the WINTER to begin on Jan. 1, thus eliminating the need for deriving load shapes for daytypes from data that overlaps the years. Also excluded are seasons that are discontinuous. PG&E's earlier specification contained such seasons. Further, the desire for comparability with GRASS excluded the possibility of seasons that do not contain an exact number of months. These constraints, together with others that came from the nature of the resource mix leave few margins for choice in the seasonal split. Specifying the seasons in this way, however, makes manipulation of the load data considerably more tractable. Given the seasons selected, the

[5] For clarity's sake, we speak only of falls in load, but an actual control program will likely cause both falls and increases in load at various times.

question becomes how to represent the variation experienced in each season by exactly 4 daytypes.[6]

Our daily load shapes are based on historical PG&E hourly area loads. Historical loads were available for the years 1979-84, although the data set for 1984 was in a different format and required some preprocessing before it could be used as a parallel input to the other data.[7] One of the reasons for our interest in the load shape inputs was the desire to respecify the model for the PG&E area, rather than its territory. This required recasting many inputs, including the system load shapes. The load data contains many shapes for different subsets of the PG&E area, but we concentrated on shape #7, and all of our shapes are derived from it. The load data documentation describes #7 as "Area Load 1983 Definition (AL-DWR)," and the equivalent definition for the 1984 format is "HLL #216 - 143."

4. Defining daytypes

The LMSTM manual describes the four daytypes as "weekend day (daytype 1)," "normal weekday (daytype 2)," "peak weekday (daytype 3)," and "extreme weekday (daytype 4)," but is somewhat vague about the exact definitions of these categories. It suggests that the frequency for each daytype should be as follows: weekend (28% of days), normal (57% of days), peak (10% of days), extreme (5% of days). The logic behind distinguishing extreme and peak daytypes with relatively few days reveals the load management orientation of LMSTM. For the purposes of evaluating load control programs, these days are likely to be of special interest, and from the point of view of financial evaluation, cost savings on these daytypes may dominate the analysis. We have maintained the principle of high load daytypes with few members since our policy case will be a load-shaving one. One of the immediate questions to arise concerns the variable by which days are ranked from extreme to normal. Four possible variables that spring to mind are daily peak GW load, a measure of daily load variation such at the variance or peak to off-peak ratio, system GWh output, or a combination index of some kind. The manual proposes ranking the days of the year in order of "energy load," a piece of schizophrenic jargon, that appears to mean GWh output. *Prima facie*, the correlation between peaks and output should be high, and inspection of the PG&E data confirms this. However, in the case of the EXTREME and PEAK daytypes, the number of days is small enough that the resulting shape could be sensitive to a selection of days with only one or two different members. As a practical matter, we decided to rank days only by GWh output, in keeping with the manual's recommendation. An interesting sensitivity case would involve a ranking on another variable, perhaps an index that included peak, output, and other variables, such as temperature. The correct ranking procedure is not at all an obvious matter, and, although the die is now cast for this study, we do not consider the issue closed by any means. The definition of daytypes is unconstrained relative to the choice of seasons, and the possibility of choosing some other kind of day other than high load ones for special scrutiny is still open.

Having ranked each annual data block in descending order of system output, the next problem is where to make the incisions in the data set that divide it into the daytypes. The definition of WEEKEND could either be based on whether or not the day in question was actually a work day or by assuming that all low load days are holidays.[8] The manual's suggestion that a fixed

---

6 In principle less daytypes could be defined with a consequent loss of detail.

7 Appendix A contains a sample of the PG&E data for 1979.

8 Note that in some rare service territories the holidays might be high load days, but this is not true of PG&E.

percentage go into each daytype leaves us uninspired. This approach has the user take the hierarchically ordered data set and sever it into blocks of equal size for the same daytype of each season. Surely better ways of grouping the days into natural sets exist.

5. Weekends

Turning firstly to the WEEKEND's, daytype 1, as mentioned above, two obvious definitions present themselves.

1. All the low output days could be assumed to be WEEKEND's, and so this daytype can be considered just the lowest block in the hierarchical listing of days. This is the manual's suggestion and it further proposes the lowest 28% of days be taken, which is very close to 2/7.
2. Actual holidays and weekends could be tagged in the data set and assigned to this daytype.

Perusal of the PG&E data shows that these two definitions result in different sets, although only by a few days, never more than 5 per season. This never constitutes a dominant fraction of the days in this daytype, which is always over 20. From a computational point of view, finding all the holidays involves the more complex algorithm. Some holidays, such as July 4, are easily found, but others, like Presidents' Day, are more difficult. Further, there would be some question about whether days such as Martin Luther King Jr.'s birthday and Columbus Day, should be considered holidays or not.

Two observations on the PG&E data lead us to a definition of WEEKEND that combines the two alternatives, and one that leads to a simple algorithm for finding them.

1. The first is that weekends and holidays are almost always low output days and they group together at the bottom of the ordered data. Outputs are generally higher on Saturdays than Sundays, but few weekdays intervene at the bottom 2/7 of the table. Looking further at actual weekend load shapes, they are qualitatively different to working weekdays, and form a natural group.
2. A second observation made is that holidays are always low output days. Even if they are above the bottom 2/7 block, they never make it above the NORMAL daytype block.[9] For example, although our WINTER season is long, and includes all of April, New Years' Day, which occurs at one of the coldest times of the year, is captured in the lowest 28% of daily outputs for every year of the PG&E data. In 1984, it falls on a Sunday and is the lowest output day of all in our WINTER season. In fact, appearance in the bottom block of the hierarchical data are a good predictor that a stray weekday is a holiday.

This suggests that all days in the bottom block be considered WEEKEND's whatever day of the week they fall on. Holidays that result in low output are categorized as WEEKEND's, and those that do not, end up in the NORMAL days. This eliminates the need for an algorithm to find all the holidays. Further, since the Saturdays and Sundays are easy to find, they are simply added to the days appearing in the lowest 2/7 block. This is the definition for daytype 1 that we adopted for all seasons; that is, 4 of the 16 daytypes are taken care of.

Also, if the researcher was attempting a bottom-up approach, holidays would be high load days for some enduses, e.g. cooking, and not for others.

[9] Definition of this daytype follows.

6. Other daytypes

Now we turn to the division of the remaining block of data into the other 12 daytypes. As noted above, we are suspicious of the manual's recommendation that a simple proportional split is adequate. At first, we believed that a standard clustering algorithm would be able to subdivide the data into natural daytype groupings and the question of how many days to place in each daytype would take care of itself. We followed up on this conjecture by generating an SPSS data set from the load data that contained only the date, the maximum load, and daily output for every day of the raw data. To account for year to year growth and weather variation, the load and output were normalized to 1.0 as the highest observation. The basic principle of clustering is simple. The objective is to subdivide a set of observations into groups such that the sum of the variances within groups is minimized while the sum of the variances between groups is maximized. Several algorithms for achieving this optimal division are available in SPSS and other statistical packages. We concentrated on minimization of the squared distances between group members approach, the most commonly used. The results of this attempt are sadly disappointing. The data simply do not cluster well for our purpose. Since the number of groups desired is 3, the approach is simply one of letting SPSS freely group the data into three clusters. If all of the years' data are used, the results are disappointing because the three clusters contain wildly different numbers of days from each year. If, for example, the extreme day cluster ends up with 30 days from 1979 and 3 from 1984, it is not clear what rule should be applied for assigning days to the extreme category. The same issue arises if the data is clustered one year at a time. The clusters, which are then only valid for one year of data, are wildly different between years, and do not provide an obvious rule for future daytype divisions. A further disadvantage of the clustering approach is that the sizes of the natural clusters do not coincide with our *a priori* belief that the PEAK and EXTREME days are ones of special interest. Natural clustering could result, for example, in an extreme daytype containing most of the days. From the point of view of expected costs, this may be the optimal choice; however, in the analysis of load management strategies, the few days of extra high output are likely to be of special interest, even though their effect on overall costs may not be great. In a sense, the qualitative analysis of results places a special premium on the extreme days. If a load shape is derived that simply averages these days in among a large category of extreme days, the special information is lost.

We finally adopted a more pedestrian but intuitively appealing method to defining the daytypes. Plots are generated that show the distribution of days by output for each season; an example appears as Figure 2. We "eyeball" the plots for each season and seek out a natural division into daytypes. The results for each season are shown in Figures 3-6. The seasons show interesting differences. The distributions for FALL and WINTER are more compact than the other two, and their modes come closer to the maximum. Given our penchant for a small number of days in the EXTREME and PEAK daytypes, the divisions are quite natural. They are certainly not, however, beyond dispute and, indeed, we do not feel we have really progressed very far at all towards defining a "correct" method for making the split.

7. Cost of averaging

However, given the split made, the numbers of days in each daytype appear Figure 1. We wrote a Fortran program, DTSPLT, that processes the PG&E data and produces a load shape for input to LMSTM. The program can produce daytype load shapes based on any one year of data or any group of sequential years of data. When the correct data are identified the days are averaged to produce a representative shape.[10] The choice of averaging to perform the reduction of

[10] DTSPLT is a simple f77 program that is available from the authors.

multiple days to one shape was, probably, an excessively hasty one. We come to this conclusion because the shapes produced by DTSPLT are flatter than those previously used by PG&E.

Figures 7-10 portray some sample load outputs from LMSTM. Note that the vertical scale varies to conveniently demonstrates the differences between the plots on that page. All of the data are actual LMSTM run output for the years and daytypes noted. The LBL shapes are the outputs that result from using LBL's average load shapes as inputs.[11] The MAN output results from using the original PG&E input with only the MANCON customer class active, while FIV has the first five classes active. For the SUMMER EXTREME day, as a general rule, the LBL shape is less extreme, i.e. lower peaks and shallower valleys, than either PG&E alternative. The difference reaches 8.5%. Qualitatively, however, the shapes are similar. The SPRING NORMAL days are more interesting. Qualitatively, the LBL shape differs in two details; the afternoon peak shows a lower bump between 13:00 and 18:00, and the LBL shape exhibits a dramatic secondary peak at 21:00. During the off-peak period, LBL appears between MAN and FIV, one of the very few occasions it does so. PG&E folklore claims that the shapes originally used in LMSTM were based on a forecast for 1983, and not actual data. However, the differences are disturbing because less extreme loads lead to less extreme marginal costs and less interesting time-dependent savings. Among the other daytypes it is generally true that LBL's average load shape exhibited less extreme behavior than the original PG&E shapes.

We are concerned that perhaps the flattening of the LBL shapes results from the averaging process used to derive them. Naturally, averaging across daytypes and years tends to eliminate extremes. This brings us to a recurring dilemma. On the one hand, average conditions are the natural ones to plan for and the ones on which to base financial estimates, and yet, on the other hand, studies of load control naturally focuses on extreme conditions. Further, over time the value of a load control program may be dominated by high returns in extreme years or on extreme days. In other words, the value of the program might be non-linearly related to load. This is a weighty issue that we have no doubt we will return to, yet the problem at hand must be faced.

8. Judgment call

Consider the messy Figure 11, which shows 7 legitimate load shape inputs to LMSTM, all derived from the PG&E data. The ones named 1979-84 are based on a single year's data, while "average" is the shape we have hitherto discussed. Qualitatively, the shapes are similar; year to year variation is not great when loads have been normalized to a peak. In fact, on this scale following any one shape among the pack is impossible. Figure 11 lulls the viewer into the belief that year-by-year variations in the shape are small and unimportant. Interestingly, 1979 and 1981 have peaks at 17:00, 1982, 1983 and average at 14:00, and 1980 and 1984 have bimodal peaks at 15:00 and 16:00. No year has a unique 15:00 peak, and 16:00 is the clear mode. Note that 1983 is something of an outlier in the off-peak hours but is in the midst of the pack on-peak. This is just the kind of load variation we are looking for, and after much soul-searching we elected to use shapes based only on 1983 data. This is purely a judgment call, although 1983 is considered a fairly typical weather year and it is not an unreasonable choice, being one of the latter years in the sample. Since we are suspicious of the averaging process, a latter version of

[11] The PG&E outputs come from a hybrid case which employs PG&E load shapes but LBL seasons and daytypes. The results are, therefore, meaningless in themselves but serve as a basis for comparison.

DTSPLT might derive representative shapes by some other criterion, such as the mode. However, the consequences of such a strategy for forecast results requires careful thought.

Table 1

Energy Use Ratios

SEASON 1, DAYTYPE 1:	0.662
SEASON 1, DAYTYPE 2:	0.744
SEASON 1, DAYTYPE 3:	0.773
SEASON 1, DAYTYPE 4:	0.803
SEASON 2, DAYTYPE 1:	0.750
SEASON 2, DAYTYPE 2:	0.854
SEASON 2, DAYTYPE 3:	0.950
SEASON 2, DAYTYPE 4:	1.000
SEASON 3, DAYTYPE 1:	0.681
SEASON 3, DAYTYPE 2:	0.764
SEASON 3, DAYTYPE 3:	0.807
SEASON 3, DAYTYPE 4:	0.843
SEASON 4, DAYTYPE 1:	0.691
SEASON 4, DAYTYPE 2:	0.804
SEASON 4, DAYTYPE 3:	0.865
SEASON 4, DAYTYPE 4:	0.904

9. Why it matters

Figure 12 is designed to show that the issues delved into here are of some importance and not merely of academic interest. Figure 12 is by no means a typical case; it is a carefully chosen example, but an important one. Two LMSTM load outputs are shown, based on two sets of load shape inputs, the average and 1983 versions. Once again, the shapes are qualitatively similar, and particularly the lower night-time load of the 1983 shape seems inconsequentially different. Below, however, the marginal cost output appears for each case and the 1983 load produces a dramatically different marginal cost off-peak, more than 2 c/kWh lower. This is exactly the kind of discontinuity effect that intuitively the analyst knows is going on in the model when loads change, yet it is type of effect very few would bother to track down through sensitivity cases to the load inputs. Of course, the change in off-peak marginal cost will alter pumping prospects markedly in this example. Further, it is important to bear in mind that there are fully 30 days of this type in our specification, so the net effect on annual results might also be noticeable.

An interesting exercise would be to rebuild from the LMSTM input load shapes load duration curves for PG&E. This is computationally not difficult, and these could be compared to parallel data from GRASS. Unfortunately, the necessary GRASS data have not been available to us to undertake this exercise.

10. Energy use weights

The daytype energy use weights follow naturally from the daytype definition because each daytype is associated with an output level derived from the original load data. The weights are shown in Table 1.

11. Conclusion

The choice of daytype definition is a critical step in the analysis conducted so far. The LMSTM manual and the general modeling literature offer little guidance on this difficult problem. We attempted to unearth a natural division of the days into groupings using a clustering algorithm but found this pursuit futile. We resorted to an eyeballing approach using the distributions of daily output for each daytype. Using PG&E's load data files, we derived representative system load shapes for each daytype by averaging across the days in the daytype, and, in one case, across all the years of the data. We settled on the shapes derived from 1983 data alone as the most sensible choice for our purposes. Many issues related to the definition of daytypes remain unclear and this should be a fertile area of future modeling research.

Figure 1

Distribution of Daytypes for all Seasons

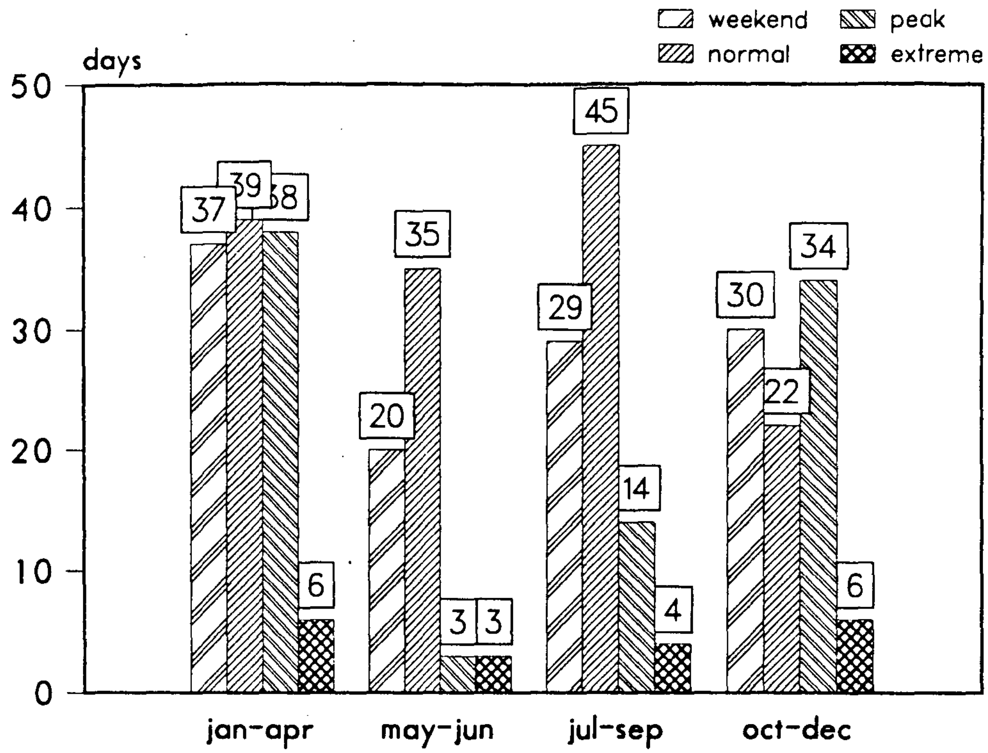

Figure 2

Distribution of Total Daily Sales, All Winters (excluding weekends and holidays)

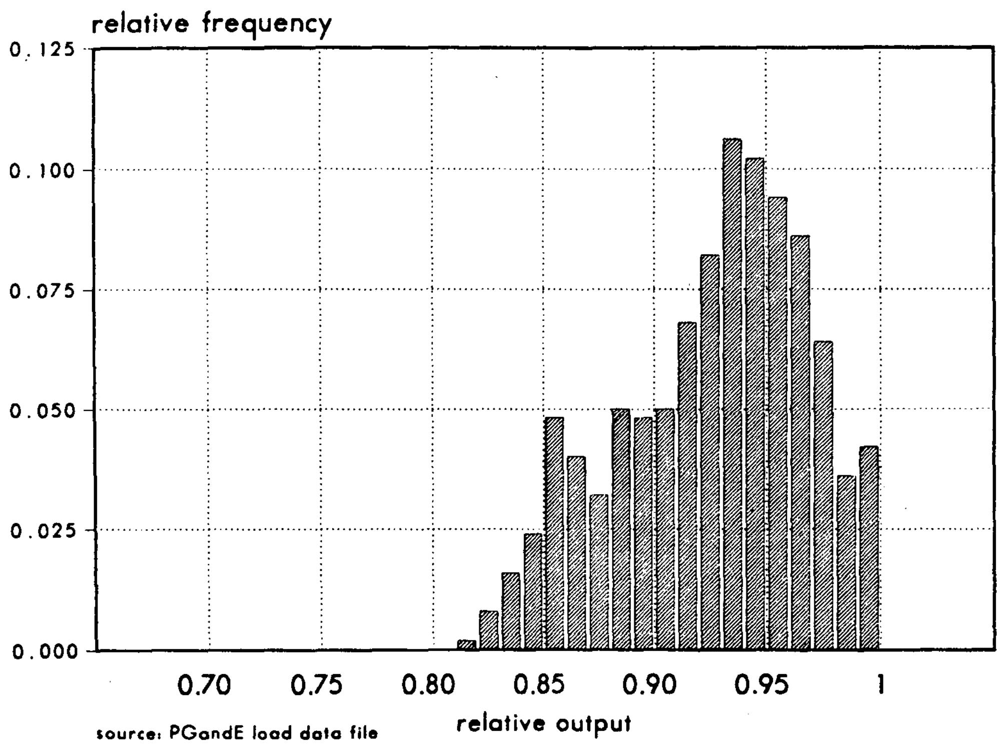

-- XBL 872-535

Figure 3

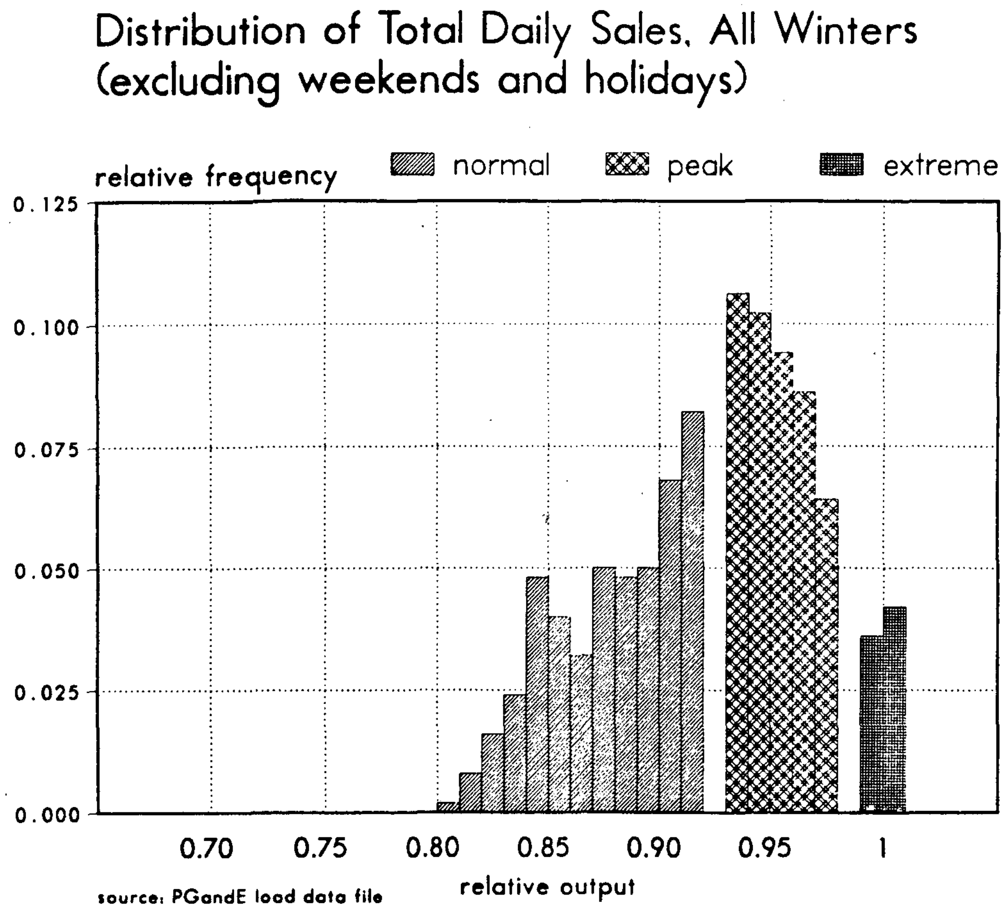

Figure 4

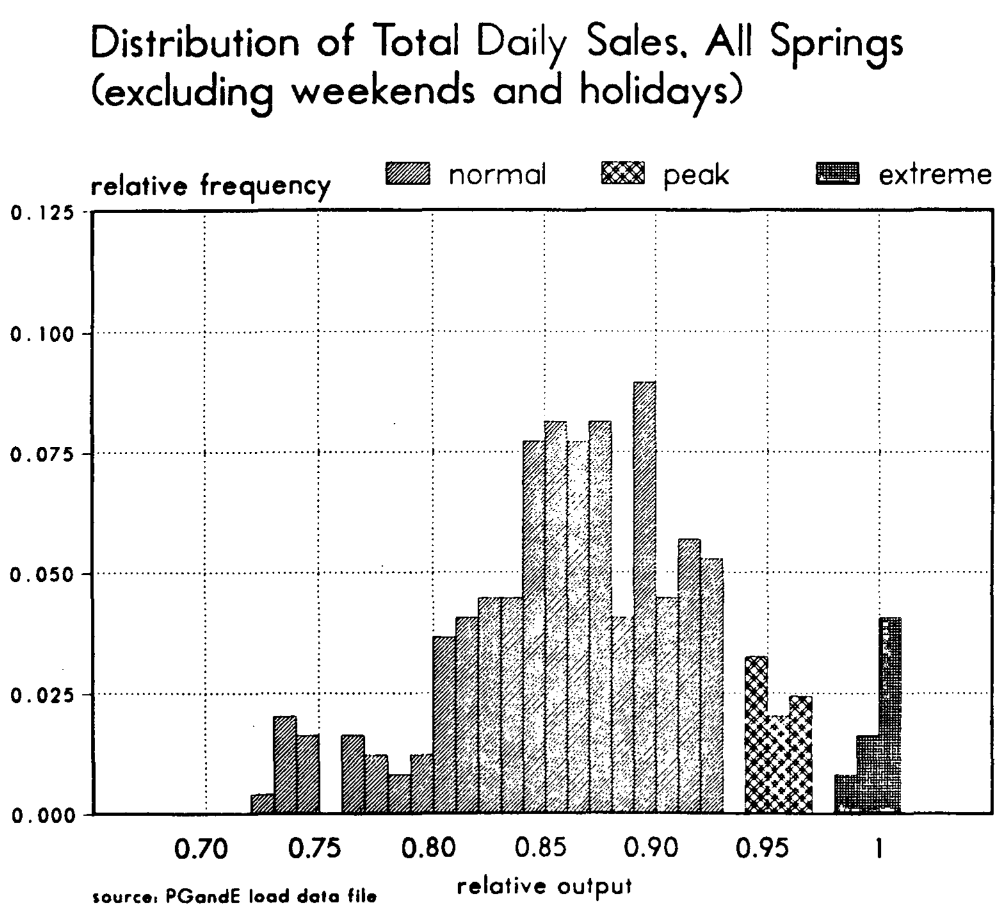

-- XBL 872-536

Figure 5

Distribution of Total Daily Sales, All Summers (excluding weekends and holidays)

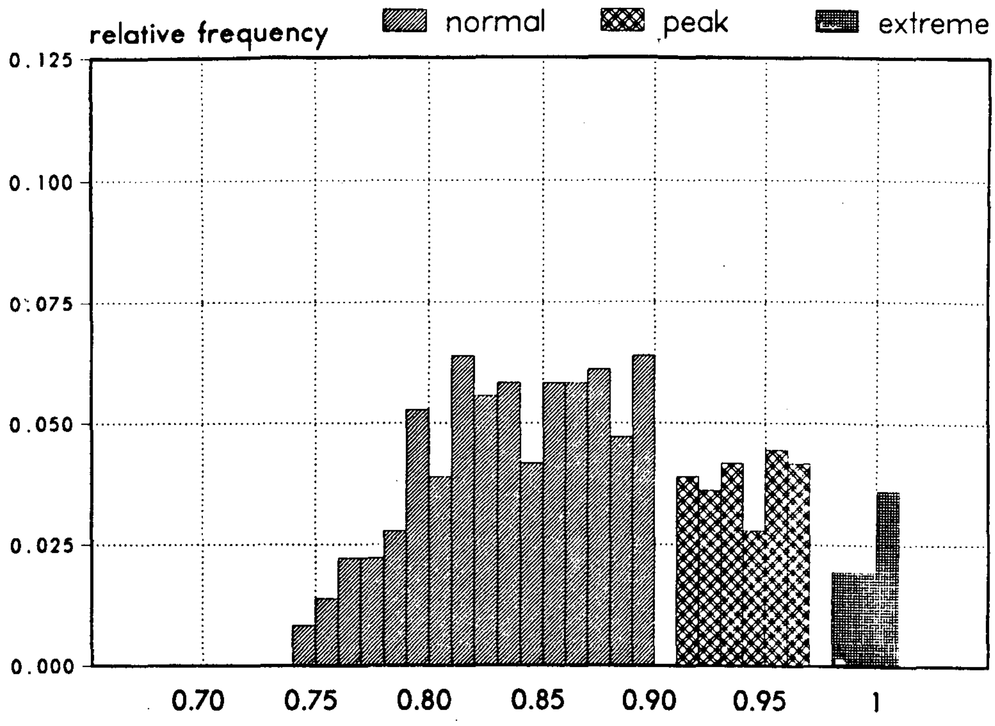

source: PGandE load data file

Figure 6

Distribution of Total Daily Sales, All Falls (excluding weekends and holidays)

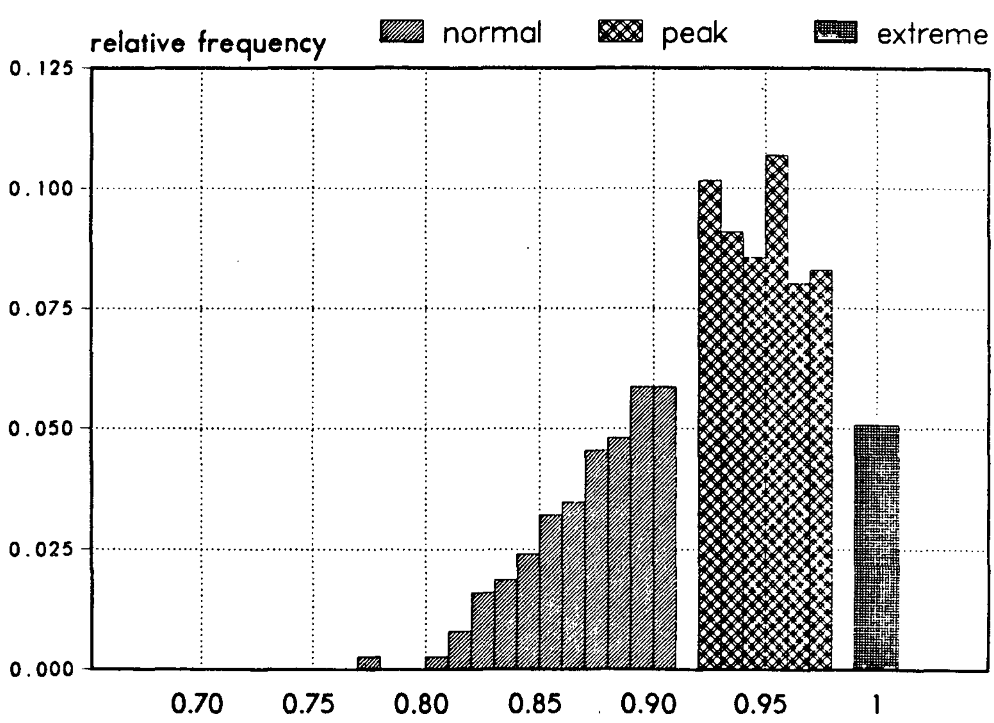

source: PGandE load data file

XBL 872-537

Figure 7

Comparison of Three PGandE Load Shape Outputs from LMSTM
(Spring Normal Days 1989)

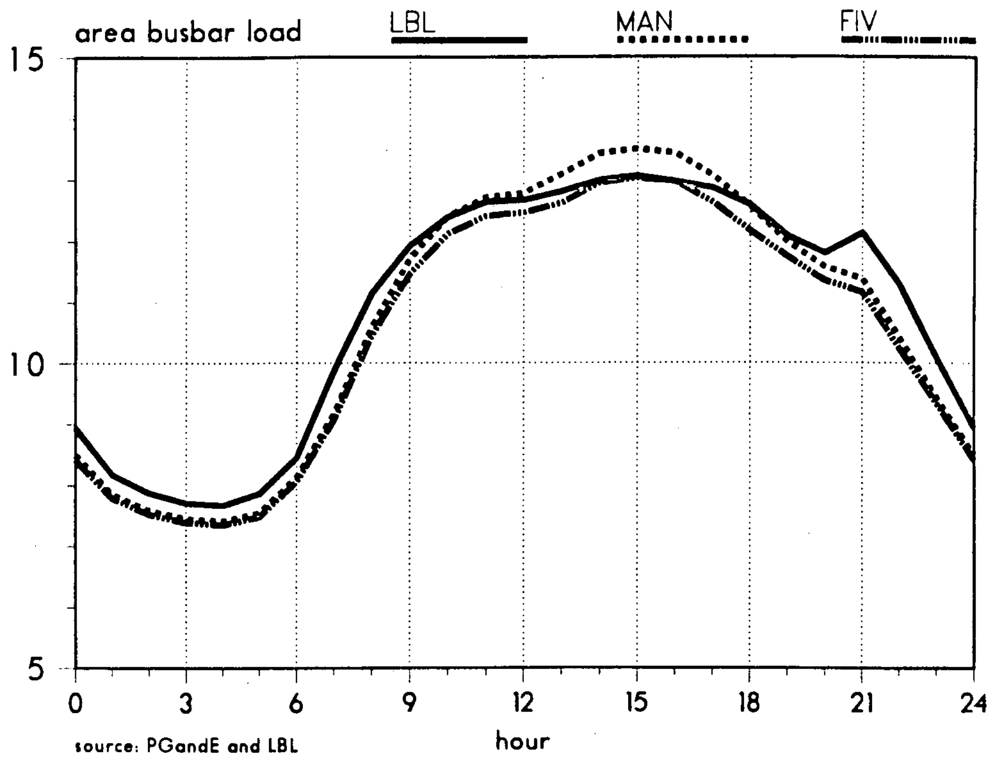

Figure 8

Comparison of Three PGandE Load Shape Outputs from LMSTM
(Spring Normal Days 2003)

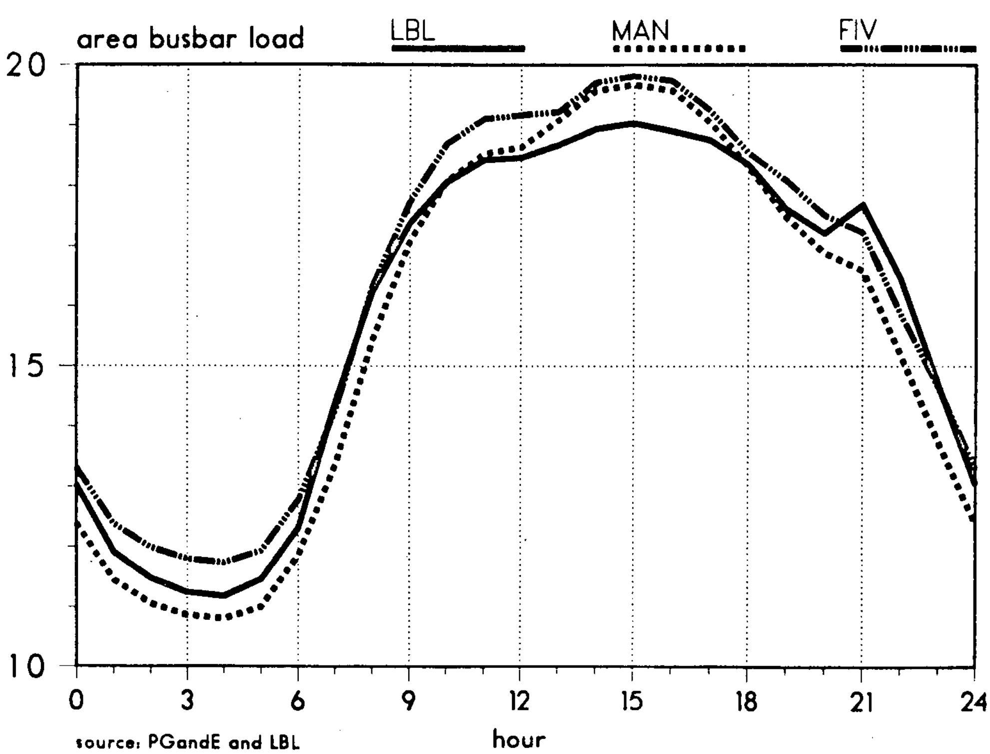

-- XBL 872-538 -

Figure 9

Comparison of Three PGandE Load Shape Outputs from LMSTM (Summer Extreme Days 1989)

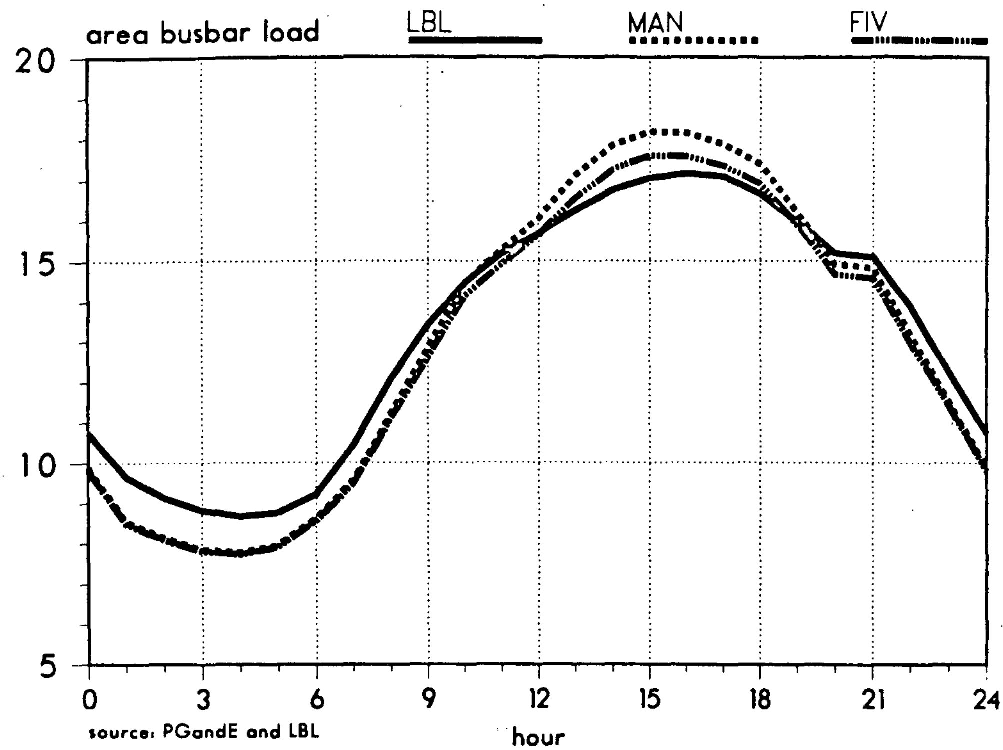

Figure 10

Comparison of Three PGandE Load Shape Outputs from LMSTM (Summer Extreme Days 2003)

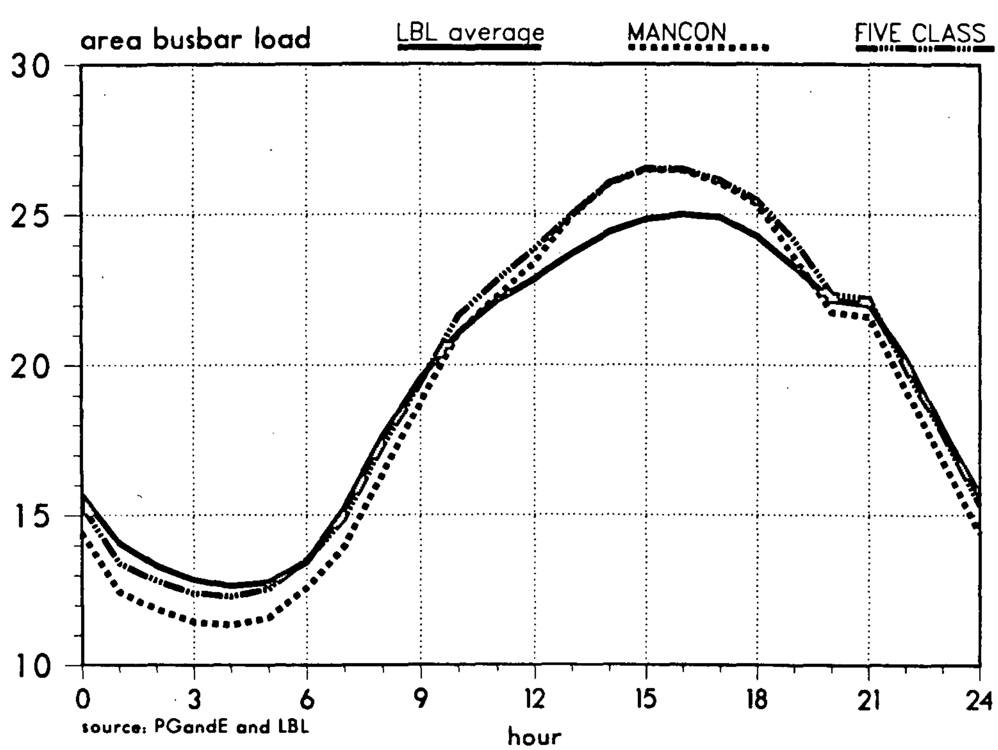

-- XBL 872-539 -

Figure 11

Seven Possible PGandE Load Shape Inputs to LMSTM (Summer Extreme Day)

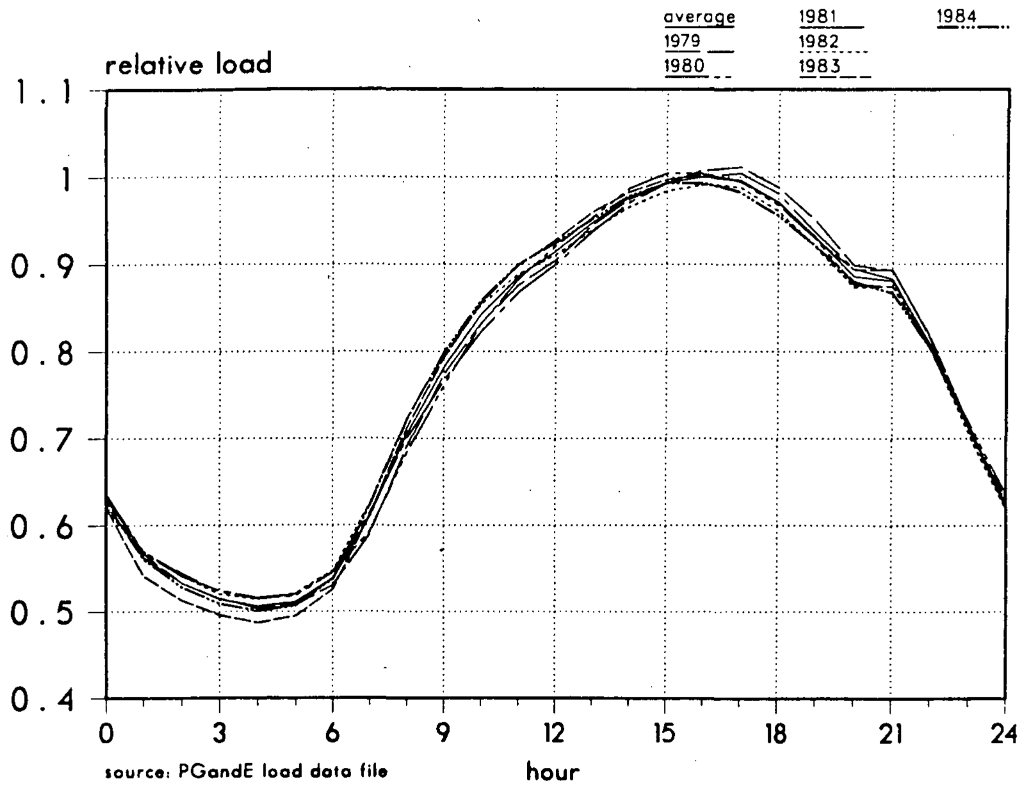

Figure 12

Extreme Example of Marginal Cost Sensititivity to Load Shape Inputs - Fall Weekend, 1989

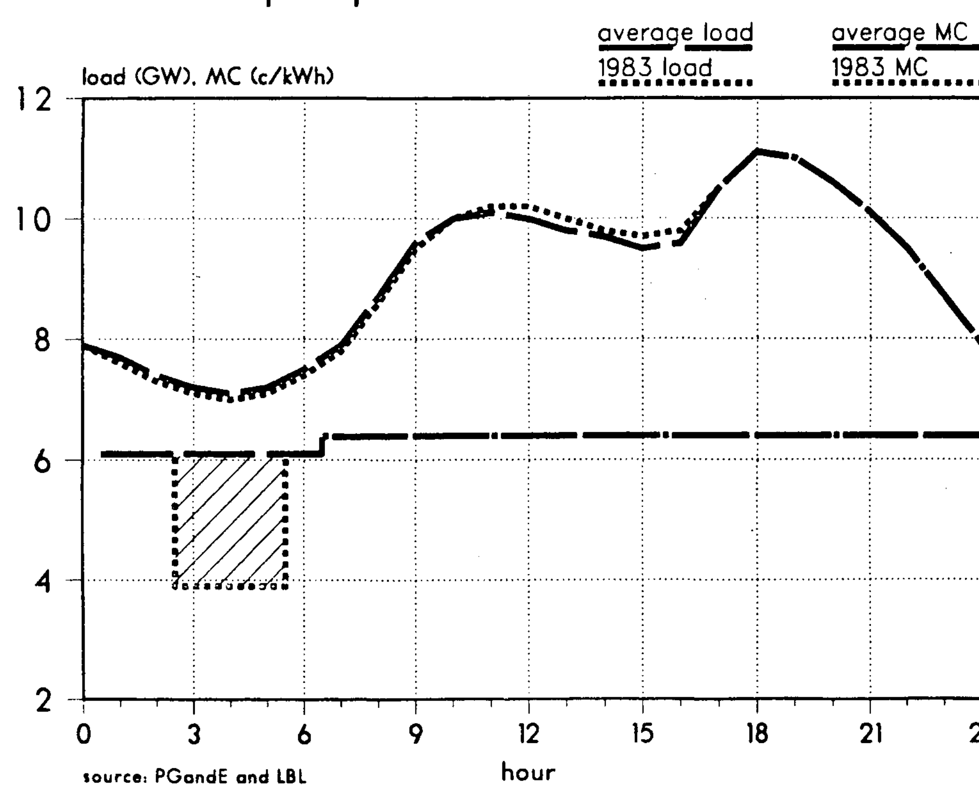

-- XBL 872-540 -

Appendix B

Translating GRASS Hydro Data into LMSTM Inputs

1. Introduction

This appendix describes a method for adapting hydro data from the GRASS model to LMSTM. Specifically, this appendix covers the following resources: PG&E's basic hydro, QF hydro, Northwest hydro, and Northwest thermal. Correct representation of PG&E's hydro system is essential for successful modeling of PG&E's supply system because seasonal variation in marginal costs is closely related to hydro availability. Unfortunately, because the format of the hydro information in GRASS is very different than the format required by LMSTM, calibrating the two models is not a trivial exercise.

This appendix describes a general method for adapting GRASS hydro data to LMSTM. Then this method is applied to the GRASS run used for PG&E's 1985 Long-Term Plan (LTP).

2. Representation of Hydro in GRASS

Tables 1 and 2 show hydro data from a typical year in GRASS.

Three aspects of the tables are strikingly different to LMSTM's inputs.

A. GRASS is able to operate a monthly resource allocation and data is available for each month, whereas LMSTM divides the year into four seasons.

B. In the example of Table 1 and 2, 22 separate categories of hydro are differentiated (7 categories are shown). This data represents the hydro resources of PG&E, qualifying facilities (QF's) selling hydro to PG&E, and Northwest firm energy that is sold to both PG&E and municipal utilities (MUNIs). In addition, there are 6 *thermal* plants included in the GRASS run that represent Northwest economy hydro energy. However, these resources are treated separately by GRASS and are discussed in section 4.2.

While LMSTM allows the user to operate a large number of individual plants, they have to be grouped into a relatively small set of GENERATING TECHNOLOGY GROUPS (24 maximum), and dispatch is done at this level.

C. For each hydro category listed in the GRASS data, the data set listed consists of minimum and maximum capacities (Table 1) and an energy output for every plant every month (Table 2).

3. General Methodology for Translating GRASS Hydro Data into LMSTM Inputs

The problem at hand is how GRASS hydro data should be translated to an LMSTM format. The following methodology has been developed to solve this problem.

A. The GRASS annual data are broken into the chosen LMSTM seasons. The nature of the GRASS data immediately shows that it is wise to define the LMSTM seasons so that they consist of whole months. Note that the definitions previously used by PG&E were not whole months, and, in fact, were not defined as continuous time periods. For the purposes of aggregating to seasons, the energy numbers are summed across the months of the season, and the capacities are averaged.

B. The number of plants is radically reduced to six technology groups as shown in Table 3. For PG&E hydro and QF hydro, the minimum capacity of the GRASS data is taken to be

Table 1

Sample GRASS Hydro Capacity Data

Pacific Gas and Electric Company Run LT85BAS1

Final SPC Values - All Modifications - Avg. Period 1985-2005

Monthly Simulation Release PCOST85F Date 85086

Avg. Year 1985 Yearly Page 3

Annual Input Summary 1985 (continued)

Hydro Generation Unit Data

NO.	Unit ID NAME	CAPACITIES (MW)												
1	BNONPGE	28	34	37	38	68	68	51	43	44	44	46	52	MAX. CAP.
		11	15	21	19	26	33	38	27	16	7	10	19	MIN. CAP.
	.	.	.	.	.	.	.	.	.	.	.	.	.	
	.	.	.	.	.	.	.	.	.	.	.	.	.	
9	NWNCPAHY	13	13	13	13	9	13	18	20	27	25	24	18	MAX. CAP.
		0	0	0	0	0	0	0	0	0	0	0	0	MIN. CAP.
10	NWPGEHY1	519	519	519	519	229	519	0	0	0	0	0	0	MAX. CAP
		0	0	0	0	0	0	0	0	0	0	0	0	MIN. CAP.
11	NWPGEHY2	0	0	0	0	0	0	452	528	819	708	674	447	MAX. CAP.
		0	0	0	0	0	0	0	0	0	0	0	0	MIN. CAP.
12	NWPGEHY3	0	0	0	0	0	0	452	528	819	708	674	447	MAX. CAP.
		0	0	0	0	0	0	0	0	0	0	0	0	MIN. CAP.
13	NWSMUDHY	41	41	41	41	13	41	78	92	150	127	120	77	MAX. CAP.
		0	0	0	0	0	0	0	0	0	0	0	0	MIN. CAP.
14	NWWAPAHY	123	123	123	123	52	123	218	225	400	259	268	215	MAX. CAP.
		0	0	0	0	0	0	0	0	0	0	0	0	MIN. CAP.
	.	.	.	.	.	.	.	.	.	.	.	.	.	
	.	.	.	.	.	.	.	.	.	.	.	.	.	
	TOTAL	1391	1687	1930	2193	2561	2198	1858	1635	1530	1198	1304	1310	MIN. CAP.
		5760	5859	5805	6034	5748	6373	6691	6799	7662	6890	6827	6365	MAX. CAP.

the ROR capacity and the remainder is taken to be the pondage. For Northwest hydro, GRASS always assumes that the minimum capacity is zero. Thus, there is only Northwest pondage hydro and no run-of-river Northwest hydro.

C. To allocate the energy to the technology groups, the following approach is taken. ROR capacity is assumed to operate at 100% load factor in every season. The remaining seasonal energy is allocated to the corresponding PON group.

Table 2

Sample GRASS Hydro Energy Data

Pacific Gas and Electric Company Run L185BAS1

Final SPC Values - All Modifications - Avg. Period 1985-2005

Monthly Simulation Release PCOST85F Date 85086

Avg. Year 1985 Yearly Page 4

ENERGIES (GWH)

1	BNONPGE	16	20	24	25	34	39	35	24	15	10	10	17	262
	.	.	.	.	.	.	.	.	.	.	.	.	.	
	.	.	.	.	.	.	.	.	.	.	.	.	.	
9	NWNCPAHY	2	2	2	2	1	2	3	3	3	4	3	3	30
10	NWPGEHY1	77	70	77	75	34	75	0	0	0	0	0	0	408
11	NWPGEHY2	0	0	0	0	0	0	67	79	111	105	97	66	525
12	NWPGEHY3	0	0	0	0	0	0	67	79	111	105	97	66	525
13	NWSMUDHY	6	5	6	6	2	6	12	14	21	19	17	11	125
14	NWWAPAHY	18	17	18	18	8	18	32	34	164	39	39	32	437
	.	.	.	.	.	.	.	.	.	.	.	.	.	
	.	.	.	.	.	.	.	.	.	.	.	.	.	
	TOTAL	1784	1861	2037	2315	2672	2474	2478	2331	2261	1929	1840	1820	25802

Table 3

LMSTM Technology Group Definitions for its Hydro Resources

Group Name	PG-ROR	PG-PON	QF-ROR	QF-PON	NW-PON	NW-THM
Group ID	20	21	22	23	24	12
Plant Type	must-run	energy limited	must-run	energy limited	energy limited	shutdown
Plant Type No.	7	2	7	2	2	4
Group Description	PG&E's run of river hydro	PG&E's pondage hydro	QF run of river hydro	QF pondage hydro	NW firm energy, treated as pondage	NW economy energy, treated as thermal

D. Associated with each LMSTM technology group is one or more TECHNOLOGY TYPES. These types specify the capacity, cost, and availability of a specific resource. An example of a technology type specification (from SUPPLY.IN) is shown in Figure 1. Because the hydro resource varies from season to season, it would at first seem necessary to specify four technology types for every technology group shown in Tables 1 and 2, a total of 24 technology types. However, a clever way around this requirement is available by the judicious use of the plant maintenance scheduling. When a plant is required to go on maintenance for a certain fraction of a season, LMSTM has the capability to downrate the plant's capacity for the season. This seasonal downrating feature can be used to obtain just the hydro capacities desired for any season.

a) *ROR*. The ROR problem is relatively easy, and is shown in Figure 2. Since a ROR plant is assumed to run at 100% load factor, it can be designated a MUST RUN facility, and be given a capacity equal to its level in the highest season C_{max}. If left in this way, LMSTM would run at full speed all year and an excess energy output of the hatched area in Figure 2, A + B + C, would result. The key is simply to prevent the model from generating too much energy by setting a maintenance schedule that diminishes the available capacity to the desired level for each season, C_1, C_2, etc. LMSTM's maintenance scheduling comes at the TECHNOLOGY TYPE level and consists of five numbers, as can be seen in Figure 1. Each column is a TECHNOLOGY TYPE. The first number of interest are SCHEDULED MAINT, the fraction of total time lost each year. This total scheduled maintenance is represented in Figure 2 as $(A + B + C) / C_{max} * h_4$, where h_4 is normally 8760 hours. The other four inputs to LMSTM, labeled SEASON FRACTION, are the seasonal fractions for each season. These numbers are shown in Figure 2 as A / (A + B + C), B / (A + B + C), etc.

b) *PON*. The same hatched area, shown in Figure 3 as A + B + C, is lost by the capacity shortage, and this can be taken care of in the same way as for the ROR. In the pondage cases, however, there is also a "loss" represented by the areas D + E + F + G, which is a result of the ENERGY LIMIT, and so the actual energy produced is represented by the cross-hatched areas H + I + J + K. The energy limit becomes binding when there is not enough water behind the dam to run the generators at full speed. LMSTM has a category of plant for ENERGY LIMITED technologies and this must be the classification used for this application. This input block, labeled ENERGY LIMIT, can be seen in Figure 1. Only if a group has this classification does an energy limit entered in the input file become binding. The energy limit inputs consist of five numbers, as shown in Figure 1. The first is the number of hours that the plant could be run at full speed before the energy limit would be reached. In this case, that is annual total energy available divided by the highest capacity level, or in the figure, $(H + I + J + K) / C_{max}$. The other four SEASON FRACTION's distribute the total energy available among the seasons. These are represented in Figure 3 as H / (H + I + J + K), etc.

4. LMSTM Hydro Input Data Using the 1985 GRASS Long-Term Plan

Using the GRASS Print4 output from PG&E's 1985 LTP, we created the necessary inputs to correctly represent PG&E's hydro system in LMSTM. This section discusses both the hydro and the Northwest thermal resources that are specified in GRASS.

a) Hydro Resources

As described in the general methodology section above, it is necessary to aggregate GRASS hydro data across units and across months. The desired numbers of this aggregation procedure are yearly capacities, seasonal maintenance schedules, and, if a pondage hydro resource, seasonal energy limits. We aggregated the hydro resource data for 1985, 1986, 1989, 1990, 1993, 1996 and 2003 using a spreadsheet written on Lotus 123. The inputs and outputs of this aggregation procedure are shown for 1989 in Figure 4. The first page of Figure 4 shows information that is copied directly from the 1985 GRASS run. The second page shows the spreadsheet's outputs. The top third of page 2 of Figure 4 shows the seasonal capacities and energies for LMSTM's hydro technology groups: PG-ROR, PG-PON, QF-ROR, etc. Note that there is never any capacity or energy listed for NW-ROR, because the hydro resource from the Northwest is treated totally as pondage or thermal energy (see below for separate description of the Northwest thermal resource). The middle third of page 2 of Figure 4 converts the seasonal capacity data into a seasonal maintenance schedule. For example, for PG-ROR hydro, the maximum seasonal capacity, C_{max}, is 2350MW and it is available in the spring. The total scheduled maintenance required in 1989 relative to C_{max} is shown in the far right column of the figure. For PG-ROR, its value is 0.2716. The fractional allocations of the the total maintenance requirement are shown to be 0.292 (WINTER), 0.00 (SPRING), 0.278 (SUMMER), and 0.431 (FALL).

If the hydro resource is an energy limited technology, the spreadsheet computes the yearly energy limit (MWh/MW, equivalent to hrs.) and the seasonal allocations of the energy limit. These numbers are shown in the bottom third of page 2 of Figure 4. For PG-PON, the yearly energy limit is 2267 hours and the seasonal fractions are 0.2857 (WINTER), 0.1753 (SPRING), etc.

If we were only interested in 1989, then we would need only one technology type for each of our technology groups. However, we are interested in the period 1983 to 2005. While technology type definitions can allow updating of capacities over time, their representation of maintenance schedules, energy limits, and prices are fixed for all times. Unfortunately, these variables show considerable change over our period of analysis for some of our technology groups. Thus, we are forced to create two or more technology types for some of our technology groups. The following describes the LMSTM technology types that we created based on the 1985 GRASS run.

PG&E Hydro. For PG&E's own hydro resources, maintenance schedules, energy limits, and prices are relatively stable over time. Thus, only one technology type is necessary for each of the two technology groups. Table 4 shows some of the most important parameter specifications for the PG-ROR and PG-PON technology types.

Table 4
LMSTM Representation of PG&E's Hydro System

Group Name	PG-ROR	PG-PON
Type Name	PG-ROR	PG-PON
Type No.	55	56
User	PG&E	PG&E
Year Range	1983-	1983-
Seasons Covered	ALL YEAR	ALL YEAR
Point-Est. Year	1996	1996
Cap. Cost ($/kWh)	2126	2126
Fixed O&M ($/kW/yr)	11.93	11.93
Fuel Name	HYDRO	HYDRO
Fuel ID	21	21
Heat Rate (Btu/kWh)	10,000	10,000
Disp. Weight Factor	1.0	1.0
Energy Limits used?	NO	YES

A few notes will make Table 4 and the tables that follow easier to understand. In Table 4, the line "User" means the utility or group of utilities that use the power. This information is important because some hydro is dispatched by PG&E but is used by the MUNIs. The "Point-Estimate Year" is the year from which the maintenance schedule and energy limit fractions were computed from GRASS. The "Dispatch Weight Factor" is a term that reduces the variable operating cost of a plant seen by LMSTM when dispatching plants. Values other than 1.0 are used if there are non-price factors (such as QF dispatch priority due to PURPA) that affect dispatch.

QF Hydro. The hydro produced by QF's does show considerable change over time with respect to its scheduled maintenance and energy limits. Because of this, we created three technology types for every technology group. As shown in Table 5, the 3 technology types cover the following periods: 1983-1988, 1989-1995, and 1996 and beyond. During each of these periods, we chose a representative single year for the type's maintenance schedule and energy limits. These years are noted in Table 5 as the "Point-Estimate Year." Note, however, that capacities are updated in every year that significant change occurs.

Table 5

LMSTM Representation of Independently-Produced (QF) Hydro

Group Name	QF-ROR			QF-PON		
Type Name	QFROR1	QFROR2	QFROR3	QFPON1	QFPON2	QFPON3
Type No.	57	58	59	61	62	63
User	PG&E	PG&E	PG&E	PG&E	PG&E	PG&E
Year Range	1983-1988	1989-1995	1996-	1983-1988	1989-1995	1996-
Seasons Covered	ALL YEAR	ALL YEAR	ALL YEAR	ALL YEAR	ALL YEAR	ALL YEAR
Point-Est. Year	1985	1989	1996	1985	1989	1996
Cap. Cost ($/kWh)	0.0	0.0	0.0	0.0	0.0	0.0
Fixed O&M ($/kW/yr)	0.0	0.0	0.0	0.0	0.0	0.0
Fuel Name	QF-COG	QF-COG	QF-COG	QF-COG	QF-COG	QF-COG
Fuel ID	24	24	24	24	24	24
Heat Rate (Btu/kWh)	10,000	10,000	10,000	10,000	10,000	10,000
Disp. Weight Factor	1	1	1	0.1	0.1	0.1
Energy Limits used?	NO	NO	NO	YES	YES	YES

NW Hydro. As noted previously, it is a convention to represent Northwest firm energy as pondage hydro. Like QF hydro, we split up Northwest pondage hydro into three technology groups. The important characteristics of these groups are shown in Table 6.

b) Northwest Thermal Resources

As noted previously, it is a convention to model Northwest economy energy as a thermal resource. Translating GRASS's representation of this resource into LMSTM's format requires an aggregation procedure similar to the one performed for the hydro data. Table 7 shows where the Northwest thermal resource appears in the 1985 GRASS run. Because of the relatively smaller number of units and because we do not have to worry about energy limits, the task of creating LMSTM inputs for Northwest hydro is easier than was the case for the other hydro data.

Figure 5 shows the inputs and outputs of a spreadsheet written to aggregate the Northwest thermal data from GRASS. The top half of pages 1 and 2 of Figure 5 shows the Northwest

Table 6

LMSTM Representation of Northwest Firm Purchases

Group Name	NW-PON		
Type Name	NWPON1	NWPON2	NWPON3
Type No.	64	65	66
User	PG&E & MUNIs	PG&E & MUNIs	PG&E & MUNIs
Year Range	1983-1988	1989-1995	1996-
Seasons Covered	ALL YEAR	ALL YEAR	ALL YEAR
Point-Est. Year	1985	1989	1996
Cap. Cost ($/kWh)	0.0	0.0	0.0
Fixed O&M ($/kW/yr)	0.0	0.0	0.0
Fuel Name	NW-PUR	NW-PUR	NW-PUR
Fuel ID	22	22	22
Heat Rate (Btu/kWh)	9,480	9,480	8,990
Disp. Weight Factor	0.1	0.1	0.1
Energy Limits used?	YES	YES	YES

thermal hydro data for all years in which there is a significant change in capacity. Page 1 shows the capacity available to PG&E and page 2 shows the capacity available to the MUNIs. Note that while we disaggregate energy supplied to PG&E and to the MUNIs so as to accurately count PG&E's costs, *all* Northwest thermal energy is dispatched by PG&E. Note also that the yearly data on Northwest thermal is split into two. This is done because the price that is paid for Northwest thermal energy (by both PG&E and MUNIs) increases in the second half of the year. Because technology type definitions do not allow for prices that vary across seasons, it is necessary to treat energy purchased in the first half of the year (WINTER and SPRING) separately from energy purchased in the second half of the year (SUMMER and FALL).

The bottom half of pages 1 and 2 of Figure 5 compute the maintenance schedules four times for each year; once for PG&E (page 1) and once for the MUNIs (page 2) and, for each of these two utilities, once for the first half of the year and once for the second half. We found an additional technology type break-out was necessary when we found that the maintenance schedules change significantly over time. In total, the distinction between PG&E and the MUNI's, seasonal changes in price, and temporal changes in the maintenance schedules results in the creation of 8 technology types for the NW-THM technology group. The important characteristics of these technology types are shown in Table 8.

Table 7

Sample GRASS Northwest Thermal Hydro Data

Pacific Gas and Electric Company Run LT85BAS1

Final SPC Values - All Modifications - Avg. Period 1985-2005

Monthly Simulation Release PCOST85F Date 85086

Avg. Year 1985 Yearly Page 7

Thermal Unit Data

Annual Input Summary 1985 (continued)

	UNIT ID--						RATING (MW)						
	NAME	JAN.	FEB.	MAR.	APR.	MAY	JUNE	JULY	AUG.	SEP.	OCT.	NOV.	DEC.
	.	.	.	.	.	.	.	.	.	.	.	.	.
	.	.	.	.	.	.	.	.	.	.	.	.	.
3	NWNCPATH	14	14	14	14	18	14	9	7	1	2	3	9
4	NWPGETH1	560	560	560	560	705	560	1	1	1	1	1	1
5	NWPGETH2	560	560	560	560	705	560	1	1	1	1	1	1
6	NWPGETH3	1	1	1	1	1	1	734	581	1	223	289	745
7	NWSMUDTH	109	109	109	109	137	109	72	58	1	23	30	73
8	NWWAPATH	277	277	277	277	348	277	182	175	1	141	132	185
	.	.	.	.	.	.	.	.	.	.	.	.	.
	.	.	.	.	.	.	.	.	.	.	.	.	.

Table 8

LMSTM Representation of Northwest Economy Energy

Group Name				NW-THM				
Type Name	NWTPG1	NWTPG2	NWTPG3	NWTPG4	NWTMU1	NWTMU2	NWTMU3	NWTMU4
Type No.	33	34	35	36	37	38	39	40
User	PG&E	PG&E	PG&E	PG&E	MUNIs	MUNIs	MUNIs	MUNIs
Year Range	1983-1994	1983-1994	1995-	1995-	1983-1989	1983-1989	1990-	1990-
Seasons Covered	WIN, SPR	SUM, FALL	WIN, SPR	SUM, FALL	WIN, SPR	SUM, FALL	WIN, SPR	SUM, FALL
Point-Est. Year	1987	1987	1995	1995	1987	1987	1995	1995
Cap. Cost ($/kWh)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Fixed O&M ($/kW/yr)	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Fuel Name	NW-PUR	NW-PUR	NW-PUR	NW-PUR	NW-MUN	NW-MUN	NW-MUN	NW-MUN
Fuel ID	22	22	22	22	20	20	20	20
Heat Rate (Btu/kWh)	11,650	14,020	12,128	14,020	12,440	12,440	12,646	12,646
Disp. Weight Factor	1.0	1.0	1.0	1.0	1.0	1.0	1.0	1.0
Energy Limits used?	NO	NO	NO	NO	NO	NO	NO	NO

Figure 1

Sample Technology Type Specification in LMSTM

(E)
TYPE ATTRIBUTES

TECHNOLOGY TYPE ID #	55	56	57	58	59	00
TECHNOLOGY TYPE NAME	PG-ROR	PG-PON	QFROR1	QFROR2	QFROR3	
TECHNOLOGY GROUP ID #	20	21	22	22	22	0
CAPITAL COST ($/KW)	2126.	2126.	0.	0.	0.	0.
" " ESCAL.ID #	0	0	0	0	0	0
DIR CST SPNT AT BGN WU YR($/KW)	0.	0.	0.	0.	0.	0.
AFUDC ACCURED AT BGN WU YR($/KW	0.	0.	0.	0.	0.	0.
FIXED O&M DIRECT ($/KW/YR)	11.93	11.93	0.0	0.0	0.0	11.93
" " " ESCAL.ID #	0	0	0	0	0	0
FIXED O&M INDIRECT ($/KW/YR)	0.00	0.00	0.00	0.00	0.00	0.00
" " " ESCAL.ID #	0	0	0	0	0	0
VAR. O&M DIRECT ($/MWH)	0.00	0.00	0.00	0.00	0.00	0.00
" " " ESCAL.ID #	0	0	0	0	0	0
VAR. O&M INDIR. ($/MWH)	0.00	0.00	0.00	0.00	0.00	0.00
" " " ESCAL.ID #	0	0	0	0	0	0
MIN. OPERATING LEVEL (.%)	1.00	0.01	1.00	1.00	1.00	0.01
MIN. OPERATING COST (.%)	0.80	0.01	0.80	0.80	0.80	0.01
STARTUP COST ($/MW)	0.	0.	0.	0.	0.	0.
" " ESCAL.ID #	0	0	0	0	0	0
HEAT RATE (BTU/KWH)	10000.	10000.	10000.	10000.	10000.	10000.
DISTRIBUTED FUEL ID #	21	21	24	24	24	21
TRANSMISSION LOSS (.%)	0.00	0.00	0.00	0.00	0.00	0.00
DISPATCH WEIGHT FACTOR (.%)	1.0	1.0	1.0	1.0	1.0	1.0
SCHEDULED MAINT. (.% TIME/YR)	0.2723	0.0620	0.3158	0.2102	0.2104	0.000
SEASON 1 (WIN) FRACTION (.%)	0.2900	0.5475	0.4588	0.0705	0.0402	0.000
" 2 (SUM) " (.%)	0.2785	0.0421	0.0607	0.3460	0.3676	0.000
" 3 (FAL) " (.%)	0.4315	0.0000	0.4735	0.5824	0.5922	0.000
" 4 (SPR) " (.%)	0.0000	0.4103	0.0000	0.0000	0.0000	0.0000
STATE 1 AVAILABLILITY (.%)	1.0	1.0	1.0	1.0	1.0	1.0
" " PROBABILITY (.%)	1.00	1.00	1.00	1.00	1.00	1.00
STATE 2 AVAILABLILITY (.%)						
" " PROBABILITY (.%)						
STATE 3 AVAILABLILITY (.%)						
" " PROBABILITY (.%)						
STATE 4 AVAILABLILITY (.%)						
" " PROBABILITY (.%)						
ENERGY LIMIT (MWH/MW)	0.	2322.3	0000.	0000.	0000.	2000.
SEASON 1 (WIN) FRACTION (.%)	0.0000	0.2954	0.0000	0.0000	0.0000	0.3288
" 2 (SUM) " (.%)	0.0000	0.1910	0.0000	0.0000	0.0000	0.2521
" 3 (FAL) " (.%)	0.0000	0.3286	0.0000	0.0000	0.0000	0.2521
" 4 (SPR) " (.%)	0.0000	0.1849	0.0000	0.0000	0.0000	0.1671
EFFICIENCY (.%)	0.0	0.0	0.0	0.0	0.0	0.0
POWER RATIO (.%)	0.0	0.0	0.0	0.0	0.0	0.0
INITIAL CAPACITY (MW)	2350.	3791.	30.	0.	0.	0.
UNIT SIZE (MW)	25.	25.	25.	25.	25.	25.
YEAR OF CHANGE	1986	1986	1986	1989	1996	
NEW CAPACITY (MW)	2350.	3801.	50.	164.	253.	
YEAR OF CHANGE	1990	1989	1989	1990	2003	
NEW CAPACITY (MW)	2361.	3919.	0.	188.	257.	
YEAR OF CHANGE		1990		1993		
NEW CAPACITY (MW)		4136.		239.		
YEAR OF CHANGE		1993		1996		
NEW CAPACITY (MW)		4205.		0.		
YEAR OF CHANGE		1996				
NEW CAPACITY (MW)		4281.				
YEAR OF CHANGE						
NEW CAPACITY (MW)						
YEAR OF CHANGE						
NEW CAPACITY (MW)						
YEAR OF CHANGE						
NEW CAPACITY (MW)						
YEAR OF CHANGE						
NEW CAPACITY (MW)						
YEAR OF CHANGE						
NEW CAPACITY (MW)						

Figure 2

Seasonal Representation of Run of River Hydro

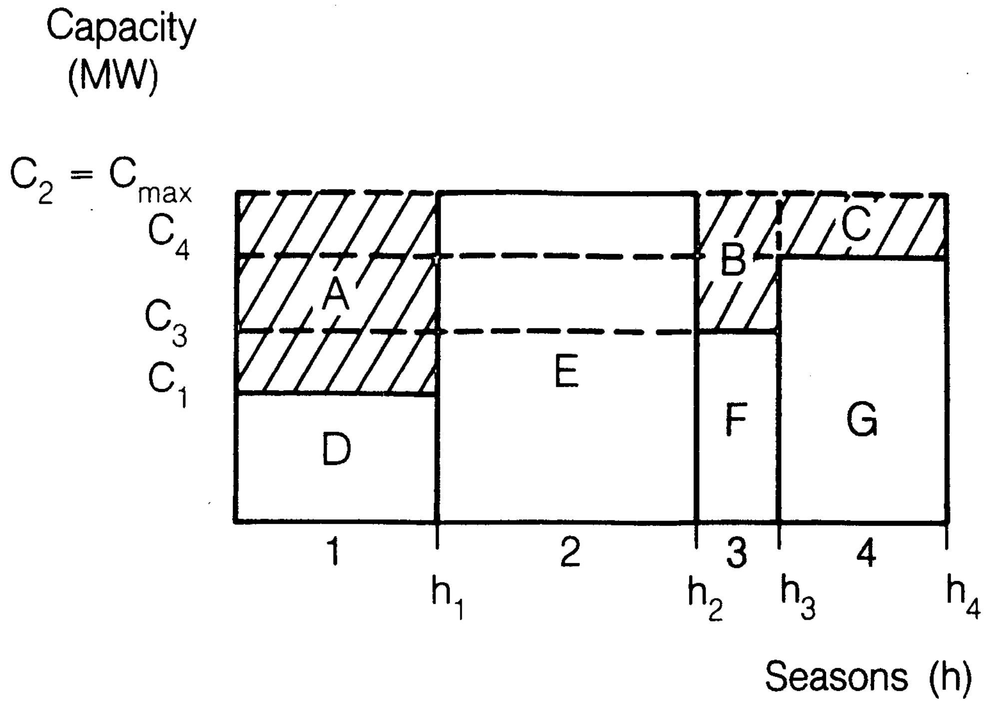

XBL 871-8803

Figure 3

Seasonal Representation of Pondage Hydro

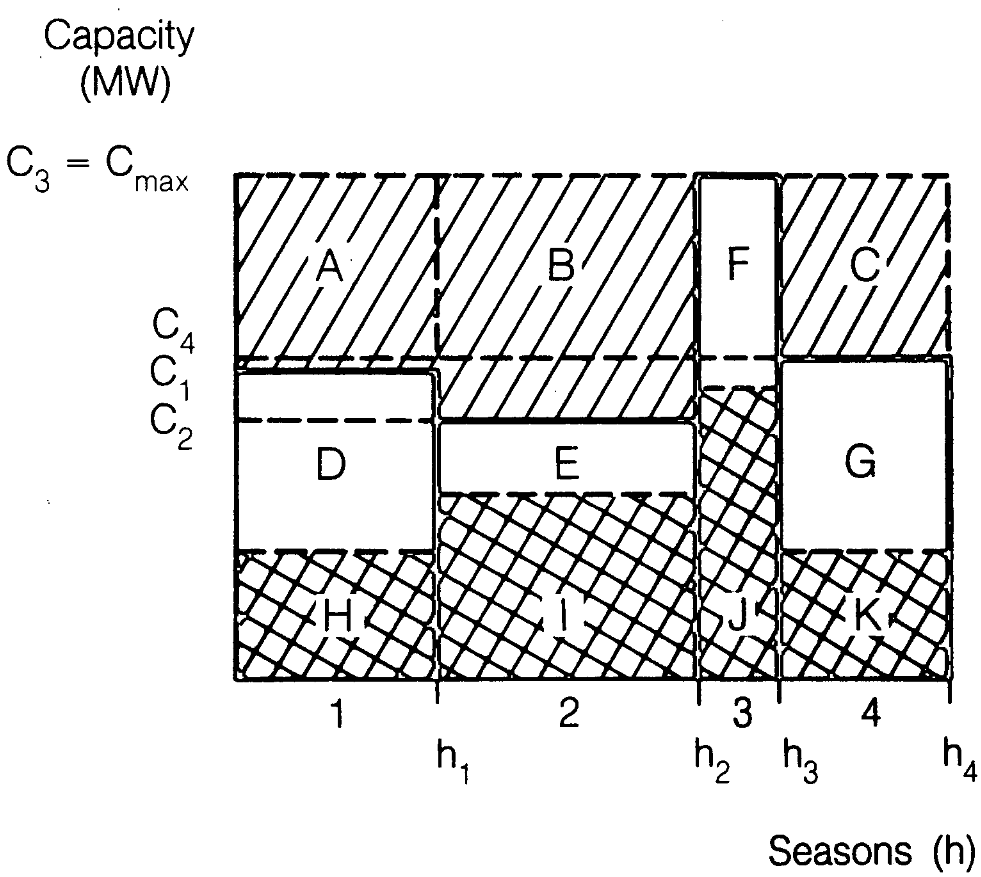

XBL 871-8804

Figure 4

Sample Hydro Data Calculations

YEAR: 1989 LMSTM HYDRO INPUT DATA SOURCE: PGandE 1985 LONG TERM PLAN GRASS RUN

UNIT NO. AND NAME		JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC
CAPACITIES (MW)													
1 QF (BNONPGE)	MAX CAP	265	302	344	334	322	313	279	236	240	213	244	254
	MIN CAP	105	143	195	183	165	163	150	104	96	60	89	104
NW HYDRO													
9 NWNCPAHY	MAX CAP	13	13	13	13	9	13	19	21	27	26	25	20
	MIN CAP	0	0	0	0	0	0	0	0	0	0	0	0
10 NWPGEHY1	MAX CAP	569	569	569	569	229	569	0	0	0	0	0	0
	MIN CAP	0	0	0	0	0	0	0	0	0	0	0	0
11 NWPGEHY2	MAX CAP	0	0	0	0	0	0	492	566	819	749	720	504
	MIN CAP	0	0	0	0	0	0	0	0	0	0	0	0
12 NWPGEHY3	MAX CAP	0	0	0	0	0	0	492	566	819	749	720	504
	MIN CAP	0	0	0	0	0	0	0	0	0	0	0	0
13 NWSMUDHY	MAX CAP	46	46	46	46	13	46	82	100	150	135	129	88
	MIN CAP	0	0	0	0	0	0	0	0	0	0	0	0
14 NWWAPAHY	MAX CAP	135	135	135	135	52	135	238	237	400	257	272	244
	MIN CAP	0	0	0	0	0	0	0	0	0	0	0	0
TOTAL NWHY	MAX CAP	763	763	763	763	303	763	1323	1490	2215	1916	1866	1360
	MIN CAP	0	0	0	0	0	0	0	0	0	0	0	0
TOTAL PGandE	MAX CAP	5124	5229	5204	5438	5511	5742	5552	5489	5530	5146	5149	5238
	MIN CAP	1380	1672	1909	2174	2535	2165	1820	1608	1514	1191	1294	1291
TOTAL AREA	MIN CAP	1485	1815	2104	2357	2700	2328	1970	1712	1610	1251	1383	1395
	MAX CAP	6152	6294	6311	6535	6136	6818	7154	7215	7985	7275	7259	6852
ENERGY (GWH)													
1 QF (BNONPGE)		124	157	214	195	190	178	161	124	113	62	91	114
9 NWNCPAHY		2	2	2	2	1	2	3	3	3	4	4	3
10 NWPGEHY1		85	76	85	82	34	82	0	0	0	0	0	0
11 NWPGEHY2		0	0	0	0	0	0	73	84	91	111	104	75
12 NWPGEHY3		0	0	0	0	0	0	73	84	91	111	104	75
13 NWSMUDHY		7	6	7	7	2	7	10	15	18	20	19	13
14 NWWAPAHY		20	18	20	20	8	20	35	34	161	39	40	36
TOTAL NWHY		114	102	114	111	45	111	194	220	364	285	271	202
TOTAL PGandE		1690	1776	1943	2267	2609	2389	2312	2147	1863	1661	1597	1625
TOTAL AREA		1928	2035	2271	2573	2844	2678	2667	2491	2340	2008	1959	1941

TOTAL YEAR NW HYDRO
1989 (GWH) (.%)
PGandE 1520 0.712

MUNI: 613 0.287

Figure 4 (cont.)

LMSTM INPUT HYDRO SUMMARY

YEAR: 1989							
SEASON:	WIN	SPR	SUM	FALL	TOTALS		
LMSTM SEASON ID:	1	4 ~~3~~	2	3 ~~4~~			
NO. OF MONTHS:	4	2	3	3			
NO. OF HOURS:	2880	1464	2208	2208	8760		
----------------------------		-------	-------	-------			
ROR-PGandE CAP. (MW)	1783.8	2350.0	1647.3	1258.7			
ENERGY (GWH)	5137.2	3440.4	3637.3	2779.1	14994.0		
POND-PGandE CAP. (MW)	3465.0	3276.5	3876.3	3919.0			
ENERGY (GWH)	2538.8	1557.6	2684.7	2103.9	8885.0		
ROR-QF CAPACITY (MW)	156.5	164.0	116.7	84.3			
ENERGY (GWH)	450.7	240.1	257.6	186.2	1134.6		
POND-QF CAPACITY (MW)	154.8	153.5	135.0	152.7			
ENERGY (GWH)	239.3	127.9	140.4	80.8	588.4		
ROR-NW CAPACITY (MW)	0.0	0.0	0.0	0.0			
ENERGY (GWH)	0.0	0.0	0.0	0.0	0.0		
POND-NW CAPACITY (MW)	763.0	533.0	1676.0	1714.0			
ENERGY (GWH)	441.0	156.0	778.0	758.0	2133.0		
"SCHEDULED MAINTENANCE" FOR THE HYDRO SYSTEM					"MAX" ENERGY (GWH)	TOTAL SCHED MAIN.	TOTAL SCHED MAIN. (.% OF MAX)
----------------------------		-------	-------	-------			
ROR-PGandE (GWH)	1630.8	0.0	1551.5	2409.7	20586	5592.0	0.2716
ROR-PGandE (.%)	0.2916	0.0000	0.2775	0.4309		1.0000	
POND-PGandE (GWH)	1307.5	940.6	94.2	0.0	34330	2342.3	0.0682
POND-PGandE (.%)	0.5582	0.4016	0.0402	0.0000		1.0000	
ROR-QF (GWH)	21.6	0.0	104.5	175.9	1437	302.0	0.2102
ROR-QF (.%)	0.0715	0.0000	0.3460	0.5824		1.0000	
POND-QF (GWH)	0.0	1.8	43.6	4.6	1356	50.0	0.0369
POND-QF (.%)	0.0000	0.0366	0.8715	0.0919		1.0000	
ROR-NW (GWH)	0.0	0.0	0.0	0.0	0	0.0	ERR
ROR-NW (.%)	ERR	ERR	ERR	ERR		ERR	
POND-NW (GWH)	2738.9	1729.0	83.9	0.0	15015	4551.8	0.3032
POND-NW (.%)	0.6017	0.3798	0.0184	0.0000		1.0000	
ENERGY LIMITS FOR THE FOR THE HYDRO SYSTEM							
----------------------------		-------	-------	-------			
POND-PGandE (MWH/MW)	2267.1	2267.1	2267.1	2267.1			
POND-PGandE (.%)	0.2857	0.1753	0.3022	0.2368	1		
POND-QF (MWH/MW)	3802.1	3802.1	3802.1	3802.1			
POND-QF (.%)	0.4067	0.2174	0.2386	0.1373	1		
POND-NW (MWH/MW)	1244.5	1244.5	1244.5	1244.5			
POND-NW (.%)	0.2068	0.0731	0.3647	0.3554	1		

Figure 5

Sample Northwest Thermal Data Calculations

	JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	TOTAL

NOTE: AVAILABLE CAPACITY IS SPLIT IN TWO IN MID YEAR TO ALLOW FOR DIFFERENT ENERGY COST FOR DIFFERENT TIMES OF THE YEAR

	JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	TOTAL
SEASON:	WIN				SPR		SUM			FALL			
LMSTM SEASON ID:	1				3		2			4			
NO. OF MONTHS:	4				2		3			3			
NO. OF HOURS:	2880				1464		2208			2208			8760
PGandE SYSTEM													
MONTHLY CAPACITIES (MW)													
1985	1120	1120	1120	1120	1410	1120							
1985							734	581	1	223	289	745	
1987	1070	1070	1070	1070	1410	1070							
1987							654	505	1	140	198	630	
1990	1294	1294	1294	1294	2098	1294							
1990							479	269	1	1	1	424	
1995	980	980	980	980	2012	980							
1995							191	5	1	1	1	21	
2000	952	952	952	952	1904	952							
2000							178	1	1	1	1	7	

LMSTM SEASONAL INPUTS

		WIN	SPR	SUM	FALL		MAX ENERGY (GWH)	TOTAL SCHED. MAINTANENCE (.% OF TOTAL)
SEASON AVG.:	1985	1120	• 1265	0	0			
SCHED. MAIN. (GWH)		418	0	2793	2793	6004	11081	0.5418
SCHED. MAIN. (.%)		0.0696	0.0000	0.4652	0.4652	1.00		
SEASON AVG.:	1985	0	0	• 439	419			
SCHED. MAIN. (GWH)		1263	642	0	43	1949	3843	0.5072
SCHED. MAIN. (.%)		0.6482	0.3295	0.0000	0.0223	1.00		
SEASON AVG.:	1987	1070	• 1240	0	0			
SCHED. MAIN. (GWH)		490	0	2738	2738	5965	10862	0.5492
SCHED. MAIN. (.%)		0.0821	0.0000	0.4590	0.4590	1.00		
SEASON AVG.:	1987	0	0	• 387	323			
SCHED. MAIN. (GWH)		1114	566	0	141	1821	3387	0.5376
SCHED. MAIN. (.%)		0.6115	0.3109	0.0000	0.0776	1.00		
SEASON AVG.:	1990	1294	• 1696	0	0			
SCHED. MAIN. (GWH)		1158	0	3745	3745	8647	14857	0.5820
SCHED. MAIN. (.%)		0.1339	0.0000	0.4331	0.4331	1.00		
SEASON AVG.:	1990	0	0	• 250	142			
SCHED. MAIN. (GWH)		719	366	0	238	1322	2187	0.6046
SCHED. MAIN. (.%)		0.5438	0.2764	0.0000	0.1798	1.00		
SEASON AVG.:	1995	980	• 1496	0	0			
SCHED. MAIN. (GWH)		1486	0	3303	3303	8092	13105	0.6175
SCHED. MAIN. (.%)		0.1836	0.0000	0.4082	0.4082	1.00		
SEASON AVG.:	1995	0	0	• 66	8			
SCHED. MAIN. (GWH)		189	96	0	128	413	575	0.7185
SCHED. MAIN. (.%)		0.4576	0.2326	0.0000	0.3098	1.00		
SEASON AVG.:	2000	952	• 1428	0	0			
SCHED. MAIN. (GWH)		1371	0	3153	3153	7677	12509	0.6137
SCHED. MAIN. (.%)		0.1786	0.0000	0.4107	0.4107	1.00		
SEASON AVG.:	2000	0	0	• 60	3			
SCHED. MAIN. (GWH)		173	88	0	126	386	526	0.7353
SCHED. MAIN. (.%)		0.4471	0.2273	0.0000	0.3256	1.00		

Figure 5 (cont.)

	JAN	FEB	MAR	APR	MAY	JUN	JUL	AUG	SEP	OCT	NOV	DEC	TOTAL
MUNI SYSTEM													
MONTHLY CAPACITIES (MW)													
1985	400	400	400	400	503	400							
1985							263	240	1	166	165	264	
1987	383	383	383	383	503	383							
1987							238	232	1	159	151	225	
1990	489	489	489	489	796	489							
1990							182	104	1	1	1	161	
1995	418	418	418	418	856	418							
1995							83	45	1	1	1	12	
2000	419	419	419	419	844	419							
2000							74	1	1	1	1	1	

LMSTM SEASONAL INPUTS

		WIN	SPR	SUM	FALL		MAX ENERGY (GWH)	TOTAL SCHED. MAINTANENCE (.% OF TOTAL)
SEASON AVG.:	1985	400	452	0	0			
SCHED. MAIN. (GWH)		148	0	997	997	2142	3955	0.5416
SCHED. MAIN. (.%)		0.0692	0.0000	0.4654	0.4654	1.00		
SEASON AVG.:	1985	0	0	168	198			
SCHED. MAIN. (GWH)		571	290	67	0	929	1737	0.5344
SCHED. MAIN. (.%)		0.6152	0.3127	0.0721	0.0000	1.00		
SEASON AVG.:	1987	383	443	0	0			
SCHED. MAIN. (GWH)		173	0	978	978	2129	3881	0.5486
SCHED. MAIN. (.%)		0.0812	0.0000	0.4594	0.4594	1.00		
SEASON AVG.:	1987	0	0	157	178			
SCHED. MAIN. (GWH)		514	261	47	0	822	1562	0.5260
SCHED. MAIN. (.%)		0.6250	0.3177	0.0573	0.0000	1.00		
SEASON AVG.:	1990	489	643	0	0			
SCHED. MAIN. (GWH)		442	0	1419	1419	3279	5628	0.5827
SCHED. MAIN. (.%)		0.1348	0.0000	0.4326	0.4326	1.00		
SEASON AVG.:	1990	0	0	96	54			
SCHED. MAIN. (GWH)		276	140	0	91	507	838	0.6048
SCHED. MAIN. (.%)		0.5436	0.2763	0.0000	0.1801	1.00		
SEASON AVG.:	1995	418	637	0	0			
SCHED. MAIN. (GWH)		631	0	1406	1406	3444	5580	0.6171
SCHED. MAIN. (.%)		0.1832	0.0000	0.4084	0.4084	1.00		
SEASON AVG.:	1995	0	0	43	5			
SCHED. MAIN. (GWH)		124	63	0	85	271	377	0.7206
SCHED. MAIN. (.%)		0.4562	0.2319	0.0000	0.3118	1.00		
SEASON AVG.:	2000	419	632	0	0			
SCHED. MAIN. (GWH)		612	0	1394	1394	3401	5532	0.6147
SCHED. MAIN. (.%)		0.1800	0.0000	0.4100	0.4100	1.00		
SEASON AVG.:	2000	0	0	25	1			
SCHED. MAIN. (GWH)		73	37	0	54	164	222	0.7380
SCHED. MAIN. (.%)		0.4455	0.2265	0.0000	0.3281	1.00		

LAWRENCE BERKELEY LABORATORY
UNIVERSITY OF CALIFORNIA
INFORMATION RESOURCES DEPARTMENT
BERKELEY, CALIFORNIA 94720