

ENERGY DIVISION

**CUSTOMER-SPECIFIC METRICS FOR THE REGULATION
AND LOAD-FOLLOWING ANCILLARY SERVICES**

BRENDAN KIRBY and ERIC HIRST

January 2000

Sponsored by
Consortium for Electric Reliability Technology Solutions
Office of Energy Efficiency and Renewable Energy
U.S. Department of Energy
Washington, DC 20585

OAK RIDGE NATIONAL LABORATORY
Oak Ridge, Tennessee 37831
managed by
LOCKHEED MARTIN ENERGY RESEARCH CORPORATION
for the
U.S. DEPARTMENT OF ENERGY
under contract No. DE-AC05-96OR22464

CONTENTS

	Page
SUMMARY	v
LIST OF ACRONYMS	vii
1. INTRODUCTION	1
2. DATA	5
DATA QUALITY	5
TEMPORAL AGGREGATION FOR REGULATION	6
TIME-AVERAGING PERIOD	10
3. REGULATION	13
SYSTEM-LEVEL METRICS	13
CUSTOMER-SPECIFIC METRICS	14
RESULTS	21
4. LOAD FOLLOWING	23
SYSTEM-LEVEL METRICS	23
CUSTOMER-SPECIFIC METRICS	25
RESULTS	27
5. CONCLUSIONS	31
ACKNOWLEDGMENTS	33
REFERENCES	35
REGULATION VECTOR-ALLOCATION METHOD	37

SUMMARY

In competitive electricity markets, the costs for each ancillary service should be charged to those who cause the costs to be incurred with charges based on the factors that contribute to these costs. For example, the amount of generating capacity assigned to the regulation service is a function of the short-term volatility of system load. Therefore, the charges for regulation should be related to the volatility of each load, not to its average demand.

This report discusses the economic efficiency and equity benefits of assessing charges on the basis of customer-specific costs (rather than the traditional billing determinants, MWh or MW), focusing on two key real-power ancillary services, regulation and load following. We determine the extent to which individual customers and groups of customers contribute to the system's generation requirements for these two services. In particular, we analyze load data to determine whether some customers account for shares of these two services that differ substantially from their shares of total electricity consumption.

We defined and applied metrics for regulation and load following. For regulation, we chose the standard deviation (MW) of the thirty 2-minute values in each hour. For load following (MW), we chose the difference between the maximum and minimum values of the 30-minute rolling-average load during each hour.

We also developed and applied methods to allocate these system-level metrics to individual customers and to groups of customers. The regulation allocation method uses a trigonometric relationship to correlate an individual customer's regulation burden with the total burden. The load-following allocation method calculates each customer's share of the total requirement on the basis of its coincident load-following requirement.

Application of these allocation methods shows that charging customers for these ancillary services on the basis of average loads can be inequitable. For one control area, a few large industrial customers account for 34% of system load, compared with 93% of the regulation and 58% of the load-following requirements (Fig. S-1). These customers disproportionately use these services but, in general, are not paying their fair share under typical utility tariffs. The subsidies inherent in today's ancillary-service pricing methods cannot, and should not, be sustained. Indeed, industrial customers with near-time-invariant loads, such as aluminum smelters and paper mills, will justifiably claim they require none of these services and, therefore, should not have to pay for them.





Fig. S-1. The shares of system regulation (top) and load following (bottom) calculated for several large industrial customers for 12 days in February 1999. Note that these customers sometimes account for more than 100% of the total, implying that the nonindustrial customers were reducing the requirements for these services.

LIST OF ACRONYMS

ACE	Area control error
AGC	Automatic generation control
CPS	Control performance standard
FERC	U.S. Federal Energy Regulatory Commission
ISO	Independent system operator
NERC	North American Electric Reliability Council
RTO	Regional transmission organization

INTRODUCTION

In its recent Notice of Proposed Rulemaking, the Federal Energy Regulatory Commission (FERC 1999) wrote, “The Commission believes that, whenever it is economically feasible, it is important for the RTO [regional transmission organization] to provide accurate price signals that reflect the costs of supplying ancillary services to particular customers.” Earlier, FERC (1996) wrote in its Order 888, “Because customers that take similar amounts of transmission service may require different amounts of some ancillary services, bundling these services with basic transmission service would result in some customers having to take and pay for more or less of an ancillary service than they use. For these reasons, the Commission concludes that the six required ancillary services should not be bundled with transmission service.”

Recognizing the economic efficiency and equity benefits of assessing charges on the basis of customer-specific costs, we investigated the requirements for two key ancillary services, regulation and load following. We determined the extent to which individual customers and subgroups of customers contribute to the system’s generation requirements for these two services, in particular whether some customers account for shares of these two services that differ substantially from their shares of total electricity consumption.

The remainder of this section defines these two ancillary services and explains how they differ from each other. Chapter 2 describes the data we obtained, the quality-control checks we applied to the data, and our data conversions to separate system-level regulation from load-following requirements. Chapters 3 and 4 present (1) the methods we developed to create system and customer-specific metrics for regulation (Chapter 3) and load following (Chapter 4) and (2) the results obtained with these methods. Chapter 5 presents our conclusions.

Because electricity is a real-time product, control-area operators must adjust generation to meet load on a minute-to-minute basis. As the electricity industry becomes deintegrated, with competitive generation separated from regulated transmission and system control, defining the requirements and responsibilities to meet time-varying customer loads is increasingly important. Regulation and load following are the two key ancillary services required to perform this function. (Energy imbalance is, depending on one’s definition, either the service or the accounting function that corrects for hourly errors in the provision of energy and the other two services.)

Loads can be decomposed into three elements (Fig. 1). The first element is the average load (base) during the scheduling period, 85 MW over the one hour shown in this case. The second element is the trend (ramp) during the hour and from hour to hour (the morning pickup

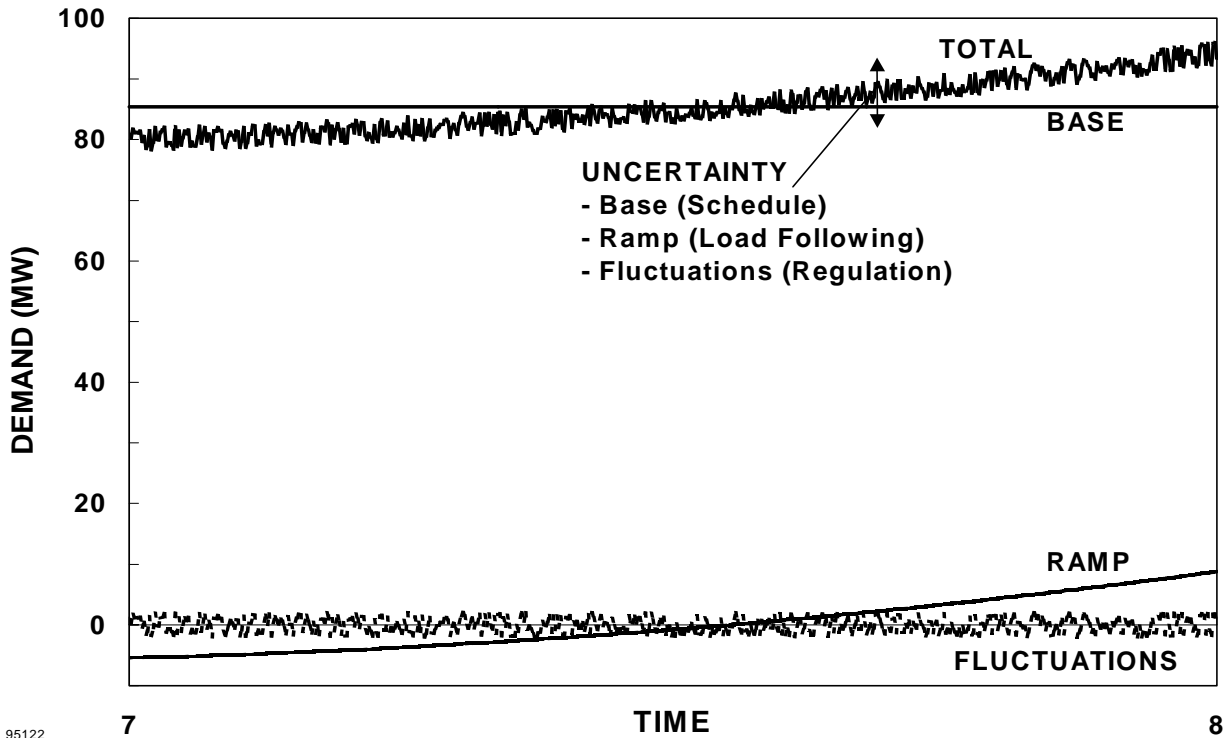


Fig. 1. Components of a hypothetical load on a weekday morning.

in this case); here that element increases from -5 MW at 7 a.m. to $+9$ MW at 8 a.m. The third element is the rapid fluctuations in load around the underlying trend; here the fluctuations range over ± 2 MW. Combined, the three elements yield a load that ranges from 78 to 96 MW during this hour.

The system responses to the second and third components are called load following and regulation. These two services (plus, perhaps, energy imbalance) ensure that, under normal operating conditions, a control area is able to balance generation to load. The two services are briefly defined below [see also FERC (1996), Hirst and Kirby (1998), and Interconnected Operations Services Working Group (1997)]:

- Regulation is the use of online generating units that are equipped with automatic generation control (AGC) and that can change output quickly (MW/minute) to track the moment-to-moment fluctuations in customer loads and to correct for the unintended fluctuations in generation. In so doing, regulation helps to maintain interconnection frequency, manage differences between actual and scheduled power flows between control areas, and match generation to load within the control area. This service can be provided by any appropriately equipped generator that is connected to the grid and electrically close enough to the local control area that physical and economic transmission limitations do not prevent the importation of this power.

- Load following is the use of online generation equipment to track the intra- and inter-hour changes in customer loads. Load following differs from regulation in three important respects. First, it occurs over longer time intervals than does regulation, 10 minutes or more rather than minute to minute. Second, the load-following patterns of individual customers are highly correlated with each other, whereas the regulation patterns are largely uncorrelated. Third, load-following changes are often predictable (e.g., because of the weather dependence of many loads) and have similar day-to-day patterns. Even when not predictable by the control-area operator, the customer can inform the control center of impending changes in its electricity use.

FERC (1996), in its Order 888, which defined six ancillary services, did not discuss load following. Perhaps because of this omission, most utilities and independent system operators (ISOs) do not include load following in their tariffs.* The absence of this service requires the California ISO to acquire much more regulation (as well as other services, such as replacement reserves and supplemental energy) than it otherwise would (Wolak, Nordhaus, and Shapiro 1998). Specifically, the California ISO buys regulation in amounts equivalent to about 5% of daily load, compared with about 1% for most utilities. Thus, the ISO is substituting an expensive service (regulation) for an inexpensive one (load following). Perhaps because of these problems, FERC (1999), in its notice on RTOs, proposed to require RTOs to operate real-time balancing markets. The primary resource for these markets is generators that can change output every five or ten minutes (i.e., to follow load).

*The Mountain West Independent Scheduling Administrator in Nevada is the only U.S. entity we know of that has proposed to create an explicit load-following service (FERC Docket No. ER99-3719-000, July 23, 1999). Most utilities provide load following through the 5- or 10-minute economic dispatch of their generating units.

DATA

We obtained 30-second data from a control-area operator on generation and load for a 12-day period in February 1999, a total of 34,560 observations.* For each 30-second interval, the data included total generation, net imports, total load, and the loads for several individual industrial customers. These large industrial customers include, among others, steel mills, oil refineries, and air-separation facilities. For confidentiality reasons, we scaled all the data shown here.

We summed the industrial loads to create a subgroup that we called *industrial load*. We called the difference between total load and industrial load *nonindustrial load*. Figure 2 shows the hourly loads for five days (Wednesday through Sunday). The total and nonindustrial loads show the expected winter patterns with morning and evening peaks, and with lower loads (by about 10%) on the weekends. The industrial load, on the other hand, is relatively constant from hour to hour. Its coefficient of variation (ratio of standard deviation to mean) is about half that of the nonindustrial load.

DATA QUALITY

This data set includes almost 3 million elements. Given its large size, it is not surprising that we found several anomalies in the raw data. Overall, one record is missing every three hours. The analysis software we developed for this project does not require all the data for a particular period to be present. We used linear interpolation to impute the missing data.

Another type of data anomaly was harder to recognize. We found 12 cases where one or more industrial loads rise for a single reading without a corresponding rise in the total load. For these load changes to be real, a simultaneous drop would have to occur in nonindustrial load for that single time step. We deleted five data points as being clearly impossible. We dropped three points because the control-area operator data-collection system identified problems at those times. We deleted one record because multiple industrial loads spiked simultaneously with no corresponding spike in the total load.

Failure to eliminate bad data results in errors in the analysis. Spikes in the data show up as abnormally high regulation requirements (high hourly standard deviation of the short-term

*We obtained and analyzed comparable data for a 12-day period in August and September 1999. We do not report these late-summer results because they are so close to those reported here for February.




Fig. 2. Hourly system load, nonindustrial load, and industrial load (the sum of several large industrial loads) for five days.

fluctuations) for one or more of the loads. To facilitate analysis, the software performs several data checks. Total generation is compared with the sum of the individual generator outputs and the difference is reported if the imbalance exceeds 10 MW. Reported and calculated area control error (ACE) is reported if the difference exceeds 10 MW. Spikes in industrial load not reflected in total load are reported if they exceed 100 MW. These anomalies are recorded by the software so that the specific data points can be examined for possible elimination. Dropping too many data points runs the risk of eliminating interesting and important events. We identified questionable data in fewer than 0.1% of the 34,560 records.

TEMPORAL AGGREGATION FOR REGULATION

Assessing the individual customer contribution to the overall regulation requirement necessarily involves generation performance. A control area is not expected to perfectly match generation and load instantaneously. Rather, generation matches load with some time lag, and, therefore, generation matches load only approximately. The North American Electric Reliability Council (NERC 1999) Control Performance Standards (CPS) 1 and 2 set statistical limits on the allowable differences between one-minute averages of the control area's difference between aggregated generation and interchange schedules relative to load (i.e., ACE). The cost of regulation is a function of the opportunity and operating costs of the generation capacity used to provide regulation. The control-area operator studied here is

satisfied with its CPS performance. It meets NERC requirements without wasting money by overcompliance.

Although the AGC systems at most utility control centers send raise and lower pulses to individual generators as frequently as every two or four seconds, generators do not follow such short-term load fluctuations. Our prior work (Hirst and Kirby 1996) suggests that generation follows load at the one- to two-minute interval.

We began this portion of the analysis with the 30-second data on generation and load regulation for the full 12 days. Regulation is here quantified as the difference between total load (generation) and the 30-minute rolling average of load (generation). We aggregated these 30-second data to create three additional datasets with 60-, 120-, and 240-second averages of the original data.

Figures 3 and 4 show, for one hour, the relationship between the regulation components of load and generation. The first graph shows that load fluctuates rapidly and substantially around its rolling average, whereas generation moves much more slowly and with a much smaller amplitude. In going from 30- to 60- to 120- to 240-second averages, the generation patterns are largely unchanged, but the load patterns become much smoother, smaller in amplitude, and slower moving. The ratios of the standard deviation of generation to the standard deviation of load increase from 0.45 for 30-second averages to 0.5, 0.6, and 0.7 with 60-, 120-, and 240-second averages. Note that even with 4-minute averages, load volatility remains 40% higher than generation volatility. A close look at the graphs showing 2- and 4-minute averages suggests that generation lags load by at least two minutes (Fig. 4).

These visual observations are confirmed by statistical analysis. We calculated correlation coefficients between generation and load for each of the four datasets for each of the 12 days. We then repeated these analyses by lagging generation 1, 2, ..., 12 time periods from load (Fig. 5). The correlation between generation and load increases in going from 30-second averages to 4-minute averages. The correlation between lagged generation and load is highest for a lag of 3 to 4 minutes.

Finally, we ran regression models of generation as a function of current and prior-period loads with 30-, 60-, and 120-second data. The models all have high explanatory power (the R^2 values are all about 0.7). The models all show that current load is a poor predictor of current generation. (Remember, we are discussing here only the fluctuations in generation and load about their 30-minute rolling averages, not the totals.) Load has a statistically significant effect on generation between 1 and 9 minutes.

We found these results surprising. We had expected to find that 30 seconds was too short a time-averaging period and that the appropriate period would fall between 1 and 2 minutes because CPS, the performance metric for the control area, is based on 1-minute averages. These results suggest that 4 minutes may be more appropriate. Although the data and




Fig. 3. Relationship between the regulation components of utility generation (solid line) and load (dashed line) with 30-second averages (top) and 60-second averages (bottom).


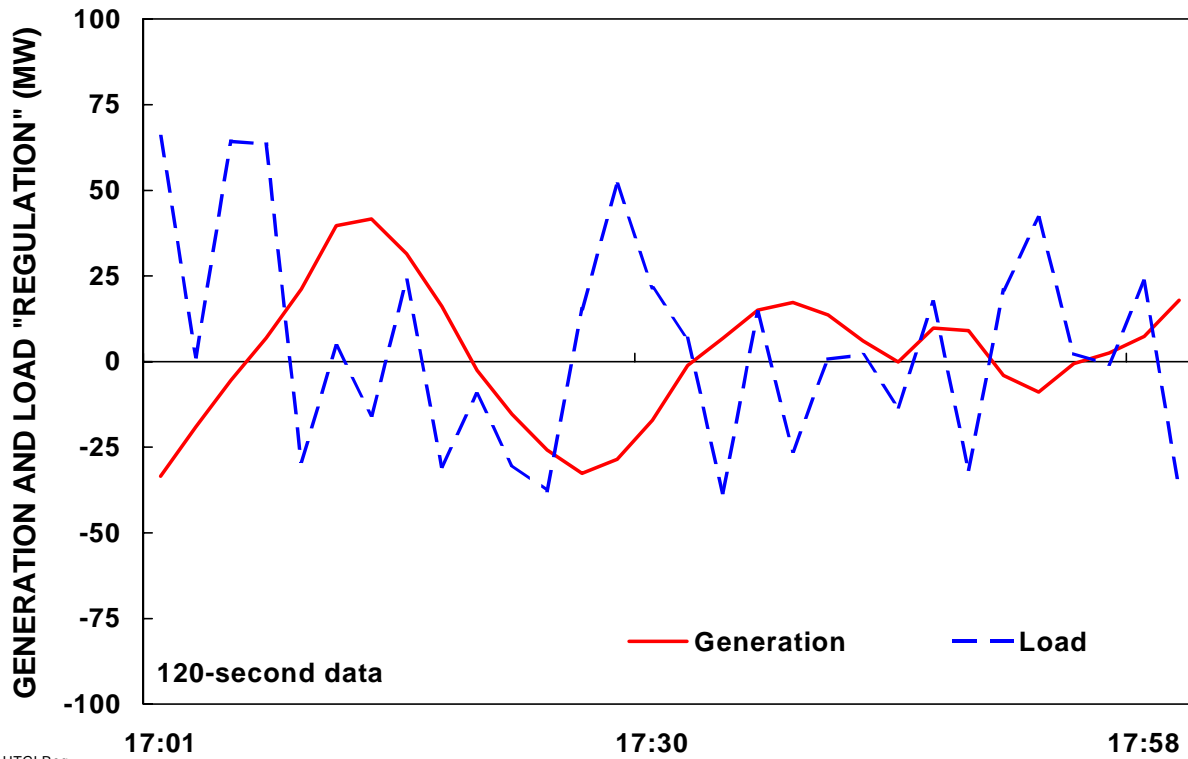
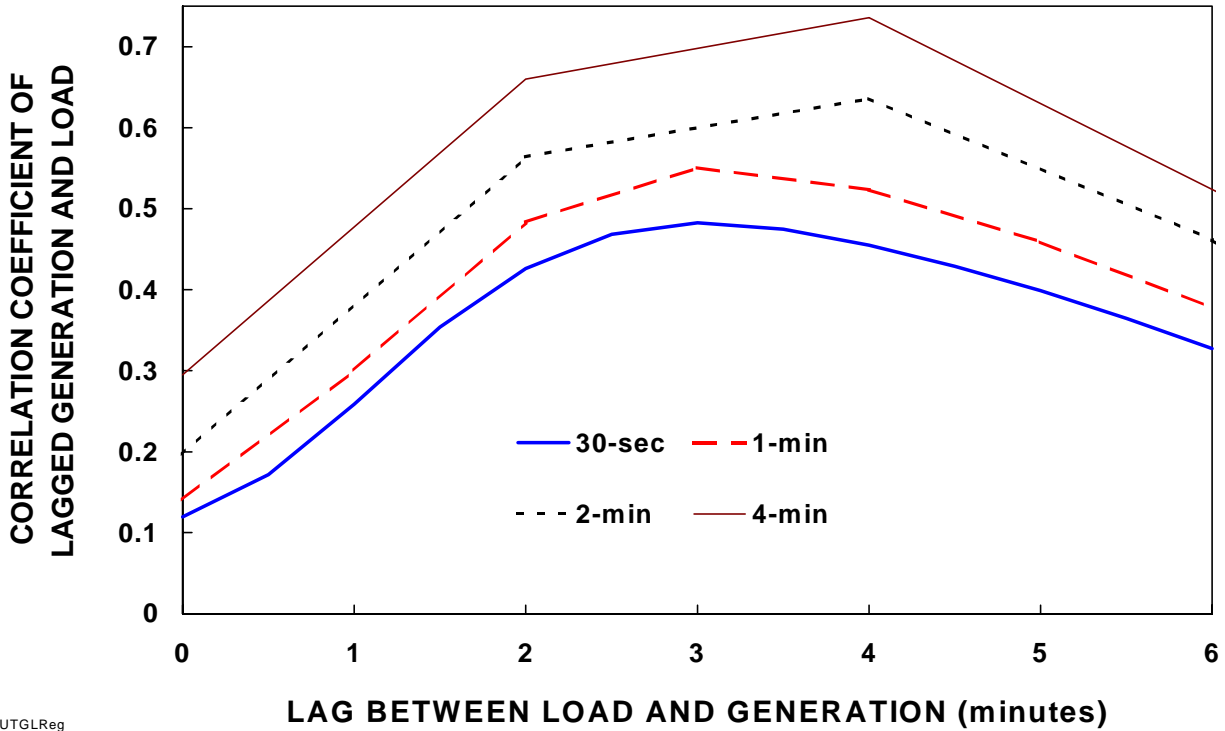


Fig. 4. Relationship between the regulation components of utility generation (solid line) and load (dashed line) with 120-second averages (top) and 240-second averages (bottom).



UTGLReg

Fig. 5. Correlation coefficient between generation and load regulation as a function of the time-averaging period (30 seconds to four minutes) and the lag between load and generation.

analysis suggest a 4-minute time-averaging period, we chose a 2-minute period because this is a conservative approach (i.e., it preserves more information) and it is consistent with the results of our analysis of regulation data for two other utilities.

TIME-AVERAGING PERIOD

There is no hard-and-fast rule to define the temporal boundary between regulation and load following. If the time chosen for the split is too short (e.g., five minutes), too much of the fluctuations will appear as load following and not enough as regulation. If the boundary is too long (e.g., 60 minutes), too much of the fluctuations will show up as regulation and not enough as load following. But in each case, the total is unchanged and is captured by one or the other of these two services.

We used rolling averages to define the boundary between the two services.* We tested 20-, 30-, 40-, 50-, and 60-minute rolling averages with 2-minute data. For the 30-minute rolling

*The use of a rolling average to separate regulation from load following is an analytical convenience, not possible in real time. System operators instead use sophisticated analytical methods to forecast loads for the next few hours, based on current and expected weather conditions, prior loads, and other factors.

average, as an example, we calculated the rolling average for each 2-minute interval as the mean value of the seven earlier values of the variable, the current value, and the subsequent seven values:

$$\text{Load following}_t = \text{Load}_{\text{estimated-t}} = \text{Mean} (L_{t-7} + L_{t-6} + \dots + L_t + L_{t+1} + \dots + L_{t+7}),$$

$$\text{Regulation}_t = \text{Load}_t - \text{Load}_{\text{estimated-t}}.$$

Selecting the appropriate rolling-average period to analyze load following depends on the factors that affect generation costs, which are specific to the generators providing load following. Several questions must be answered to determine the appropriate rolling-average period. Does the predictability of load following reduce costs by enabling economic dispatch for the units that provide load following? Are costs reduced for all slower movements or only for predictable movements? What is the lead time required to enable load following?

We analyzed the impact of changing the rolling-average period. The top part of Fig. 6 shows a shift in load-following responsibility (but no similar shift in the regulation responsibility) between the industrial and nonindustrial customers as the rolling-average period changes from 20 to 60 minutes. As discussed in Chapter 3, we use the standard deviation of regulation as the key metric for this service. For purposes of this comparison with load following, we multiply the regulation standard deviation by 3, which captures 99% of its variation.

We also compared the regulation and load-following magnitudes. The bottom part of Fig. 6 shows that the sum of total load-following and regulation requirements remain relatively constant, varying from 161 MW at 20 minutes to 156 MW at 60 minutes. Nonindustrial regulation requirements remain constant, as well, at 7 MW. Nonindustrial load following rises as the time-averaging period is extended, from 23 MW at 20 minutes to 31 MW at 60 minutes. Lengthening the rolling-average period from 20 to 60 minutes shifts 22 MW of the industrial variability from load following to regulation.

Changing the rolling-average duration has two major impacts. Shortening the period shifts the industrial variability from regulation to load following and shifts load following from the nonindustrial to the industrial customers. Selecting the correct rolling-average duration should be based on a break point in the generation cost drivers to provide each service, probably based on the costs of economic dispatch and the time required to implement it. For this analysis we used a 30-minute rolling average.

Because this control-area operator is satisfied with its CPS performance, we do not deal with the *amount* of each service that should be provided. This project deals only with the *allocation* among customers of the existing services and their costs.





Fig. 6. The effects of changes in the rolling-average period on the allocation of load variation between regulation and load following by type of load.

REGULATION

SYSTEM-LEVEL METRICS

Selection of an appropriate metric for regulation should be based on the underlying cost drivers to provide the service (i.e., the units that are on AGC). For example, if the costs are dominated by ramp-rate (MW/minute) considerations, then the speed of regulation units might be a better metric than the amount of regulation capacity. In addition, the chosen metric should be amenable to mathematical manipulation that permits definition of the relationships between individual loads and the total load. In this project, we analyzed load data only; we did not consider the types and costs of generating units used to provide the service. (We did, as explained in Chapter 2, examine the dynamics of generation vs loads.)

We examined four possible metrics for regulation, using the 2-minute regulation values defined above:

- Standard deviation (MW) of the 30 values in each hour;
- Average of the absolute values (MW) of the 30 values in each hour;
- Average regulation rate (MW/minute), calculated as one-half the average of the absolute values of the 29 differences between adjacent regulation values;* and
- Maximum regulation rate (MW/minute), calculated as one-half the maximum value of the absolute values of the 29 differences between adjacent regulation values.

Table 1 summarizes the average, maximum, and minimum values of these four metrics for total system load. Overall, regulation (as measured by the standard deviation) is 1.3% of total load. The two magnitude metrics (standard deviation and average of the absolute values) are highly correlated with each other. These correlation coefficients are all above 0.95 for total load, nonindustrial load, and industrial load. These very high correlations suggest that we need use only one of these two magnitude factors; we chose the standard deviation.

*We multiply the average by one half to convert from the 2-minute differences to 1-minute differences. This metric can also be calculated as the total path length (the sum of the absolute values of all the MW movements) during an hour divided by 60.

Table 1. Regulation metrics for system load for 12 days in February 1999 (288 hourly observations)

	Standard deviation (MW)	Average of absolute values (MW)	Average rate (MW/minute)	Maximum rate (MW/minute)
Average	31	25	14	41
Maximum	50	46	21	84
Minimum	16	13	8	18

On the other hand, the correlations between either of these two magnitude variables and the average regulation speed variable are lower (0.7 for nonindustrial load and 0.6 for industrial load). These lower correlation coefficients suggest that we may need two variables to accurately capture regulation, a magnitude (MW) factor and a ramprate (MW/minute) factor. Without information on the costs of regulation ramprate, we did not pursue this second metric.

Finally, the correlation coefficients between load itself and regulation are very low, suggesting that load is a poor predictor of regulation requirements.

Figure 7 shows the hour-to-hour patterns in regulation magnitude for weekdays and weekends. Figure 8 shows the variations from day to day in average regulation burden. These graphs show that the industrial loads have much greater volatility than do the nonindustrial loads. Indeed, as a share of total load, the industrial loads require about six times as much regulation as do the nonindustrial loads.


CUSTOMER-SPECIFIC METRICS

Having established system-level metrics for regulation, we turn our attention to the development of metrics for customer-specific assignment of the total regulation amount. This customer allocation is especially important for utilities that have nonconforming loads (e.g., steel mills).


Because regulation is the short, minute-to-minute fluctuations in load, the regulation component of each customer's load is largely uncorrelated with those of other customers. If each customer's load fluctuations (σ_i) is completely independent of the remainder of the system, the total regulation requirement (σ_T) would equal

$$\sigma_T = \sqrt{\sum \sigma_i^2} , \quad (1)$$

where i refers to an individual customer and T is the system total.



UtRegLF



UtRegLF

Fig. 7. Average hourly regulation requirement for weekdays (top) and weekends (bottom) for total load, nonindustrial load, and industrial load.




Fig. 8. Average daily regulation requirement.

In this idealized case, the share of regulation assigned to each customer would equal

$$\text{Share}_i = (\sigma_i / \sigma_T)^2, \quad (2)$$

and there would be no need to analyze interactions among customer loads in calculating the total regulation burden.

If, on the other hand, the loads are completely correlated with each other [i.e., the correlation coefficient (r) between each pair of loads equals 1], the total regulation requirement is the simple sum of the individual requirements:

$$\sigma_T = \sum \sigma_i. \quad (3)$$

In this idealized case, the share of regulation assigned to each customer would equal

$$\text{Share}_i = \sigma_i / \sigma_T. \quad (4)$$

Table 2 shows the total regulation requirement and customer shares for two customers, one with a standard deviation of 3 MW and the second with a standard deviation of 4 MW. If the two loads are independent of each other ($r = 0$), the total regulation requirement is 5 MW.

If, however, the two loads are completely and positively correlated with each other ($r = 1$), the total regulation requirement is 7 MW. Finally, if the two loads are completely and negatively correlated with each other ($r = -1$), the total regulation requirement is 1 MW.

Table 2. Total regulation requirement (and % of total) for two loads as a function of the correlation between the two loads

Load	Standard deviation (MW)	Regulation amount (% share)		
		Uncorrelated	Positively correlated	Negatively correlated
1	3	1.8 (36%)	3 (43%)	-3 (-300%)
2	4	3.2 (64%)	4 (57%)	4 (400%)
Total	--	5	7	1

Figure 9 shows results from an analysis of data for 19 large industrial customers from another control area. As expected, the actual value of the total regulation requirement is slightly (9%) higher than the total calculated as if the loads were completely uncorrelated. Also as expected, the actual value is much (63%) less than that calculated as if the loads were completely correlated. In this case, the loads exhibit a slight positive correlation with each other.

Figure 10 shows geometrically the possible relationships between two loads (A and B) and the total regulation requirement. (Each element could be represented by its standard deviation.) In the top panel, the two loads are uncorrelated, and the total regulation requirement is the square root of the sum of the squares of the individual loads (Eq. 1). In the middle panel, the same two loads are negatively correlated with each other, yielding a total regulation requirement less than that shown in the top panel. The bottom panel illustrates a situation opposite to that shown in the middle panel; in this case, the two loads are positively correlated with each other, and the total regulation requirement is more than that shown in the top panel.

The question is how to allocate fairly the total regulation requirement between these two loads (and, by extension, among several loads). The allocation method should yield results that are independent of any subaggregations. In other words, the assignment of regulation to load L should not depend on whether L is billed for regulation independently of other loads or as part of a group of loads. In addition, the allocation method should reward (pay) loads that reduce the total regulation burden.* In the middle panel of Fig. 10, load A offsets some of the

*A third criterion for choosing an allocation method could be independence of the order in which loads are added to the system. This objective overlaps with the first one discussed above.

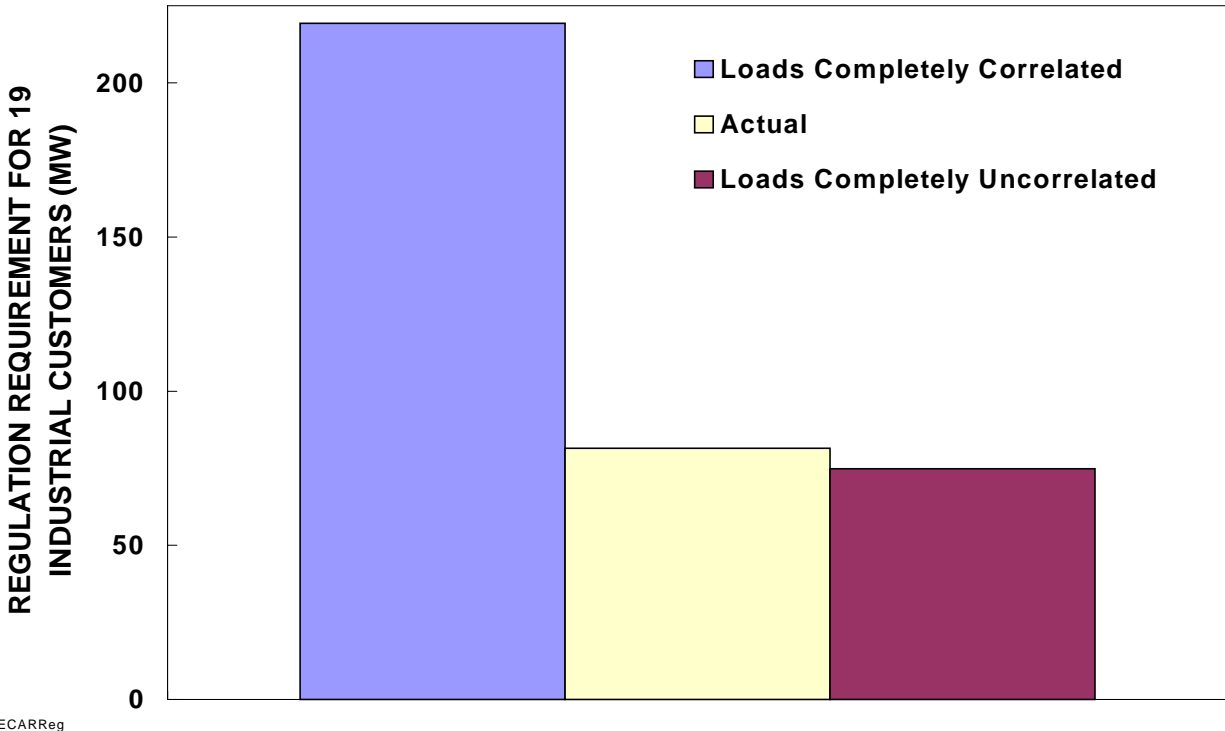


Fig. 9. Regulation requirement for 19 customers, showing the relationships among the actual value and those requirements that would occur if the loads were completely uncorrelated or were perfectly correlated.

regulation requirement of B , yielding a total requirement less than that of B alone; A should be paid for its contribution to reducing regulation.

Figure 11 illustrates schematically the method that we developed for such allocations; see the Appendix for details.* This method works for the two extreme situations discussed above, when loads are either completely uncorrelated or perfectly correlated. More important, this method yields reasonable results for the intermediate cases when loads are partially correlated with each other. Consider two loads A and B and the *Total*, with the regulation requirement of each based on the standard deviation of the short-term fluctuations. We propose a geometric approach to calculating the contribution of A to the *Total*, based on the projection of A onto the *Total* (shown as X in Fig. 11):

$$X = (Total^2 + A^2 - B^2)/(2 \times Total) . \tag{5}$$

*Alternative methods can probably be applied to this problem, but they may not scale appropriately, may not handle loads that are neither completely uncorrelated nor completely correlated with each other, and may be sensitive to subaggregations. The final section of the Appendix describes an alternative approach that we tested but found unsatisfactory.



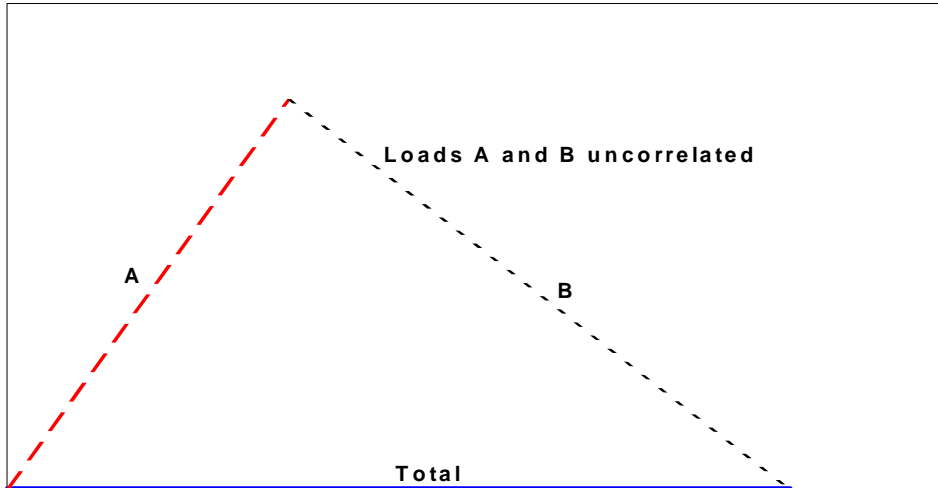


Fig. 10. The relationships between the standard deviations of two loads and the total regulation requirement, assuming the two loads are uncorrelated, negatively correlated, or positively correlated.




Fig. 11. Geometric allocation of individual loads *A* and *B* to regulation *Total*. *X* is *A*'s share of the total. *B*'s share, by subtraction, is *Total - X*.

The contribution of *B* to the *Total* is then equal to $Total - X$ or

$$Total - X = (Total^2 + B^2 - A^2) / (2 \times Total) . \quad (6)$$

This method can be extended to three or more loads through disaggregation of the total into various components. The only computational requirement is to calculate the standard deviation of each component and of each subtotal (total minus load *i*). Consider, as an example, a utility that wants to assign regulation charges separately for the residential class, the commercial class, five industrial customers, and the remainder of the industrial class, eight groups in all. The utility would calculate, for each hour, the standard deviation of eight subtotals (total - residential class, total - commercial class, and so on) as well as the standard deviations of each group of customers and the total, 17 values in all.

Table 3 illustrates numerically how the allocation process works with more than two loads. The table shows the standard deviations for four industrial loads and their total. The final column shows the method's allocation of the 26.3-MW regulation total to the four loads. If these four loads were completely uncorrelated, the total regulation requirement would be 30 MW, 14% higher than the actual. (The actual regulation requirement is lower than what would occur if the loads were uncorrelated because the loads are slightly negatively correlated with

each other.) If, on the other hand, the loads were completely and positively correlated, the total requirement would be 58 MW, 121% higher than the actual.

The method proposed here can accommodate a mix of individually metered loads and subaggregations, such as several large industrial loads that are metered separately and aggregations of thousands of residential and commercial customers. The subaggregations of the nonmetered residential and commercial loads will have the correct share of regulation assigned to them; any cost shifting will occur within the subaggregations and not between the subaggregations and the individually metered loads. This desirable property greatly reduces the need to meter any but the most nonconforming loads.

Table 3. Application of regulation-allocation method to four industrial loads

Load	Standard deviation (MW)	Standard deviation of total minus load (MW)	Regulation allocation (MW)
A	20.0	21.2	12.2
B	12.5	25.7	3.6
C	10.1	26.0	2.2
D	15.5	22.3	8.3
Total	26.3		26.3

RESULTS

During the 12 days studied, the hourly regulation standard deviation for the system as a whole ranged between 16 and 50 MW, with a mean of 31 MW (see Figs. 7 and 8). The nonindustrial and industrial standard deviations averaged 10 and 31 MW. (For comparison, the average nonindustrial load was 1280 MW, and the average industrial load was 670 MW, yielding a total system load of 1950 MW during these 12 days.)

The allocation method assigned the industrial customers 93% of the regulation total, almost triple their 34% share of system load. As shown in Fig. 12, there were several hours when the industrial customers were assigned more than 100% of the regulation requirement (reaching 132% one afternoon). During the hours that the industrial share exceeded 100%, the nonindustrial customers would have received a credit for regulation, offsetting their regulation costs during the other hours.

We applied the same method to allocate the industrial load among its components. Interestingly, two of the loads are negatively correlated with the others, yielding small negative regulation requirements.




Fig. 12. The share of hourly system regulation requirement assigned to several large industrial customers, which averaged 93% over these 12 days.

LOAD FOLLOWING

SYSTEM-LEVEL METRICS

We examined two possible metrics for load following:

- Load-following magnitude (MW) measured as the difference between the maximum and minimum values of 30-minute rolling-average load during each hour and
- Load-following rate (MW/minute) measured as the ratio of the first metric divided by the number of minutes between the highest and lowest load values.

Unlike regulation, load following is a signed quantity, positive if it is rising during the hour and negative if it is falling. Table 4 summarizes the average, maximum, and minimum values of these two load-following metrics for total system load. Figure 13 shows the hour-to-hour pattern in load following for weekdays and weekends. Unlike regulation, there is a clear diurnal pattern, reflecting the morning and early-evening peaks and the late-evening dropoff shown in Fig. 2. The nonindustrial loads track this diurnal pattern closely, while the industrial load is much more erratic in its load following. Figure 14 shows the variations in average load-following requirement from day to day. Overall, load following is 3.3% of total load.

Table 4. Absolute values of load-following metrics for system load for 12 days in February 1999 (288 hourly observations)

	Magnitude (MW)	Ramp rate (MW/minute)
Average	64	1.7
Maximum	181	6.3
Minimum	11	0.2

The load-following magnitude and ramprate metrics are highly correlated. The correlation coefficients between the magnitude and ramp rate are above 0.95 for total load, nonindustrial load, and industrial load (Fig. 15). These high correlation coefficients suggest that load-following requirements are adequately captured by only one factor. Physically, the high correlation coefficients mean that the load-following ramping requirements are greatest during the morning and early-evening pickups and late-evening dropoff.





Fig. 13. Average hourly load-following requirement for weekdays (top) and weekends (bottom).




Fig. 14. Average daily load-following requirement.

As with regulation, the correlation coefficients between load and load-following magnitude are very small, suggesting that load itself is a poor predictor of load-following requirements.

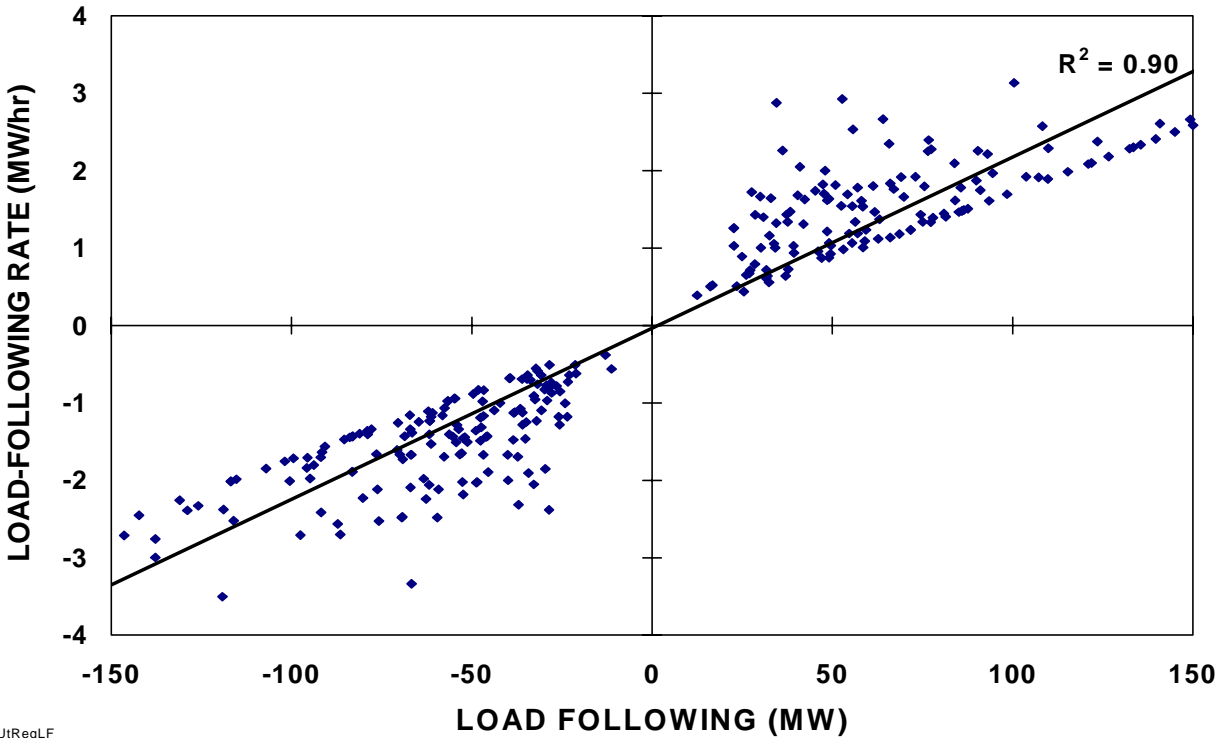
CUSTOMER-SPECIFIC METRICS

We calculate each customer's share of load following (or that of each group of customers) as the ratio of the customer's *coincident* load-following amount to the total load-following amount:

$$\text{Share}_i = (\text{Load}_{i, T_{max}} - \text{Load}_{i, T_{min}}) / (\text{Load}_{T_{max}} - \text{Load}_{T_{min}}),$$

where i refers to a customer or group of customers, T_{max} is the time within the hour that the system reaches its maximum load, and T_{min} is the time within the hour that the system reaches its minimum load. Note that T_{max} and T_{min} refer to the times of the maximum and minimum *system* loads, not those for the individual components.

Figure 16 illustrates how this method works. The graph shows the rolling averages for one hour for a hypothetical system, with residential, commercial, and industrial loads. The residential load in this example increases monotonically throughout the hour, from its minimum



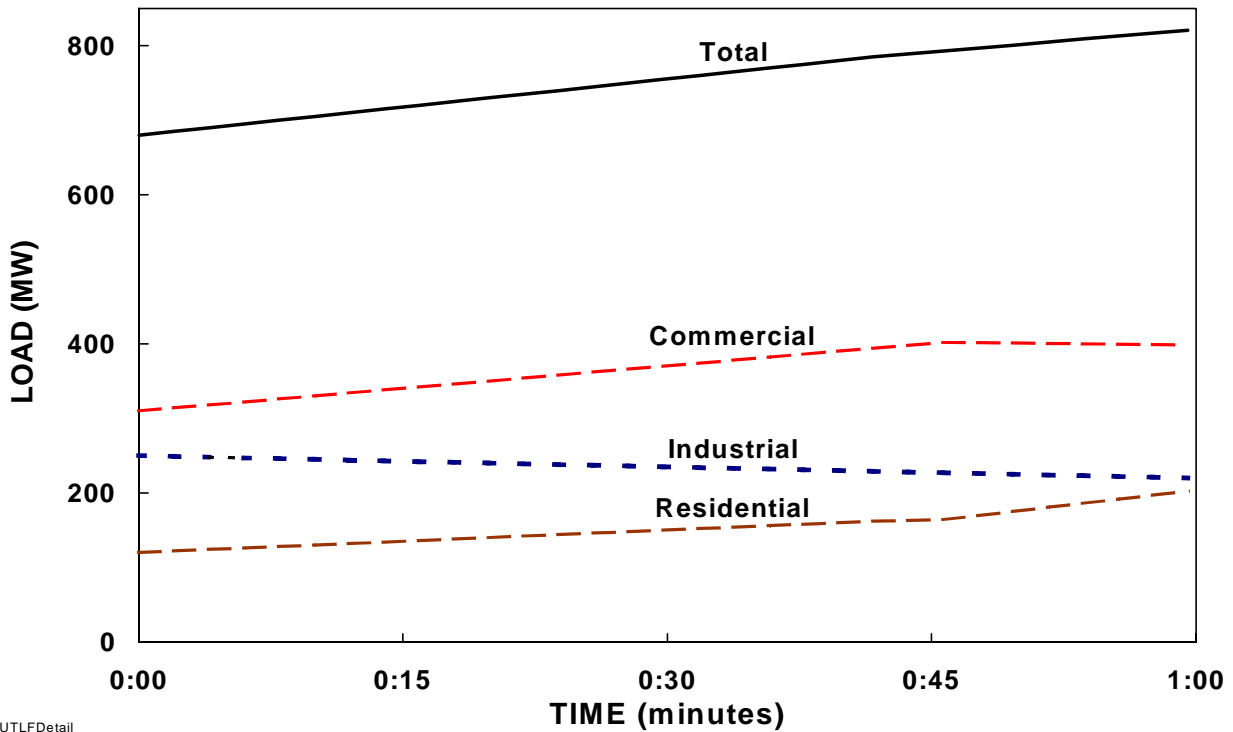
UIRegLF

Fig. 15. Load-following ramprate as a function of load-following magnitude for total load.

value of 120 MW at 0:00 to its maximum value of 203 MW at 1:00. The commercial load also increases during the hour, from 310 MW at 0:00 to 402 MW at 0:46, followed by a slight decline to 399 MW at 1:00. The industrial load, unlike the residential and commercial loads, decreases during the hour, from 250 MW at 0:00 to 220 MW at 1:00. The sum of these changes yields the system load, which increases from 680 MW at 0:00 to 821 MW at 1:00; this difference implies a total load-following requirement of 141 MW with a ramp rate of 2.35 MW/minute [= (821 - 680)/60].

A simple sum of the load-following requirements for each component alone yields a total of 205 MW (= 83 + 92 + 30), far higher than the 141 MW actual. This discrepancy is a function of the signs of the different components as well as the noncoincidence of these changes with the movement of the system as a whole. Allocating the total load-following burden on the basis of the noncoincident movements of each component would unfairly penalize loads that are noncoincident. Even worse, such an allocation scheme would charge customers for load following even if their load moved counter to the system load (as the industrial load does in this example).

In this example, allocating the total load-following requirement according to coincident loads assigns 59% of the total to the residential class and 63% to the commercial class and gives a 21% credit to the industrial class. Allocating the requirement on the basis of



UTLFDetail

Fig. 16. Hypothetical system load and its components for one hour. Although the system load-following requirement is 141 MW, the sum of the components is 205 MW, primarily because the industrial load decreases during the hour while the residential and commercial loads increase.

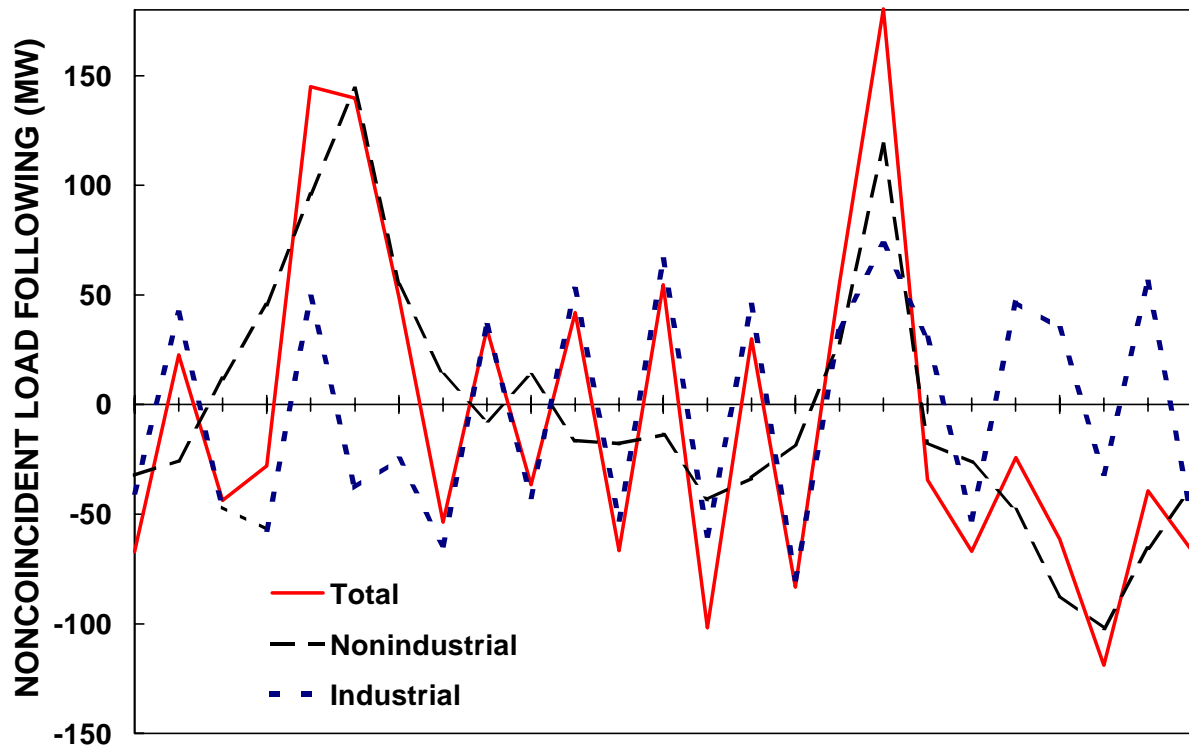
noncoincident loads would charge the industrial customers for 15% of the total even though their loads were moving counter to the system load during this hour.

RESULTS

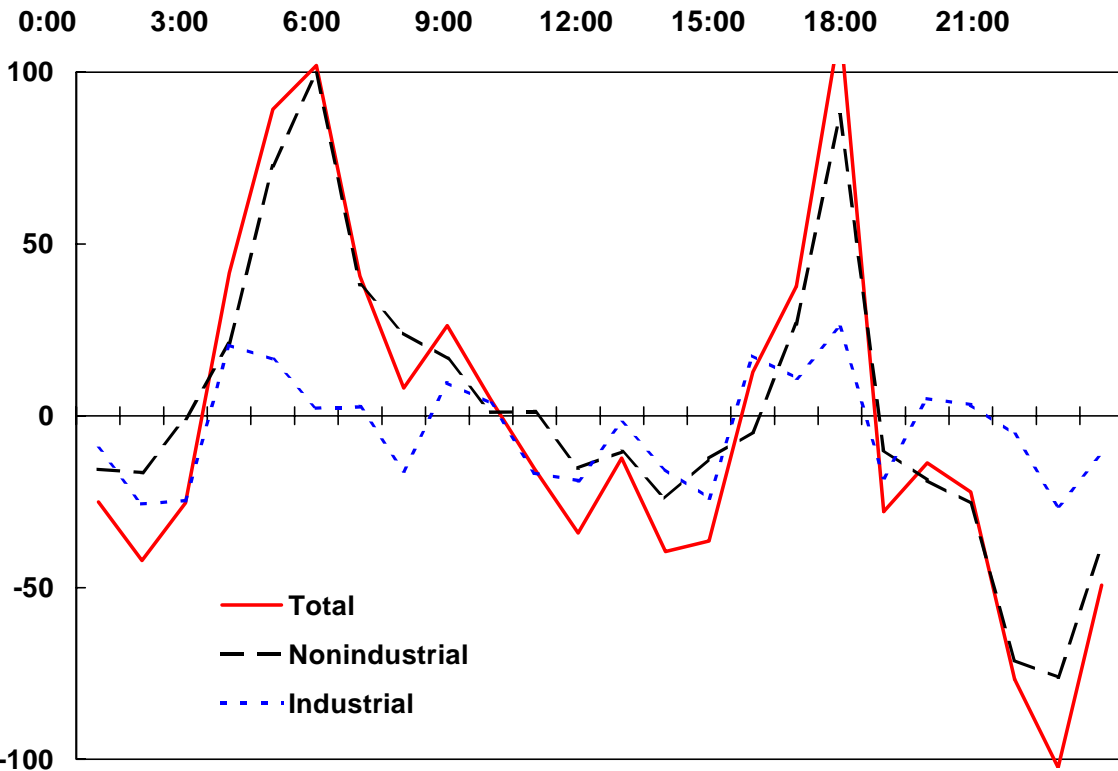
The top part of Fig. 17 shows load-following requirements for one of the 12 days studied in this project. The dramatic hour-to-hour variations are quite different from the pattern we had expected to see. The data show large positive load-following requirements at 4 and 5 a.m. and again at 5 p.m., with smaller peaks at 6, 8, 10 a.m., and noon. The data also show large negative values at 7 and 11 a.m. as well as at 1, 3, and 10 p.m.

We had anticipated an early-morning peak, an early-evening peak, and a late-evening dropoff (Figs. 2 and 13). The hourly averages across all the days show just such a pattern (bottom of Fig. 17). Averaged across all 12 days, load-following requirements peak at 4 and 5 a.m. and again at 5 p.m. Requirements then drop sharply at 9 and 10 p.m.

The averaged data show that the nonindustrial load contributes most to the total load-following requirement during the morning and early-evening ramp up and again during the late-evening ramp down. However, during the other hours of the day, the industrial load dominates. Whereas the nonindustrial load-following pattern is consistent from day to day, the industrial



UTFL



UTFL

Fig. 17. Load-following requirement for one weekday (top) and averaged over 12 days in February 1999 (bottom).

pattern is not. Indeed, the top part of Fig. 17 shows the industrial load-following requirement changing substantially from hour to hour, often swinging from positive to negative and back again. The nonindustrial load, on the other hand, shows the expected winter diurnal pattern.

Figure 18 shows the absolute value of system load following and the coincident contributions from the two components. It shows clearly the importance of the industrial load during the hours of mild load-following changes. In particular, during hours 0 through 4, 7 through 17, 19, 20, and 24, the industrial load accounts for more of the total load-following requirement than does the nonindustrial load. Unlike the nonindustrial load, the industrial load's load-following pattern is not predictable from day to day. (When averaged over several days, the industrial load-following requirement appears to be much smaller than it actually is.)

Because of the patterns shown in Figs. 17 and 18, the industrial loads account for much more of the load-following requirement than we had anticipated: 58%, far above their 34% share of total load (Fig. 19). Because the *cost* of load following is likely to vary from hour to hour and be more expensive during peak-demand periods, the industrial share of load-following costs is likely to be lower than 56%. Correspondingly, the nonindustrial shares of load following and energy are 42% and 66%. Given this substantial difference between shares of load and load following, customer-specific assignment of load following is probably warranted for these large industrial customers.

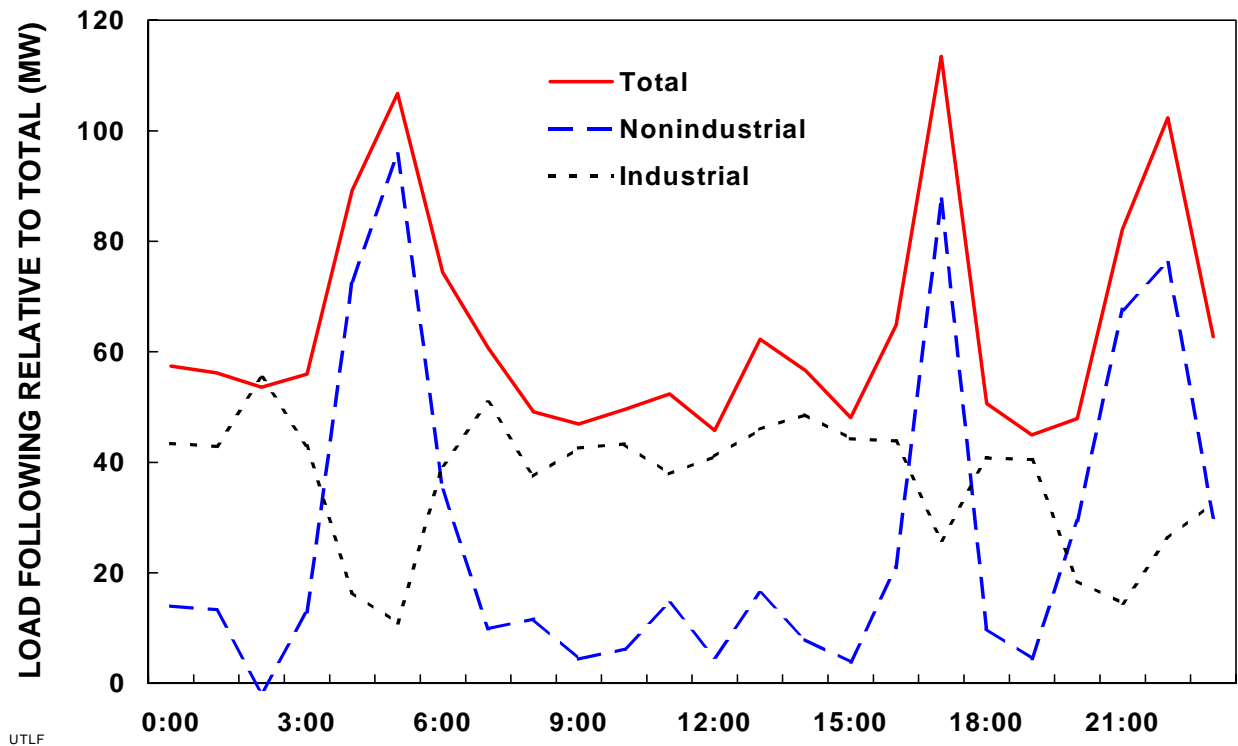


Fig. 18. System load-following magnitude by hour for 12 days in February 1999 and the contribution to these totals by nonindustrial and industrial customers.




Fig. 19. The share of hourly load-following requirement assigned to several large industrial customers for 12 days in February 1999, which averaged 58% over these 12 days.

The share of load following assigned to these industrial customers varies substantially from hour to hour, as shown in Fig. 19. The industrial share exceeds 100% for 26% of the hours and is less than zero for 10% of the hours.

We also examined the individual industrial loads and their relationship to the total industrial load. Here, too, the shares of load-following requirement vary considerably, both in absolute terms and relative to the energy shares. For example, one customer accounted for 22% of the industrial energy use but 40% of the industrial load-following requirement. On the other hand, another accounted for 33% of the energy share but none of the load-following requirement.

CONCLUSIONS

Utilities typically have several large industrial customers with unusual time signatures (called nonconforming loads). Because these loads may contribute disproportionately to a company's load-following and regulation requirements, we developed and applied methods to quantify system-level and customer-specific requirements for these two ancillary services.

Regulation and load following (along with, perhaps, energy imbalance) are the real-power ancillary services that a control area uses to maintain the necessary real-time balance between generation and load. The NERC Control Performance Standards 1 and 2, both measures of area control error, determine whether the amounts and use of the generation resources devoted to regulation and load following are adequate. Because this control-area operator's CPS performance meets NERC criteria, this project focused on the allocation of existing requirements rather than on a determination of the appropriate amounts of each service to provide.

Electricity consumption varies with time. If the consumption data are examined at the 4-second level, they will show considerable random movement. If, on the other hand, they are examined at the hourly level, they will show only smooth and consistent day-to-day patterns. The data themselves do not determine how best to disaggregate the total into the energy, load-following, and regulation components. This disaggregation requires judgment as well as analysis of the control area's mix of generation and loads. Because this control-area operator's generation follows load at only the 2- to 4-minute level, we aggregated the 30-second data to 2-minute averages. We used a 30-minute rolling average to characterize load following; regulation is then the difference between actual load and the rolling average. A longer rolling-average period would shift some load following from industrial to nonindustrial customers and would shift some industrial load following to industrial regulation.

We defined and applied two metrics for regulation and two for load following:

- Standard deviation (MW) of the thirty 2-minute values in each hour;
- Average regulation rate (MW/minute), calculated as one-half the average of the absolute values of the 29 differences between adjacent regulation values;
- Load-following magnitude (MW), equal to the difference between the maximum and minimum values of 30-minute rolling-average load during each hour; and

- Load-following rate (MW/minute), equal to the ratio of the load-following-magnitude metric divided by the number of minutes between these two (highest and lowest) load values.

The amount of generating capacity provided for regulation is a multiple of the regulation standard deviation to ensure sufficient probability of meeting these temporal variations in load. Multiplying the standard deviation by two provides 95% coverage, and multiplying by three provides 99% coverage.

The two load-following metrics (magnitude and ramprate) are very highly correlated; as a consequence, we think it is sufficient to consider only the magnitude metric for this service. The two regulation metrics are also positively correlated, but not to the extent that occurs for load following; therefore, it may be useful to consider both regulation metrics.

We also developed and applied methods to allocate these system-level metrics to individual customers and to groups of customers. The regulation allocation method uses a trigonometric relationship to determine the amount of each customer's regulation requirement that is correlated with the total requirement. The load-following allocation method calculates each customer's share of the total requirement on the basis of its coincident load-following requirement.

Table 5 summarizes results on system and customer-specific assignments. Although the industrial customers as a group account for 34% of the total energy, they account for 93% of the regulation requirement and 58% of the load-following requirement. Within this group of large customers, the allocations are also disproportionate. For example, one customer accounts for 8% of total energy but 38% and 22% of regulation and load following, respectively. On the other hand, another customer accounts for much smaller shares of regulation and load following than of energy.

Current U.S. utility practice (i.e., the tariffs filed with FERC as required by Order 888) typically charges customers for these ancillary services on the basis of average load (i.e., energy). Assume, for purposes of an example, that the hourly costs of regulation and load following are \$10/MW and \$5/MW, respectively, and that the amounts of generating capacity to provide these services average 93 MW (three times the 31 MW in Table 5) and 64 MW, respectively. Given these assumptions, the average cost to provide these two services would be \$1220/hour over this 12-day period. The traditional method of assigning customer-specific charges for these services would bill nonindustrial customers \$800/hour and industrial customers \$420/hour. Using the allocation methods developed here, \$600/hour would be shifted from the nonindustrial bill (a 75% cut) to the industrial bill (a 140% increase) for these two ancillary services. Similar shifts would occur among the individual industrial customers.

Table 5. Characteristics of total load and its components for 12 days in February 1999

	Energy		Regulation		Load following	
	(MW)	Share (%)	(MW)	Share (%)	(MW)	Share (%)
Total load	1954	—	31.2	—	63.9	—
Nonindustrial load	1284	65.7	2.2	7.2	27.0	42.3
Industrial load	670	34.3	29.0	92.8	36.9	57.7
1	264	13.5	16.4	52.6	20.9	32.7
2	33	1.7	2.8	9.0	4.2	6.6
3	77	3.9	6.4	20.5	5.9	9.2
4	10	0.5	0.5	1.7	0.4	0.6
5	10	0.5	2.9	9.2	5.4	8.5
6	275	14.1	-0.1	-0.2	0.0	0.0

The kinds of subsidies identified here among customers and customer classes likely exist for other electricity products, including installed-capacity requirements. For installed capacity, industrial customers (with high load factors) are probably subsidizing other customers with lower load factors.

The results presented here are consistent with anecdotal evidence from other control areas. The regulation requirements for one utility are 50% higher when a single metal-fabrication customer operates than when that customer is offline. Another utility has two steel mills that account for 3% of total load, but over 50% of regulation and load-following requirements. A steel mill in a third utility's service area accounts for 1% of load and 22% of regulation requirements; a paper mill in the same service area accounts for 5% of load but only 1% of regulation. Finally, these results for February 1999 are very similar to those obtained with analysis of late-summer 1999 data from the same control area.

ACKNOWLEDGMENTS

We thank Ziad Alawayan, Malcolm Bertsch, Joseph Eto, Joanne Horalek, Stephen Hoffman, John Hughes, Laura Manz, Gregory Pakela, Brian Scott, and Robert Taylor for their helpful comments on a draft of this report. We thank Fred O'Hara for editing this report.

REFERENCES

E. Hirst and B. Kirby 1996, *Ancillary-Service Details: Regulation, Load Following, and Generator Response*, ORNL/CON-433, Oak Ridge National Laboratory, Oak Ridge, TN, September.

E. Hirst and B. Kirby 1998, *Unbundling Generation and Transmission Services for Competitive Electricity Markets: Examining Ancillary Services*, NRRI 98-05, National Regulatory Research Institute, Columbus, OH, January.

Interconnected Operations Services Working Group 1997, *Defining Interconnected Operations Services Under Open Access*, EPRI TR-108097, Electric Power Research Institute, Palo Alto, CA, May.

North American Electric Reliability Council 1999, *NERC Operating Manual*, "Policy 1 — Generation Control and Performance," Princeton, NJ, November.

U.S. Federal Energy Regulatory Commission 1996, *Promoting Wholesale Competition Through Open Access Non-Discriminatory Transmission Services by Public Utilities; Recovery of Stranded Costs by Public Utilities and Transmitting Utilities, Final Rule*, Docket Nos. RM95-8-000 and RM94-7-001, Order No. 888, Washington, DC, April 24.

U.S. Federal Energy Regulatory Commission 1999, *Regional Transmission Organizations, Notice of Proposed Rulemaking*, Docket No. RM99-2-000, Washington, DC, May 13.

F. Wolak, R. Nordhaus, and C. Shapiro 1998, *Preliminary Report on the Operation of the Ancillary Services Markets of the California Independent System Operator*, Market Surveillance Committee, California Independent System Operator, Folsom, CA, August 19.

REGULATION VECTOR-ALLOCATION METHOD

We can think of regulation as a vector and not just a magnitude. For example, start with load A . It might be a single house or an entire control area with a regulation burden of 8. Consider another load B with a regulation burden of 6 that we want to combine with A . If loads A and B are perfectly correlated positively, they add linearly, as shown in the top of Fig. 20. If the two loads are perfectly correlated negatively, their regulation burdens would add as shown in the middle of Fig. 20. Typically, loads are completely uncorrelated and the regulation requirement for the total is the square root of the sum of the squares, or 10 in this case (bottom of Fig. 20).

Multiple uncorrelated loads are always at 90 degrees to every other load. They are also at 90 degrees to the sum of all the other loads. This characteristic requires adding another dimension each time another load is added, which is difficult to visualize beyond three loads. Fortunately, the math is not any more complex. The fact that each new uncorrelated load is at 90 degrees to every other load and to the total of all the other loads is quite useful. The analysis of any number of multiple loads can always be broken down into a two-element problem, the single load and the rest of the system.

Return to the two-load example but consider the more general case where loads A and B are neither perfectly correlated nor perfectly uncorrelated. We may know the magnitude of A and the magnitude of B , but we do not know the magnitude of the total without measuring it directly (i.e., we do not know the *direction* of each vector). We can, however, measure the total regulation requirement and use this vector method to *allocate* the total requirement among the individual contributors.

We know the total regulation requirement because we meter it directly as the aggregated regulation requirement of the control area. We can know the regulation requirement of any load by metering it also. We can know the regulation requirement of the entire system less the single load we are interested in by calculating the difference between the system load and the single load at every time step, separating regulation from load following, and taking the standard deviation of the difference signal. Knowing the magnitudes of the three regulation requirements, we can draw a vector diagram showing how they relate to each other (Fig. 21).

How much of the total regulation requirement is the responsibility of load A ? We can calculate the amount of A that is aligned with the total and the amount of B that is aligned with the total. We can do this geometrically (as shown below) or with a correlation analysis.




Fig. 20. The relationships among the regulation components (*A* and *B*) and the total if *A* and *B* are positively correlated (top), negatively correlated (middle), or uncorrelated (bottom).

Y is perpendicular to the total regulation *T* (uncorrelated). *X* is aligned with *T* (correlated). *A*'s contribution to *T* is *X*. Knowing *A*, *B*, and *T*, we can calculate *X*. (We could also calculate *Y*, but there is no need to do so.) We can write two equations relating the lengths of the various elements:

$$A^2 = X^2 + Y^2 \tag{A-1}$$

$$B^2 = (T - X)^2 + Y^2 \tag{A-2}$$

Subtract Eq. A-2 from Eq. A-1 to get

$$A^2 - B^2 = X^2 - (T - X)^2 + Y^2 - Y^2 ;$$




Fig. 21. The relationship among the regulation burdens of loads *A* and *B* and the total (*T*) when *A* and *B* are neither perfectly correlated nor perfectly uncorrelated.

$$A^2 - B^2 = X^2 - (T^2 - T \times X - T \times X + X^2) = -T^2 + 2T \times X .$$

Solving for *X* (load *A*'s contribution to the total *T*) yields

$$X = (A^2 - B^2 + T^2)/2T . \tag{A-3}$$

Recall that we can decompose a collection of any number of loads into a two-load problem consisting of the load we are interested in and the rest of the system without that load (Fig. 22). We can solve Eq. A-3 for as many individual loads as we wish. Variable *T* remains the total regulation requirement, variable *A* becomes each individual load's regulation requirement, and variable *B* becomes the regulation requirement of the total system *less* the specific load of interest.

This allocation method works well with any combination of individually metered loads and load profiling for the remaining loads. The load profiling can be as simple as making the usual assumption that the other loads' regulation requirements are proportional to their energy requirements. Or measurements of a sample set can be taken to determine the magnitude of their regulation burdens. This vector-allocation method is used to determine the regulation burden of each of the metered loads. The residual regulation burden is then allocated among the remaining loads, assuming they are perfectly uncorrelated.

ALTERNATIVE APPROACH

We initially tested an alternative approach. That approach calculated the *incremental* regulation requirement for each load by calculating the system regulation requirement without that load and then defining the allocation for the individual load as the difference between the total regulation requirement and the regulation requirement without that single load.




Fig. 22. Application of vector-allocation method to the case with more than two loads.

$$\sigma_{Incremental-j} = \sqrt{\sum \sigma^2} - \sqrt{\sum \sigma^2 - \sigma_j^2}$$

$$\sigma_{Incremental-j} / \sum \sigma_{Incremental} = \frac{\sqrt{\sum \sigma^2} - \sqrt{\sum \sigma^2 - \sigma_j^2}}{\sum_k \left[\sqrt{\sum \sigma^2} - \sqrt{\sum \sigma^2 - \sigma_k^2} \right]}$$

This method works well in many cases. However, when a load *reduces* the overall regulation requirement, this method blows up. Consider an example in which load *A* has a regulation requirement of 40, load *B* has a requirement of 22, and the total is 30. The incremental burden for *A* is +10, and for *B* it is -8. The total incremental burden is +2, which means that *A* is charged for 150 (= 30 × 10/2) and *B* receives a credit for 120 (= 30 × -8/2). These results are unappealing because they are so sensitive to small changes in the incremental contributions of any single load. In this case, load *A* would get charged for 150 MW, even though its total was only 40 MW, a completely unrealistic result. (Even if loads are 100% correlated, no load should get charged for more regulation than it would require if it was the only load in the control area.) The method developed in the body of the report, and which we recommend, assigns a regulation requirement of 33.6 to *A* and -3.6 to *B*, yielding results that are much more reasonable.