MEASURING GENERATOR PERFORMANCE IN PROVIDING REGULATION AND LOAD-FOLLOWING ANCILLARY SERVICES

Eric Hirst and Brendan Kirby
Consulting in Electric-Industry Restructuring
Oak Ridge, Tennessee 37830

January 2001

Prepared by
OAK RIDGE NATIONAL LABORATORY
P.O. Box 2008
Oak Ridge, Tennessee 37831-6285
managed by
UT-Battelle, LLC
for the
U.S. DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

CONTENTS

List of Figures	v
Acknowledgments	vii
Executive Summary	ix
Abbreviations and Acronyms	xi
1 INTRODUCTION	1
2 DATA	3
3 LOAD FOLLOWING	5
Metrics	5
Performance Relative to System Requirements	5
Performance Relative to Control-Center Requests	5
Results	6
4 REGULATION	11
Metrics	11
Performance Relative to System Requirements	11
Expected Values of Generation	12
Performance Relative to Expected	14
Results	17
5 CONCLUSIONS	25
REFERENCES	

LIST OF FIGURES

1	Hourly system load and interchange schedule for five days and 2-min regulation requirements for system load and interchange schedule for two hours for February 1999	4
2	Average hourly load-following requirement for winter weekdays	7
3	Average winter load-following contributions from generation and ACE	8
4	Supply of load-following services from generation and ACE as functions of total hourly load-following requirement	8
5	Winter load-following contributions from three generators	9
6	Winter performance of generation in providing load following relative to control-center requests for the service	10
7	Regulation can be split into orthogonal components, one aligned with AGC requests and the other independent of those requests.	16
8	Area control error for 12 hours on one winter day	17
9	Average winter hourly regulation requests (expected) relative to actual output and unit error or actual output correlated and uncorrelated with AGC requests	20
10	Average winter hourly regulation requests (expected), output (actual), and error (SCE) for unit A and unit C	21
11	Regulation performance at the 2-min level for units A and C for two hours on one winter day	22
12	Average winter weekday regulation performance for three generators: production correlated with expected and performance metrics	23

ACKNOWLEDGMENTS

The work described in this report was coordinated by the Consortium for Electric Reliability Technology Solutions and was funded by the Assistant Secretary for Energy Efficiency and Renewable Energy, Office of Power Technologies, Transmission Reliability Program of the U.S. Department of Energy

We thank Stephen Hoffman, Brian Scott, and Paul Spicer for their very helpful comments on the draft of this report. We thank Gregory Pakela and Paul Huseman for their assistance during the course of this project. We also thank Carolyn Moser for editing the report; and we thank Gwen Scudder for managing the clearance, publication, and distribution process.

EXECUTIVE SUMMARY

Independent system operators manage ancillary-service markets in four parts of the United States: California, the Mid-Atlantic, New York, and New England. Regional transmission organizations are likely to operate similar markets in other parts of the country within the next few years.

Because these ancillary services are expensive and are bought and sold in competitive markets, the electricity industry needs methods to measure the services delivered, generally in real time. The prices for the regulation service range from less than \$10/MW-hr in New England to almost \$50/MW-hr in the PJM Interconnection. Currently, the lack of performance metrics requires these system operators to pay for the amount of capacity reserved rather than for the service delivered.

This project developed and applied metrics for two real-power ancillary services, load following and regulation. The application is based on a small control area, using system and generator-specific data for two 12-day periods, one in February 1999 and the other in August–September 1999. These data include 30-second values for system load, net exports, area control error, Interconnection frequency, control-center requests to each generator, and the output from each generator.

To be consistent with emerging competitive bulk-power markets, we defined load following as an hourly (rather than a daily) service. The load-following requirement each hour is the signed difference between the highest and lowest smoothed values during the hour of retail load plus net exports. Using these two times ($T_{\max}$ and $T_{\min}$), we computed the contribution of each generator to the overall requirement. In a similar fashion, we defined the control-center request for load-following service as the signed difference between the highest and lowest smoothed values of requested output levels during the hour. Each unit's performance relative to this unit-specific expectation is based on its change in output between these two times.

Analyzing regulation performance is much more complicated than analyzing load following because of the frequent changes in direction. We developed and applied several metrics to measure the hour-to-hour regulation performance of individual generators relative to the control-center requests. These metrics use the standard deviation as the measure of volatility in control-center requests, generator output, the components of generator output [aligned with the request and orthogonal to (independent of) the request], and the supplier control error. Our preferred metric for regulation is equal to the ratio of the standard deviations of unit output relative to the expected output multiplied by the correlation coefficient between these two values:

$$\text{Correlation (Actual, Expected)} \times [\text{StDev (Actual)} / \text{StDev (Exp)}] .$$

The individual units differed substantially in their regulation performance, both in the amount of regulation provided and in their performance metric. Figure ES.1 shows the average hourly

performance of three generators. Unit C consistently performed well, while Unit B was erratic in its output, sometimes moving in the opposite direction from the control-center requests.

Fig. ES.1. Real-time performance metrics for three generators for 12 days in February 1999.

On this chart, 1.0 signifies perfect performance, 0 means the unit completely disregarded the system-operator requests, and negative performance means the unit worsened performance and perhaps should be charged for the regulation service.

Although these metrics should be tested in other utility settings, the results developed here suggest that these metrics can be implemented today. Such metrics are especially important in competitive markets, where the system operator would pay suppliers for real-time performance in delivering the requested regulation and load-following services.

ABBREVIATIONS AND ACRONYMS

ACE	area control error
AGC	automatic generation control
CPS	Control Performance Standard
NERC	North American Electric Reliability Council
RTO	regional transmission organization
SCE	supplier control error

INTRODUCTION

In an earlier project, we analyzed data on total system load as well as the loads of eight large industrial customers (Kirby and Hirst 2000). We analyzed these data in terms of system and customer-specific requirements for two real-power ancillary services, regulation and load following. We conducted these analyses using 12 days of data from February 1999 plus 12 days of data from August and September 1999.

The project discussed here focused on the supply side (provision) of these two services. Specifically, we examined the output of this control area's generation resources, in aggregate and by unit. We analyzed the performance of these generating units in two ways. First, we analyzed the contribution of these generators to overall system performance [generally relative to the North American Electric Reliability Council (NERC) standards]. Second, we analyzed performance relative to what the control center requested of the generators.

The two services examined here are briefly defined below. See also 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/min) to track the moment-to-moment fluctuations in customer loads and to correct for unintended fluctuations in generation. Regulation helps to maintain Interconnection frequency, manage differences between actual and scheduled power flows among 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 min 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 load following is not predictable by the control-area operator, the customer can inform the control center of impending changes in its electricity use.

The U.S. Federal Energy Regulatory Commission (FERC), in its Order 888, which defined six ancillary services, did not discuss load following (FERC 1996). However, its Order 2000 on regional transmission organizations (RTOs) requires RTOs to operate real-time balancing markets (FERC 1999). The primary resource for these markets is generation that can change output every 5 or 10 min (to follow load).

Because the restructuring of the U.S. electricity industry is still a work in progress, the industry has not yet agreed how best to develop and apply metrics for load following and regulation.* The basic functions of the two services are clear, however: regulation follows the short-term, largely random fluctuations, and load following compensates for the larger, slower, generally correlated and predictable swings associated with the daily load cycle. Developing metrics that quantify the supply of these two services by individual generators is proving to be difficult for the electricity industry. Should the metrics compare the performance of individual generators to overall system performance [e.g., as measured by the NERC (1999) Control Performance Standards (CPS)]? Or should the performance of an individual generator be compared only to the system-operator instructions sent to that generator? Both types of metrics are important, and both are developed here. Finally, such metrics will probably be applied differently within a vertically integrated utility assessing the performance of the generators it owns and operates versus a competitive market in which an RTO acquires (buys on behalf of customers) the output of certain generators.

Chapter 2 discusses the data provided for this project. Chapter 3 presents the metrics we developed for load following and the performance of the generators providing this service with respect to these metrics. Chapter 4 covers the regulation service. And Chapter 5 summarizes our key findings and conclusions.

*NERC's Interconnected Operations Services Subcommittee (formerly a Task Force) continues to work on the technical and policy issues associated with such metrics. The Task Force published draft compliance templates, including one for regulation and load following, in June 2000.

DATA

The control area provided 30-second data on generation and load for winter 1999 and summer 1999. For each 30-second interval, the data include total generation, net exports, total load, area control error (ACE), and Interconnection frequency. In addition, the data include the output from 11 of the control area's generating units (accounting for 90% of the system's total capacity) and the desired output from each unit. To improve readability, this paper presents results for only three of these units.

We split each observation into its load following and regulation components using the same method we developed for our prior load-analysis project (Kirby and Hirst 2000). Specifically, we aggregated the 30-second data to 2-min averages and used a 30-min rolling average to separate regulation from load following. We calculated the rolling average for each 2-min interval as the mean value of the seven earlier values of the variable, the current value, and the subsequent seven values:^{*}

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

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

We applied this allocation method to all the data used in these analyses.

As discussed in our previous report (Kirby and Hirst 2000), this method is somewhat arbitrary and imperfect. It is arbitrary in that the time-averaging period (30 min in this project) and the temporal aggregation of raw data (2 min) cannot be predetermined. In principle, the control-area characteristics (dynamics of generation and load) should determine these two factors. The method is imperfect in that it does not fully separate load following from regulation. In particular, the 30-min rolling average contains some short-term fluctuations that are more related to regulation than to load following.[†]

As discussed in Chapter 3, we used the sum of system load plus interchange schedules (essentially, net exports for this control area) as the reference against which to assess generator performance in providing the two ancillary services analyzed here. As shown in the top chart of Fig. 1, in the winter energy exports follow an hour-to-hour pattern similar to the native-load pattern. The schedules generally have the same morning pickup and late-evening drop-off that load does; however, the schedules do not show the same drop-off and re-pickup that load does in the late afternoon. In other words, interchange schedules contribute to this control area's load-following requirements.

^{*}In practice, system operators cannot know future values of load. They generally produce short-term forecasts of these values to aid in generation-dispatch decisions.

[†]Alternative methods, such as fast Fourier transforms, might more fully separate the two functions. However, the selection of the frequency (time interval) used to split the two services would still be arbitrary.

In the summer, the correlation between the load-following components of system load and interchange schedule is not as strong as in the winter. Because schedules ramp from one hour to the next during a 10-min period starting at the top of each hour, interchange schedules also impose a regulation burden, as shown in the bottom of Fig. 1.

Fig. 1. Hourly system load and interchange schedule for five days (*top*), and 2-min regulation requirements for system load and interchange schedule for two hours (*bottom*), both for February 1999.

LOAD FOLLOWING

Defining and using metrics for load-following performance is simpler than the comparable tasks for regulation. This situation is a consequence of the much simpler performance required for load following (large, generally predictable monotonic movements each hour) than for regulation (frequent, unpredictable small movements in both directions). This chapter discusses the analysis conducted with the load-following (smoothed) component of the time-series data, as discussed in Chapter 2.

METRICS

Performance Relative to System Requirements

Because the pattern of electricity exports is similar to that of its loads, it is important to consider interchange schedules along with system load in defining the total load-following requirement (top of Fig. 1). This control area generally sells increasing amounts of power during the morning ramp-up and decreasing amounts during the evening ramp-down. In essence, it sells both energy and the load-following service at wholesale.

As a consequence of these exports, we defined the total load following requirement (in megawatts) each hour as the signed difference between the maximum value of the sum of system load plus net exports and the minimum value of the sum of system load plus net exports:

$$\text{Load Following} = [\text{Load} + \text{Net Exports}]_{T_{\max}} - [\text{Load} + \text{Net Exports}]_{T_{\min}},$$

where $T_{\max}$ is the time within the hour that the system reaches its maximum value of load plus exports, and $T_{\min}$ is the time within the hour that the system reaches its minimum value.

We calculated the performance of generators, either individually or as a portfolio, on the basis of the change in output between $T_{\max}$ and $T_{\min}$ for each hour. Because the two times are the same for each generator, the total contribution to load following from generation is equal to the sum of the contributions of the individual units.

Performance Relative to Control-Center Requests

To measure the performance of a unit relative to the control-center request to that unit, we adopted the same metric used to measure the contribution of individual loads to the total load-following requirement. We calculated the load-following request based on the signed difference between the highest and lowest 2-min MW requests (i.e., the expected output from each generator) during an hour. The unit's load-following contribution is

$$\text{Generation}_{T_{\max}} - \text{Generation}_{T_{\min}},$$

where $T_{\max}$ is the time within the hour when the load-following component of the request to that specific generator reached its maximum value and $T_{\min}$ is the time within the hour when the load-following component of the request to that generator reached its minimum value.

The performance of Generator i relative to its expected value is then

$$\text{Performance}_i = (\text{Generation}_{i, T_{\max}} - \text{Generation}_{i, T_{\min}}) / (\text{Expected}_{i, T_{\max}} - \text{Expected}_{i, T_{\min}}) .$$

This metric can have values above 1.0 (implying better than perfect performance) and less than zero (implying a load-following movement that hurts, rather than helps, the system). As a practical matter, such values are rare.

RESULTS

In the winter, the average hourly load-following requirement for system load plus interchange schedule was 83 MW, of which load accounted for 60 MW and interchange schedules for 23 MW (Table 1). On average, generation provided 70 MW of the load-following service, and ACE (the rest of the Eastern Interconnection) provided the remaining 13 MW. In other words, this control area obtains a nontrivial portion (16%) of its load-following service from other control areas in the Eastern Interconnection; this issue is discussed further below. The summer load-following requirement was about 10% higher, and the share from schedules was slightly less (Table 1).

Figure 2 shows the weekday hourly averages of the total, system-load, and interchange-schedule requirements for winter load following. During the early-morning ramp-up hours, the average load-following requirement approached 250 MW. During the late-evening ramp-down period

Table 1. Hourly load-following requirements and supply

	Winter		Summer	
	Average (MW)	% of total	Average (MW)	% of total
Load-following requirements				
Load	59.7	72	71.5	79
Schedules	23.5	28	19.4	21
Total	83.2	100	90.9	100
Load-following supply				
Unit A	12.4	15	9.4	10
Unit B	11.5	14	10.6	12
Unit C	18.6	22	16.7	18
Total generation	70.2	84	77.2	85
ACE	13.2	16	13.8	15

Note: These values are calculated relative to the total hourly load-following requirement of the system load plus net interchange schedules (generally exports), with the signs adjusted so the load-following requirement is always positive.

(10 P.M.), the requirement reached ! 250 MW. The interchange schedules add to the load-following requirement because this control area sells increasing amounts of power during the morning ramp-up and decreasing amounts during the evening ramp-down. Loads and interchange schedules over this 288-hour period are modestly correlated, with an r of 0.42. The summer weekday patterns are slightly different, with two morning peaks in load following of about 140 MW each and a late evening drop-off with a load-following requirement of about ! 140 MW. The correlation coefficient of the load and schedule load-following requirements is 0.19, about half that for the winter.

Figure 3 shows the average winter weekday contributions of generation and ACE to the total load-following requirement. Clearly, generation dominates during the major ramp-up and ramp-down periods. ACE contributes to load following in a much smaller and seemingly random fashion. On average, however, ACE contributes to load following for most hours. The summer patterns are very similar to the winter ones.

Figure 4 shows the contributions of generation and ACE to the winter load-following requirement as a function of that requirement. Clearly, generation responds closely to the requirement ($R^2 = 0.94$) and ACE does not ($R^2 = 0.03$). ACE contributes most to the load-following requirement during those hours when the requirement is lowest (left side of Fig. 4), which are typically the mid-day hours when the industrial load dominates the load-following requirement. The patterns are almost identical for the summer period.

Fig. 2. Average hourly load-following requirement for winter weekdays.

Fig. 3. Average winter load-following contributions from generation and ACE.

Fig. 4. Supply of load-following services from generation and ACE as functions of total hourly load-following requirement.

Figure 5 shows the average winter weekday hourly contributions to load following for three generators. Although the hour-to-hour patterns for each of these units correspond closely to those for generation as a whole (compare Figs. 3 and 5), there are differences among these units. In general, these differences reflect the requests from the control center to each unit. The contributions of individual units are less well correlated during the summer period.

Figure 6 shows the performance of total generation relative to the control-center requests to the units (rather than relative to the system's requirements). Clearly, total generation closely follows load-following requests. The load-following performance of generation in aggregate is as good in the summer as in the winter. Actual performance matches expectations with an accuracy of over 90% for both the winter and summer periods (Table 2).

Fig. 5. Winter load-following contributions from three generators.

Fig. 6. Winter performance of generation in providing load following relative to control-center requests for the service.

Table 2. Load-following performance of three generators relative to unit-specific expected values

Unit	Winter			Summer		
	Expected (MW)	Actual ^a (MW)	Actual relative to expected (%)	Expected (MW)	Actual ^a (MW)	Actual relative to expected (%)
A	18.5	17.0	92	16.1	14.2	88
B	17.3	15.3	89	16.4	14.3	87
C	24.4	24.1	99	23.1	21.8	85
Total generation	80.5	74.0	92	87.0	79.1	91

^aThese values are calculated relative to the hourly expected values with the signs adjusted so that the expected load following is always positive.

REGULATION

METRICS

We examined the performance of generators providing regulation in three different ways:

1. Relative to the short-term volatility of system load or ACE. For example, to what extent does generator output reduce the standard deviation of ACE?
2. Relative to NERC's CPS1 and 2. For example, to what extent does generator output improve compliance with CPS1 and CPS2?
3. Relative to control-center requests for regulation up and down movements. For example, to what extent does generator output conform to the AGC requests from the control center?

Performance Relative to System Requirements

We considered two options for the standard against which to measure the performance of generation: system load or ACE, or CPS compliance. Load volatility is important to the control area only as it affects CPS performance and the costs of using generators to maintain CPS compliance within NERC bounds. Similarly, reducing the magnitude of ACE is of little benefit to the control area unless that reduction moves the CPS metrics from out of compliance to compliance. Therefore, we focused on generator contribution to CPS performance.

Focusing on CPS compliance raises three complications.

1. Measuring generator performance is complicated by the fact that both CPS1 and CPS2 are pass-fail measures; either the control area is in compliance (i.e., CPS1 is greater than 100% *and* CPS2 is greater than 90%), or it is not in compliance. Improving CPS2 from (for example) 92% to 94% is of little value to the control area except, perhaps, as insurance against possible later poor CPS performance. Thus, generator AGC movements benefit the control area only when they bring a noncompliant CPS metric into compliance.
2. CPS compliance is measured over long times—monthly for CPS2 and annually for CPS1. How do we assess the hourly or daily performance of a generator when any problems it might cause for CPS could be offset during other hours?
3. How should performance on these two metrics, CPS1 and 2, be weighted? What happens if the output of a particular generator improves CPS1 but worsens CPS2?

Because this control area's CPS compliance is well above the NERC requirements, both with and without the regulation component of generation, we did not address these complications.

The overall measure of generation performance relative to system requirements is based on two sets of ACE values, with and without the generator(s) in question. We then calculate hourly values for CPS1 and 2 based on these two sets of ACE values. In both cases, we assume that the generator continues to perform its load-following function, if any. That is, we analyzed only the regulation portion of each generator's output.

First, the generator's regulation performance can be *subtracted* from the actual ACE values, and the resultant 2-min ACE values can be used to calculate new hourly values for CPS1 and 2. The differences between these new CPS values and the original (actual) values show the effects of this generator on the CPS metrics.

Second, the regulation component of all the generation can be subtracted from the actual ACE values, yielding a set of ACE values and associated CPS metrics as if the units within that control area provided no regulation service. Then, the regulation performance of the unit in question can be *added* to ACE and the change in CPS performance associated with this generator determined. The two sets of metrics yielded results that were qualitatively similar.

Attributing the poor performance of a single generator to the system's monthly (CPS2) and annual (CPS1) performance will be very difficult. For example, the poor performance of a generator may lead to a CPS2 violation during one hour. However, if the system otherwise would have fewer than 13 such violations during the rest of the day, the generator's poor performance will have no effect on the system's performance; it still meets the 90% compliance requirement for CPS2.

Whether and how to use regulation metrics related to system performance is unclear. In the discussions leading up to NERC's proposed Policy 10 (NERC 2000), some task force members felt that if a generator behaves in a way that affects the overall performance of the control area in meeting the CPS requirements, the generator should be paid or charged for its effects on that performance. Generators that help reduce ACE and frequency deviations should be rewarded regardless of whether or not those actions were taken in response to control-center requests. Other task force members felt that individual suppliers should not "look over the shoulder" of the system operator at either system ACE or Interconnection frequency. To do so, they argued, would permit individual suppliers of the regulation service to second-guess the system operator's AGC dispatch instructions.

Although there is merit to both arguments, we believe generators should respond to control commands from the system operator. While CPS, ACE, and frequency are important reliability indicators, they are not the only factors the system operator controls. The system operator may direct a specific generator up or down for other reasons, such as to reduce transmission-line loading or accommodate another generator that is about to come on- or off-line. Turning the operator's directions into mere "suggestions" could undermine system control and reliability.

Because the system operator is ultimately responsible to NERC for CPS performance, it is inappropriate for other entities to independently try to meet these standards. For example, it is reasonable to charge a generator for the costs of CPS noncompliance if the unit's failure to respond contributes to the control area's noncompliance. But what if the noncompliance results from the system operator's failure to ask for enough resources or response? It may be better to leave the responsibility for CPS violations with the system operator and burden the generators only with failure to follow the operator's directions. For these reasons, we conclude that generator performance should be measured against the system operator's instructions only.

Expected Values of Generation

The data we received included overall and daily estimates of the AGC ramp-rate capability (in MW/min) of each unit. In addition, we received 30-second data for each unit on its actual output,

the upper and lower limits on that output (which determine the regulation range for the unit), and the AGC request from the control center. This *desired* value of output for each unit reflects the system operator's request for generator movement unconstrained by the physical limitations of the unit. [AGC systems can be (and are) designed and tuned to request movements consistent with unit constraints.] That is, the desired value assumes that the unit has no upper or lower limits and can respond with an infinite ramp rate to AGC requests.*

In principle, we should also constrain the expected values by the acceleration limits of the unit. Ignoring the acceleration limits of generators worsens, perhaps greatly, *apparent* generator performance (personal communication from P. Spicer, Wisconsin Public Service Corp., Green Bay, Wis., Nov. 13, 2000). Because data on unit-specific acceleration limits were not available, we ignored these constraints in this project. Thus, the expected values we derived are optimistic, and the performance metrics we derived are worse than they actually are.

We developed a new variable called *expected* output intended to reflect the physical characteristics and limits of each generator. These values are computed as follows:

1. Compare the desired output at time t with the expected output at $t-1$. If the absolute value of this difference exceeds the unit's stated ramp rate, set the expected output at time t to be consistent with the unit's stated ramp rate. The only exception to this rule occurs if the unit, during this time interval, moved at a rate faster than its stated ramp rate, in which case the expected output is limited by the actual ramp rate.
2. Compare the desired output at time t with the upper and lower limits. If the desired output falls outside this range (either higher than the high limit or lower than the low limit), the expected output is set to either the upper or lower limit, depending on whether the desired output at time t is greater or less than the expected (not desired) output at time $t-1$. The only exception to this rule occurs if the unit, at this time, is operating beyond its stated upper and lower limits, in which case the expected output is limited by the actual output.

Mathematically, this computation can be stated as

$$\text{Expected}_t = f(\text{Desired}_t, \text{Expected}_{t-1}, \text{RR}, \text{UL}, \text{LL}, \text{Actual}_t, \text{RR}_a),$$

where

Expected	=	physically feasible output (MW) that can be expected from the unit at time t
Desired	=	raw AGC request to the unit (MW)
Actual	=	actual output of the unit (MW)
RR	=	stated ramp rate for the unit (MW/min)
RR _a	=	actual ramp rate for the unit (MW/min) between times $t-1$ and t , calculated as $ \text{Actual}_t - \text{Actual}_{t-1} / \Delta t$
t	=	time (min) and $\Delta t = t - (t-1)$ (min)
UL	=	stated upper limit for the unit (MW)

*For this control area, the sum of the desired signals across all the generators results in a total generation request that matches ACE in shape and is about three times larger than ACE. If the generators accurately followed these desired signals, ACE would double in size and many CPS violations would occur. This hypothetical situation demonstrates that this control area sends AGC requests to its generators that reflect the fact that generators will not respond fully to those requests. Thus, the system operator, its AGC system, and the generators operate in a closed loop.

LL = stated lower limit for the unit (MW)

The value of $Expected_t$ is calculated using a two-step process: (1) check against UL or LL depending on whether the request is an increase or a decrease; and (2) check against the RR ability of the unit. In both steps, if the unit is operating outside its stated range or is moving faster than its stated RR, then the actual performance is the limit against which the Expected is calculated.

If $Desired_t > Expected_{t-1}$:

$$Expected_t = \text{Min}\{\text{Max}(UL, \text{Actual}_t), \text{Max}[(Expected_{t-1} + RR \times \Delta t), \text{Actual}_t], \text{Desired}_t\}.$$

If $Desired_t < Expected_{t-1}$:

$$Expected_t = \text{Max}\{\text{Min}(LL, \text{Actual}_t), \text{Min}[(Expected_{t-1} - RR \times \Delta t), \text{Actual}_t], \text{Desired}_t\}.$$

As shown in Exhibit 1, the values of expected generation for each unit should be calculated using short-interval (e.g., 30-second) data.

Exhibit 1. Time Intervals for Calculating Expected Values

We computed values of $Expected_t$ at each 30-second interval and averaged these to generate the 2-min values. We also calculated $Expected_t$ directly from the 2-min averages of the desired and actual signals and found that this yields expectations that are much more volatile than those based on the 30-second values. We believe that the expectations based on the 30-second values are much more reasonable than those based on the 2-min averages. The 30-second values better reflect the control-center AGC requests.

Performance Relative to Expected

We explored several possible metrics for individual generators to measure their performance against the expected values of generation discussed above. (Exhibit 2 lists several desirable characteristics for such metrics.) These generally follow the supplier control error (SCE) concept developed by NERC in its proposed Policy 10. SCE is equal to $\text{Output}_{\text{Actual}} - \text{Output}_{\text{Expected}}$. Two possible metrics are the following:

Metric 1: $1 + [\text{StDev}(SCE_t)/\text{StDev}(Exp_t)]$.

Metric 2: $\text{Correlation}(\text{Actual}_t, \text{Expected}_t) \times [\text{StDev}(\text{Actual}_t)/\text{StDev}(Exp_t)]$.

In these formulations, t refers to the thirty 2-min values in an hour, and the standard deviations are calculated with these thirty 2-min values. If the standard deviation of the actual is greater than the standard deviation of the expected value, the second metric yields a value greater than 1.0, implying a performance greater than 100%. In actuality, this over-response is as bad as an under-response and should be treated the same way. For example, a value of 1.2 should be considered equal to 0.8 in terms of providing compensation for that hour. Our definition of expected automatically values overperformance when such overperformance is consistent with the desired request.

Exhibit. 2. Desirable Characteristics of Generator Performance Metrics

Linearity

The metric should reflect the generator's performance in a manner that is proportional to the useful work the generator did. If the generator supplied half of what it should have, the metric should indicate 50%.

Harm vs Nonperformance

The metric should *appropriately* distinguish between failure to provide the desired response and introduction of undesirable fluctuations. Fluctuations that are exactly opposite of the desired behavior are worse than random fluctuations. Random fluctuations are similar to those imposed by loads.

Useful Enhanced Performance

The metric should not penalize a generator that exceeds its ramp rate or high/low limits *if* doing so helps the system operator. Our use of "Desired" and "Expected" helps by allowing the Expected value to exceed the generator's stated capabilities if the generator chooses to do so *and* if the extra generator movement is consistent with the Desired value.

Capacity vs Performance

The metric should appropriately recognize that the system operator reserves a certain amount of regulating capacity ahead of time, whether or not the operator uses that capacity in real time. For example, the system operator might reserve 100 MW of regulation capacity beforehand and then call on only 60 MW in real time. If the unit delivers 40 MW of regulation service, should it be compensated for the 40 MW it provided, for 80 MW (because it failed to provide 20 MW of the 100 offered to the system operator), or for 67 MW (two-thirds of the 100 MW offered)?

Adjusting Performance Compensation for Small Requests

When asked to provide zero or small amounts of regulation, the generator may through natural random motion swamp the desired response. Judging and paying for performance on the basis of percentage compliance with such small requests would be unfair and inappropriate. It may be appropriate to establish a deadband in the performance metric(s) near zero regulation request so that generator performance in this region is not even judged.

Regulation Delivery vs Regulation Consumption

A generator is simultaneously a provider and consumer of regulation. Specifically, a generator can be providing the desired regulation but also having a random regulation component. These two components should be measured separately.

Support of NERC Standards (CPS 1 and 2)

An alternative (or additional) consideration could tie generator performance and compensation to the control area's compliance with the NERC Control Performance Standards. It may, however, be difficult for a control-area operator to demonstrate unambiguously that a generator's poor performance harmed the control area in some way (e.g., caused a violation of NERC requirements or

Both metrics discussed above should have a maximum value of 1.0 (100% compliance, or perfect performance), but a minimum value can be negative. A value of zero implies complete disregard of control-center requests, while a negative value implies generator movement counter to that requested (which would hurt system performance).^{*} When a unit is not on AGC, we set the expected output of the unit equal to zero, and the regulation metrics described above are not computed.

An alternative normalizing factor could be the amount of regulation capacity provided by each generator. This factor has two advantages. First, it is simple to compute. Second, it does not suffer from a problem that could occur when the control center requests very little AGC movement during an hour and the unit's control system contains small amounts of error, such that the percentage error is large. In such cases, a unit could be unfairly penalized for failing to follow very small requested movements.

Finally, we developed and tested a pair of metrics, the first of which measures the contribution to the provision of the regulation service and the second of which measures the load-type use of regulation by the generator. Both metrics rely on the correlation coefficient between actual and expected output, as described above:

$$\text{Regulation contribution (MW)} = \text{StDev(Actual)} \times \text{Correlation(Actual, Expected)} .$$

$$\text{Load requirement (MW)} = \text{StDev(Actual)} \times \sqrt{1 - \text{Correlation}^2} .$$

As suggested by Fig. 7, the actual generator movements are split into orthogonal components, one corresponding to the control-center requests and the other independent of those requests. These two metrics, unlike the ones discussed above, have units of capacity. The prior metrics are dimensionless and are scaled so that 1.0 implies perfect regulation performance.

Two additional factors need to be considered before a metric is suitable for use in commercial markets for the regulation service. First, it may be necessary to include a deadband in the metric because a generator cannot respond exactly to requested movements. That is, the AGC systems at a generator may have some error that prevents the unit from following precisely the AGC request. In addition, generators are

Fig. 7. Regulation can be split into orthogonal components, one aligned with AGC requests and the other independent of those requests.

^{*}In principle, the chosen performance metric would be used to determine the hourly payment to each generator providing the regulation service (or the hourly charged if the generator's performance is sufficiently poor). As a practical matter, it probably makes sense to limit payments and charges to the range $\pm 100\%$ of the hourly price of regulation.

sufficiently complicated machines that, with or without AGC, they cannot maintain desired output exactly. A unit should not be penalized for small errors in its performance.*

More important, the metric may need to account for the fact that the amount of regulation *delivered* during an hour may be less than the amount *requested* during an hour, which, in turn, may be less than the amount *reserved* for that hour. How should the metric be used to set the regulation payment for a unit in that case? For example, consider the situation in which the system operator reserves 100 MW of regulation capacity, the system operator requests 50 MW during the hour, and the unit delivers 30 MW that hour. Does the unit get paid that hour for the 30 MW provided, for 80 MW (the 100 reserved minus the 20 not delivered) or for 60 MW (because 60% of the requested amount was provided)? This is less an analytical issue and more a policy one. In today's ancillary-services markets, managed by independent system operators in California and the Northeast, suppliers are paid primarily for the amount of regulation capacity reserved rather than for the amount of service delivered in real time.

RESULTS

During the 12-day winter period, this control area had an average of 0.16 CPS2 violations per hour (the maximum number of violations allowed is 0.6 per hour). Its CPS1 performance averaged 161% (the NERC requirement is 100% or more on an annual basis). Thus, this control area easily met the NERC requirements. Figure 8 shows ACE for 2-min intervals for 12 hours on one weekday; clearly, ACE is generally well within the CPS2 bounds. Although the summer performance (0.31 CPS2 violations per hour and CPS1 equal to 118%) was not as good as the winter performance, the control area still complied fully with the NERC requirements.

*"Small" can be defined relative to the amount of capacity reserved for regulation each hour or relative to the generator's maximum regulation generating capability. Alternatively, it could be defined relative to the amount of regulation service requested from the control-area operator. The former type of definition may be more practical.

We then removed the regulation component (but not the load-following component) of generation from ACE and recalculated compliance with CPS1 and 2. CPS1 performance improved slightly in the winter, from 161% to 168%. On the other hand, CPS2 performance declined slightly, from 0.16 to 0.35 violations per hour (i.e., from 97% to 94%, with a minimum NERC requirement of 90% on a monthly basis). For the summer, both CPS1 and CPS2 declined slightly (from 0.31 to 0.47 CPS2 violations per hour, and from 118% to 115% for CPS1) when we removed the regulation component of generation.

Thus, as a practical matter, the regulation component of generation had little effect on CPS performance and no effect on CPS compliance, for both the winter and summer periods. Because the total generation portfolio had so little effect on CPS compliance, we did not examine the contributions of individual generators to CPS compliance.

These results suggest two conclusions. First, the contribution to CPS performance from the fast movements of the generators are minor and ambiguous. Second, the regulation component of generation is not needed to maintain compliance with NERC's CPS requirements. That is, the AGC systems could be disconnected from all the control area's generators and CPS compliance would still be above the minimums required. However, the generation *is* needed for load following. These conclusions are slightly overstated because our separation of load following from regulation uses 15 min of *future* values of generation, load, and ACE. In reality, uncertainty about these future conditions leads to the use of some regulation to provide the load-following service.

Turning to the performance of generation relative to the AGC requests, Table 3 summarizes the performance of three generators and the total portfolio over the two analysis periods. The amount of time on AGC differed across the units and between the two seasons. As noted above, the performance metric was calculated only for those times a unit was on AGC; i.e.,

Fig. 8. Area control error for 12 hours on one winter day.

$$\text{Correlation}(\text{Actual}_i, \text{Expected}_i) \times [\text{StDev}(\text{Actual}_i)/\text{StDev}(\text{Exp}_i)] .$$

Overall, generation provided 52% of the expected regulation service in the winter (Table 3). Among the three units, Unit C performed best, with a metric of 86%, and Unit B performed worst (25%). In the summer, the overall performance was lower (37%), primarily because Unit C did not follow its regulation requests as well as it had in the winter.

Figure 9 shows the winter performance of the generation portfolio, hour by hour for weekdays. Both parts of the figure show expected regulation output. The top part of the figure also shows actual output and SCE, while the bottom part shows the amounts of actual output that are correlated and uncorrelated with the expected values. The summer data show a regulation requirement about 15% higher and an actual performance somewhat poorer than for the winter data.

Figure 10 shows the winter performance for two generators, A and C. Unit A had average performance (48%, as shown in Table 3), whereas Unit C had excellent performance (86%). A comparison of the two parts of this figure shows that the actual standard deviation for regulation output is much closer to the expected value for C than for A. Similarly, the SCE for C is much smaller relative to its actual value than for A. Figure 11 shows the details of the winter regulation performance for these two units for a 2-hour period. For the first hour shown, A had a 68% performance and C had a 94% performance; for the second hour, the comparable metrics were 26% and 80%. As can be seen in Table 3, for the summer period performance at both units declined (from 48 to 22% for unit A, and from 86 to 45% at unit C).

Figure 12 summarizes the average winter hourly weekday performance of three generators. Unit C dominates the production of regulation, both in terms of output and performance metric. Unit B contributes little to regulation and has poor performance metrics. For the summer period, the differences among the units in their regulation contribution were much smaller.

Table 3. Regulation performance (standard deviations in MW and dimensionless metric) of individual generators

Unit	Time on AGC (%)	Expected	Actual	SCE	Actual × correlation coefficient	Performance metric (1.0 = perfect)
Winter						
A	91	7.4	5.6	5.4	3.8	0.48
B	85	4.9	3.4	5.0	1.2	0.25
C	96	8.9	9.2	4.3	7.9	0.86
Total generation	—	25.0	18.8	17.6	13.4	0.52
Summer						
A	73	5.5	3.7	5.6	1.4	0.22
B	89	5.3	3.9	5.3	1.4	0.27
C	96	7.7	6.9	6.9	3.8	0.45
Total generation	—	29.2	20.6	25.0	11.2	0.37

UReg

CONCLUSIONS

Using 2-min data on generator output and control-center requests, we analyzed the performance of generation resources in delivering regulation and load following, two key real-power ancillary services. To conduct these analyses, we first developed suitable performance metrics that can be applied to individual generators as well as to the entire resource portfolio. We developed two types of metrics. The first focuses on performance relative to system requirements, taken here to be compliance with NERC Control Performance Standards. The second type focuses on performance relative to the control-center requests to each generator. We conducted these analyses using 12 days of data from February 1999 and 12 days of data from August and September 1999.

We defined load following as an hourly (rather than a daily) service to be consistent with emerging competitive bulk-power markets. The load-following requirement each hour is the signed difference between the highest and lowest smoothed values during the hour of retail load plus net exports. Using these two times ($T_{\max}$ and $T_{\min}$), we computed the contribution of each generator to the overall requirement. In a similar fashion, we defined the control-center request for load-following service as the signed difference between the highest and lowest smoothed values of requested output levels during the hour. Each unit's performance relative to this unit-specific expectation is based on its change in output between these two times.

Overall, this control area's generation accounted for about 85% of the average hourly load-following requirement (83 MW in the winter and 91 MW in the summer). The rest of the Eastern Interconnection (reflected in the ACE values) contributed the remaining 15%. The ACE contribution to load following occurs primarily during those hours when the industrial load dominates the load-following requirement. Unit C was the most important contributor to load following in the winter, accounting for 22% of the total (Table 1). Units A and B each contributed more than 10% of the total. In the summer, unit C contributed about 20% of the total. The individual units generally performed very well relative to the control-center requests to each unit; performance was 85% or better for all the units analyzed.

Analyzing regulation performance is much more complicated than analysis load following because of the frequent changes in direction. Our analysis of regulation performance relative to system requirements focused on compliance with CPS1 and 2 rather than offsets to system-load volatility or reductions in ACE. We found that CPS compliance was good enough that removal of the entire regulation component of generation had almost no effect on these compliance values. Thus, we conclude that the contribution of generation regulation output to CPS performance is minor and ambiguous. However, the load-following output of these generators is essential to maintain compliance with CPS1 and 2.

We developed and applied several metrics to measure the hour-to-hour regulation performance of individual generators relative to control-center requests. These metrics use the standard deviation as the measure of volatility in control-center requests, generator output, the components of generator output [aligned with the request and orthogonal to (independent of) the request], and the supplier control error.

The individual units differ substantially in their regulation performance, both in the amount of regulation provided and in their performance metric. In the winter Unit C provided the most regulation—8 MW on average (Table 3). (This was also the highest value for all the data examined.) Overall, generation provided an average of 19 MW of regulation in the winter, of which 13 MW was aligned with the expected value, yielding a score of 52%. In the summer, Unit C again provided the most regulation (4 MW) . Overall, generation provided an average of 21 MW in the summer, of which 11 MW was aligned with the expected value, yielding a score of 37%. Thus, regulation performance was better in the winter than the summer.

In summary, we defined and applied metrics to measure the real-time performance of generators, in aggregate and individually, in delivering the regulation and load-following services. Although these metrics should be tested in other utility settings, the results developed here suggest that these metrics can be used by traditional, vertically integrated utilities and in RTO competitive-market settings. Such metrics are especially important in competitive markets, where the RTO would pay suppliers for real-time performance in delivering the requested regulation and load-following services.

REFERENCES

- Hirst, E., and B. Kirby. 1998. *Unbundling Generation and Transmission Services for Competitive Electricity Markets: Examining Ancillary Services*. NRRI 98-05. National Regulatory Research Institute, Columbus, Ohio, January.
- Hirst, E., and B. Kirby. 2000. "Time-Averaging Period for Regulation Service." *IEEE Power Engineering Review* **20**(7), 10–12, July.
- Interconnected Operations Services Working Group. 1997. *Defining Interconnected Operations Services under Open Access*. EPRI TR-108097. Electric Power Research Institute, Palo Alto, Calif., May.
- Kirby, B., and E. Hirst. 2000. *Customer-Specific Metrics for the Regulation and Load-Following Ancillary Services*. ORNL/CON-474. Oak Ridge National Laboratory, Oak Ridge, Tenn., January.
- NERC (North American Electric Reliability Council). 1999. "Policy 1—Generation Control and Performance." *NERC Operating Manual*. Princeton, N.J., November.
- NERC (North American Electric Reliability Council). 2000. *Interconnected Operations Services: Draft 3.1 of Proposed Policy 10*. Princeton, N.J., February.
- FERC (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, D.C., April 24.
- FERC (U.S. Federal Energy Regulatory Commission). 1999. *Regional Transmission Organizations*. Docket No. RM99-2-000, Order No. 2000. Washington, D.C., December 20.

INTERNAL DISTRIBUTION

1.	V. D. Baxter	14	P. N. Leiby
2.	L. G. Berry	15.	J. M. MacDonald
2.	D. J. Bjornstad	16.	V. C. Mei
3.	M. A. Brown	17.	C. I. Moser
4.	J. E. Christian	18.	A. C. Schaffhauser
5.	T. R. Curlee	19.	M. Schweitzer
6.	P. D. Fairchild	20.	R. B. Shelton
7.	R. G. Gilliland	21.	J. J. Tomlinson
8.	S. W. Hadley	22.	B. E. Tonn
9.	S. G. Hildebrand	23.	J. W. Van Dyke
10.	L. J. Hill	24.	T. J. Wilbanks
11.	P. J. Hughes	25.	Central Research Library
12.	B. J. Kirby	26.	Document Reference Section
13.	R. M. Lee	27.	Laboratory Records (RC)

EXTERNAL DISTRIBUTION

29. Dr. Lilia A. Abron, President, PEER Consultants, P. C., 1460 Gulf Blvd. Apt. 1103, Clearwater, Florida 33767.
30. Dr. Douglas C. Bauer, Executive Director, Commission on Engineering and Technical Systems, National Research Council, Harris 280, 2001 Wisconsin Ave. NW, Washington, D.C. 20007
31. Dr. Susan L. Cutter, Professor and Chair, Director, Hazards Research Lab., Department of Geography, University of South Carolina, Columbia, South Carolina 29208.
32. Mr. P. Richard Rittelmann, FAIA, Executive Vice President, Burt Hill Kosar Rittlemann Associates, 400 Morgan Center, Butler, Pennsylvania 16001-5977.
33. Dr. Susan F. Tierney, The Economic Resource Group, Inc., One Mifflin Place, Cambridge, Massachusetts 02138.
34. Dr. C. Michael Walton, Ernest H. Cockrell Centennial Chair in Engineering, Department of Civil Engineering, University of Texas at Austin, Austin, Texas 78712-1076.
- 35.-36. OSTI, U.S. Department of Energy, P. O. Box 62, Oak Ridge, TN 37831 Office of Assistant Manager for Energy Research and Development, DOE/ORO, P.O. Box 2001, Oak Ridge, TN 37831-8600.
37. External Electric Industry Policy Studies group distribution mailing list and extra copies to Brendan Kirby, 4500N, Rm. 102, MS-6070.