MIT Libraries | DSpace@MIT

MIT Open Access Articles

Energy Storage for Use in Load Frequency Control

Massachusetts Institute of Technology

Energy Storage for Use in Load Frequency Control

Olivia Leitermann, James L. Kirtley
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
77 MASSACHUSETTS AVE.
CAMBRIDGE, MASSACHUSETTS
EMAIL: kirtley@mit.edu

Abstract—*Certain energy storage technologies are well-suited to the high-frequency, high-cycling operation which is required in provision of load frequency control (LFC). To limit the total stored energy capacity required while reducing the cycling burden on traditional thermal generators, the LFC signal may be split between thermal generators and energy storage units. To evaluate the dispatch of energy storage units in concert with thermal generators, this paper presents energy-duration curves and ramp-rate-duration curves as graphical tools. The energy storage requirement and thermal ramping requirement may also be graphically compared to provide insight for cost evaluations.*

I. INTRODUCTION

Traditional thermal generators are limited in their ability to provide load frequency control (LFC) because of restrictions on power ramp rate. [1]–[3]. By contrast, certain energy storage devices are well suited to such high-frequency, high-power cycling operation [4]–[6].

The approach described in this work is to split the LFC signal between two dissimilar sets of assets: nimble but lower-capacity energy storage units, and slower traditional thermal generators. The goal is to decrease the fuel and maintenance requirements of the thermal generators and to enable the integration of variable generation resources that can increase LFC requirements [7]. Further, this approach may enable better provision of LFC through the use of storage to track fast fluctuations without increased cost. This paper proposes graphical tools for use in evaluating the suitability and dispatching of an energy storage unit for LFC duty. Also proposed is a broad strategy for incorporating the storage in dispatch.

The use of energy storage units for LFC has been limited by the concern that they will unexpectedly be completely filled or emptied and hence be made unavailable. The graphical tools suggested here, called energy-duration curves and ramp-rate-duration curves, seek to manage and inform the dispatch of energy storage for LFC. Using these tools on a representative data set, it is easy to see the net energy storage and ramp rates that LFC requires. Unit outage rates due to insufficient energy capacity may be predicted based on historical data and the system may be designed to avoid or mitigate any such outages. Different methods for dividing the power signal between energy storage and thermal assets may then be easily compared using the curves.

II. BACKGROUND

Load frequency control, or the minute-to-minute adjustment of generated power on the grid to follow fluctuations in load, is traditionally provided by baseload, mid-merit, or peaking thermal plants running at part load [1]. This is an expensive mode of operation, because in addition to the efficiency penalty imposed by part-load operation, the plants also suffer an additional efficiency penalty when throttling [8], [9]. The varying output power can also increase maintenance costs and increase wear and tear on the plants [3]. Even so, the LFC performance of these plants may not be sufficient for good grid control, as response times and ramp rates are limited [2].

By contrast, energy storage plants can be very well suited to the provision of LFC. Some energy storage technologies (such as flywheels and some battery chemistries) are very nimble and can rapidly change power settings with virtually unlimited ramp rates [4]. Unlike in arbitrage applications like load-shifting, in this application a small energy capacity is not a major difficulty. Some energy storage plants have successfully been incorporated into the electric grid for frequency control [5], [6], [10]. Still, the use of energy storage for this application is not widespread.

The use of energy storage units for LFC requires some different analytical tools from those associated with traditional thermal generating units. Because of the mixture of time scales in this problem, time-series graphics of LFC power requirements offer little insight. For thermal units, the load-duration curve is a useful tool for examining use patterns [11]. However, load-duration curves provide no information on ramp rates or required net energy delivery. Fourier decompositions also fail to provide insight into net energy and ramp rate because the magnitude and phase decompositions are not easily interpreted. When using an energy storage unit in concert with thermal units, two related metrics become more important. The first we shall call the ramp-rate-duration curve, and the second the energy-duration curve.

III. DIVIDE REGULATION BURDEN

The main approach taken in this paper is to divide the burden of LFC between fast energy storage units and slower traditional thermal generators [12]. In this way, the energy storage can assume the fastest-cycling portion of the required

---

[1]Paper No. 631

LFC and allow thermal generators to be operated at steadier conditions. This approach is illustrated in Fig. 1.
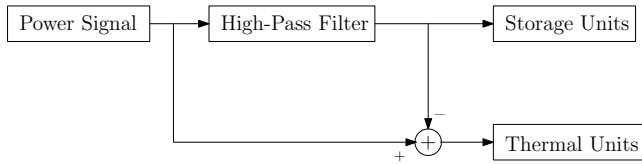

Fig. 1. Block diagram of scheme to partition load frequency control signal between thermal generators and energy storage units.

To demonstrate this concept, this paper uses a data set from a United States balancing area which includes 10-second total control area load. The data set runs for 9 non-consecutive days representing different load conditions and times of year. All data processing is done on each day individually to avoid artifacts due to the discontinuity between days. The raw load data is initially processed through a 5-point median filter to remove anomalous data spikes. This minimally processed load data is included as Fig. 2. This data is used in the following sections to illustrate the techniques which are described.
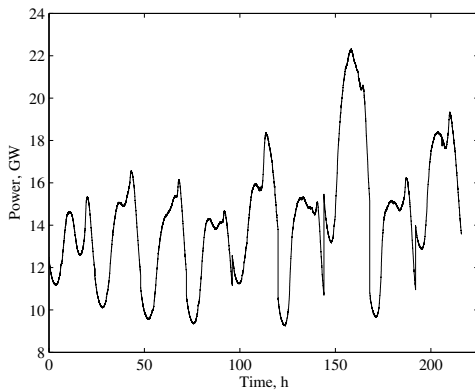

Fig. 2. Total load power in a balancing area on 9 non-consecutive days. Data points are sampled every 10 seconds. Note that discontinuities between days do not represent actual step changes in load and each day is processed individually in figures that follow.

IV. RAMP-RATE-DURATION CURVE

The ramp-rate-duration curve displays the use of ramping capability in much the same way as a load-duration curve displays the use of (thermal) power capacity. It is a visual representation of the fraction of time that a certain total ramp rate is required of a generating system. A ramp-rate-duration curve can be constructed by first determining the ramp rate by taking the derivative (or finite differences) of the dispatched power curve. The ramp-rate-duration curve is then created by tallying the fraction of time the ramp rate is at or below a
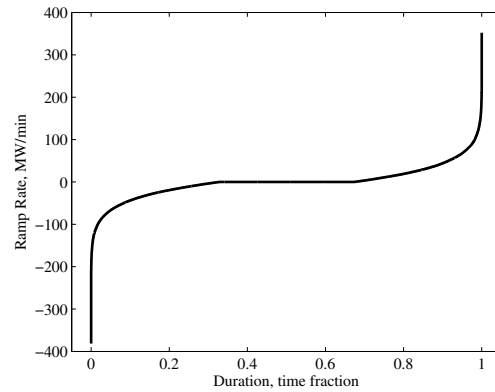

Fig. 3. Ramp-rate-duration curve for total area load power (corresponding to Fig. 2). The large ramp rates are difficult and expensive for thermal generators to produce.

certain level, for example by sorting the ramp rate curve data [13].

The ramp-rate-duration curve corresponding to Fig. 2 is in Fig. 3, where a small number of outlying points have been removed.[2] The flat portion in the center of the curve is due to the use of the median filter. This figure clearly indicates the potentially large ramp rate requirement of the regulation signal, compared to the capabilities of thermal units. Although some of the larger ramp rates indicated by this plot may be due to data collection errors, the data indicate an average absolute ramp rate of about 24 MW/min, out of a total generating capacity of 10-20 GW [1].

The reason for using a ramp-rate-duration curve stems from a performance difference between traditional thermal assets and energy storage units. While in general energy storage units are able to ramp from one power level to another very quickly, ramping with thermal units is slow and is more expensive than steady state operation [2]–[4]. From the ramp-rate-duration curve of an LFC signal sent to a thermal unit, both the maximum ramp rate to be required of the unit and the fraction of time the unit is ramping at any given rate are clear.

V. ENERGY DURATION CURVE

The energy-duration curve is similar to the ramp-rate-duration curve, but tallies net energy required at each instant. It is primarily of interest for resources that cannot deliver nonzero average power. Because a nonzero average power value will lead to a ramp in energy upon integration, the low frequency components of a power signal must be removed before an energy-duration curve is created. As an illustration, a Chebyshev type I high-pass filter of order 3 is used to

[2]Ramp rate points outside of 5 standard deviations were removed if their absolute values exceeded 4 times the absolute value of neighboring points on both sides. This intentionally conservative method for detecting telemetry errors resulted in the removal of less than 0.1% of data points. Some remaining points may also be the result of telemetry errors. Balancing Area engineers indicated that sustained load ramps do not generally exceed about 40 MW/min, although ramp rates for shorter changes may be higher.

separate high-frequency and low-frequency components. The break frequency of this filter is approximately 1/60 minutes. The high-frequency and low-frequency components of the signal are illustrated in Fig. 4.
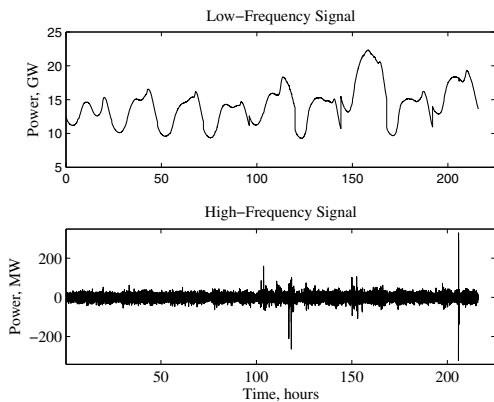

Fig. 4. High- and low-frequency portions of the total power signal of Fig. 2 using a Chebyshev type I filter of order 3, with a cutoff frequency of approximately 1/60 min, in a scheme like that illustrated in Fig. 1. The sum of the upper and lower graphs is equal to the total signal.
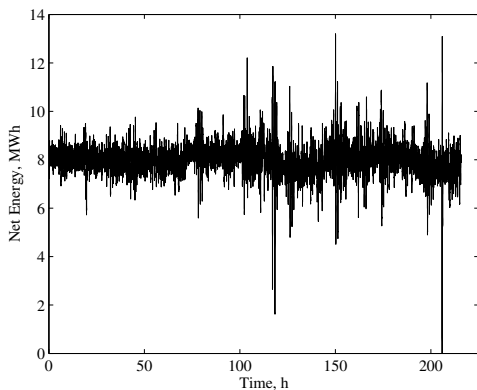

Fig. 5. Integral of high-frequency (bottom) portion of Fig. 4, the total power signal of Fig. 2 processed using a Chebyshev type I high-pass filter of order 3 with a cutoff frequency of approximately 1/60 min, in a scheme like that illustrated in Fig. 1. The majority of low-frequency content has been removed.

Once a power signal with only high-frequency content has been generated, the net energy as a function of time is obtained by taking the integral of the power curve. This is pictured in Fig. 5. However, if low frequency components of the power signal have been effectively removed, it is difficult to draw further conclusions about energy storage requirement from this graph.

The energy-duration curve is created by tallying the percent of time that the net energy requirement is at or below a certain level. (For example, this can be done by sorting the data points.) Figure 6 shows the energy-duration curve corresponding to the high-frequency portion of the control signal of Fig. 4. Note that the absolute energy values in the
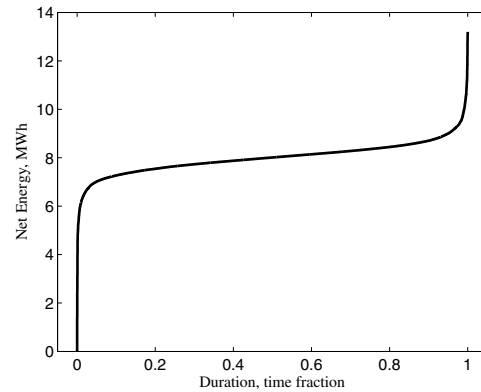

Fig. 6. Energy-duration curve for Fig. 5, output of a high-pass filter on total area load power. It can be seen that for the high frequency portion of the signal, a limited amount of stored net energy is required.

energy-duration curve are of less interest than the total change, so the data has been shifted vertically to set the low energy point to zero.

The advantage of the energy-duration curve is that it simplifies the evaluation of whether a given energy storage unit is well-suited to following a given power signal. First, the maximum required energy storage capacity is readily available from the two ends of the graph. Furthermore, it is easy to determine how often the storage capacity is being used. As can be seen from Fig 6, an energy storage unit of only 14 MWh would be sufficient to follow the high-frequency regulation signal of Fig. 4.

VI. Example Filter Evaluation

To illustrate the utility of the energy- and ramp-rate-duration curves, consider again the task of partitioning the load power signal of Fig. 2 between fast-acting, limited-energy storage units and slower, limited-power thermal generators. The goal of the partition is to limit the total required energy storage and the maximum and average ramp rate of the thermal units while adequately responding to the entire signal. The ultimate aim is to produce a method to partition the power requirement in real-time. Hence, only causal candidate filters are investigated.

To demonstrate the use of this technique, the effects of a class of simple filters will be explored. All filters seek to partition the signal between the thermal and energy storage assets to more effectively take advantage of the strengths of each unit type. The Chebyshev type I high-pass filter used in Section V was selected for its good attenuation of low frequencies and its fast transition band. A filter of order 3 was found to offer a good compromise between fast rolloff and the increased delay produced by additional poles.

A range of cutoff frequencies was selected with periods from 20 minutes to 90 minutes. Because filters with these cutoff frequencies lead to moderate energy storage requirements, higher frequencies are not included in these plots. Figure 7 illustrates the energy-duration curves which result

from these filters. The energy capacity requirement for these filters ranges from about 12 to about 130 MWh. The power requirement for the high frequency component of these filters is included as Fig. 8. Each filter requires about ±400 MW and thus the 20 minute cutoff filter requires about 2 minutes of energy storage, while the filter with a 90 minute period cutoff requires about 20 minutes of energy storage. The substantial improvement in thermal unit ramp rate which these filters provide may be seen in Fig. 9. It is notable that both the energy-duration curves and the ramp-rate-duration curves have long "tails," indicating a requirement for high energy storage capacity and high ramping capability which are used only infrequently.

Fig. 9. Ramp-rate-duration curve of low frequency portion of load power signal, obtained by subtracting from the original signal the output of Chebyshev type I high-pass filters of order 3 with several different cutoff frequencies. Compared to Fig. 3, both the maximum and the mean ramp rates required of the thermal units have been substantially reduced.
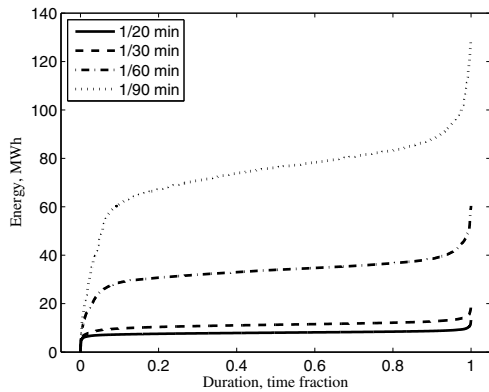

Fig. 7. Energy-duration curve of high frequency portion of load power signal using Chebyshev type I high-pass filters of order 3 with several different cutoff frequencies.
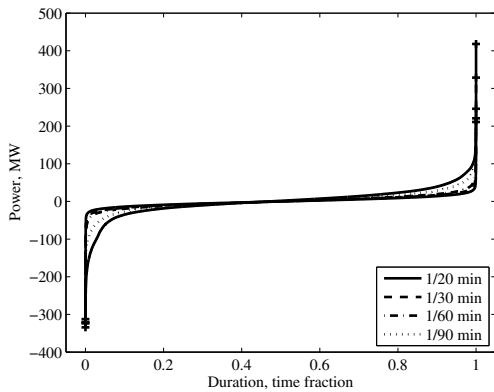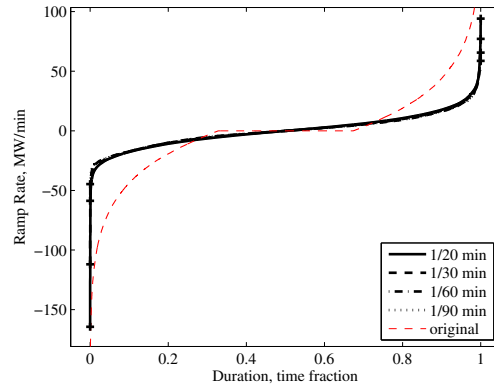
average ramp rate versus maximum energy storage requirement. Instead of plotting the percent of time spent at each ramp rate or at at each stored energy level, the average of the absolute ramp rate is used as a ramping cost metric and plotted against the maximum energy storage required. In this way several filters or dispatch methods can be compared, and an economically and operationally appropriate solution may be selected by comparing the cost of energy storage with the cost of ramping. An example of this plot including the Chebyshev filters is shown as Fig. 10.
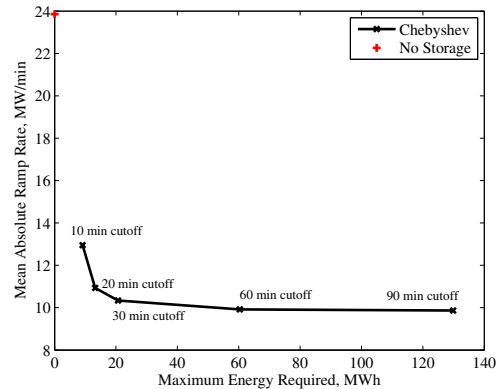

Fig. 10. Comparison of different filters in terms of average absolute ramp rate versus maximum energy storage requirement. The filters of Figs. 7–9 are included.

Fig. 8. Power-duration curve of high frequency portion of load power signal using Chebyshev type I high-pass filters of order 3 with several different cutoff frequencies. Displays the power requirement corresponding to the energy requirement of Fig. 8.

VII. RAMP RATE VERSUS ENERGY STORAGE

Another way to visualize the interaction of fast energy storage with traditional thermal units is with of a plot of

VIII. CONCLUSIONS

This paper has demonstrated some of the advantages of dividing the burden of LFC between fast energy storage units and slower traditional thermal generators. The use of fast energy storage can lead to substantial reductions in the ramp rate requirement of the thermal units and thus to reduced costs.

Energy-duration curves and ramp-rate-duration curves are useful tools for evaluating the performance of dispatch methods. Slope-duration curves in particular may also prove useful in other applications which focus on ramping at different time scales, such as for economic dispatch.

REFERENCES

[1] N. Jaleeli, D. N. Ewart, L. H. Fink, L. S. VanSlyck, and A. G. Hoffman, "Understanding automatic generation control," *IEEE Transactions on Power Systems*, vol. 7, pp. 1106–1122, Aug. 1992. A report of the AGC Task Force of the IEEE/PES/PSE/System Control Subcommittee.

[2] Y. V. Makarov, J. Ma, S. Lu, and T. B. Nguyen, "Assessing the value of regulation resources based on their time response characteristics," Tech. Rep. PNNL-17632, Pacific Northwest National Laboratory, Richland, WA, Jun 2008.

[3] S. A. Lefton, P. M. Besuner, and G. P. Grimsrud, "Managing utility power plant assets to economically optimize power plant cycling costs, life, and reliability," in *Fossil Plant Cycling Conference*, pp. 195–208, 1994.

[4] Electric Power Research Institute, I. Gyuk, and S. Eckroad, "EPRI-DOE handbook of energy storage for transmission and distribution applications," Tech. Rep. 1001834, U.S. Department of Energy, Palo Alto, CA and Washington, D.C., Dec 2003.

[5] M. L. Laszarewicz and J. A. Arseneaux, "Status of pilot projects using flywheels for frequency regulation," in *IEEE Power Engineering Society General Meeting*, 2006.

[6] P. A. Taylor, "Update on the Puerto Rico Electric Power Authority's spinning reserve battery system," in *11th Annual Battery Conference on Applications and Advances*, 1996.

[7] B. Parsons, M. Milligan, B. Zavadil, D. Brooks, B. Kirby, K. Dragoon, and J. Caldwell, "Grid impacts of wind power: A summary of recent studies in the United States," *Wind Energy*, vol. 7, pp. 87–108, 2004.

[8] D. C. H. Prowse, "Improvements to a standard automatic generation control filter algorithm," *IEEE Transactions on Power Systems*, vol. 8, pp. 1204–1210, Aug. 1993.

[9] I. A. Erinmez, D. O. Bickers, G. F. Wood, and W. W. Hung, "NGC experience with frequency control in England and Wales—provision of frequency response by generators," in *IEEE Power Engineering Society Winter Meeting*, pp. 590–596, 1999.

[10] B. Roberts and J. McDowall, "Commercial successes in power storage," *IEEE Power and Energy Magazine*, pp. 24–30, March/April 2005.

[11] R. R. Booth, "Power system simulation model based on probability analysis," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-91, pp. 62–69, Jan 1972.

[12] E. Hirst and B. Kirby, "Defining intra- and inter-hour load swings," *IEEE Transactions on Power Systems*, vol. 13, pp. 1379–1385, Nov. 1998.

[13] P. Sørensen, N. A. Cutululis, A. Vigueras-Rodriguez, L. E. Jensen, J. Hjerrild, M. H. Donovan, and H. Madsen, "Power fluctuations from large wind farms," *IEEE Transactions on Power Systems*, vol. 22, pp. 958–965, Aug. 2007.