



Benefits of Fast-Response Storage Devices for System Regulation in ISO Markets



Prepared For: The AES Corporation
By: KEMA Incorporated, Fairfax, VA

June 2008

Copyright © 2008, KEMA, Inc.

The information contained in this document is the exclusive, confidential and proprietary property of KEMA, Inc. and is protected under the trade secret and copyright laws of the U.S. and other international laws, treaties and conventions. No part of this work may be disclosed to any third party or used, reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, without first receiving the express written permission of KEMA, Inc. Except as otherwise noted, all trademarks appearing herein are proprietary to KEMA, Inc.



Table of Contents

1. Introduction	2
2. Simulation Description	5
3. Scenarios Studied.....	7
4. Control Algorithm Development	8
5. System and Storage Performance Metrics	13
6. Analysis of Results.....	14
7. Regulation Revenues.....	17
7.1 Price Duration Curves	18
7.2 Settlements.....	19
7.3 Revenues and Gross Margin.....	19
8. Conclusions – Storage Device Effectiveness for ISO Regulation Services and Economics	20

List of Exhibits:

Exhibit 1: High-level Representation of the Model	4
Exhibit 2: Close-Up of Various Control Signals for Two Hours of a Peak Day	6
Exhibit 3: Regulation and Storage Scenarios for CA, NY, and PJM ISOs	8
Exhibit 4: Behavior of Storage Parameters over a Typical Peak Day - CAISO.....	10
Exhibit 5: Close-Up showing effects of Generation Ramping during Morning Hours.....	11
Exhibit 6: Behavior of Storage Parameters Over 21-hour Period - CAISO	11
Exhibit 7: Close-Up showing effects of Generation Ramping during the Morning Hours	12
Exhibit 8: Storage Settlements for Peak Day - CAISO	14
Exhibit 9: System Performance with Power Spectral Density of ACE as Index.....	16
Exhibit 10: System Performance with Average of Absolute Value of ACE as Index.....	16
Exhibit 11: Storage Unavailability (% of Time).....	17
Exhibit 12: Sample Regulation Price Duration Curve for CAISO	18

Benefits of Fast-Response Storage Devices for System Regulation in ISO Markets

1. Introduction

There has been considerable interest in the application of High Performance Energy Storage Systems (HPESS) for use in system regulation services – responding to Automatic Generation Control (AGC) signals – particularly in an Independent System Operator (ISO) environment where significant revenues can be collected. As plans for rapid and large deployment of renewable generation resources, especially wind, may create a greater requirement for regulation due to the inherent volatility of the renewable resource, advances in storage technologies have led to serious consideration of using fast-acting energy storage to provide regulation service and help to solve the potential problem.

However, prospects of widespread use of storage for regulation raises a number of complex questions about effectiveness and economics. This paper reports on an ambitious attempt to demonstrate the effectiveness of HPESS for AGC regulation service and to explore answers to some of the more compelling questions around how best to apply storage for regulation. One such question is whether HPESS is more effective than conventional generation in meeting regulation service requirements.

For our study, a detailed dynamic simulation of several different ISO systems and their generation resources, including AGC and real time markets, was used to examine:

- the performance of HPESS used for regulation,
- the impact on system performance, control strategies for AGC and storage, and
- the energy market economics of storage in this application.

Based on this modeling, we reached a number of conclusions about the desirable energy-to-power ratio of the storage device and the control schemes used to operate it. We also found that the ISO AGC may need to be adapted to best exploit the new storage technologies. In effect, this will require developing and demonstrating a concept of a "fast regulation service" integrated with existing AGC and real time dispatch paradigms. Finally, we also provided an initial analysis of the overall (and generally favorable) regulation and energy economics of using storage devices for this application.

Problem Statement

When regulation is provided by thermal units, there are costs imposed in terms of lowered heat rate, the foregone opportunity to sell energy with the capacity utilized for regulation, and additional wear and tear on the generating unit. Further, there may be reason to believe that regulation service additionally causes degraded emissions performance. In control areas where regulation is largely sourced from thermal generation, therefore, the use of storage to provide regulation service may help to reduce such costs.

HPESS technologies, such as advanced batteries, flywheels, and compressed air, are all candidates to join pumped hydro as viable storage solutions for bulk power systems. Several organizations have announced plans to deploy advanced storage technologies for various system applications, and at least one is attempting to make a business of using storage to provide regulation services.

A number of questions arise when considering storage for system regulation.

- What capacity and duration is needed for storage devices to respond to persistent calls for regulation service in each direction?
- What effects will large amounts of storage have on overall system performance,
- Can the AGC algorithms be designed to take advantage of the “fast response” capabilities of the emerging storage technologies?

The first two concerns largely revolve around the ratio of energy storage capacity-to-power capacity (storage duration¹ at full power) of the storage device. Understanding how this ratio interacts with device performance and overall system performance when storage is used for regulation is a key to the overall economics of storage for regulation. Another question which strongly drives the overall economics of storage for regulation is how the charge – discharge cycle efficiencies affect the energy consumption of the HPESS in the real time energy markets.

These questions are all interrelated. The ISO AGC algorithm and general system dynamic performance will dictate how the storage device is used for regulation. The storage duration period of the device will influence its performance and overall economic and the control scheme used by the storage operator to maintain a charge state of the device.

In attempting to answer these questions, it was decided to use a detailed dynamic simulation of three different representative ISO environments, including AGC systems and real time markets, to explore the performance of HPESS when used for regulation. The paper focuses one of these areas and provides information on others in the appendix.

Methodology

For this study, KEMA adopted a dynamic system model / simulation of power systems (generation, load, and interconnected control areas), grid operator / utility control systems, and real time dispatch that has been used to assess the impact of wind generation on the national

¹ Storage Duration Times used in the report and subsequent tables assume maximum charge rates, based on the storage device’s maximum power and energy capacities.

system of the Netherlands. This model was used as a basis for conducting this study. The California ISO, (CAISO), PJM RTO, and New York ISO systems were chosen as target markets and systems for representation and analysis. (Note: publicly available data was used for this study. The three regional systems mentioned above did not participate in or review the results.)

The model fully represents power plant dynamics, system frequency response, load frequency response, hourly scheduling, and real time dispatch and AGC in sufficient fidelity for the purposes of demonstrating the effectiveness of storage resources and analyzing the balancing energy requirements of storage inefficiencies. It also is able to demonstrate the overall system performance when using storage in lieu of conventional generation and to help document how the storage resource performs as a regulating resource as a function of the duration or storage time (energy-to-power ratio) of the device.

Exhibit 1 shows an overview of the model at a high level. The storage system is modeled as one large storage resource available for ancillary services at the system level. The storage model also includes provisions for calculating the net real time energy settlements as well as a variety of outputs and statistics for reporting and analysis.




Exhibit 1: High-level Representation of the Model

The storage model incorporates energy and power limitations, parasitic losses, and charging/discharging inefficiency as a function of power level.

A number of alternative ISO AGC schemes were examined to take advantage of one of the characteristics of storage that point the way towards development of a "fast regulation service" protocol. The potential of such developments bear directly and strongly on the economics of storage-supplied regulation, particularly with the energy-to-power capacity ratios (duration times) that are effective with different control regimes.

These issues are discussed along with general conclusions and recommendations in Section 4 of this report.

2. Simulation Description

Our first assessment focused on the California ISO (CAISO), New York ISO (NYISO), and PJM territories. For the scenarios examined, key simulation inputs are the hourly energy schedules for generation and interchange together with the actual load curve for the simulated day. The simulation then adds auto-correlated noise to the actual day load data to generate the second by second simulated loads.

The generation company schedules are derived from the actual load and interchange, with a typical load forecasting error added to the hourly schedules to replicate the real world process. The interchange and generation schedules as used by the AGC and the Generation company subsystems also are "ramped" from one hourly value to the next from 10 minutes before the hour to 10 minutes after the hour.

The Real Time Market (or Balancing Energy) clearing function uses a "stack" of bids / offers for real time energy from the generation companies to dispatch balancing energy every 5 minutes. The Real Time Market or Balancing system looks at the schedules, the schedules at 5 and 10 minutes forward, the actual generation, and determines the needed balancing energy from the generating companies via a market-clearing algorithm. Energy instructions (increments (inc) and decrements (dec)), are generated and issued to the Generating Companies (or "GenCos"). These result in real time generation schedules for the GenCos. There is also a corresponding Real Time Market Clearing Price (MCP). GenCos allocate the inc and dec instructions (increment, decrement) to their units. (Note: regulation services are unit based, not portfolio based in these markets. However, our GenCo allocation of regulation to the units dealt with this.)

The AGC system computes the Area Control Error (ACE) and processes it to determine regulation signals – a real time control signal – for the ancillary service "regulation" which is provided by units who were awarded regulation services. Typical ACE behavior for the CAISO is shown in Exhibit 2. On this plot we also show the Filtered ACE – an internal AGC signal which has been filtered to remove high frequency noise that the system is not expected to respond to – and the AGC Control signal which typically includes an integral term².

² The integral term refers to a component of the control signal which is generated by taking the integral of the control signal. It acts to force the process to return the error signal to zero




Exhibit 2: Close-Up of Various Control Signals for Two Hours of a Peak Day

ACE=Area Control Error as computed according to the NERC standard;

Filtered ACE = filtered version of Area Control Error to reduce high-frequency noise;

Evaluating the Effectiveness of Storage

When we add storage as a regulation resource to the mix, there are a number of questions we wish to answer:

- *Does AGC performance degrade when storage is substituted for traditional generation providing the same amount of regulation services?*

We address this question by comparing ACE metrics such as the average of the absolute value, and by comparing the Power Spectral Density of the ACE signal.

- *Is the storage capacity adequate to the regulation service contracted? That is, does the device ever hit minimum or maximum charge and remain unable to respond for significant periods of time?*

We address this question by examining plots of the charge/discharge and energy storage level of the storage device; and by computing a metric of the percent of time during a period that the device is unable to provide the instructed regulation.

- *What are the energy economics of the storage device?*

We address this question by computing and evaluating the real time energy usage of the device (net in and out) in each 5-minute dispatch interval and computing the storage device energy revenues using the Market Clearing Price (MCP) just as is done for a generator. We later computed the expected regulation services revenues to develop a "gross margin" from operations for the device.

- *Are the remaining generators providing regulation being impacted by the use of storage for part of the regulation service? That is, are they being subject to greater frequency of movement and/or amplitude within their committed regulation bandwidth?*

This question is important because generators price regulation services in terms of their opinion of the opportunity cost of not being able to provide other energy or ancillary services (which can be computed from market data) and their estimate for and decreased efficiency from the unit because it is regulating.

We examine the impact of using storage on remaining generation by examining the AGC control signal – the same metrics as used for ACE can be applied to the control signal. Thus the average of the absolute value of this control signal over a period is a metric, as is the Power Spectral Density. We can also look at the actual generation for movement. However, the regulation movement is dominated on the hourly boundaries by ramping activity as the GenCo schedules go up and down.

3. Scenarios Studied

Three systems were modeled - the California ISO, the PJM RTO, and the NYISO. Each of the three systems was represented as a number of generating companies with a distribution of steam, CCGT, CT, hydro, and other generation sources. The total capacity of each type of resource was arranged to match the total capacities by type in the ISO control area. This approach was used for all territories. There are important differences in these systems.

- The PJM system relies heavily on coal fired steam generation (which is slower dynamically than other generation types) and combustion turbines.
- The California system has no coal units, steam units are fired by natural gas, and many combined cycle combustion turbines as well as significant hydroelectric resources. KEMA received feedback from the ISOs that ACE value was representative of typical signals for PJM and California.
- New York also has significant hydroelectric resources. For purposes of comparison, the load volatility in the New York model was set to be consistent with a moderate ACE so that storage effectiveness in that regime could be seen. Exhibit 3 shows typical characteristics of these three systems, the total amounts of regulation required, and the assumed storage device storage size and control signals.

Scenario	Peak Load (MW)	Installed Coal Capacity (MW)	Installed Steam Capacity (MW)	Installed CCGT Capacity (MW)	Installed CT Capacity (MW)	Installed Hydro Capacity (MW)	Total Generation (MW)	Total Regulation Procured (MW)	Storage Capacity (MW)	Storage Control Signal
California	57,000	0	18,248	19,946	18,248	13,097	74,112	300	100	AGC Control
								300	100	Filtered ACE
								500	100	AGC Control
								500	100	Filtered ACE
New York	32,000	15908	0	7,873	6,072	6,073	42,886	200	50	AGC Control
								200	50	Filtered ACE
PJM	110,000	97281	11,280	22,870	26,042	6,545	164,988	1000	200	AGC Control
								1000	200	Filtered ACE

Exhibit 3: Regulation and Storage Scenarios for CA, NY, and PJM ISOs

Actual peak days from 2007 were used to develop 24-hour load and interchange schedules for use in the evaluation of storage effectiveness and economics. Published actual and forecast load data from the ISO web sites were used in establishing the schedules and the load forecast errors. Because the published data does not include second by second load values – in fact such are only known to the ISO via intertie and generation data – colored noise was added to the actual load data to represent second by second load behavior. The amplitude and frequency or autocorrelation behavior of the load noise was "tuned" so that a base case simulation without storage matched representative ISO ACE behavior in terms of amplitude, excursions, and zero crossing.

The real time prices reported for the peak day and the published bid stacks of inc/dec bids were then used in setting up the simulated inc/dec offers into the real time market. These resulted in the settlements calculation for real time energy usage by the storage device.

In KEMA's analysis, it is possible to run simulations with varying sizes of storage systems. At small sizes the system is basically a "price taker" in the market and the key question is whether it can respond to AGC raise / lower commands. At larger sizes, the device (or aggregation of devices) begins to influence system performance and impact the amount of regulation acquired from conventional generation.

Different scenarios for the size of the storage resource related to the system peak load and system normal regulation volumes procured / utilized were established and simulated. For each scenario, base cases (no storage) were simulated to establish a baseline and then the chosen storage scenario was simulated over a range of energy-to-power ratios or storage duration, ranging from 6 minutes to (i.e. 6 minutes - 10 MWH and 100MW) to 2 hours.

4. Control Algorithm Development

AGC systems typically have two or three important signals in terms of the processing of intertie and frequency values into control signals sent to the plants. The first is the instantaneous Area Control Error (ACE) computation per NERC standards, which is:

$ACE = 10 * \text{Bias} \times (\text{Frequency Error}) + \text{Deviation from Interchange schedule in MW.}$

where Bias is expressed in MW / 0.1 Hz

Because the ACE signal has a certain amount of (relatively) high frequency noise which is below the typical threshold of response and would simply cause unnecessary oscillation of control signals, the ACE is filtered to reduce the ACE to a signal more tuned to what the control area wishes to control. We can call this the filtered or "smoothed" ACE. The filtered or smoothed ACE then goes to a control algorithm that implements a linear PID³ to determine the total control signal or "AGC Control" which is used to control the regulating units.

As a first step in investigating different control strategies for exploiting storage as a regulation resource, the study explored whether to use the ACE, the filtered ACE, or the AGC control signal as the control signal for the storage device.

- The ACE demonstrated an obvious aspect of integrating a very fast storage device - such as a battery. The storage device could follow the noisy ACE signal better than any conventional generator, but to what purpose? It adds cycling duty to no benefit. Therefore, it was not considered further as a storage control signal.
- The filtered ACE was a way to make use of a fast resource as well as to avoid having the storage device participate in load following during ramping as described above. It is an avenue of further investigation for the development of "fast regulation ancillary services" – and can be expected to lead to a different filtering and control algorithm for the purpose. It turns out in simulated practice that using the Filtered ACE signal has significant overall advantages to both the storage device and system performance.
- The AGC control signal has the advantage for the purpose that it represents the signal for which a storage device would be required to respond to under existing market protocols and control regimes. It does not recognize the need for a storage device to be "paid back" by being allowed to restore its energy level to a targeted range, nor is it likely to cross zero during a ramping period. Thus a storage device that is controlled by a derivative of that signal is likely to need a higher energy-to-capacity ratio than one which is controlled by a filtered ACE signal.

³ PID: Acronym for linear combination of the error signal, its integral, and its derivative

Example of Storage Performance for Two Control Strategies

Exhibit 4 shows the storage performance with a storage device of 10 MWH capacity and 100MW power rating. This storage device, responding to the filtered ACE, comes close at times to reaching its stored energy capability but never does actually hit its limits. Towards the end of the 21st hour, the energy level is hitting maximum capability (meaning it has charged, supplying negative regulation, to the maximum) but not quite. Exhibit 5 is a close up of the 6 AM hour when excessive ramping is driving system ACE beyond acceptable limits and the units and the storage device are providing maximum regulation.

Exhibit 7 shows the same case and same storage device, but this time the AGC control signal incorporating the integral term (AGC Control) is used to control the storage device. In this case, the storage device performance is not as good. It can be seen that shortly after 7 AM the storage device charge state is at a limit and the charging power goes to zero and remains there until the AGC asks for energy reversal. Therefore, in this case the storage device would have failed to comply with regulation instructions.

In some markets, this results in the resource being declared ineligible to provide regulation services; in others it results in the resource being financially penalized. Investigating the way that the storage device energy-to-power ratio and AGC control paradigms interact to both effectively regulate ACE and make good use of the storage device resource is a major thrust of this study.




Exhibit 4: Behavior of Storage Parameters over a Typical Peak Day - CAISO
(Control signal is Filtered ACE.)




Exhibit 5: Close-Up showing effects of Generation Ramping during Morning Hours.




Exhibit 6: Behavior of Storage Parameters Over 21-hour Period - CAISO
(Control signal is AGC, i.e., Filtered ACE plus typical AGC control algorithm).




Exhibit 7: Close-Up showing effects of Generation Ramping during the Morning Hours.

The study examined the implications of both of the above control approaches and found that shorter duration times were acceptable when a filtered ACE signal was used as the control but that when the AGC control (as exists today) is used, longer duration time is required. Storage performance for regulation purposes is acceptable in either case, provided that appropriate energy-to-power ratios are utilized.

Another area that was explored was whether overall AGC performance can be improved via any special techniques. One that proved valuable was to "feed forward" the actual storage power level into the AGC control signal; that is, to decrease the signal being sent to the conventional units to reflect the fast response of the storage unit. This was more valuable during "normal" periods than during ramping periods, however. Note that the feed forward effect plus the use of the entire Filtered ACE signal as a storage device control signal implies that for many ACE excursions the storage device will perform as much of the total regulation as it is able – avoiding control to the plants altogether. During ramping periods the integral term in AGC control becomes large and the conventional plants are used as they have to be to match ramping needs.

Another possible scheme is to separate the ACE or the AGC control signal into high and low frequency components. Then the conventional units can respond to the low frequency signal and the storage to the high frequency signal. This seemed to give good results when no large ramping is present. However, when large ramps are present the "filtered ACE with feedforward" as described above gave better performance.

5. System and Storage Performance Metrics

Several metrics are used to assess overall system and storage performance in all the studies described below. All the metrics are computed and shown for 3 hour time periods, dividing the simulated day into seven segments from midnight to 9 PM. This is done because the behaviors in the morning ramp and evening ramp are distinctly different from the behavior at midnight, and of course, load is at peak in the afternoon. With a single metric, that differentiator would be lost.

System Performance Metrics

- Average (Absolute value (ACE)) (ACE is computed as in Equation (1), Section 4). This is a metric that describes the average magnitude of ACE over a time-period.
- Power Spectral Density⁴ of ACE. This metric shows both the total "energy level" in ACE and also the distribution of it across frequency. It is potentially useful as a way to compare simulated ACE to actual. It also shows whether the overall system performance is altered in terms of the frequency characteristics of ACE.

Storage Performance Metrics

Another useful index is a computation of the percent of time when the storage is "unavailable" meaning that it is unable to respond to AGC control because it is fully discharged / charged and cannot provide more power as lower / raise respectively. As might be expected, for higher energy-to-power ratios this metric is essentially zero; for lower ratios and in response to the AGC control (instead of the filtered ACE) it becomes a significant factor.

Storage settlement dollars can be shown in total for a time period or a day or graphically across the day. As might be expected, there will be periods when the storage is "ahead" financially based on net energy delivered/consumed and the relative pricing of that energy, but overall the storage should be expected to "owe" some settlements due to the inefficiencies and losses. *It would be possible to simulate this for different values of storage efficiency but this was not done for this study; this study considered only representative efficiency of a general storage device. Energy settlements will be sensitive to that efficiency*

⁴ Power Spectral Density is describes how the variance of a signal or a time series is distributed with frequency. If $f(t)$ is a finite-energy (square integrable) signal, the spectral density $\Phi(\omega)$ of the signal is the square of the magnitude of the continuous Fourier transform of the signal (here power is taken as the integral of the square of a signal, which is the same as physical energy if the signal is a voltage applied to a 1-ohm load). The actual formula is shown below:

$$\int_{-\infty}^{\infty} |f(t)|^2 dt = \int_{-\infty}^{\infty} \frac{F(\omega) F^*(\omega)}{2\pi}$$

A sample of the storage device settlement economics is shown in Exhibit 8 for two cases that differ only in the AGC control signal applied to the storage device. As can be seen in these plots, there is considerable difference in the storage device economics between these two cases.




Exhibit 8: Storage Settlements for Peak Day - CAISO

In Exhibit 8, the usage of the storage device as a fast regulation service provider and the storage device responding to the current regulation paradigm control signals are shown for storage settlements on a peak day. In this figure, a positive settlements value indicates that the storage device is "paying" the ISO for balancing energy; a negative figure means that the ISO is paying the storage device for balancing energy. These figures show the cumulative net settlements for the period shown. In these two cases, the storage device comes out ahead financially on real time energy overall. That is not always the case; depending upon the relative real time prices of periods when the storage device is asked to charge / discharge, the storage device may be a financial winner or loser. The probability is that the overall settlements should be favorable. However, the simulations demonstrate that net energy profitability of the storage device can vary widely across time-periods and scenarios such that it is hard to develop generalizations.

6. Analysis of Results

For different scenarios, a number of different storage configurations (different storage duration times or energy-to-power ratios) were simulated and assessed. The results from these simulations are summarized in Exhibits 9, 10, and 11.

In Exhibit 9, the performance of the power system with each storage configuration and control option is compared to the base case (no storage). This index is derived from the Power

Spectral Density (PSD) of ACE. The average of the three-hour PSD is computed for each case and then normalized to the base case – so all cases are ultimately summarized as an index where the base case is 1, and values larger than 1 represent a degradation and less than 1 an improvement. Two sets of scenarios were run with the California model – one where the ISO procures 300 MW of regulation total (which is typical) and one where it procures its tariff maximum at 500 MW (more appropriate for a peak day). One case is presented for the New York model. Only one set was run for PJM, representing anecdotally typical conditions for a peak day.

Exhibit 10 shows the same scenarios and cases, but with the index computed using the metric average of the absolute value of ACE which is indicative of the rough "bandwidth" or typical magnitude of ACE. To compare the two metrics – a steady state or "DC" value of ACE of 100 MW would produce a PSD of zero but an average absolute value of 100. A sine wave of amplitude 100 at a specific frequency would produce a PSD of the energy of that sine wave (function of the frequency) and an average absolute of the $100 / \text{sq root of } 2$.

Exhibit 11 shows the same scenarios and cases, but this time the index is the percentage of time that the storage device is "unavailable" - meaning fully charged when it is being asked to supply "down" regulation, or fully discharged when it is asked to supply "up" regulation.

Yellow cells indicate an unacceptable level of storage availability – when the storage is unavailable or unable to respond to AGC signals more than 1 % of the time. We draw several strong conclusions from this data:

- 1) Use of the Filtered ACE signal as a separate control indicator to the storage rather than the AGC control output gives better performance in all scenario comparisons, and better performance than the base case with no storage. Remember that in this scheme, the actual storage output is "fed forward" to the AGC control or subtracted from the control signal sent to the plants. So the AGC is consciously taking advantage of the faster response of the storage in these cases, to apparent advantage.
- 2) A storage duration 6 minutes or 12 minutes is not generally acceptable in many of these scenarios. However, a duration of 30 minutes is generally acceptable which is encouraging for the overall storage economics.
- 3) The relative AGC performance is not necessarily monotonic with the energy-to-power ratio or storage duration times. This is due to the complexities of the storage control algorithm, which is always biasing the storage charging power to restore the charge level to 50% of range at a 10-minute rate in these cases. With different storage duration capacities, this results in different biases and thus different storage response to AGC.

These results indicate that the charge maintenance control needs to be "tuned" to the specifics of the ISO environment, the storage duration capacity, and the regulation signal being utilized.

Another metric could be the NERC CPS1⁵ metric. As noted above, the CPS1 metric is >160 for all time periods in all the various simulations for the California case. Similarly, the other cases also show fully acceptable CPS1 metrics across all the storage applications. And equally importantly, it does not degrade as a result of any storage configurations across all the cases. For that reason the plots are uninteresting and are not presented. This result was obtained for all the scenarios and for all systems simulated. Thus, the use of storage resulted in no degradation of NERC CPS1 performance as simulated.

Scenario	Total Regulation Procured (MW)	Storage Capacity (MW)	Storage Control Signal	Ratio of PSD with Storage to PSD in Base Case					
				2 Hours	1 Hours	30 min	15 min	12 min	6 min
Markets									
California	300	100	AGC Control	0.957	0.864	0.953	-	0.961	0.97
	300	100	Filtered ACE	0.902	0.88	0.898	-	0.929	0.931
	500	100	AGC Control	1.244	1.241	1.236	-	1.238	1.204
	500	100	Filtered ACE	0.809	0.709	0.712	-	0.818	0.893
New York	200	50	AGC Control	1.091	0.954	1.014	-	1.149	0.982
	200	50	Filtered ACE	0.483	0.539	0.495	-	0.613	0.723
PJM	1000	200	AGC Control	0.988	0.963	-	0.892	-	0.897
	1000	200	Filtered ACE	0.725	0.722	-	0.721	-	0.76

Exhibit 9: System Performance with Power Spectral Density of ACE as Index

Scenario	Total Regulation Procured (MW)	Storage Capacity (MW)	Storage Control Signal	Ratio of Average Absolute Value of ACE with Storage to Base Case Average					
				2 Hours	1 Hours	30 min	15 min	12 min	6 min
Markets									
California	300	100	AGC Control	1.110	1.110	1.108	-	1.137	1.125
	300	100	Filtered ACE	0.990	1.003	0.997	-	0.973	1.011
	500	100	AGC Control	1.025	0.964	1.023	-	1.023	1.043
	500	100	Filtered ACE	0.912	0.896	0.904	-	0.896	0.915
New York	200	50	AGC Control	1.226	1.134	1.303	-	1.166	1.150
	200	50	Filtered ACE	0.783	0.822	0.816	-	0.826	0.923
PJM	1000	200	AGC Control	0.978	0.971	-	0.953	-	0.967
	1000	200	Filtered ACE	0.817	0.866	-	0.822	-	0.849

Exhibit 10: System Performance with Average of Absolute Value of ACE as Index

⁵ CPS1 (in %) = 100*(2-(C)*(Δf)*ACE)) where Δf and ACE are computed for each "clock-minute" interval in the period over which CPS1 is being calculated.(normally an hour) and the constant C is chosen for each interconnection based on target frequency bounds

Scenario	Total Regulation Procured (MW)	Storage Capacity (MW)	Storage Control Signal	Storage Unavailability - Percent (%) of Time					
				2 Hours	1 Hours	30 min	15 min	12 min	6 min
Markets									
California	300	100	AGC Control	0.002	0.002	0.002	-	11.14	31.045
	300	100	Filtered ACE	0.013	0.020	0.020	-	0.012	3.737
	500	100	AGC Control	0.007	0.008	0.006	-	0.007	6.301
	500	100	Filtered ACE	0.018	0.001	0.018	-	0.020	0.020
New York	200	50	AGC Control	0.008	0.007	1.40	-	34.75	41.536
	200	50	Filtered ACE	0.034	0.039	0.037	-	0.041	19.712
PJM	1000	200	AGC Control	0.02	0.023	-	0.018	-	13.693
	1000	200	Filtered ACE	0.042	0.012	-	0.012	-	0.015

Exhibit 11: Storage Unavailability (% of Time)

These tables also show an interesting phenomenon. When the storage is "unavailable" a significant percentage of the time the overall AGC performance does not degrade that much if at all from the base case, at least when filtered ACE is used as the control signal. The reason for this is that the feedforward effect, as described above, then automatically assigns the full AGC control signal to the conventional units, up to the amount of regulation procured. Thus the algorithm is largely "self correcting" when the storage capacity is exhausted. This is not the case when the AGC control is used as the signal to the storage device.

7. Regulation Revenues

The storage owners are paid for the *regulation service*, an ancillary service that acts in the California market as a forward option, independent of energy delivered. The successful provider of regulation services bears several costs (which are not borne by a storage device) such as wear and tear, increased fuel costs, emissions, and the opportunity cost of not being able to provide energy or reserves.

Regulation markets are cleared daily for the day-ahead provision of hourly regulation capacity. Successful bidders are all paid the clearing price for the amount of MW (really MW / minute) awarded.

For a storage owner of a successful regulation provider, the regulation services revenue calculation is straightforward: The revenue is calculated by multiplying the hourly price of regulation by the scheduled regulation power. For the CAISO runs, the storage device equals 100 MW in each case.

The CAISO pays for both regulation up and down, it simply uses the cleared price for bids, no compensatory payments for opportunity or other costs. In PJM, regulation payments include an opportunity cost for the inability to supply reserves or energy with the capacity committed to regulation. NYISO also computes compensatory payments.

7.1 Price Duration Curves

The simulations are done for 24 hour periods (peak days). Except for unusual circumstances, the ISO typically buys the same amount of regulation from one day to the next (300 MW normally, in the California case, but up to 500 MW on peak days allowed). This study used a "price duration curve" approach to assess the number of hours in a year for which a given price level for regulation was observed (in 2007) and then estimated the annual revenues by multiplying the integral of the price duration curve, in effect, by the 100 MW.

A sample price duration curve is shown in Exhibit 12.

A separate price duration curve is derived for each month of the year (Weekdays only are reported). Note that on average, regulation services in California are worth about \$9/MW, but vary considerably in the "tail" of higher prices associated with peak conditions when the day ahead spot energy price and reserve price are also high.




Exhibit 12: Sample Regulation Price Duration Curve for CAISO
 RU = Regulation Up (left y-axis); RD = Regulation Down (right y-axis).

This approach allows a straightforward calculation of the revenues that accrue at monthly prices for 100 MW of regulation services. This approach does not take into account the inevitable depression in the price that would occur when an additional 100 MW of regulation capacity is on the market. Presumably, the average price would not drop much as there still would have to be 200MW of traditional sources – but the peak price would come down. Though this case is being shown for the California market, New York and PJM were also run and showed similar results.

7.2 Settlements

Settlement dollars from charging and discharging the storage device were calculated in the KEMA model. When positive, this number represents a cost to the operator. When negative, it is a benefit. Earnings vary across the cases depending upon whether the AGC Control or Filtered ACE is used as the control signal. The duration time of the storage device affects settlements costs as does the total regulation the ISO is employing.

The settlements also vary with the real time energy price differential from one 5-minute dispatch period to another. The interaction that arises from the presence of a storage device, the storage device configuration, and the regulation paradigm is complex and can lead to counterintuitive results. Also, the simulations show that lower storage duration times when AGC Control is used will be "unavailable" or non-responsive for a significant percent of the time. This results in no net settlements effects but would in some markets lead to financial penalties against the regulation services revenues.

7.3 Revenues and Gross Margin

The combination of regulation revenues and settlements revenues / costs result in a gross margin (or gross expense) related to regulation and energy operations. This is only part of the profitability calculation, of course – the capital recovery and all other operating costs have to be considered as well, but this study only focused on whether significant positive revenues can be achieved.

Revenues were calculated using price duration curves shown in Exhibit 12 for four representative months in 2007 - February, May, August and November. Duration curves were made separately for weekdays and weekends. Average prices were determined for each weekday and weekend hour of each month, and multiplied by the 100 MW (amount of storage used in the particular case) capacity of the storage device. Each month's revenue was then multiplied by three (3) to calculate revenue of each of the four seasons. The four seasonal revenues were then added together for the annual revenue.

Based upon this methodology, evaluations were conducted for CAISO and NYISO as those markets vary the real-time price every five (5) minutes. For the CAISO runs and a storage device that equals 100 MW, annual regulation revenues showed a positive value on the order of \$100,000 per MW or more. Similar positive results were seen in NYISO as well.

Though there were differences seen in gross margins for devices with various storage durations, the focus of the study was to show that positives revenues could be generated by used of a fast-response storage device. Hence, the calculations showed that markets could provide the incentive necessary to implement a fast-storage solution.

8. Conclusions – Storage Device Effectiveness for ISO Regulation Services and Economics

As a result of these studies, we can draw a number of conclusions regarding the use of storage devices to provide regional regulation services:

- In some cases, the fast response of a storage device makes it more effective, MW for MW, for AGC purposes than a conventional unit. This is shown by the decreased ACE metrics when the same amount of regulation is provided from a storage device, rather than from a mix of conventional units. What is extremely interesting is that the effectiveness of the storage device increases if the AGC is adapted to take advantage of a fast-acting device that also needs to be "repaid" to restore its energy levels in a finite period of time. The improvement in ACE itself is not the major benefit but the ability to achieve regulation services using storage devices with lower energy-to-power ratios (shorter duration periods) may be very important. The reduced capital cost implied by the lower energy-to-power ratio translates into lower overall costs for regulation services from the storage device.
- Storage devices responding to the filtered ACE signal are more effective than conventional generation at providing AGC regulation, in terms of system ACE performance, when used as described in this paper. This strongly implies that a lesser MW amount of regulation procured from storage and utilized as the primary resource for responding to ACE can be used as compared with the volume of regulation procured today from conventional resources.
- It is clear from the results presented that a storage device with an appropriate duration period energy-to-power capacity ratio is at least as effective as conventional generation for the supply of regulating services. Storage devices with very low duration periods (translating into very low energy-to-power ratios < 0.2 or 12 minutes) may be unsuitable, depending upon the characteristics of ACE in a particular system.
- Depending on the relative extent of storage used for regulation, the use of storage also results in decreased regulation signals to conventional units, which provides benefits in terms of incremental maintenance and possibly reduced adverse emissions effects.
- The energy operating economics of a storage device are favorable at today's price levels for regulation and real time energy. The overall economic effectiveness of a storage device for this application also depends on the capital costs which will vary with different technologies and which are expected to come down over time.