Analysis of Independent Electricity Storage Opportunities on the California Grid

Michael Berger and Eric Munsing

May 9th 2014

Dr. Scott Moura

CE 290:002 Energy Systems & Control

1. ABSTRACT

This project studies the revenue potential associated with using electricity storage system to arbitrage energy prices on the electricity grid operated by the California Independent System Operator (CAISO). A dynamic program (DP) produces optimal charge/discharge patterns for 466 individual pricing nodes on the CAISO grid, assuming operation of the storage system as a price-taker, no discharge power constraints, and no concurrent operation in ancillary services markets. The simulation is analyzed over a range of round-trip efficiencies from 40% to 100%, representing the full range of current storage technologies. The distribution of revenues for these nodes is found to be highly skewed with a long tail of high-value nodes offering revenues twice as high as those of median nodes. This understanding can help policy makers and energy storage companies understand the suitability of specific storage technologies to particular nodes on the grid, and can be used to help forecast the penetration of storage at different price points.

2. INTRODUCTION

2a. Motivation and Background

As wind and solar power become significant sources of electricity generation both globally and in the U.S., grid operators are challenged to balance loads with these intermittent, non-dispatchable resources. While many grid operators are increasing their reserve margin in response to the introduction of intermittent generation, energy storage systems have been highlighted as a way of smoothing generation and matching intermittent renewables to daily variations in price. This is attractive to owners of renewable generation capacity who have to sell their power at off-peak hours or curtail their generation due to transmission constraints, to power consumers who are interested in a reliable source of renewable energy, and to grid operators who are working to ensure that changes in demand can be met by concurrent changes in supply. This shift of power from low-value times to high-value times is characterized by the storage operator's ability to capture revenue from these price differences, and it is this arbitrage value which will ultimately dictate the degree to which energy storage can be used to shift generation.

In addition to the primary energy market, ancillary services markets hold much potential for storage operations, and in the near term there may be more value from storage in the ancillary services market than in the primary energy market. However, if regulators view the shifting of generation and load as an important step to enabling the penetration of renewable generation, it is the intertemporal arbitrage of energy which will ultimately fill this goal.

Current Uses of Storage

95% of current energy storage capacity is found in pumped-hydro storage systems, which are typically constrained by their power throughput and are operated on a diurnal cycle (pumping water uphill during cheap nighttime hours and releasing the water through turbines during peak daytime hours). The limited availability of sites suitable for pumped-hydro storage and the environmental impacts of new pumped-hydro developments has effectively stopped the development of these systems in the United States.

Thermal energy storage (pre-cooling or pre-heating a thermal mass) has received considerable investment as a method for shifting demand, but is not considered here as it does not store electricity and thus cannot participate in energy markets.

Novel electricity storage systems (mechanical, electro-mechanical, and magnetic) have begun to see limited deployment in niche applications. Renewable energy farms are being sited at wind farms to capture energy that otherwise would be curtailed (Ivanpah, Solana) or to smooth out wind supply and provide a “firm” energy supply for a power purchase agreement (AES, Duke). Grid operators are experimenting with energy storage systems in pockets of high demand or supply flux (Tehachapi Battery Storage Project, New York ISO, Portland General Electric). And behind-the-meter energy storage is being explored to help customers avoid paying high time-of-use or block demand pricing (Stem, Tesla/SolarCity). These diverse applications share one characteristic: they are all responding to economics which are not captured on wholesale energy markets, but are rather attempting to capture loss from fringe cases.

Much hope is being held out for the ability of storage to shift large amounts of energy from low-value to high-value times of day in order to facilitate the large-scale integration of renewable energy sources into this electricity system. For this to happen, energy storage will need to become economical not just to prevent curtailment or fulfill specific contractual stipulations, but as part of the wholesale energy market.

This is ultimately what will dictate the large-scale adoption of storage technologies, and enable the temporal shifting of intermittent renewables to high-value times. This idea has currently only been realized by pumped hydro storage, a technology with incredible economies of scale but limited suitability for most applications. This research explores the potential of using other storage systems for energy arbitrage, by assessing where on the California power grid there might be the highest arbitrage value.

2b. Relevant Literature

Policy And System Optimal Operations

Energy storage programs are being considered by both federal and state electricity regulators as an important tool for providing a “smarter” power grid, and a recent ruling by the California Public Utilities Commission set out requirements for utilities to purchase storage contracts by 2020.

Both Sandia 2010 and Sioshani 2012 review potential applications for energy storage, including energy arbitrage, ancillary services, increased reliability, and distribution services. Sandia 2010 forecasts potential market value and capacity for these services based on a heuristic framework for market analysis developed by Eyer 2004, which takes into account the size of the electricity market but does not examine actual control strategies. Sioshani 2012 contrasts work on this potential value with the actual revenue a storage operator can capture in the current energy markets.

NREL 2013 addresses this by simulating the operating costs of a simplified transmission grid with storage, modeled at one-hour intervals using the commercial PLEXOS energy modeling. The study finds that the system benefits of storage on the grid ranged from $25-75/kW but the market value of storage (the value which could be captured by an independent storage operator) was in $7-27/kW, suggesting significant spillovers which can not be captured by the storage operator in the current energy markets.

Multi-Market Optimization

Because the ancillary services market poses a higher potential value of storage than the primary energy market (NREL 2013), a number of studies have examined how to optimally schedule storage to capture opportunities in both ancillary services and primary energy markets. Walawalkar approaches the estimation of regional estimation of storage value by aggregating nodes based on statistical price distributions and evaluating the value of storage using a fixed time horizon. Kazempour Et Al (2009) formulate the multi-market problem as a mixed-integer nonlinear programming problem for a pumped-hydro storage system of fixed size attached to a hydroelectric power plant, and Kirby (2012) approached a similar problem for a pumped-hydro system using average Californian grid prices using a similar mixed-integer approach.

Storage for Firming Renewables

Korpaas et al. explored the market and technology dynamics of combined wind and storage facilities in their 2003 article. Their approach outlines the necessary information to fully model a storage technology and optimize for revenue using a dynamic program. Their system description includes efficiencies, ramp rates, and capacities. The model takes a day's worth of wind speed data as an input and returns the amount of energy stored as well as the daily revenue as a function of time.

Energy Arbitrage Market

The research described here is most similar to Sioshani et al 2009, which analyzes the energy arbitrage value of a price-taking storage system over a number of nodes on the PJM grid, and explores sensitivity to power throughput, efficiency, natural gas prices, and forecast accuracy. Sioshani goes on to relax the the price-taking assumption and explores the potential surplus value (economic welfare) created by storage through its impact on market price, a topic not covered in the current research.

These other studies have typically studied power-limited storage systems, in which the energy reservoir is too large to be filled or dispatched in one time step. While this is typical of large reservoir systems such as pumped hydro and CAES, this may not be true for new storage technologies. This assumption leads to an overstatement of the value per kW of dispatch capability, but an understatement of the value per kWh of energy storage.

The current study focuses on understanding and characterizing the variation in storage value over an electricity grid, where storage operators using modular storage systems may choose their location in order to maximize profits. Rather than seeking to study a specific technology with associated round-trip efficiency and power limits, we are interested in understanding the opportunities which different storage technologies may be able to profitably take advantage of.

2c. Focus of the Study

This study employs a dynamic program (DP) to calculate the optimal revenue for a storage technology at different nodes on the California grid. A less computationally intensive heuristic model is also employed as a validation step. Identifying the highest optimal revenue nodes will indicate the sites at which profitable grid-scale storage installations will first become feasible as storage costs drop.

3. TECHNICAL DESCRIPTION

3a. Energy Price Data

Data on the day-ahead locational marginal price (LMP) of power in the California power grid for calendar year 2013 was collected from the online API of the California Independent System Operator (http://oasis.caiso.com/). The hourly day-ahead market was used, as this market sees the highest volume of contracts, captures both seasonal, diurnal, and hour-to-hour price fluctuations, would be the key market in which energy might be shifted between different times, and its resolution is coarse enough to be computationally tractable for processing 1-year blocks of data.

Data was collected for the hourly pricing at all aggregated pricing nodes (APnodes) on the CAISO grid. These nodes represent interties with other grids, generation buses, and distribution buses. Because these aggregate several underlying nodes (on average 10 individual price nodes for each APnode), they will not show as much volatility or price fluctuation as individual pricing nodes, and thus provide a conservative estimate of the value of storage.

Finer temporal resolution use of data from the hour-ahead (15-minute interval) or real-time market (5-minute interval) would capture slightly more arbitrage opportunities, but the amplitude of the price fluctuations is typically much smaller than diurnal price swings, and so would add little incremental value for the storage operator. Because of the limited degradation associated with rapid battery charge/discharge, it is likely that storage operators would exploit these short-term arbitrage opportunities.

3b. Mathematical Problem Statement

The first step of any optimization problem is to state the goal of the program. This dynamic program will take, as inputs, the round-trip efficiency of a storage system, the storage capacity of the system, and the nodal price data described in section 3a, and will maximize revenue, outputting the optimal charging and discharging pattern for a given time horizon. The parameters and variables used in this analysis are:

k = time index, from 0 to N+1, where N = time horizon

Δt = timestep size [hours]

$P_{grid}(k)$ = Power sent to the battery in timestep k [kW]

$c_{grid}(k)$ = Spot market clearing price [$/kWh]

η = round trip efficiency of storage process

η_{in} = one-way trip efficiency during charging

η_{out} = one-way trip efficiency during discharging

E_{max} = storage system capacity [kWh]

Dynamic Program

A dynamic program was written to capture the anisotropic property of the storage system efficiency. A linear program cannot accurately represent a storage system with a round trip efficiency of less than 100%. The dynamic program maximizes the profits from the energy storage system given market price data, system efficiency, and storage capacity as inputs.

Defining Value Function

The value function is simply the total profits accrued between time step k and the end of the time horizon, N+1, where the profits accrued in timestep N+1 are zero (this is the necessary boundary condition).

$$V(k) = max(cgrid(k)*Pgrid(k) + V(k+1))$$

Constraints

Transmission constraints around the physical node location were not considered, since these will vary in time and can be extremely complex to model. Ramping rates also are not considered for this project. This is because most storage systems have sufficiently large ramp rates, making their inclusion on a 15-minute time interval superfluous. This approximation may be less accurate for storage technologies such as old pumped-hydro plants, which can have ramp rates of a few minutes. The lack of transmission and ramping constraints means that when buying or selling power is optimal, it will either completely charge or completely discharge the storage capacity. With these limitations in mind, the active constraints on the DP can be represented as:

$$\eta out Emin - Egrid(k) \Delta t \leq Pgridk \leq Emax - Egrid(k) \Delta t \; \eta in$$

Finally, the solver iterates backwards in time from k = N+1 to N=0, solving over the two possible states of charge (E_{max} and 0), and the three possible power outputs (Full discharge, no output, and Full charge) to identify the most valuable charge/discharge cycle.

Validation

To check that the DP is behaving as expected, a small sample set of data was run and plotted in Figure 1. The green function is the state of the battery's charge, ranging from 0 to 100 kWh, and the blue function is the price data for this node over the same time period. One can see that the state of charge of the battery jumps from 0 to 100 at periods of locally low prices, and then falls from 100 to 0 during periods of high prices. This bang-bang behavior is expected given the constraints.

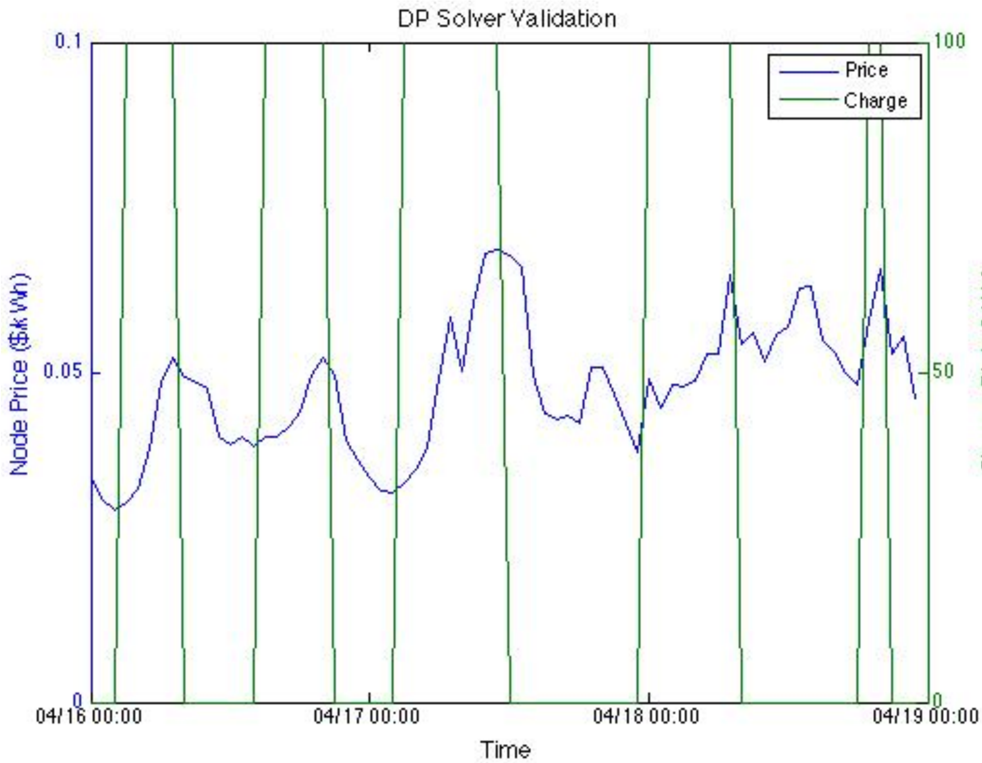

Figure 1: Validation of DP over 3-day test data.

3c. Comparison With Heuristic Algorithm

Under the assumptions which were made in this report, the arbitrage behavior can be captured by a simple heuristic: sell whenever a peak price P_{max} will be followed by a trough of $P_{max}*\eta$ or lower, and buy whenever a minimum price P_{min} will be followed by a rise in price of P_{min}/η or higher.

This algorithm was implemented in Matlab as an additional method for validating the results, and was found to have the benefit of being dramatically faster than the DP solver (taking less than 1/300 of the DP solver time), and produced results identical to those of the DP approach. This demonstrates that the assumption of perfect price foresight is not necessary to reach optimal behavior, and that a day-ahead market forecast would likely be sufficient for a storage operator to capture the full opportunity from diurnal arbitrage.

4. RESULTS

The results of running the solver on all 466 nodes with a round-trip efficiency of 85% are shown in figures 2-4, represented as histograms.

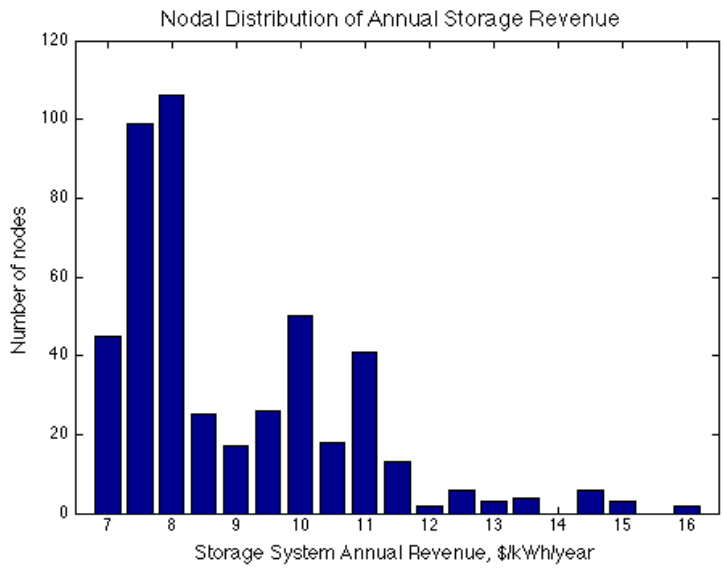

Figure 2:Distribution of Nodal Annual Revenue

Figure 2 shows a skewed distribution for the annual revenue, with some nodes making as much as two times the median revenue. These nodes are important, since they would be the most valuable nodes at which to install storage. The higher-value nodes are likely at areas of high demand or transmission constraints, whereas the strong clustering of low-value nodes likely represents the low congestion on California's grid.

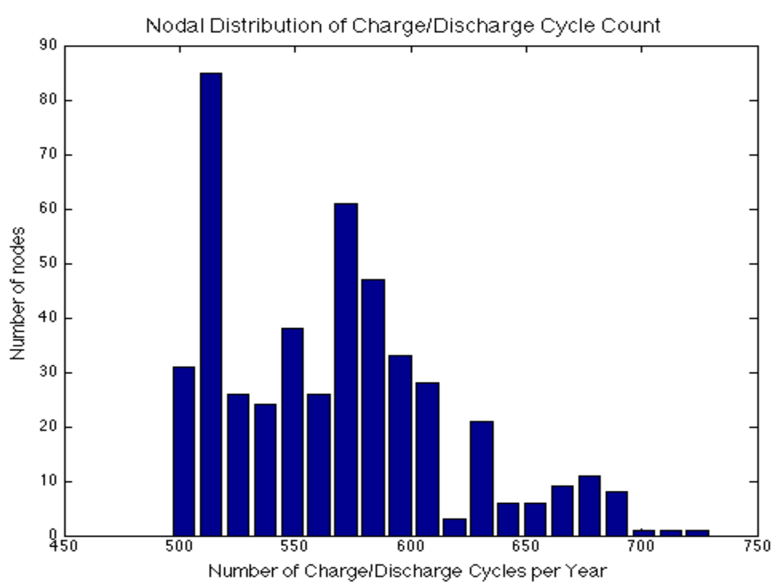

Figure 3: Distribution of cycle count per year.

Figure 3 shows the distribution of how many cycles a node undergoes in its profit maximizing charge/discharge pattern. The number of cycles varies from 500 to ~730 per year. This information could be used to help identify the highest value nodes with the lowest optimal cycles/year, which would reduce wear and tear on the batteries and power electronics.

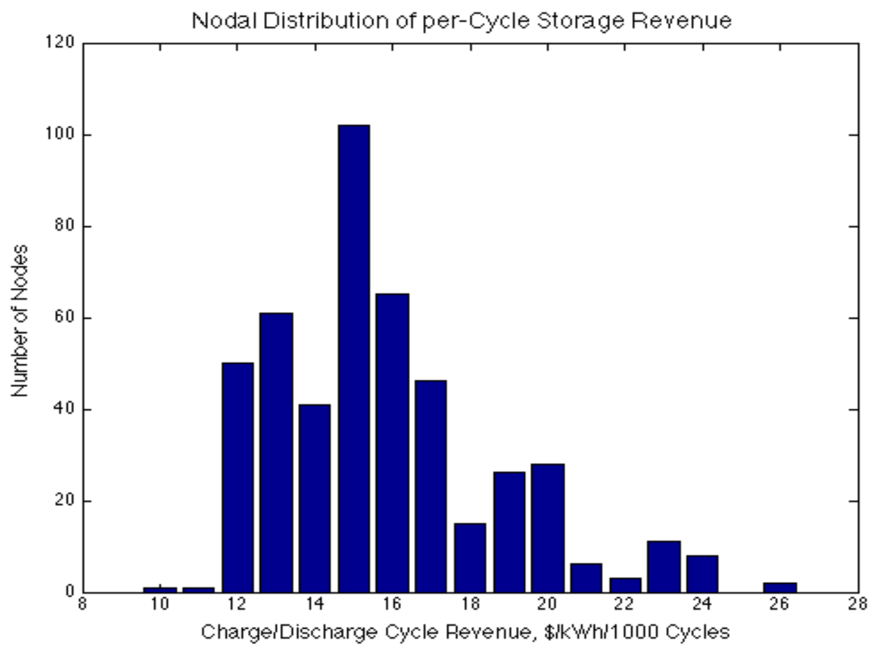

Figure 4: Distribution of annual revenue per 1000 cycles per kWh.

Figure 4 integrates the information from figures 2 and 3, and displays the revenue per 1000 cycles. The nodes with the highest values (to the right on this graph) will have the highest profits, and it is worth noting that again there is a skew to the right, with a few nodes having as much as twice the revenue per 1000 cycles as the grid average. These nodes would be ideal for technologies that are sensitive to the number of cycles.

Sensitivity Analysis

In order to understand how different storage systems might capture value from energy arbitrage, we ran our simulation over a range of round-trip efficiencies, from 40% to 100% round-trip efficiency. As efficiency increases, the storage system has fewer energy losses, increasing revenue. Simultaneously, the storage operator is able to profitably trade smaller price fluctuations and make more transactions, but the incremental value of each transaction goes down because the high-magnitude price swings can still be captured by lower-efficiency designs. These effects combine to create a slightly nonlinear increase in revenue with increasing efficiency.

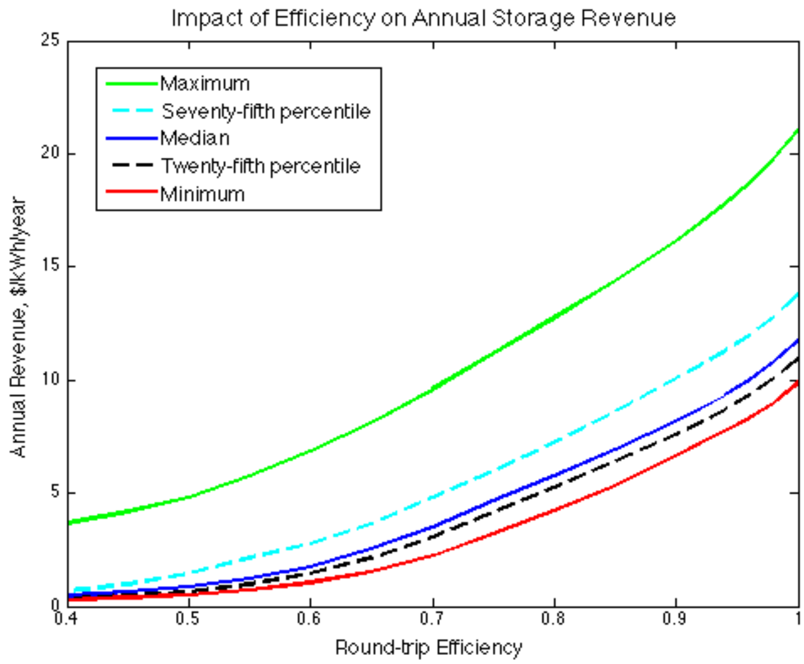

Figure 5: Sensitivity of annual revenue to round-trip efficiency.

The number of charge/discharge systems increases more strongly with increasing efficiency, as the operator takes advantage of smaller hourly price fluctuations. This effect would be much more noticeable if the study had also included hour-ahead and real-time power markets, where high-efficiency systems may be able to take advantage of small but frequent price fluctuations.

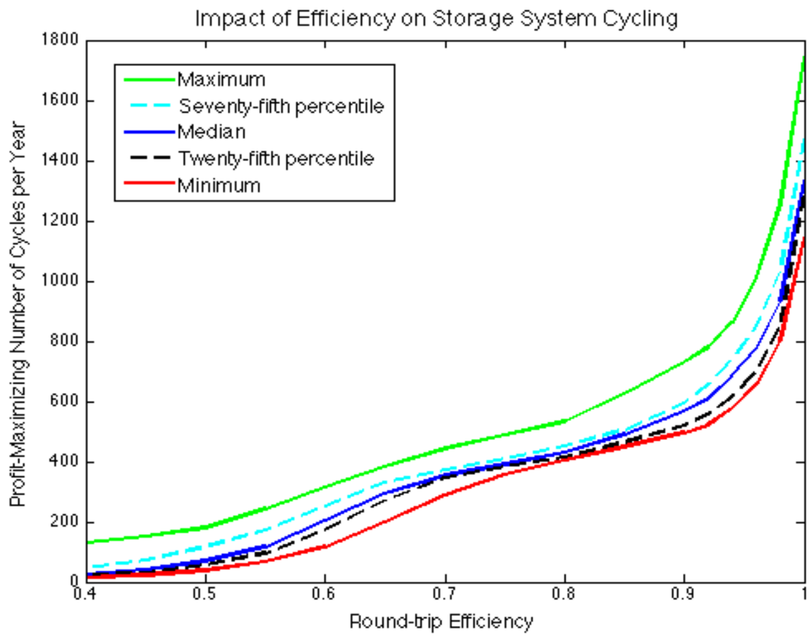

Figure 6: Sensitivity of optimal number of cycles to round-trip efficiency.

These two effects combine to create a general downward trend in the per-cycle revenue of the storage operator, which starts off very high as the operator takes advantage of infrequent but high-magnitude price swings, and tapers off dramatically at high efficiencies as the operator chases the small profits associated with frequent hourly price swings.

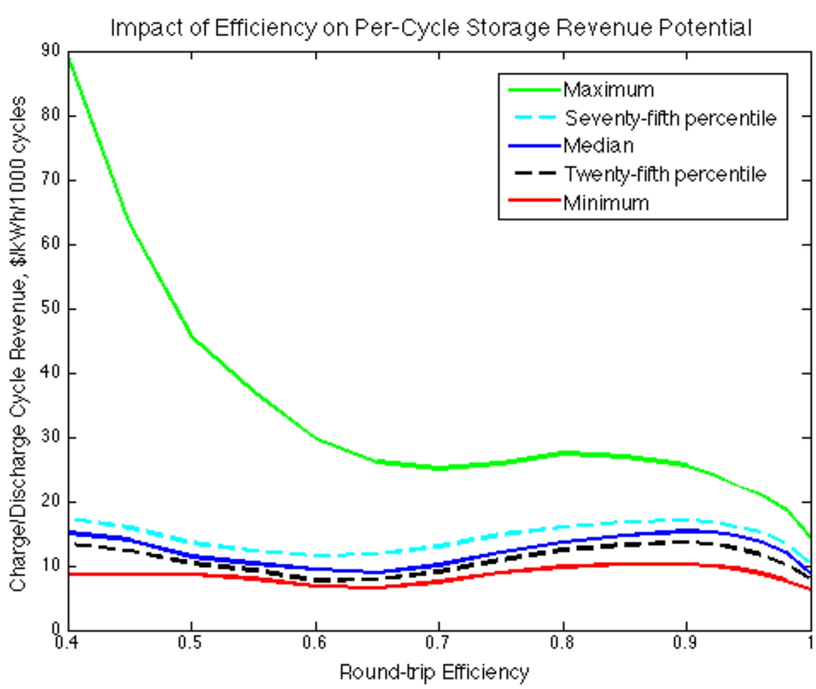

Figure 7: Sensitivity of per-cycle revenue to round-trip efficiency.

There appears to be a “sweet spot” which maximizes median per-cycle revenue, found when the operator assumes 80-90% round-trip efficiency. At this level, the storage system is cycling roughly twice a day, capturing both the morning and afternoon peak prices, but is not pursuing the smaller hourly variations in price.

It should be noted that while these results depict revenue-maximizing behavior, actual profit-maximizing strategies will likely act more conservatively in order to respect technical constraints and operational costs of a given technology (power charging/discharging limits, degradation in battery due to rapid charging, etc.).

High- and Low-Value Nodes

The nodes with highest and lowest revenue potential were examined across the range of efficiencies studied. As efficiency changes, a general reordering of the nodes was observed as the price-maximizing charge pattern changes, though some nodes were found to be outliers over a wide range of efficiencies, and when measured both by annual revenue and revenue per cycle.

These results were interpreted by mapping the node’s ID to a CAISO’s published list of price resources, though only a portion of the nodes can be identified through published data, as the full node mapping is classified.

Table 1: Extreme Value Nodes and Physical Interpretation

High-Value Nodes for Storage:	
CAISO APnode ID	*Interpretation*
VESTAL_2_KERN-APND	Kern River Hydroelectric Unit 3
POD_EASTWD_7_UNIT-APND	Eastwood Pumped-Hydro Storage Facility
BIGCRK_2_EXESWD-APND	Big Creek Hydroelectric Project
GILRPP_1_PL1X2-APND	Gilroy Energy Center Units 1&2
PANDOL_6_UNIT-APND	Delano Energy Company Aggregate Node
Low-Value Nodes for Storage:	
BUCKCK_7_PL1X2-APND	Bucks Creek Aggregate Node

Surprisingly, a number of the nodes that show the highest value of storage were associated with large hydroelectric facilities in relatively remote areas (Eastwood Pumped Hydro Storage Facility on Shaver Lake, the Big Creek Hydroelectric complex around Huntington and Shaver lakes, and the Kern River generator near Kernville). For all of these nodes, price repeatedly approaches $1000/MWh, and even during low-demand season exhibits relatively large daily price swings, making storage valuable. These facilities are all operated by Southern California Edison company, and may be operated counter to profit-maximizing behavior to meet other stakeholder needs (optimizing operations of a full hydroelectric network, minimum water flow regulations, requirements to supply water to irrigators, or operated to provide ancillary services rather than generation). 2013 was also a very low water year in Southern California, and it is possible that these high prices would not be replicated with higher reservoir levels. However, it is interesting to note that one of the nodes with lowest revenue potential was associated with the Bucks Creek hydroelectric project, a large hydroelectric complex in Northern California operated by Pacific Gas & Electric.

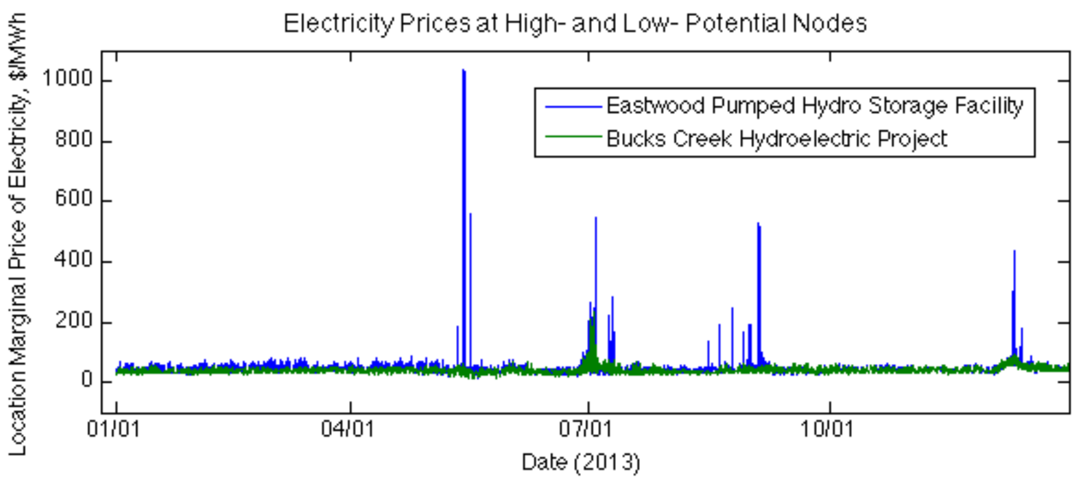

Figure 8: Example electricity prices at high- and low-value nodes.

5. DISCUSSION

This study identifies the revenue potential for independent storage operations on the Californian grid, and as such can be used to guide the identification of locations and technologies that can profitably arbitrage power on the wholesale market. If energy storage is to help shift intermittent renewable generation from low-value times to times of peak demand in order to reduce the use of fossil fuel generators, it will be necessary for energy arbitrage to be profitable on the wholesale energy market.

A comparison with other studies of energy arbitrage value in the California market shows that our study produces slightly higher values than those found in previous literature. Eyer assessed storage value based only on expected growth in demand, and not in replacement of current generation demand, and did not optimize bid strategy using actual price data. Both Byne and Silva-Monroy's study was most similar to the current study, examining a relatively fast-discharging system with high capacity, and producing results consistent with those of the current study. Kirby studied pumped hydro power stations with limited discharge capacity, which leads to a lower value of arbitrage, and may have studied a low-value node.

Study	Annual Value	Value/kWh/year	Assumptions
Eyer et al 2004	$49/kW	$4.90	10 hour, 90% efficiency
Byne and Silva-Monroy 2012	$25-41/kW	$6-10	4 hour, 90% efficiency
Kirby 2012	$46/kW	$2.9	16 hour, 75% efficiency

In our study, we find that there is a "long tail" of nodes with high revenue potentials from storage, and it is at these nodes where storage facilities will first be profitable. As the costs of storage technologies decrease, the number of nodes which are profitable will rapidly increase, likely creating a "boom" in storage development, particularly if development of storage systems results in cost reductions through learning-by-doing or economies of scale in production.

Comparison with Current Storage Costs

Table 2: Summary of Current Storage Technology Characteristics

Technology	Reference Year	Round-Trip Efficiency (%)	Storage System Installed ($/kWh)	Projected Cycle Life (Cycles)	$/kWh/ 1000 Cycles
Advanced Lead-Acid Battery	2011[1]	80	330	2000	165
Sodium-Sulfur	2011[1]	75	350	3000	117
Lead-Acid with Carbon Electrodes	2011[1]	75	330	20000	17
Zn/Br Batteries	2011[1]	70	400	3000	133
Vanadium Redox batteries	2011[1]	65	600	5000	120
Lithium-Ion Batteries	2011[1]	85	600	4000	150
Compressed Air Storage	2011[1]	70	5	25000	0.2
Pumped Hydro	2011[1]	85	75	25000	3
Flywheels	2011[1]	95	1600	25000	64
Supercapacitors	2011[1]	95	10000	25000	400

DOE Near-Term Target	2013[2]	75	250	4000	63
DOE Long-Term target	2013[2]	80	150	5000	30
Model S (estimate)	2014[3]	85	260	4000	65
Model E (forecast)	2014[4]	90	166	4000	42

[1]"Energy Storage Systems Cost Update", Sandia National Labs, 2011

[2]"Grid Energy Storage", Department of Energy Report, December 2013

[3]http://insideevs.com/tesla-battery-in-the-model-s-costs-less-than-a-quarter-of-the-car-in-most-cases/

[4]http://cleantechnica.com/2014/03/04/flawed-study-exaggerates-cost-tesla-model-e/

The revenues presented in our study are represented in $/kWh, but it is important to note that the costs of an actual storage system would include not just storage capacity but also other costs which would not scale with capacity: power rectifying and transmission equipment, facilities, and overhead- which could combine to cost as much as the storage capacity itself (Sandia 2011). A profitable storage operator would need to be able to procure all of these for a levelized cost (in $/kWh/year or $/kWh/cycle) less than the revenue potentials outlined in our analysis.

Technological Prognosis

Pumped Hydro and underground CAES systems are able to capture great economies of scale, and thus CAES may become the first new storage technology to be widely economical. The current price of batteries makes storage uneconomical over the charge lifetimes listed above for all technologies except lead-acid batteries with carbon electrodes. Widescale application of energy storage may require significant technological innovation to lower costs, increasing battery lifetime, or intelligently control charge patterns to minimize battery degradation.

While not studied explicitly in this paper, we can expect that the efficiency requirements required for trading short-term fluctuations in energy prices on the hour-ahead or real-time markets will require the use of flywheels, ultracapacitors, or other high-efficiency, long cycle-life systems. Fortunately, the short duration of these trades should not lead to a significant cost from the leakage that characterizes these systems.

Limitations:

Because this study does not consider power limits, internal losses, or costs associated with particular charging patterns, the results of this study cannot be used to evaluate the suitability of a specific storage technology for deployment on a generation node.

As the elasticity of node marginal price with respect to power supply is not considered, the current approach cannot evaluate the total value that could be extracted by a storage operator.

Previous studies have found that the ancillary services market offers a higher marginal value for services provided by storage, but this has not been considered here because the authors were most interested in understanding the price at which storage could be economically used to shift generation.

The assumption of perfect market foresight, necessary to be able to solve the full year using a linear programming approach, should be largely appropriate in the day-ahead market but would need to be revised for shorter-term markets. This assumption could be relaxed in ??

Future Steps

The full value of storage could be better understood by simultaneously optimizing bidding in the ancillary services and primary energy markets, which could be approached using a mixed-integer nonlinear programming or trough an expanded DP. Modeling the benefits of participating in the hour-ahead and real-time markets could be approached through use of a stochastic dynamic program.

The dynamic programming approach used here will eventually allow for the relaxation of some of the other assumptions, particularly the elasticity of the energy price, linear costs of storage, lack of power constraints, and nonlinear operational costs associated with impact of the charging pattern on the battery health.

We would expect that as these additional costs become integrated, the battery will more often operate with partial reserves, as its ability to take full advantage of arbitrage opportunities becomes limited and the operator seeks to hedge profits in the face of uncertain outcomes.

6. SUMMARY

Using actual price data for 466 aggregated price nodes on the California electricity grid, we utilized a dynamic program to find revenue potentials in the range of $7-17/kWh/year or $10-27/kWh/1000 cycles, with a long tail of high-value nodes which are of significantly higher value than have been explored in previous studies. By calculating revenue potential over a variety of nodes, we are able to provide both more accurate estimates of storage value, and guide policy decisions towards supporting storage systems which are well-suited for a particular application. A sensitivity analysis was performed by sweeping the round-trip efficiency of the storage technology over a range of 40-100% and measuring the changes to annual revenue, optimal annual cycles, and revenue per cycle.

Looking forward, there are two major subjects to address. The first is that this analysis did not include any revenue streams from ancillary services markets, which can be a major contributor to a storage facility’s profits. These markets could be addressed in a similar manner to what was presented in this paper and synthesized to produce a more robust and accurate assessment of the true break-even cost of various storage technologies. The second is that the analysis has provided useful information about the storage opportunities at nodes in the grid, however there is no mapping of node IDs to physical locations, which could be informative to both future investment in storage facilities and the academic literature.

7. REFERENCES

Denholm, P., Jorgenson, J., Hummon, M., Jenkin, T., Palchak, D., Kirby, B., Ma, O., and O'Malley, M. "The Value of Energy Storage for Grid Applications" National Renewable Energy Laboratory. May 2013.

Eyer, J., and Corey, G. "Energy Storage for the Electricity Grid: Benefits and Market Potential Assessment Guide" Sandia National Laboratories Report 2010-0815. Feb. 2010

Eyer, J. M., Iannucci, J. J., Corey, G. P., December 2004. Energy storage benefits and market analysis handbook. Tech. Rep. SAND2004-6177, Sandia National Laboratories.

Kazempoour, S.J; Hosseinpour, Majid; Moghaddam, Mohsen Parsa: "Self-Scheduling of a Joint Hydro and Pumped-Storage Plant in Energy, Spinning Reserve, and Regulation Markets" IEEE Power & Energy Society General Meeting, Calgary AB July 2009.

Kirby, Brendan: "Co-optimizing Energy and Ancillary Services from Energy Limited Hydro and Pumped Storage Plants" HydroVision July 2012

Korpaas, M., Holen, A., and Hildrum, R., "Operation and sizing of energy storage for wind power plants in a market system" 2003.

Sioshani 2010: Sioshani, Ramteen; Denholm, Paul; Jenkin, Thomas: "Market and Policy Barriers to Deployment of Energy Storage", Economics of energy and Environmental Policy 2012

Sioshani, Ramteen; Denholm, Paul; Jenkin, Thomas; Weiss, Jurgen: "Estimating the value of electricity storage in PJM: Arbitrage and some welfare effects", Energy Economics 31 (2009) 269-277

Walawalkar, Rahul; Apt, Jay, Mancini, Rick: "Economics of Energy Storage for Energy Arbitrage and Regulation in New York", Energy Policy 35 (2007) 2558-2568

8. APPENDICIES

Appendix A: Wind Power and Storage

We entered this project with an interest in understanding the application of energy storage used in conjunction with a wind farm, an area where most storage system deployments to date have been made (Tehachapi storage project, Duke advanced lead-acid wind storage project, AES wind storage project, etc.).

To this end, we collected data on the power produced by a large (60-turbine) wind farm near Mojave, CA and located close to the large Alta wind development. Power output for each turbine at 10-minute intervals was received from the manufacturer of the turbines, which collects the data for maintenance and management purposes. The power production data is collected over SCADA communications, which were found to be rather unreliable: only about 51% of all data points were available, and there were long periods (i.e. the entirety of June 25-June 29th) during which turbines would not report their data. Further, there was significant variability in the wind production data, particularly at low power output levels- i.e. when a few turbines at the farm might be turning, but others are stationary or swaying listlessly.

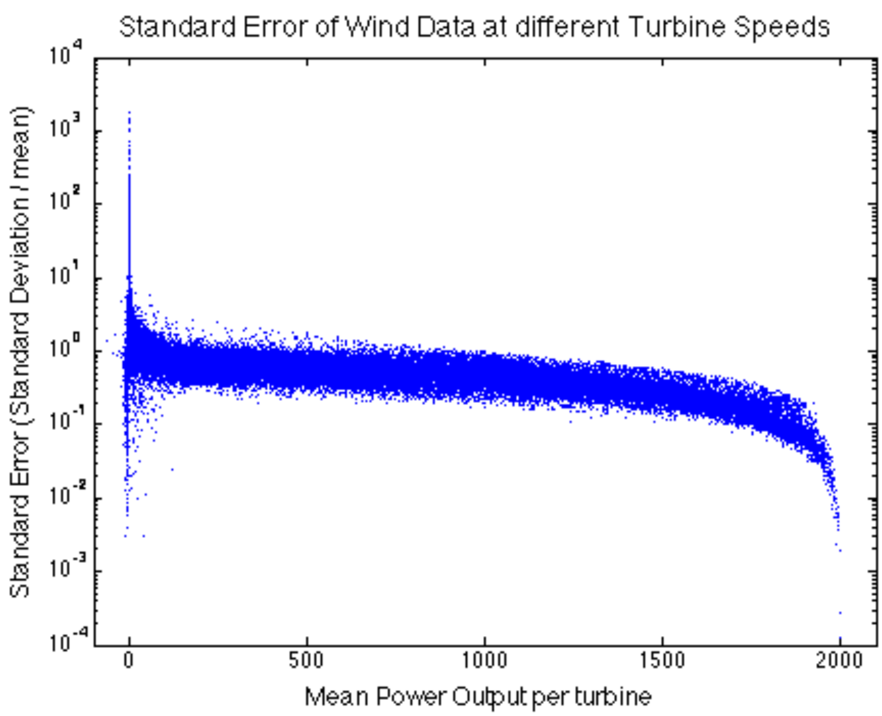

Figure 9: Error in wind data as turbine output increases. Note the high uncertainly at low output.

We characterized this potential error by looking at the cumulative energy production error over the course of the year, yielding a cumulative error graph (see below).

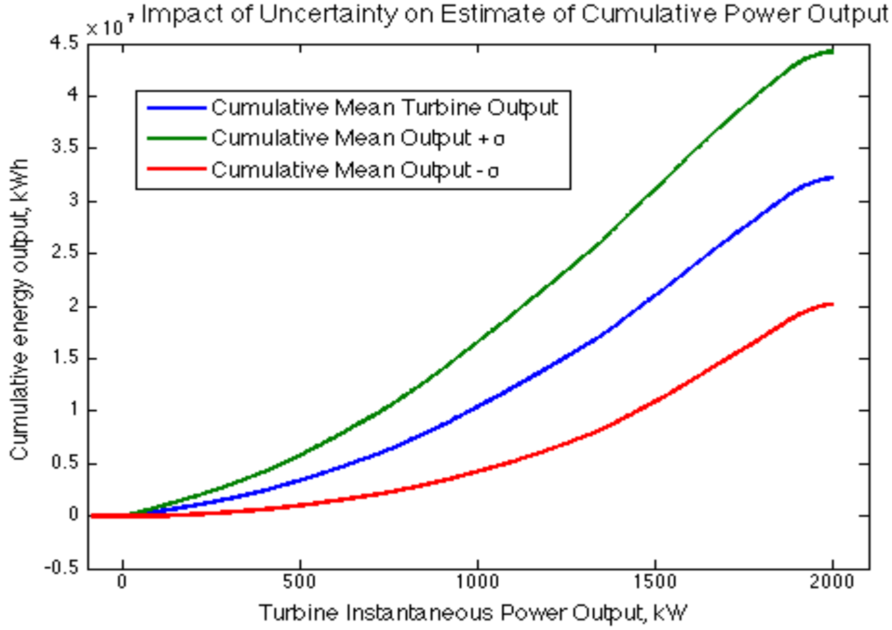

Figure 10: Cumulative impact of uncertainty on power output.

While this cautions against the accuracy of using the mean production data given our available dataset, we moved forwards with a mean wind production value until we could better formulate the problem.

We approached the problem by modeling a simple power system problem in the wind generation and electricity price were external inputs, storage capacity was fixed, and the energy dispatch (and implicitly the amount of energy in storage) was our control variable. Solving this as a linear program, we were surprised to find that the pattern of energy storage was dependent on locational marginal price alone, and not on the wind production. To maximize profits, the generation operator would fill his storage reservoir when power was cheap, and sell the stored energy when power was expensive, regardless of the production of the wind farm.

Upon further inspection, we came to realize that unless power production is being curtailed (i.e. turbines are being shut off or power is being grounded) due to severe congestion constraints, there is no need to treat storage and wind production as linked systems in order to maximize profit. The example below illustrates this:

Nighttime LMP: \$20/MWh. 1 MWh wind produced

Daytime LMP: \$40/MWh. 0 MWh wind produced

Scenario 1: Wind and storage integrated.

Nighttime: Store 1 MWh

Daytime: Sell 1 MWh

Net System Profits: \$40

Scenario 2: Wind and storage separated

Wind Operator:

Nighttime: Sell 1 MWh for \$20/MWh

Daytime: Sell 0 MWh for \$40/MWh

Profits: \$20

Storage Operator:

Nighttime: Buy 1 MWh for \$20/MWh

Daytime: Sell 1 MWh for \$40/MWh

Profits: \$20

Net system profits: \$40

As long as there is an efficient market for power, there is no difference in strategy between the storage system integrated with a renewable energy operator, and an independent storage operator buying energy from the grid. In the integrated case, the market price of power is the opportunity cost of the storage operator, and so there should be no change in behavior patterns- even if the price of energy goes negative.

However, if production is being curtailed due to severe transmission constraints, then the market is not operating efficiently, and the wind farm would be throwing away a resource which could be stored for future use- this is sometimes the case in highly transmission-constrained regions, such as Northeastern Texas, some desert solar projects, and in future Californian grid scenarios where ramping constraints limit the ability of solar to feed into the grid. Further, storage may make sense if an operator is not just optimizing profits from energy sales, but is also interested in other profits- i.e. if a wind farm is able to achieve a higher price in a power purchase agreement by offering “firm” rather than intermittent power, if the wind operator is seeking to reduce variability in profits (profitability of storage and profitability of wind are negatively correlated, as shown above), or if the wind operator is facing other costs associated with the intermittency they are introducing to the grid. However, data on the incidence and value of these cases (i.e. data on wind curtailment, models for the cost of intermittency, etc.) are not readily available, and we chose to continue to focus on the value of simple energy arbitrage, and thus of our study of simple energy arbitrage we have not considered these fringe cases, although they have led to the initial energy storage projects in the U.S. Based on this decision, we removed wind energy production from our problem formulation.

Appendix B: Price-finding by Binary Search

When formulating our problem to find the optimal size of storage for a wind farm at a given cost of storage, we found that the optimal solution was always to either build the maximum size of storage possible within the transmission constraints, or build no storage at all. We recognized that the cost of storage which separated these two prices would be the break-even price for the storage operator, below which the storage system would be profitable and should be built as large as possible to maximize profits, or above which the system would not be profitable and the investment would not be made.

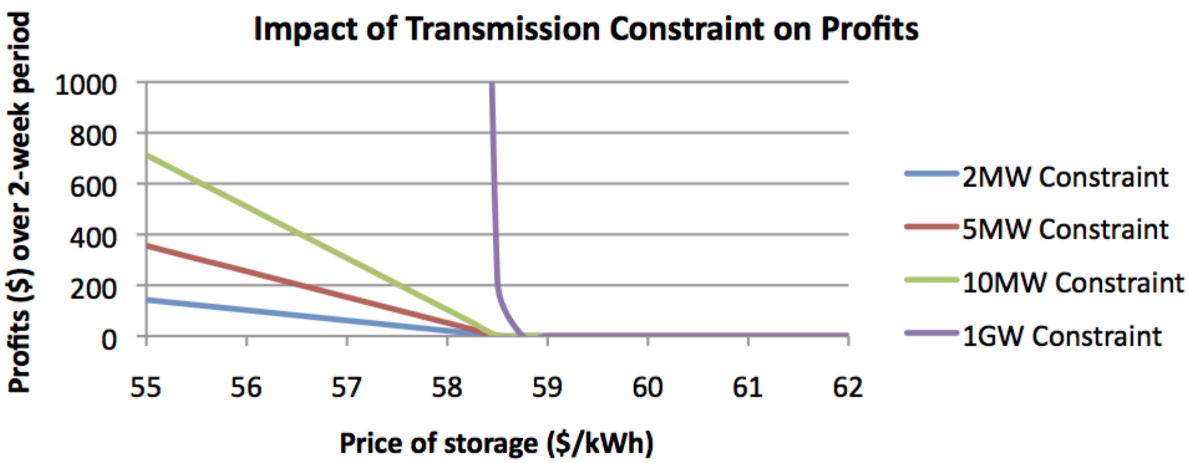

Figure 11: Visualization of profits as transmission constraints are loosened.

Still using the same LP formulation, we developed a wrapper function that conducted a binary search over a range of storage prices to identify the optimal price of storage, as follows:

1) The algorithm first evaluates the storage profits at the midpoint of the price range.

2) If storage was profitable at this price, then the price is too low, and the algorithm recursively calls itself on the upper half of its price range.

3) If the storage is not profitable at this price, the price is too high and the algorithm recursively calls itself on the lower half of its price range

4) If the price range being studied is below a target price resolution, then the algorithm has found the approximate solution and quits.

The algorithm was slow, requiring 9 evaluations of the linear program to find a solution price to within $0.50/kWh, which resulted in a solution time of ~15 minutes/node.

We realized this could be simplified by fixing the amount of storage and then computing the revenue of the system. The break-even price is then the price of storage that neutralizes these profits.

Appendix C: Linear Program

This project began as a linear program (LP) that would incorporate wind production data, market price data, storage costs, and storage efficiencies. As the project moved forward, it was identified that the wind production data was irrelevant for identifying the correct storage cost, as mentioned in Appendix A. It was also discovered that the anisotropic nature of the charging and discharging efficiencies of the storage system makes an LP solution incorrect for any efficiencies other than 100%. It was this reason that caused us to move to a DP that can solve the problem despite the anisotropy. This all being said, the initial formulation of our LP can be found below and is correct if n = 1.0.

Cost Function

The cost function being minimized for the LP consists of the sum of the revenue from all time periods where the system is buying or selling power, as well as a amortized unit cost of storage multiplied by the optimal storage capacity.

$min: -\sum_{i=0}^{N-1} P_i C_i + C_{stor} E_{max}$

The decision variables are E(i) from 0 to N, P(i) from 0 to N,-1 and Emax.

$$x = \begin{bmatrix} E_0 \\ E_1 \\ \vdots \\ E_N \\ - \\ P_0 \\ P_1 \\ \vdots \\ P_{N-1} \\ - \\ E_{max} \end{bmatrix}$$

Constraints

The constraints on this problem include both equality and inequality constraints.

$$E(i) \le E_{max} \text{ for all } i = 0,\dots N$$

$$E(i) \ge 0 \text{ for all } i = 0,\dots N$$

$$E(k+1) = E(k) - P_g(k) \text{ all } k = 0,\dots N-1$$

$$0.95\,E(0) \le E(N) \le 1.05\,E(0)$$

$$P_g(k) \le P_{max}$$

$$P_g(k) \ge -P_{max}$$

The equality constraints ($A_{eq}x = b_{eq}$) can be represented in matrix form as such:

$100\cdots|000\cdots|-0.5-110\cdots|nT00\cdots|00-11\cdots|0nT0\cdots|000-1\ddots|00nT\cdots|0 \vdots \vdots \vdots \ddots| \vdots \vdots \vdots \ddots| \vdots x=$

$nTPwindE0$

The inequality constraints (Ax = b) can be represented in matrix form as such:

$100\cdots|000\cdots|-1010\cdots|000\cdots|-1001\cdots|000\cdots|-1 \vdots \vdots \vdots \ddots| \vdots \vdots \vdots \ddots| \vdots x\leq 000 \vdots$