Battery Storage Optimization and Deployment

Anthony Duer, Applied Energy Group

ABSTRACT

Energy storage has long been a part of many utilities' energy procurement strategies. However, unlike a typical power plant where power is produced at the same time it is consumed, energy storage facilities take in energy when there is a lull in demand and then discharge that energy when demand is high. While worldwide grid connected energy storage capacity is currently fairly small, it is poised to grow quickly over the coming decade. Information Handling Services (IHS) recently predicted that energy storage capacity would grow from 340 MW in 2013 to 6 GW in 2017 and exceed 40 GW by 2022 (Wilkinson, 2014) – a compound growth rate over 70%. There are a number of technologies available to meet this expected growth, such as pumped hydro and thermal storage, but the one technology receiving the most attention recently is electro chemical, or battery, storage. In addition, there has been a growing trend of deploying storage 'behind the meter' in customers' homes and buildings. With this shift, however, the perspective of the benefits and costs suddenly shift from the utility to the customer, who will want to use the batteries to maximize their benefit. However, with this change in perspective comes new challenges such as noisier load profiles and lower accuracy load forecasts that prove a challenge to systems trying to maximize demand reduction. This paper explores setting up a multivariate regression model for a building and then builds discharge strategies that can be utilized to maximize return on available battery storage.

Introduction

Spurred by the automotive industry's adoption of hybrid and fully electric cars, considerable research has been dedicated to making batteries smaller, lighter, cheaper, and capable of storing more energy. This has resulted in, approximately, a 6% improvement in energy density per year as old chemistries are reformulated and new chemistries are discovered and commercialized (Anderson, 2009). Currently, lithium-ion batteries are the leader in market share, but new chemistries, such as lithium-sulphur and lithium-oxygen, are in development that may offer much higher energy densities than those available today (Van Noorden, 2014).

When it comes to the energy industry, batteries are hardly a new thing. However, only recently has interest in building scale, rather than utility scale, solutions really started to take hold in the minds of consumers and investors. With Tesla beginning to sell, and immediately sell out, of 3.3 and 6.4 kWh home-scale batteries and development continuing on 100 kWh commercial-sized 'powerpacks' capable of scaling 'infinitely', battery storage is beginning to come to the forefront (Tesla, 2016). With this shift, however, incentives and conditions are suddenly changing. From an operational perspective, the owner of the battery is no longer as concerned with things such as reserve capacity or grid stability. The primary concern is extracting as much financial value from the battery as possible before it begins to degrade. Likewise, the battery management system must become much smarter as the deployment strategy no longer has a nice, steady demand curve to work with but must now, instead, deal with a potentially erratic load as various pieces of equipment at the site cycles on and off.

The following figure best illustrates this concept. The figure shows three load profiles for what should, in theory, be the hardest day of the year for a utility: the annual peak day. In this case, for the California Independent System Operator (CAISO), that day occurred on September 15th, 2014 where total demand peaked at over 45 GW of instantaneous power (CAISO 2016). The three load profiles shown are that of the CAISO, Pacific Gas & Electric's (PG&E) service territory (covering most of northern and central California), and the building of interest, which will be discussed below:

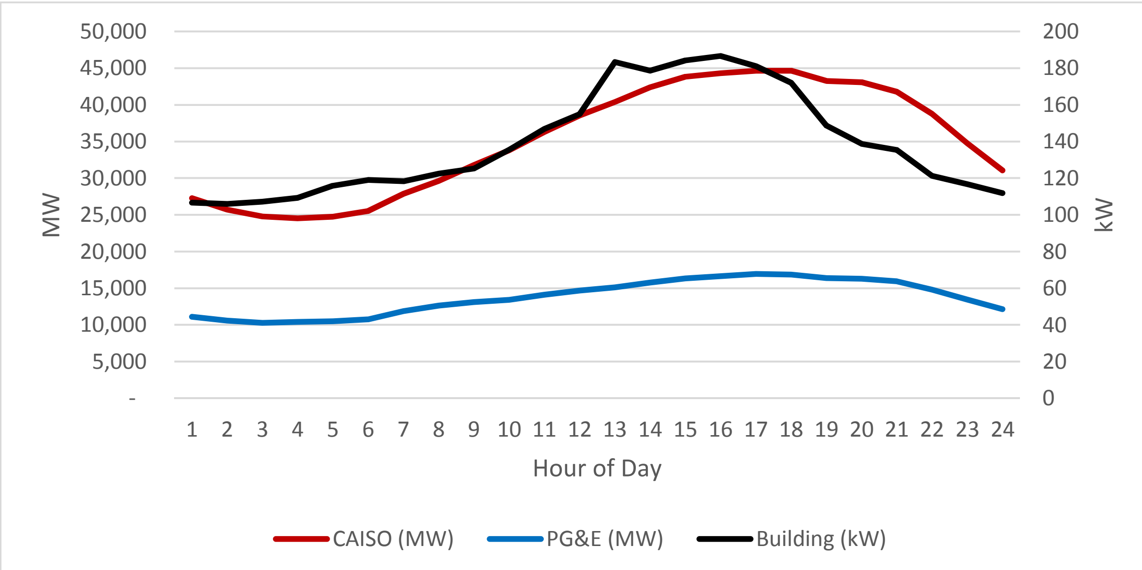

Figure 1. CAISO, PG&E, and Building Peak Day Profiles

The graph illustrates the concept core to the paper: as the scope of usage increases, the load profile tends to be smoother as customer load profiles average together and random effects within one location cancel out random effects in another location. This makes for easier deployment scheduling as demand is more predictable. However, with building level storage, this comfort is taken away and the model must then deal with a load profile that will tend to be much noisier than a model can compensate for. This is critically important as demand charges are a ratchet charge: once a value is exceeded, the charge is set for that month. Therefore, the model must leave room for forecasting error as getting things wrong once potentially negates savings for the entire month.

The Building

When beginning any modeling exercise, it's generally best to establish what data are wanted, what data are needed, and what data are available. Rarely do the three overlap. For this exercise, the data that ended up being used was a surprisingly short list:

- Interval Data
- Hourly weather data
- Tariff Data

These three pieces of data gave enough information to bootstrap the model and begin development. Along the way, different abstractions of this data (such as cooling degree days) were added and any inconsistencies within the data were smoothed out[1]. Once suitable datasets were located, the modeling process began and the three options discussed below were developed. However, first, a brief discussion of the data used:

Interval Data

The interval data was the first piece of data that was sought out. This ended up being the most difficult part as the data had to be publicly available and of a long enough duration to model a full year. Thankfully, Lawrence Berkeley National Lab published (OpenEI 2015) the electric (and gas) interval data for their Building 74 – a 45,382 square foot LEED Platinum research facility. This facility has a fairly normal load profile: a morning ramp to a midday peak to an evening ramp to a lull in demand overnight. The interval data itself was fairly high quality 15 minute interval data. The data provided covered the entirety of 2014 and the first half of 2015. The analysis and modeling period selected was calendar 2014.

Weather Data

Once the interval data and, thus, the location of the facility was established, the next step involved getting the weather data for the facility. For that, NOAA's National Climatic Data Center (NCDC) was utilized. This service provides historical weather data for various weather stations around the globe. For the facility, Downtown Oakland[2] was selected as the closest location to pull weather data from. In a real world scenario, a building would likely be able to utilize its own historical weather data to improve results.

The data received from NCDC was found to be of high quality – NOAA puts the data through various screens and the flags any erroneous results so they can be eliminated. Since the data is fairly raw, however, some additional modeling work was required to align it with the interval data. Specifically, since the reads weren't aligned on the hour, they had to be prorated for each affected hour and then summed to provide a weighted-average temperature for any given hour in the analysis year (2014).

Tariff Data

Finally, the last piece of the puzzle developed was modeling an appropriate tariff for the building. With a maximum demand in 2014 of 236 kW, PG&E's Medium General Demand-Metered Service (A-10) tariff was selected. When performing economic analysis of the battery, care was taken to make sure that the rules within the tariff were followed (Peak vs. Part-Peak, etc.) and that the appropriate values, such as Maximum Demand, were captured for each billing cycle. For ease of modeling, bills were assumed to cover from the first to the last of the month. All of the methodologies support different read cycles, however. Table 1 contains the results of the tariff analysis and lays out the base case for the building.

[1] Most of the smoothing out involved missed intervals that were caught up at the next reading (e.g. 25, 25, 0, 50…). In these instances, the average value was applied to all missing intervals.
[2] While it'd be best to use the weather observed at the building, this station was selected as the closest geographically to the building of interest.

Table 1. Tariff Analysis for Building 74

Month	Usage		Cost		
	kWh	kW	Energy	Demand	Total
January	86,371	179	$10,433	$1,871	$12,304
February	77,722	171	$9,391	$1,792	$11,183
March	85,800	184	$10,369	$1,929	$12,299
April	87,558	223	$10,635	$2,340	$12,975
May	92,365	228	$14,402	$4,065	$18,467
June	90,544	217	$14,116	$3,876	$17,992
July	93,670	226	$14,682	$4,036	$18,719
August	91,636	216	$14,288	$3,858	$18,146
September	94,511	236	$14,824	$4,211	$19,034
October	99,210	233	$15,607	$4,152	$19,759
November	88,424	189	$10,612	$1,979	$12,591
December	85,180	170	$10,298	$1,782	$12,079
Total	1,072,990	236	$149,656	$35,890	$185,546

Battery Data

All of the models assume a 100 kWh battery pack capable of discharging at 25 kW[3]. Roundtrip efficiency is assumed to be 92% (market claims run between 85% and 96%). No assumption was made for standby losses, which are assumed to be negligible within a day and there was no value in attempting to preserve the battery by not running it on weekends or holidays or where a new peak was unlikely to occur.

Finally, once all of the pieces of data were assembled, the last thing done before creating the deployment strategies was to model the building's energy consumption. As previously mentioned trying to model a specific location proves to be more difficult than modeling a system (CAISO's or PG&E's) as a whole. Any random disturbance at the site (say a fan or rooftop unit turning on ahead of what is predicted) would cause a jump in usage that is hard to capture via a regression model. Figure 2 shows how utilizing a few pieces of data such as weather, day of week, and hour can produce strong predictive power. However, since no model is perfect, we must accept that it will occasionally be wrong from time-to-time. For example, even though the grid model has predictive power of 92%, it still misses the peak day by over 10% of the total. As is illustrated later, this is not a fatal flaw for the model but, as with most things, is something to work around.

[3] Roughly equivalent to one Tesla PowerPack of 100 kWh which can provide continuous discharge 4 hours.

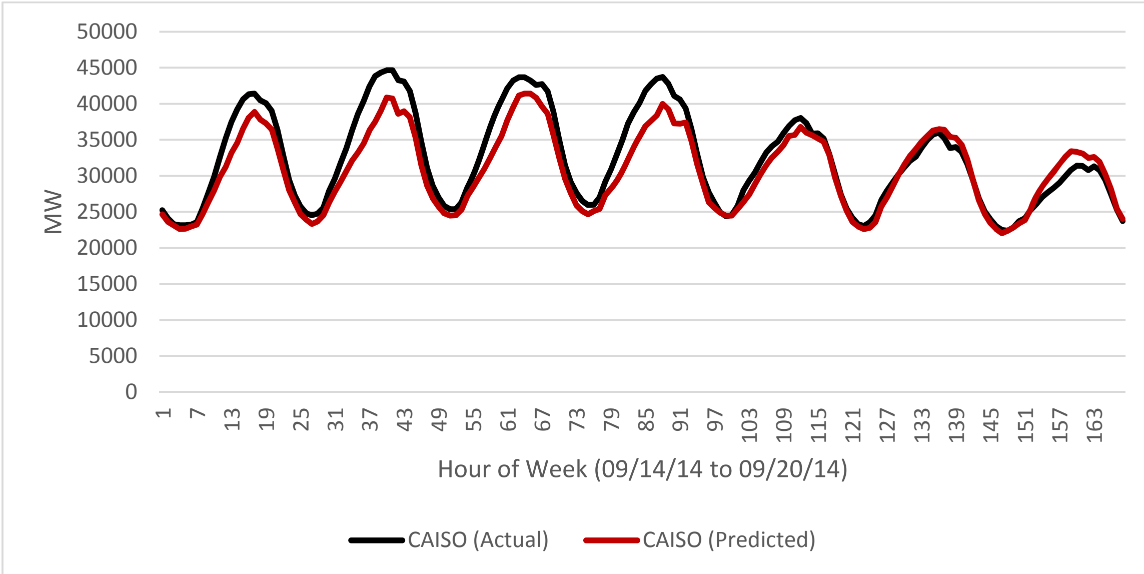

Figure 2. CAISO Actual vs. Predicted Load

As mentioned previously, formulating a comprehensive model for a building is more difficult as there are many more variables than can be captured at the '30,000 foot level'. However, only utilizing the following data, we are able to explain 82% of the observed energy usage: temperature, cooling degree hours, a month flag, a weekday flag, a holiday flag, the maximum demand for the month, whether or not the building is open (presumed operating hours), and an accumulator for that day's cooling degree hours. Using the same date range as the CAISO (system peak week) produces:

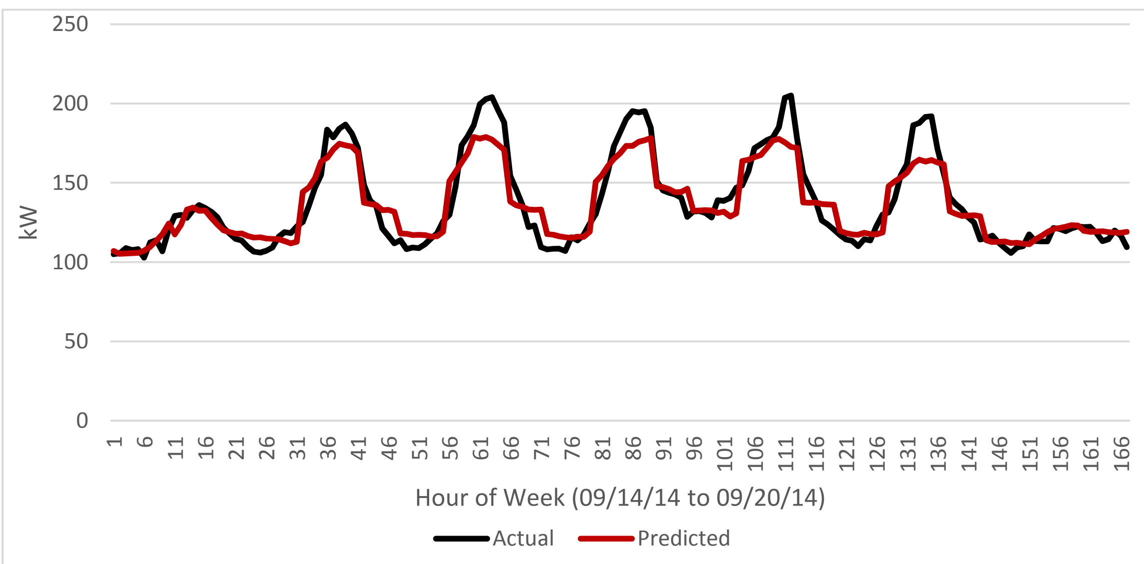

Figure 3. Actual vs. Predicted Building Load

While the system did a fairly good job of predicting the overall shape of the load on 3 out of 7 days and correctly picked up a higher-than-usual overnight load between days 4 and 5, the system did miss the overall peak. Model 3 attempts to accommodate this issue by calculating the standard deviation of the estimate and allowing the user to select a risk threshold that the system then guards against.

Model 1: Trigger Based

Perhaps the most simplistic model to begin with when developing a battery deployment strategy is that of a trigger based approach – for example, 'If the kW load in the building is less than the current maximum demand for the month and it's an off peak hour then charge the battery'. This methodology defines a series of rules about when to charge and then discharge the battery. While this approach is quick to develop and deploy, there are some drawbacks: if the rules aren't setup correctly, the battery may discharge itself preventing peaks before the hottest parts of the day when demand is likely to be highest, potentially negating any savings other than the off-peak to peak energy arbitrage. Likewise, charging the battery must be carefully set up as it does the owner no good to charge the battery during off-peak hours if that's also the time of maximum demand. The tables below provide a brief synopsis of a trigger based approach's rules and results for the building:

- Charging: Between 2am and 8am (Off Peak), attempt to charge the battery until it is full. Do not set a new maximum peak[4]. If we're in partial peak or on peak hours, demand is less than the current maximum demand, and the battery isn't full, attempt to charge it[5].
- Discharging: If the building is about to set a new peak, attempt to discharge the battery to push the peak down to the last maximum with the constraint that the battery can only discharge so much energy in a 15 minute interval. If the battery does not have sufficient energy, then completely discharge it. If we're approaching the end of the day (Off-peak transition) and the battery will need to discharge at its maximum rate in order to reach 0 kWh by the transition, begin discharging.

The following graph depicts the results of running the model during the week of the building's annual peak. Entering the week, the building had a peak of 166 kW. The first two days were a Sunday and holiday (Labor Day), which lowered load. Business resumed Tuesday and the following day, Wednesday, was when the building set its annual peak at 236 kW at 2:45pm. Because the model is not forward looking, the battery ended up trying to defend a maximum peak that was not going to stand. While the battery successfully reduced billed demand by 25 kW for the annual peak, it ended up failing to do so in 4 out of 12 months (see Table 2). There are some improvements that could be made to this model to eke out more savings by discharging the battery during on-peak hours, though that runs the risk of not having energy available for a late evening peak.

[4] The model was allowed to charge up to 145 kW total demand, even if that month's demand was lower. This avoided issues where the meter had just been read and the maximum demand at the beginning of the month was very low during the first day.
[5] There is no arbitrage opportunity at this point (buy on peak, sell on peak), but it does keep the battery available in case demand surges later in the day. In this scenario, on peak charging increased energy costs by $19 but reduced demand by $103, a net savings.

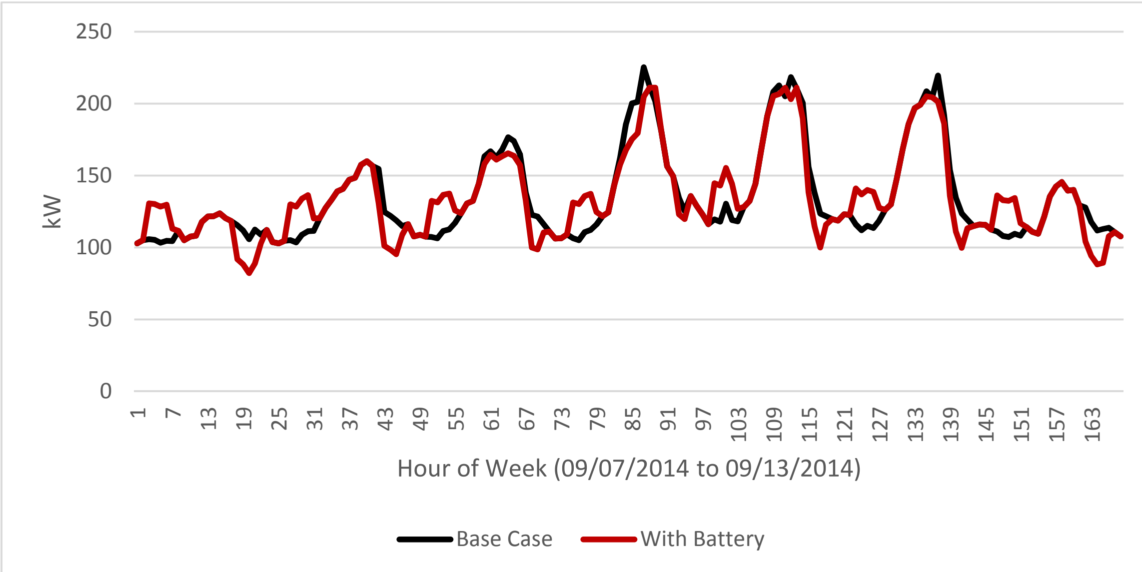

Figure 4. Trigger Based Model: Peak Week

Billing analysis shows a savings of over $4,000 in one year for the battery, which went through 365 cycles - most batteries are rated for 2,000-5,000 cycles before degrading to 80% capacity (Shahan 2015). While $3,800 came from demand savings, there was also energy cost savings even though overall energy usage was higher as the battery, essentially, arbitraged between off-peak and partial peak/on peak prices. At a cost of $250/kWh, or $25,000 for our assumed battery, this would provide for a simple payback period of just over 6 years, longer if there are associated maintenance or operational costs (NRECA 2015).

Table 2: Model 1 Savings Results

Month	Usage		Cost		
	kWh	kW	Energy	Demand	Total
January	86,645	154	$10,431	$1,609	$12,040
February	77,980	146	$9,390	$1,531	$10,920
March	86,075	159	$10,367	$1,668	$12,035
April	87,823	198	$10,630	$3,541	$14,171
May	92,637	211	$14,355	$3,771	$18,126
June	90,810	192	$14,059	$3,430	$17,489
July	93,948	201	$14,618	$3,590	$18,209
August	91,908	200	$14,241	$3,563	$17,804
September	94,780	211	$14,768	$3,765	$18,533
October	99,482	210	$15,548	$2,198	$17,745
November	88,686	182	$10,613	$1,906	$12,519
December	85,459	145	$10,294	$1,520	$11,814
Total	1,076,232	211	$149,314	$32,090	$181,403
Savings	(3,243)	25	342	$3,801	$4,143

Model 2: The Predictive Model

The next evolution in the model was to incorporate a forward looking aspect to the decision to deploy the battery storage. While the trigger model does a fine job of responding to the building as-is, there were several occasions where, early on in the month, the battery would deploy trying to preserve a peak that would be easily surpassed later in the day, no matter what the battery did. If we look at the battery's performance on November 5th, 2014, there is a clear problem:

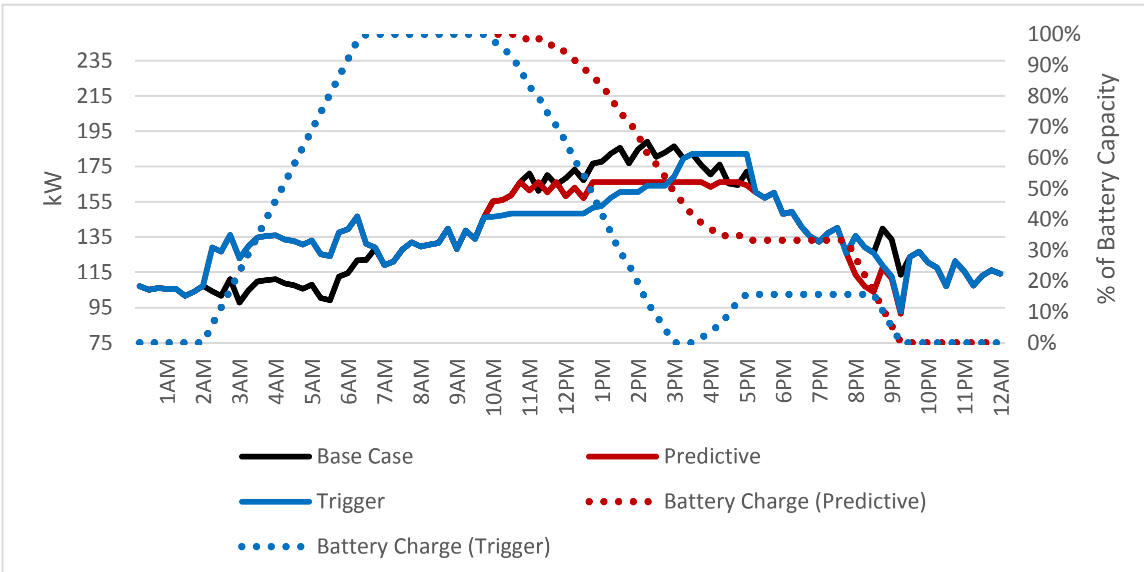

Figure 5: Model 1 and 2 Comparison for 11/05/2014

Entering the day, the building's maximum kW load for the month was 145 kW, which the trigger based model attempted to defend from interval 37 to 61. However, by interval 62, the battery had depleted and was unable to deploy, which reset the maximum demand higher, to 182 kW (which explains the recharge in intervals 63-70). So, the predictive model added an additional rule to the mix:

- Discharging: If the predicted demand is over 25 kW higher than our current new demand or if the predicted energy demand is over our current capacity of the battery, then do not deploy and wait for a higher usage period.

This additional rule causes the battery to remain in standby mode even if a new peak is occurring if it believes that a larger peak is coming. While the forecast is by no means perfect, it does provide a decent enough 'feel' for the building's load profile that the model is able to rely on it to improve its decision making process. Overall, this new step improved savings by $246, or 5.9% over the trigger based model:

Table 3: Model 2 Savings Results

Month	Usage		Cost		
	kWh	kW	Energy	Demand	Total
January	86,645	155	$10,431	$1,628	$12,058
February	77,980	146	$9,390	$1,531	$10,920
March	86,075	159	$10,367	$1,668	$12,035
April	87,824	198	$10,630	$2,078	$12,708
May	92,635	203	$14,346	$3,619	$17,965
June	90,810	192	$14,059	$3,430	$17,489
July	93,948	201	$14,618	$3,590	$18,209
August	91,908	200	$14,241	$3,563	$17,804
September	94,780	211	$14,768	$3,765	$18,533
October	99,481	209	$15,544	$3,726	$19,271
November	88,690	166	$10,613	$1,739	$12,352
December	85,459	145	$10,294	$1,520	$11,814
Total	1,076,234	211	$149,302	$31,855	$181,157
Savings	(3,244)	25	354	$4,035	$4,389

Model 3: The Probabilistic Model

The final version attempts to correct for one underlying issue with the predictive model: no forecast is perfect. Without having perfect knowledge of the future or the ability to directly control the building's load, the predictive model's energy forecast will likely be wrong by some factor. While, on average, this deviation will be small, it can have significant consequences for demand charges as they are a one-time event (i.e. there is no opportunity to correct a mistake in the prediction once the event passes).

First, though, to digress a bit, when one looks at demand and how it interacts with a battery, there are similarities that can be drawn between that system and a financial instrument: a call option. By investing money (energy) now, we purchase (store) the ability to swap the current price (demand) for the strike price (max demand when we stored the energy). What option pricing does is answer the question: how much should one value this ability? In terms of a battery storage system, if the premium that is placed on the chance future intervals will exceed the current peak demand, then the battery should not deploy, even though it may be setting a new peak.

Looking once again at the November peak day, the probabilistic model was able to capture the last 2 kW that the predictive model missed by assuming the tail end of the day was going to have heavier usage that it actually did. The probabilistic model was able to predict a lower likelihood of a new peak being developed towards the end of the day and was more aggressive in its deployment, carrying less energy into the end-of-day flush-out of energy prior to the transition into the overnight off peak period:

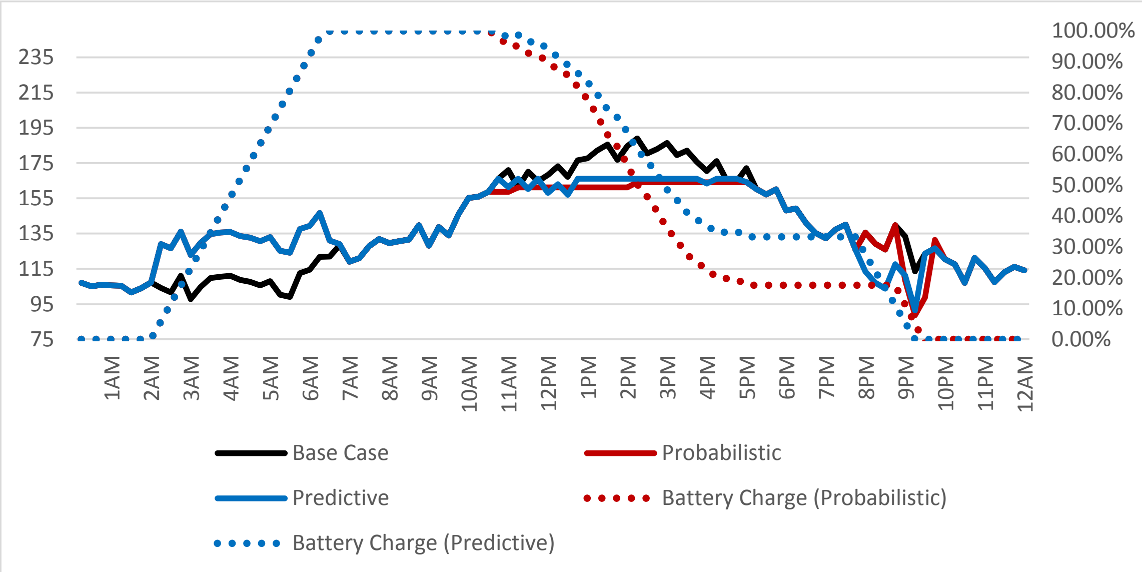

Figure 6: Model 3 vs. Model 2 November 5th Comparison

The results show that the model was able to maximize demand savings (25 kW) in 11 out of 12 months, compared to 8 out of 12 months for the predictive model. While energy cost savings were lower, which impacted the total amount, this problem would be easy to fix by incorporating peak vs. off-peak pricing into the energy valuation algorithm. If energy costs were fixed, the probabilistic model would outperform the predictive model by 1.3%, or 7.4% above the original trigger based approach.

Table 4: Model 3 Savings Results

Month	Usage		Cost		
	kWh	kW	Energy	Demand	Total
January	86,642	154	$10,432	$1,609	$12,041
February	77,972	146	$9,391	$1,531	$10,921
March	86,071	159	$10,369	$1,668	$12,036
April	87,821	198	$10,632	$2,078	$12,710
May	92,635	203	$14,361	$3,619	$17,980
June	90,810	192	$14,074	$3,430	$17,504
July	93,945	201	$14,633	$3,590	$18,223
August	91,909	200	$14,256	$3,563	$17,819
September	94,777	211	$14,783	$3,765	$18,548
October	99,481	208	$15,562	$3,706	$19,268
November	88,687	164	$10,615	$1,717	$12,332
December	85,455	145	$10,296	$1,520	$11,816
Total	1,076,205	211	$149,403	$31,795	$181,198
Savings	(3,216)	25	252	$4,096	$4,348

Conclusion

Battery storage promises to revolutionize several aspects of the utility industry. From moving solar power from daytime to evening to creating net zero buildings, there are many aspects to battery storage that are exciting and new. However, what matters most to a business is whether or not the storage system will provide value and how quickly that value can be captured. By using intelligent algorithms that not only look at what the building is doing in its current state but also what the building will (or will not) be doing in its future state, that value can be measurably improved. Each of the models presented increased in complexity but also in savings achieved. While the simplest model is also the easiest one to implement and understand, more advanced models have their advantages when it comes down to what end users care about most: results.

References

Anderson, David. 2009. “An Evaluation of Current and Future Costs for Lithium-Ion Batteries for use in Electrified Vehicle Powertrains” Accessed February 24. http://dukespace.lib.duke.edu/dspace/bitstream/handle/10161/1007/Li-Ion_Battery_costs_-_MP_Final.pdf;jsessionid=11C13DC558A77B21CF48E579CD3A4FFE?sequence=1

California Independent System Operator (CAISO). 2016. “California ISO Peak Load History 1998 through 2015”. Accessed March 11. www.caiso.com/Documents/CaliforniaISOPeakLoadHistory.pdf

NRECA. 2015. “Technology Advisory – Tesla Energy PowerWall” September 2015. www.nreca.coop/wp-content/uploads/2013/06/Technology-Advisory_Tesla-Energy.pdf

OpenEI. 2015. “Long-term Electricity and Natural Gas Consumption for LBNL Building 74” Accessed December 10. en.openei.org/datasets/dataset/lbnl-building-74

Shahan, Zachary. 2015. “Tesla Powerwall & Powerpacks Per-kWh Lifetime Prices vs Aquion Energy, Eos Energy, & Imergy” May 9. cleantechnica.com/2015/05/09/tesla-powerwall-powerblocks-per-kwh-lifetime-prices-vs-aquion-energy-eos-energy-imergy/

Tesla, 2016. “Tesla Energy” Accessed February 25. www.teslamotors.com/presskit/teslaenergy

Van Noorden, Richard. 2014. “The Rechargeable Revolution: A Better Battery” *Nature*. March 5. www.nature.com/news/the-rechargeable-revolution-a-better-battery-1.14815

Wilkinson, Sam. 2014. “The grid-connected energy storage market is set to explode, reaching a total of over 40 GW of installations by 2022.” *Market Insight, IHS*, January 15. technology.ihs.com/483008/the-grid-connected-energy-storage-market-is-set-to-explode-reaching-a-total-of-over-40-gw-of-installations-by-2022.