

Integrating Energy Storage in Electricity Distribution Networks

Aditya Mishra*, Ramesh Sitaraman*, David Irwin*, Ting Zhu‡, Prashant Shenoy*, Bhavana Dalvi§, and Stephen Lee*

*Univ. of Massachusetts Amherst, ‡Univ. of Maryland Baltimore County, §Carnegie Mellon University
{adityam,ramesh,irwin}@cs.umass.edu, zt@umbc.edu, shenoy@cs.umass.edu, bbd@cs.cmu.edu, stephenlee@cs.umass.edu

ABSTRACT

Electricity generation combined with its transmission and distribution form the majority of an electric utility's recurring operating costs. These costs are determined, not only by the aggregate energy generated, but also by the maximum instantaneous peak power demand required over time. Prior work proposes using energy storage devices to reduce these costs by periodically releasing energy to lower the electric grid's peak demand. However, prior work generally considers only a single storage technology employed at a single level of the electric grid's hierarchy. In this paper, we examine the efficacy of employing different combinations of storage technologies at different levels of the grid's distribution hierarchy. We present an optimization framework for modeling the primary characteristics that dictate the lifetime cost of many prominent energy storage technologies. Our framework captures the important tradeoffs in placing different technologies at different levels of the distribution hierarchy with the goal of minimizing a utility's operating costs. We evaluate our framework using real smart meter data from 5000 customers of a local electric utility. We show that by employing hybrid storage technologies at multiple levels of the distribution hierarchy, utilities can reduce their daily operating costs due to distributing electricity by up to 12%.

Categories and Subject Descriptors

J.7 [Computer Applications]: Computers in Other Systems—*Command and control*

Keywords

Energy; Battery; Electricity; Grid; Peak shaving

1. INTRODUCTION

Nearly 40% of energy in the U.S. is consumed in the form of electricity [28]. Increasing the percentage of electrical energy is an important part of creating a clean and sustainable energy supply, as “green” energy, e.g., from solar and wind, is generally consumed in the form of electricity. In addition, transmitting and distributing electricity is significantly more efficient than transmitting and distributing other captive energy sources, e.g., via oil and gas pipelines or trucks. However, electricity transmission and distribution (T&D)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

e-Energy' 15, July 14–17, 2015, Bangalore, India.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-3609-3/15/07 ...\$15.00.

<http://dx.doi.org/10.1145/2768510.2768543>.

costs are non-trivial, and, in some cases, such as New York and southern California, now dominate generation costs [34]. The cost and carbon footprint to generate electricity is a complex function of the electricity demand patterns, mix of generators and fuel sources, penetration of renewable energy, and T&D efficiency.

A significant fraction of these costs are determined by the electric grid's peak power demand. The peak demand influences capital costs by dictating the capacity (and number) of transmission lines, substations, transformers, etc., since utilities must size these to service the peak. In addition, since the “peaking” generators utilities activate to satisfy demand peaks are significantly less efficient and more expensive to operate than baseload generators that are continuously active, peak power demands also influence operational costs. Thus, satisfying even brief peak demand periods has a disproportionate affect on capital and operational expenses. For example, recent estimates attribute as much as 20% of the grid's generation costs in the U.S. to servicing only the top 100 hours of peak demand each year [32]. Finally, since energy lost in transmission and distribution is a function of the square of current, rising peak demand results in quadratically higher transmission losses.

The importance of reducing peak demand is one of the primary motivations for Demand Response (DR) programs, which attempt to coerce consumers into actively shifting their load from peak to off-peak periods. Since requiring consumers to actively change their behavior to shift load is often not effective, recent work has explored the use of energy storage to automatically shift load in the background, i.e., by storing energy during off-peak periods and using it during peak periods [22, 29, 30]. Prior work in this area has generally examined deploying energy storage devices (ESDs) in individual homes, where the approach can potentially reduce a consumer's electricity bill if electricity prices vary over time, e.g., such that peak prices are higher than off-peak prices. In fact, such energy arbitrage is an explicit use-case cited by Tesla for its new Power-Wall battery, which is designed for deployment in homes [19].

While prior research, and now commercial products, target energy storage for homes, such storage can be deployed at any level of the grid's hierarchy from the lowest level (at homes) to the medium level (at distribution transformers) to the top level (at distribution and bulk power substations). The choice of where to deploy energy storage presents interesting tradeoffs. For example, using energy storage in individual homes to reduce the home's peak demand requires more aggregate storage capacity than employing storage at a higher level of the grid hierarchy, since each home's peak demand does not occur at the same time yielding some smoothing from statistical multiplexing at higher levels. Since prior research largely focuses on deploying energy storage in homes, it also generally focuses on only a single type of storage technology: in particular, batteries [22, 30]. However, while batteries are the only

small-scale energy storage appropriate for homes at current price points, other ESDs become more feasible at higher levels of the grid hierarchy. Energy storage technologies differ in their cost, lifetime, energy-efficiency, etc. For example, flywheels exhibit a high energy-efficiency and lifetime, but have a high self-discharge rates and cost, while lead-acid batteries exhibit a low self-discharge rate and cost, but have a shorter lifetime and lower energy-efficiency.

Thus, our hypothesis is that intelligently employing hybrid combinations of different energy storage technologies at multiple levels of the grid’s hierarchy has the potential to reduce costs relative to deploying only a single storage technology at a single level of the hierarchy. In evaluating our hypothesis, we make the following contributions.

- **ESD and Grid Modeling.** We extensively model important ESD operational characteristics, including energy density, self-discharge rate, cycle lifetime, power ramp time, etc., to capture their tradeoffs. We examine the deployment of different ESDs using a simple model of the grid’s electricity distribution hierarchy, which includes the various costs associated with generating, transmitting, and distributing electricity.
- **Optimization Framework.** Using our above models, we develop an optimization framework that enables us to examine the benefit of using different combinations of ESDs at different levels of the hierarchy. The goal of our optimization framework is to minimize generation costs, including the capital, operational, and storage costs, for different configurations of ESDs.
- **Implementation and Evaluation.** We implement our optimization framework and then use it to evaluate in simulation the cost and benefit of different storage configurations using smart meter data from 5000 customers of a local utility. In doing so, we identify key insights into the benefits of different ESD technologies at different levels of the grid. We find that deploying hybrid ESDs at an individual level typically improves savings over any single-technology ESD deployment, while deploying multi-level hybrid ESDs typically provides the best savings. Overall, we find that ESDs can reduce distribution-related capital and operational costs by up to 12%.

2. BACKGROUND

2.1 Electric Grid

The electric grid is an interconnected network for delivering electricity from suppliers to consumers. Electricity is generated at power plants, often far from population centers, using different types of generators and fuels with different operational characteristics. Generated electricity exits the power plant and is stepped up to high voltages for long-distance transmission, since high voltages reduce transmission losses. At a substation near the final destination, a step-down transformer reduces the transmission voltage for distribution to both industrial and residential customers. At this point, distribution lines deliver electricity from the substation to end-consumers. In this work, we focus primarily on the large number of small-scale residential consumers in the grid, since they represent the vast majority of end-points in the distribution network.

2.1.1 Distribution Network

Figure 1 highlights the basic structure of electricity distribution in the grid. Electricity is fed into a bulk power substation, or a subtransmission station, which service a few “load areas” of customer demand. The bulk power substation routes the electricity to distribution substations. A distribution substation may then route the power to thousands of homes [1, 2]. Before being delivered to

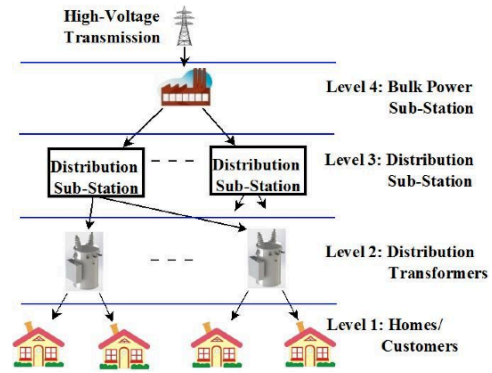


Figure 1: Typical electricity distribution hierarchy.

a building, distribution transformers near the building steps down the voltage of electricity. The number of consumers fed by a single distribution transformer varies: several homes may be fed off a single transformer in urban areas, or rural distribution may require one transformer per consumer [9].

In general, multiple distribution transformers may be connected in parallel. However, due to a lack of access to the distribution graph of an existing network and, for simplicity, in this paper, we assume the topology of the distribution network as shown in Figure 1. We base this simple model on information that is available in public domain [1, 2, 9], and use it in our experimental evaluation. Here, we assume each distribution transformer supports five homes, each distribution sub-station serves 500 transformers, and two distribution sub-stations are served by one bulk power substation. While our absolute results are specific to this simple model of a distribution network, we believe that many of our key insights are applicable to a range of real topologies, since we base our topology on publicly-available information. Importantly, our methodology and analyses extends to other types of distribution networks.

2.1.2 T&D Losses in the Grid

A fraction of electricity is lost in transmission and distribution. In the US, nationally, roughly 6% to 6.5% of the total electricity is lost each year [25]. Losses are generally divided equally between transmission and distribution. For example, in New York, transmission losses accounted for a total of 3.18% loss, while distribution losses accounted for the loss of 3.3% of the total annual electricity [5]. We use these loss values in our evaluation.

2.2 Electric Utility’s Generation Costs

An electric utility generates, transmits, and distributes electricity for sale in the electricity market [10]. A consumer’s electric bill is generally divided into three categories related to electricity’s generation, transmission, and distribution, as listed below [18, 11, 7].

Energy Charge. Consumers are charged based on the total amount of energy, in kilowatt-hours (kWh), they consume over a billing period. This charge incorporates the cost for a utility to generate the energy or buy the energy on the open market.

Distribution Charge. Consumers are charged a fee to enable utilities to recover the cost of operating and maintaining the distribution system. This charge typically has two components: an energy component, based on the amount of kWh consumed over the billing period, and a peak power component, based on the highest peak power demand in kilowatts (kW) over the billing period [18].

Transmission Charge. Consumers are charged a fee to enable utilities to recover the costs related to the delivery of electricity over high-voltage transmission lines. This energy is generally purchased from a third-party and not generated by the local utility. As with

the distribution charge, this charge has an energy component and a peak power component [18].

In some cases, consumers are not charged for energy, distribution, and transmission individually, but rather, the charges are included as part of the electricity rate. In addition to these costs, utilities also have expenditures related to the cost of materials and supplies and capital (including depreciation).

Expenses that are dependent on the total energy consumption are dictated by the pattern of end-user consumption, which cannot be controlled by a utility. However, the generation, transmission, and distribution costs incurred as a result of demand peaks can be reduced by curtailing the peaks. In addition, reducing demand peaks enables utilities to gain savings from avoided electricity costs, which include the marginal cost to produce and deliver one more unit of electrical energy. The avoided cost consists of two components—avoided energy costs (\$/MWh) and avoided capacity costs (\$/kW-month) [4]—which represent lower generation costs and the need for less peak capacity from lower peak demands.

2.3 Energy Storage to Lower Utility’s Costs

Energy storage devices can be used to store energy during low demand periods, which can then be used later to satisfy customer demands during peak demand periods, thereby reducing the net peak on the higher levels of the grid. As capital and operational expenses of the grid are largely determined by the peaks, energy storage can cut these expenses for the grid.

2.3.1 Energy Storage Technologies

We examine the potential for the following energy storage technologies to reduce an electric utility’s distribution costs.

Compressed Air Energy Storage (CAES): With Compressed Air Energy Storage, off-peak grid power is used to compress air underground. Later, when the energy is needed, this compressed air is released to power an electric generation and produce electricity. These systems are typically large, often requiring significant real-estate for storing compressed air [23], e.g., underground tanks.

Ultra-capacitors (UC): Ultra-capacitors operate similarly to electrostatic capacitors, except that they can hold significantly more energy in a size similar to that of conventional capacitors [21]. UCs are now often being used for large-scale uninterruptible power supplies (UPS) in data centers, hospitals, industrial buildings, etc. [17].

Flywheels (FW): Flywheels store kinetic energy in rotating discs. These discs are made to turn a generator for producing electricity. Flywheels can be very efficient in storing energy over short durations; however, they have high self-discharge rates due to losses from friction [21, 12]. One example of a flywheel energy storage plant is the Beacon Power plant in New York [8].

Lead Acid batteries (LA): Lead-acid batteries are one of the most widely used energy storage devices. They have long been the primary technology for stationary energy storage at both grid-scales and in off-grid homes [15].

Lithium-Ion batteries (LI): Lithium Ion batteries are the most popular type of rechargeable batteries; they are known for their relatively high efficiency and energy density [16]. Lithium Ion batteries are the primary technology in mobile systems, e.g., electric vehicles, due to their light weight in comparison with lead acid batteries. However, these batteries are now being deployed in conjunction with renewable energy to provide energy storage for homes, as evidenced by Tesla’s recent introduction of the PowerWall home battery based on lithium-ion technology [19].

While diesel generators, and other captive sources can also be considered energy storage devices, we do not consider them separately here. Pumped hydroelectric is another widely used storage

ESD	CAES	UC	FW	LA	LI
Efficiency (%)	68	95	95	80	85
Discharge:Charge Rate	4	1	1	10	5
Self-discharge (%per day)	low	20	100	0.2	0.1
Energy Density (Wh/L)	6	30	80	80	150
Power Density (W/L)	0.5	30000	1600	125	450
Ramp Time (sec)	600	0.001	0.001	0.001	0.001
Max DoD (%)	100	100	100	80	80
Energy Cost (\$/kWh)	50	500	1000	200	525
Cycle Lifetime	15000	100000	12000	2000	5000
Expected Lifetime (Years)	20	20	15	4	10

Table 1: ESD Parameters [23] [21] [39] [35]

technology in the grid; however, since it requires significant infrastructure, it is not readily deployable in the distribution networks.

2.3.2 Energy Storage Characteristics

Below we list key characteristics of the energy storage devices that are relevant to our optimizations. Table 1 list the default values of the parameters used in our study, which we derive from various scientific studies. Note that these parameters are inputs to our framework and while they vary significantly across technologies, we do not further consider the impact of varying them for a particular storage technology. We are specifically interested in how these characteristics yield different trade-offs when placing various storage technologies at different levels of the distribution hierarchy to minimize a utility’s distribution costs. We model the following characteristics:

Energy Storage Capacity: The energy storage capacity represents the total amount of energy that a device can store. Generally, the capacity is expressed in kilowatt-hours (kWh).

Maximum Charge and Discharge Rates: Usually expressed as E-rate, the maximum charge and discharge rates are a measure of the rate at which a battery can be charged or discharged relative to its total capacity [3]. For example, a 2E discharge rate is the discharge rate necessary to fully discharge the battery in half an hour.

Efficiency: Use of energy storage results in energy loss due to energy conversion. We employ a constant efficiency factor for each energy storage technology to capture these losses; e.g., typical lead-acid batteries are 80% efficient.

Self-Discharge Rate: The self-discharge rate is a phenomenon in energy storage by which the ESD loses stored energy merely with passage of time. The self-discharge rate can be significant for some technologies, such as flywheels. For energy storage technology k we model its self-discharge per unit time as a constant factor μ_k .

Cycle Lifetime: A ESD’s lifetime is usually expressed in terms of number of charge-discharge cycles. Typically, ESD lifetime is measured based on the number of cycles as a function of the depth of discharge (DoD). For a given energy storage technology, we limit its DoD and the number of charge-discharge cycles at the given DoD over a given time horizon; thereby, we control the lifetime of an energy storage device, and capture its amortized per unit energy storage cost over its lifetime in our model.

Energy Density: The energy density is the nominal battery energy per unit volume (Wh/L). The energy density determines the battery size required to achieve a given energy output [3].

Power Density: Power density is defined as the maximum available power per unit volume (W/L). The power density determines the battery size required to achieve a given power output [3].

Power Ramp Time: The Power Ramp Time is the start-up latency associated with a given energy storage technology before it can

start delivering its maximum power. This ramp-up is similar to the start-up acceleration in vehicles. Ramp times of most storage devices are very low, however for compressed air storage the ramp time may be several minutes.

3. PROBLEM STATEMENT

Although energy storage can reduce peaks and cut costs, the problem of storage deployment presents several interesting trade-offs. Peak reduction at a given level of the grid’s hierarchy enables provisioning the infrastructure at that level, as well as higher levels, for a lower peak. Therefore, storage deployment at lower levels of the hierarchy appears more beneficial than reducing the peak only at higher levels. However, in general, peaks at higher levels of the hierarchy are smaller than the sum of individual peaks at the lower levels, such as at homes; this occurs because individual homes peak at different times. The statistical multiplexing gains due to the spreading of individual peaks over time makes aggregate peaks at higher levels smaller. Therefore, deploying energy storage at the higher levels would require much less energy storage capacity, and hence lower aggregate energy storage costs.

In addition to deployment choices, the choice of storage technologies also presents tradeoffs in their cost, lifetime, efficiency, energy density, etc. For example, compressed air energy storage has a low energy cost and long expected lifetime, but a low energy efficiency and requires significant space for deployment. In contrast, lead-acid batteries have a higher energy-efficiency and a smaller form-factor, but also much higher energy costs and much shorter expected lifetime. Furthermore, different storage technologies are suitable for shaving different types of peaks: compressed air storage is suitable for wide peaks, lead-acid batteries work well for less frequent medium-width peaks, and ultra-capacitors are best for very narrow peaks (up to a minute). Since the nature of the peak demand depends on the level of the grid hierarchy—medium-width peaks are more likely at homes, whereas wide peaks are frequent higher in the hierarchy—the best choice may differ at each level.

In this work, we address the problem of deploying energy storage across a distribution grid hierarchy to cut a utility’s distribution costs. As there are a number of variables involved, such as large distribution hierarchies, time varying demand profiles, different types of storage technologies, and a range of electricity pricing plans, it is not easy to formulate a heuristic solution for this problem. Therefore, we frame it as an optimization problem. We define the problem as follows: given an electricity distribution network, an estimate of power demands, and a set of available storage technologies, the problem is to find an optimal sizing and placement of energy storage devices across the distribution hierarchy so as to minimize a utility’s expenses (Section 2.2) for distributing electricity. Our framework is general enough to provide storage provisioning solutions for a range of consumption profiles, electricity pricing plans, storage technologies, and distribution networks.

4. ENERGY STORAGE PROVISIONING AND CONTROL FRAMEWORK

We now present our optimization framework for energy storage provisioning. We intend the framework to provide storage deployment solutions for a distribution hierarchy with the goal of optimizing a utility’s cap-ex and op-ex. Inputs to the framework include power demand, the distribution network topology, and cap-ex and op-ex costs. The framework solution then outputs the optimal choice, placement, and size of energy storage devices across the hierarchy, along with optimal energy storage control patterns.

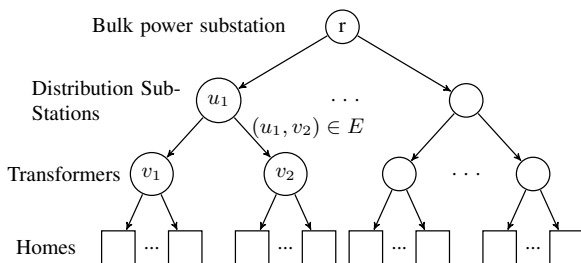


Figure 2: Illustrative graph depicting distribution hierarchy

Throughout the formulation we assume the grid distribution network is a directed graph $G = (V, E)$ —as shown in Figure 2—where V is the set of nodes (i.e., homes, distribution transformers, substations, etc.) and E represents the directed edges between these nodes. In Figure 2, circles and squares represent the nodes, and arrows represent the edges. Two types of nodes are shown in the figure: squares are the leaf nodes and circles are the non-leaf nodes. In the distribution network, leaf nodes typically represent homes. In our formulation, all leaf nodes are represented by the set labeled LeafNodes. Node r (Figure 2) represents the root node. Finally, across the formulation, k represents the k -th energy storage technology (out of K) and t is the time interval (between 1 to T).

4.1 Inputs

Power Consumption (Demand): Broadly, there are two types of problems associated with energy storage deployments in the grid’s distribution network. The first problem is determining the proper energy storage capacity and where in the hierarchy to deploy it. The second problem is determining how to charge-discharge the energy storage device to clip the peaks and realize cost savings. In this paper, we solve the first problem—the energy storage sizing and deployment problem, i.e., figuring out how much energy storage should be deployed and where. Therefore, we assume that historical power consumption traces are available, and we use these for future provisioning of storage in the distribution network. We assume that prior work can be employed to derive an accurate power demand time-series at each home (e.g., [20] [37]). Further, utilities have extensive power consumption logs over time for their customers, which can also be used as an input to our optimization framework. We divide time into T slots, each of length I . For home u we assume its power demand to be a time series $U_{sr}Dmnd_{u,t}$, where $t \in [1, T]$. We present results for both real and synthetic consumption power time series.

Capital Expenditure (Infrastructure Cost): We model two types of infrastructure costs: first, maintenance and upgrade cost, second, avoided (or marginal) capacity costs. These costs vary significantly between utilities and between locations within utilities: ranging from \$2.51/kW/month to \$46.34/kW/month [27, 26, 6]. In our experiments, we consider a cap-ex saving in the range of \$6/kW/month to \$30/kW/month resulting from peak reduction. These savings are obtained by reducing a watt of consumption from the maximum power draw $Peak_u^{max_orig}$. At any time t , power draw at node u is given by the sum of power draws of all the nodes in the sub-tree with root at the node u at time t and net energy drawn by the energy storage devices at its vicinity; the sum is denoted by $Demand_{u,t}$. The corresponding size of the tallest shaved peak at u is denoted by $Peak_u^{shvd}$.

Operational Expenditure (Tariffs): Electricity tariffs are good indicators of the operational costs incurred in distribution of electricity. Most prevalent electricity tariff models charge customers for their total energy consumption, i.e., customers have to pay a

flat $\$C_e^{unit}/\text{kWh}$ of their consumption. Recently, to shave the peak demand on their grids, utilities have introduced a peak penalty on the tallest consumption peaks across the billing cycle. Typically, the peaks are computed as a sliding window of 30 minutes over the billing period. End users then pay a penalty of $\$/\text{kW}$ based on the tallest peak. For a utility, this peak penalty translates to the energy charged in generating the peak power and the cost incurred in routing the electricity to the distribution network (as in [14]). Note that our model for the value of peak reduction derives directly from the way electric utility companies charge for the peaks, therefore the marginal value of peak reduction is constant.

Our model for capital and operational expenditure for the utilities is derived from the information available in electricity bills and reports published by the utility companies, as reported in [4, 6, 7, 11, 18]. As utilities pass on their costs to the customers, we believe utility bills closely model the actual expenses of electric utility companies. In addition, several real-world factors, such as resource availability and market price, affecting utility expenses are accounted for while computing the avoided costs. For example, among other factors, [6] accounts for wholesale electric energy price, projections of natural gas prices, generation capacity costs, cost of controlling CO_2 emissions, and the effect of implementation of anticipated federal regulations.

4.2 Optimization Problem Formulation

Decision Variables: All notations used in the framework are summarized in Table 2. Our decision variables capture both the sizing and placement of energy storage and how to operate it to minimize the peak demand. The energy storage capacity of a storage technology of type k deployed at node u is denoted by $C_{k,u}$. The average power fed into and drawn out of an energy storage device at u during time slot t is denoted by $S_{k,u,t}$ and $D_{k,u,t}$.

Optimization Objective: Our optimization objective is:

$$\text{Minimize}(\text{CapEx} + \text{OpEx} + \text{StorageCost}) \quad (1)$$

The objective function has three components: capital expenditure (*CapEx*), operational expenditure (*OpEx*), and *StorageCost*. Each component is normalized to our experiment's time horizon.

CapEx includes the capital expenses due to infrastructure deployment for electricity distribution. Assuming α_u is the maintenance, upgrade, and marginal capacity costs for each watt of infrastructure provisioning at node u , *CapEx* is given by equation (2).

$$\text{CapEx} = \sum_{u \in V} \text{Peak}_u^{shvd} * \alpha_u \quad (2)$$

OpEx is the expected utility operational costs in electricity distribution and can be represented as in (3). The *OpEx* has three components, respectively: peak surcharge paid on the tallest demand peak served by the utility, electric energy cost paid on the total electricity served to the customers, and additional avoided costs of electricity incurred as a result of inefficiencies in energy storage devices.

$$\begin{aligned} \text{OpEx} = & \text{Peak}_{root}^{shvd} * b + \sum_t \text{Demand}_{root,t} * I * a_t \\ & + \sum_{k,u,t} (S_{k,u,t} - D_{k,u,t}) * I * \gamma \end{aligned} \quad (3)$$

In the above, b is the per unit surcharge ($\$/\text{kW}$) on the tallest peak and a_t is the unit cost of electricity at time t . $\text{Peak}_{root}^{shvd}$ is the tallest peak seen at the root node, which incurs the peak surcharge. γ is the avoided electric energy cost (AEEC in $\$/\text{MWh}$), as some energy is lost in the energy storage conversion process, this lost energy incurs extra avoided costs which is added to the utility operational costs.

Note that $\text{Demand}_{root,t}$ captures the total load including losses in transmission and storage charge-discharge.

StorageCost is the cost of energy storage deployment across the grid, given by (4), where $\beta_{k,u}$ is the amortized cost of the energy storage device k at node u per unit energy adjusted to its lifetime.

$$\text{StorageCost} = \sum_{u,k} C_{u,k} * \beta_{k,u} \quad (4)$$

The lifetime of an energy storage device depends on several factors such as the depth of discharge (DoD)—a battery lasts longer for smaller DoD. The value of $\beta_{k,u}$ is an input and is determined by the DoD and the set number of charge-discharge cycles over the time horizon. In this paper, in addition to the storage costs, $\beta_{k,u}$ includes the cost of the power conversion system, balance of plant, operation and maintenance [21].

Constraints: We assume that the state of charge in all storage devices at the end of the time horizon is same as their state at the beginning, as stated in (5).

$$E_{k,u,1} = E_{k,u,T+1}, \forall k, u \quad (5)$$

At any time, an energy storage device can only store energy between a lower threshold dictated by its allowed depth of discharge and a maximum capacity; this is captured by (6).

$$(1 - \text{DoD}_k^{max}) * C_{k,u} \leq E_{k,u,t} \leq C_{k,u}, \forall k, u, t \quad (6)$$

For each storage device, the rate at which energy can be drawn from and fed into the device is bounded by its discharge (r_k^{disch}) and charge (r_k^{charge}) rates, as determined by the underlying storage technology. This is captured in equations (7) and (8).

$$0 \leq D_{k,u,t} \leq C_{k,u} * r_k^{disch}, \forall k, u, t \quad (7)$$

$$0 \leq S_{k,u,t} \leq C_{k,u} * r_k^{charge}, \forall k, u, t \quad (8)$$

Equation (9) is the energy conservation constraint, which states that the total energy drawn out of the energy storage is never greater than the energy charged to the battery multiplied by the storage efficiency (e_k).

$$\sum_{t=1}^T D_{k,u,t} \leq e_k * \sum_{t=1}^T S_{k,u,t}, \forall k, u \quad (9)$$

$$\begin{aligned} \text{Demand}_{u,t} = & \frac{1}{\eta} \left(\sum_{v:(u,v) \in E} \text{Demand}_{v,t} \right) + \sum_k S_{k,u,t} \\ & - \sum_k D_{k,u,t}, \forall u \in (V - \text{LeafNodes}), t \end{aligned} \quad (10)$$

Net power consumption at any non-leaf node u at time t , denoted by $\text{Demand}_{u,t}$, is determined by the sum of net power consumption at all its child nodes and the net power drawn/delivered by the energy storage devices deployed at the node, equation (10). (η takes care of the electricity lost due to transmission inefficiencies between node u and its children.) For example, in Figure 2, if transmission efficiency is 100%, the net power draw at node u_1 is given by the sum of net power drawn at its child nodes v_1, v_2 and the storage devices at u_1 . On the other hand, net power draw at the leaf nodes, i.e., homes, is given simply by the sum of home's electricity demand and net electricity drawn/delivered by energy storage at the home, as in equation (11).

$$\begin{aligned} \text{Demand}_{u,t} = & \text{UsrDmnd}_{u,t} + \sum_k S_{k,u,t} \\ & - \sum_k D_{k,u,t}, \forall u \in \text{LeafNodes}, t \end{aligned} \quad (11)$$

Symbol	Notation
$C_{k,u}$	Capacity of the k -th energy storage device (ESD) at node u of the grid hierarchy
C_e^{unit}	Unit cost of energy
a_t	Unit cost of electricity in interval t
$D_{k,u,t}$	Average power drawn from the k -th energy storage at node u of the grid hierarchy in time interval t
DoD_k^{max}	Maximum depth of discharge for k -th type of storage
$E_{k,u,t}$	Energy stored in k -th storage device at node u of the grid hierarchy in interval t
e_k	Efficiency of storage type k
I	Length of each time interval
$Demand_{u,t}$	Net power demand on grid at node u in interval t
$U_{srDmnd}_{u,t}$	User consumption at home node u in interval t
$Peak_u^{shvd}$	Maximum shaved peak seen by the node u
$Peak_u^{max_orig}$	Maximum original peak seen by the node u
b	\$/kW penalty on the tallest consumption peak
r_k^{charge}	Storage charging E-rate for the k -th energy storage
r_k^{disch}	Storage discharge E-rate for the k -th energy storage
R_k^{ramp}	Output power ramp up time of storage k
$S_{k,u,t}$	Average power fed into the k -th storage at node u in interval t
T	Total number of time intervals
ϕ_k^{energy}	Energy density of k -th storage technology; nominal energy per unit volume
ϕ_k^{power}	Power density of k -th storage technology; nominal power per unit volume
V_u^{max}	Maximum volume available for energy storage deployment at node u
α_u	Cost savings for each watt of under-provisioning at u
$\beta_{k,u}$	Amortized cost of storage k per unit energy adjusted to its lifetime
γ	Avoided electric energy cost (\$/MWh)
η	Electric transmission efficiency in the distribution
μ_k	Self discharge rate of storage technology type k

Table 2: Optimization framework notations

Equation (12) bounds the tallest peak ($Peak_u^{shvd}$) seen by u .

$$0 \leq Demand_{u,t} \leq Peak_u^{shvd}, \forall u, t \quad (12)$$

We also model the following energy storage characteristics in our framework, which are presented in the Appendix: the rate at which output power of a battery can increase, as some energy storages (such as compressed air) may take up to few minutes before delivering maximum rated power; battery self-discharge, as batteries lose some energy simply with passage of time; lifetime of the storage, as it affects the costs in the long term; volume needed for deploying energy storage, as some form of storages may need significant space, e.g., flywheels.

5. EXPERIMENTAL EVALUATION

5.1 Experimental Setup and Methodology

Configuration and Parameters: Our evaluation uses the grid distribution hierarchy shown in Figure 1. As explained in Section 2, we assume each distribution transformer supports five homes, each distribution substation serves 500 transformers, and two distribution substations are served by one bulk power substation. Storage devices can be placed at any of the levels in Figure 1. For simplic-

Avoided Energy Cost (γ)	\$3.53/MWh [6]
CapEx(Low)	\$6/kW/month
CapEx(Medium)	\$15/kW/month
CapEx(High)	\$30/kW/month
Energy cost (Contract pricing)	\$ 0.05/kWh
Peak Penalty (Transmission + Energy)	\$20/kW/month
Volume V_{homes}^{max}	10L (~Car Battery)
Volume $V_{transformer}^{max}$	25L (~2.5X Car Battery)

Table 3: Experiment Parameter Values.

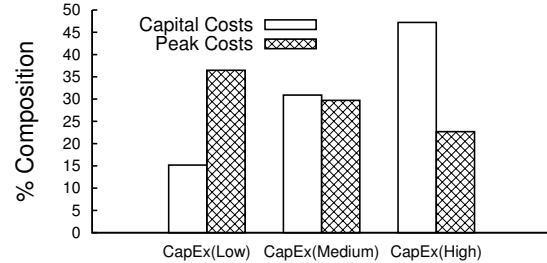


Figure 3: Composition of capex and peak penalty costs for Long-Term Contract.

ity, we present results at three levels: Homes, Transformers, and Substations, including distribution and bulk power substations.

We evaluate our framework with two op-ex cost models: a *long-term contract* and *day-ahead* model. *Long-term contract* represents the scenario where the utility either owns most of its generation or buys its energy from third parties under contracts. We adapt a real utility contract, available at [14], for evaluation. Since we include distribution costs as part of the distribution cap-ex costs, we subtract the distribution costs from the peak penalty and use the final values in Table 3. *Day-ahead* represents the case where a utility buys all of its electricity in day-ahead markets. However, the utility still incurs the peak penalty due to transmission. We use the day-ahead prices for March, 2014 from ISO New England [13]. We consider the cap-ex costs ranging from low to high, where *low* = \$6/kW/month, *medium* = \$15/kW/month, and *high* = \$30/kW/month ([27, 6, 26]).

The exact distribution of cap-ex, i.e., infrastructure cost (α), for the grid’s distribution hierarchy is not available. Thus, in this paper, we assume these costs are equally split across all the levels. As the space available at homes and distribution transformers is limited, we use a conservative value of 0.01m^3 (or roughly the size of a car battery) for the energy storage volume at homes; for transformers we set the volume to 0.025m^3 . Substations are built on large areas, so we do not constrain the available volume at substations. All results have been amortized to daily costs and savings, which includes the amortized cost of storage over its lifetime. We present results for the five ESDs, i.e., $K = 5$, discussed earlier.

We use the terms hybrid or multi-level energy storage to imply a combination of different storage technologies at a given node. Note that by using real-world day-ahead market prices and real utility consumption traces, we experiment with fine-grained time-varying prices and power consumption. Also, as we are solving an optimization problem, the computed storage values can be fractional. Although the computed numbers could differ from the actual storage capacity deployed in practice, we do not expect significant deviations from the computed values.

Workloads: For empirical evaluations, we use power consumption traces obtained from a local electric utility collected over one month (March 2014). Our traces are representative of consumption in a real electric grid. The traces contain power consumption data from 5000 homes at five minute granularity. In aggregate, we have




Figure 4: Electricity demand on a representative weekday.

more than 1.4 million unique power measurements. The average daily energy consumption of individual homes in our traces range from 15 kWh to 73 kWh.

Figure 4 shows the aggregate grid demand on a representative weekday. The figure shows that the homes peak between 6AM to 9AM (breakfast peak) and 6PM to 9PM, i.e., at dinner time. The pattern is expected based on typical work patterns, where home’s electrical activity is concentrated after office hours. Throughout this paper, unless specified otherwise, we compute the peaks at a 30 minute granularity. To gain insights into the results, we present a detailed analysis of our results on a randomly picked weekday. Later, we present the results on traces for March, 2014.

5.2 Potential Savings from Storage

Can Energy Storage Reduce Distribution Costs? We first evaluate the savings from deploying only lead-acid batteries at a single level, i.e., either at homes, or transformers, or substations. Figure 5(a) shows the percentage distribution cost savings corresponding to a low, medium, and high capital expenditure for the long-term contract pricing plan. Figure 5(a) depicts the daily percentage cost savings for *long-term contract*, which shows that even a lead-acid deployment only at homes under a low cap-ex can cut costs by 3.75%. Savings increase as cap-ex increases. Also, note that for all single-level deployments, deployment at homes shows the maximum savings. This happens because peak shaving at the lowest level (homes) provisions the infrastructure at all levels for a lower peak, thereby saving significantly in cap-ex. In contrast, savings from deployments at the transformer-level are the lowest because of the limited volume availability, which is much smaller than at homes and substations.

Result: Deploying energy storage, in this case lead-acid batteries, at a single level of the hierarchy modestly reduces costs. Deploying at the lowest level, i.e., in homes, shows the greatest savings, since it also affects peak demands at higher levels.




Figure 5: Savings from deploying lead-acid battery storage at (a) single level and (b) multiple levels under the long-term contract pricing plan.

Is Multi-Level Energy Storage Deployment Beneficial? Since related work suggests deploying lead-acid batteries only at homes, we next evaluate the impact of deploying lead-acid batteries across multiple levels of the hierarchy on savings. Figure 5(b) compares the savings from multi-level lead-acid deployment with its deployment only at homes (single-level). Savings are shown correspond-

ing to low, medium, and high capital expenditures under the long-term contract pricing plan.

For all cap-ex values, savings from a multi-level deployment surpasses that of a single-level deployment at homes. In addition, for high cap-ex, the daily cost savings from multi-level lead-acid energy storage deployment shows an increase of more than 60% compared to single-level deployment at homes.

Result: Deploying one ESD type, in this case lead-acid batteries, at multiple levels of the grid’s hierarchy further increases the cost savings up to an additional 60%.

Is Hybrid Energy Storage Deployment Beneficial? We next evaluate the additional savings possible from deploying multiple, i.e., hybrid, storage technologies at any given (single) level over a corresponding lead-acid storage deployment.

Figure 6 compares the percentage cost savings from hybrid energy storage deployment at single levels with the corresponding lead-acid deployment. Savings are shown for storage deployment at homes, transformers, and substations. Figure 6(a) and (b) shows results for the long-term contract, 6(a) is with low capex and 6(b) is with high capex.

We find that deploying hybrid energy storage boosts savings compared to lead-acid deployments, e.g., in Figure 6(a) hybrid energy storage at substations increases savings by 103%. Note that as opposed to our observation in Figure 5(a), in Figure 6(a), savings for hybrid deployment at substations is higher than that of homes. This occurs because under low cap-ex, substations can save more from greater peak shaving with hybrid energy storage, in large part, because there is no volume constraint at substations. However, as cap-ex increases, it becomes a greater component of the cost (as shown in Figure 3); as a result, the savings from deployment at homes is more than the substation’s savings for high cap-ex.

Result: Employing multiple ESD technologies in hybrid further increases savings relative to only using lead-acid batteries at any single level to as much as ~10%. Hybrid deployments are able to best match the usage pattern at any given level with the characteristics of the energy storage device.




Figure 6: Savings from deploying hybrid storage technologies at a single-level for low and high cap-ex costs under long-term contract.

Is Multi-Level Hybrid Energy Storage Deployment worth it?

Figure 7 shows how a multi-level multi-technology storage deployment can further increase savings over, first, any single level multi-technology deployment (7(a)), and second, any multi-level lead-acid (single storage technology) deployment (7(b)). Savings are shown for low, medium, and high values of cap-ex. Figure 7(a) shows that multi-level hybrid solution outperforms all the single-level hybrid solutions under all cap-ex values. For instance, 52% improvement over best single level solution under high cap-ex. We further find that a multi-level multi-technology storage deployment saves more than multi-level lead-acid deployment, as shown in Figure 7(b): for example, 83% increase in savings under low cap-ex. This increase in savings results from stringent volume constraints at lower levels, where multi-technology solutions gain an advantage

by including storage technologies with higher power and/or energy density, as opposed to lead-acid storage.

Result: A hybrid, multi-level deployment results in the greatest savings, by as much as 12%, since it is able to exploit different energy storage device characteristics at each level of the grid hierarchy, which exhibits different usage patterns.




Figure 7: Multi-Level Hybrid ESDs v/s: (a) Single-level Hybrid ESDs and (b) Multi-level lead-acid batteries under long-term contract.

How do the savings change under other pricing plans? While the above results depict savings for the long-term contract pricing plan, we have repeated all of the above experiments for the day-ahead pricing plan. In each case, we find similar cost savings for the day-ahead pricing when compared to the long-term contract plan. For example, Figure 8(a) and (b) depicts savings from a multi-level LA deployment and a multi-level hybrid deployment, respectively. As can be seen, the corresponding savings under long-term contract pricing, depicted in Figure 5(b) and 7(a), are similar to that depicted in Figure 8(a) and (b) under day ahead pricing. Since all experiments show similar results, we omit the remaining graphs for brevity (see [31] for detailed results). **Result:** Overall our experiments show that the savings due to energy storage are not specific to a pricing plan and hold for both long-term contract and day ahead pricing.




Figure 8: Savings under day-ahead pricing for multi-level lead acid batteries and multi-level hybrid storage.

Peak Reduction: Figure 9 shows the aggregate percentage peak reduction across the grid with lead-acid only, and multi-technology storage deployments. For each of the cases, we present results for both single-level deployments at homes, transformers, substations, and multi-level deployment across the hierarchy. Results are presented for long-term contract (9(a)) and day-ahead pricing (9(b)). Only medium cap-ex numbers are shown, as the numbers for other cap-ex are similar. Figure 9(a) shows even a lead-acid deployment at homes achieve a peak reduction of 11.6%, which is further increased to 16.7% by a hybrid energy storage deployment.

As we have seen, hybrid solutions have the advantage of choosing storage technologies with greater power/energy density and discharge rates. In addition, note that peak reductions achieved by substation and multi-level deployments (both hybrid and lead-acid) are very close; however, the savings achieved by them are much different (e.g., 32% difference in Figure 7(a)). Although they get similar aggregate peak reductions, the multi-level approach saves more in cap-ex by deploying energy storage devices at the lower levels. Figure 9(b) shows equivalent results for day-ahead pricing.




Figure 9: Aggregate Peak Reduction from Lead-Acid only and Hybrid ESDs for CapEx(Medium).

Result: Hybrid ESD deployments at multiple levels results in the greatest reduction in peak demand (by as much as 25%) compared to deploying hybrid ESD at individual levels or using only lead-acid batteries as the ESD.

Optimal Energy Storage Placement and Configuration: To give insight about the different types and configurations of storage technologies selected by our framework, Table 4 presents the energy storage configuration under the contract pricing for medium cap-ex. Configuration for the other cap-ex values and pricing are similar. To give an idea about absolute numbers, we include dollar savings and cost values. First, we see that if we are to use a single-level solution—just lead-acid batteries—deployment at homes does provide the best savings, because of the cap-ex gains at all levels. Second, volume constraints play an important role in limiting the benefits of lead-acid batteries in the lower level of the hierarchy; therefore, a single-level hybrid storage solution is able to increase savings by deploying a higher energy and power dense storage device, such as lithium-ion batteries, e.g., 42.5% increase at the homes. In addition, at substations where there is no volume constraint hybrid solutions employ a combination of lithium-ion, ultra-capacitors, and compressed air energy storage and further increases the savings by 15%.

Compressed air is the cheapest form of storage, however, it has a long start-up delay; ultra-capacitor and lithium-ion can be used to bridge this delay; ultra-capacitors have a very high power density, which helps in shaving tall narrow peaks, and their low energy density is complemented by lithium-ion energy storage. Finally, with the freedom of hybrid storage for multi-level deployment, we get maximum savings by deploying lithium-ion at lower levels and compressed air storage at the top level of the hierarchy.

Result: Using different ESDs at different levels of the grid’s hierarchy result in significant differences in costs and savings.

Energy Storage Costs: As the numbers in Table 4 show, the cost of energy storage is a small fraction of the total daily costs without storage devices. For example, a hybrid storage solution at the substations is only 1.83% of the total daily cost. A hybrid storage deployment at lower levels costs more than deploying lead-acid

	LA Single (\$ save, \$ cost)	Hybrid Single (\$ save, \$ cost)	Hybrid MultiLevel (\$ save, \$ cost)
Substations	LA (880.25,416.74)	LI + UC + CAES (1544.73,394.27)	CAES
Transformer	LA (558.87,125.0)	LI (841.70,236.25)	LI
Homes	LA (940.06,249.77)	LI (1340.36,472.20)	LI (2039.53,766.21)

Table 4: Storage Configuration and Placement (Long-term Contract, Medium capex): (Savings\$/day), Storage costs (\$/day)). Total cost without storage is \$21.5k/day.




Figure 10: Average daily savings for March (Day-Ahead).

batteries because lithium-ion batteries are more expensive. However, due to their higher energy and power density they also save more. Even the most expensive energy storage deployment, i.e., multi-level hybrid, incurs less than 3.6% of the total costs; most of its costs are from the lithium-ion deployment at the lower levels.

Result: The cost of energy storage is a small fraction of the total daily distribution costs without any energy storage capacity.

5.3 Longer-term Savings

For computational tractability, so far, we have presented results on a single day. However, to show that the savings hold over longer periods, we conducted experiments over an entire month. Figure 10 shows the average daily cost savings for the month of March, 2014 from our traces. Due to space constraints, savings are shown only for low and high cap-ex for day-ahead pricing; three types of deployments are shown: lead-acid at homes, multi-level lead-acid, and multi-level hybrid. Here, hybrid multi-level can save up to 11.7% for high cap-ex, and up to 9.8% for low cap-ex, which outperforms the multi-level lead-acid (low cap-ex) by 190%. The general trends in the figure are similar to what we have already seen. As peaks become taller, there is a greater opportunity for savings.

6. RELATED WORK

Much of the work on DR in the grid using ESDs has focused on cutting costs for customers with storage in presence of variable electricity pricing. For example, [24] presents an optimization approach to cut costs using ESDs in presence of spot electricity prices. Similarly, in [29] authors propose the use of energy storage in homes to cut their electricity bills under a variable prices [38], which they model as a Markov Decision Process. However, none of these approaches specifically look at cutting the costs for the utility. In fact, as noted in [22], such approaches can increase the peak demand on the grid and thereby increase its op-ex and cap-ex.

In contrast to the work above, the authors in [30] propose the use of ESDs for cutting peak demand on the grid and reducing its generation costs. Prior work has also proposed renewable energy integration to reduce consumption from the grid, e.g., [40, 36].

However, all of these consider ESD/renewable deployment only at customer premises (homes), and they evaluate their solutions only for a specific ESD technology. Finally, there has been considerable work in cost-aware provisioning and DR for datacenters, e.g., [33, 39]. The closest to our work is that done by Wang et al. [39]; here, the authors present a framework for modeling different ESDs, and the tradeoffs of placing them at different levels of datacenter power hierarchy. The authors evaluated the proposed framework using traces from real datacenters. As opposed to this, we have formulated and evaluated the solution for an electricity distribution network. We model several distribution network features which are absent in a datacenter, e.g., power losses in distribution.

7. CONCLUSIONS

In this paper, we study the novel problem of ESD deployment across distribution grid hierarchy for enabling automated Demand Response (DR). We present a generalized optimization framework for ESD deployment and control across the distribution grid hierarchy. Our framework can provide ESD provisioning solutions for a range of consumption profiles, electricity pricing plans, ESD technologies, and distribution networks. We showed that ESD provisioning can save up to 12% daily costs in distribution for the utility companies. In addition, we also present several key insights regarding ESD deployment in the distribution network.

Our work has some limitations, which we plan to address as part of future work. For example, our current model assumes the marginal value of reducing peak usage is constant, whereas in practice the marginal value varies with the magnitude of the peak. We also do not consider the impact of renewable generation, which may alter the cost of reducing peak demand. Our models assume linearity to keep the problem tractable, although there are many characteristics of ESDs, and particularly batteries, that are non-linear, e.g., capacity as a function of discharge rate due to Peukert’s law. Finally, while our capital and operational cost estimates are based on publicly available sources, and we evaluate our system over a wide range of possible costs, e.g., high, medium, and low cap-ex, these estimates may vary widely across utilities, which would effect the possible savings in the real world. However, our methodology is general and can be applied to utilities with different costs and distribution hierarchies.

8. ACKNOWLEDGEMENTS

We thank our shepherd, Mathijs de Weerd, and reviewers for their comments. This research was supported by NSF grants CNS-1413998, CNS-1253063, CNS-1503590.

References

- [1] Overview of CON EDISON System and LIC Network. <http://bit.ly/1SB6YXA>.
- [2] Utilities. http://www.nyc.gov/html/sirr/downloads/pdf/final_report/Ch_6_Utillities_FINAL_singles.pdf.
- [3] A Guide to Understanding Battery Specifications. http://web.mit.edu/evt/summary_battery_specifications.pdf, 2008.
- [4] Biennial Report of the North Carolina Utilities Commission. <http://www.ncuc.net/reports/CostAllocationReport09.pdf>, 2009.
- [5] Assessment of Transmission and Distribution Losses in New York. Technical Report PID071178 (NYSERDA 15464), EPRI, Palo Alto, CA, 2012.

- [6] Avoided Energy Supply Costs in New England: 2013 Report. <http://bit.ly/1LF0SzV>, July 2013.
- [7] Electric Bill Breakdown. <http://bit.ly/1KyBhLT>, 2013.
- [8] Beacon Power. <http://beaconpower.com/>, 2014.
- [9] Distribution Transformer. Wikipedia, 2014.
- [10] Electric Utility. Wikipedia, 2014.
- [11] Explanation of Unbundled Electricity Charges Residential Customers. <http://bit.ly/1KyBhLT>, July 2014.
- [12] Flywheel. <http://bit.ly/1FQbqwa>, 2014.
- [13] ISO New England - Hourly Data. http://www.iso-ne.com/markets/hrly_data/index.html, 2014.
- [14] Large General Service Contract Schedule. <http://www.hged.com/customers/payment-billing-rates/rates/electric-rates/183E%20-%20Large%20General%20Service.pdf>, 2014.
- [15] Lead-acid Battery. http://en.wikipedia.org/wiki/Lead%E2%80%93acid_battery, 2014.
- [16] Lithium-Ion Battery. Wikipedia, 2014.
- [17] Maxwell Ultracapacitor Uninterruptible Power Supply (UPS) Solutions. <http://bit.ly/1Brz01K>, 2014.
- [18] Summary of Rates for Public Service. <http://bit.ly/1eA6QaZ>, 2014.
- [19] PowerWall: Tesla Home Battery. <http://www.teslamotors.com/powerwall>, June 2015.
- [20] D. Akay and M. Atak. Grey prediction with rolling mechanism for electricity demand forecasting of turkey. *Energy*, 2007.
- [21] K. Bradbury. Energy Storage Technology Review. http://people.duke.edu/~kjb17/tutorials/Energy_Storage_Technologies.pdf, August 2010.
- [22] T. Carpenter, S. Singla, P. Azimzadeh, and S. Keshav. The impact of electricity pricing schemes on storage adoption in ontario. In *ACM e-Energy*, 2012.
- [23] H. Chen, T. N. Cong, W. Yang, C. Tan, Y. Li, and Y. Ding. Progress in electrical energy storage system: A critical review. *Progress in Natural Science*, 2009.
- [24] B. Daryanian, R. Bohn, and R. Tabors. Optimal Demand-side Response to Electricity Spot Prices for Storage-type Customers. *Power Systems, IEEE Transactions on*, 1989.
- [25] U. S. EIA. How much electricity is lost in transmission and distribution in the united states? - FAQ, 2014. [July-2014].
- [26] B. Horii and E. Cutter. Energy Efficiency Avoided Costs 2011 Update, December 2011.
- [27] K. Knapp, J. Martin, S. Price, and F. Gordon. Costing Methodology for Electric Distribution System Planning, November 2000.
- [28] L. L. N. Laboratory. U.S. Energy Flowchart 2013. <https://flowcharts.llnl.gov/>, August 2014.
- [29] A. Mishra, D. Irwin, P. Shenoy, J. Kurose, and T. Zhu. Smartcharge: Cutting the Electricity Bill in Smart Homes with Energy Storage. In *ACM e-Energy*, 2012.
- [30] A. Mishra, D. Irwin, P. Shenoy, and T. Zhu. Scaling Distributed Energy Storage for Grid Peak Reduction. In *ACM e-Energy*, 2013.
- [31] A. Mishra, R. Sitaraman, D. Irwin, T. Zhu, P. Shenoy, B. Dalvi, and S. Lee. Integrating energy storage in electricity distribution networks. Technical Report UM-CS-2015-012, UMass Computer Science Tech Report, May 2015.
- [32] A. T. C. on National Public Radio. Interview with Dan Delurey, President Demand Response Smart Grid Coalition. *How Smart is the Smart Grid?*, July 7th, 2010.
- [33] S. Pelley, D. Meisner, P. Zandevakili, T. F. Wenisch, and J. Underwood. Power Routing: Dynamic Power Provisioning in the Data Center. In *ACM Sigplan Notices*, 2010.
- [34] W. Pentland. Perverse Economics of the Electric Grid: As Generation Gets Cheaper, Transmission Costs Soar. <http://onforb.es/1HJM5AP>, 2013.
- [35] S. M. Schoenung and W. V. Hassenzahl. Long-vs. Short-Term Energy Storage Technologies Analysis. A Life-Cycle Cost Study. A Study for the DOE Energy Storage Systems Program.
- [36] J. Taneja, D. Culler, and P. Dutta. Towards Cooperative Grids: Sensor/Actuator Networks for Renewables Integration. In *SmartGridComm*, October 2010.
- [37] G. K. Tso and K. K. Yau. Predicting electricity energy consumption: A comparison of regression analysis, decision tree and neural networks. *Energy*, 2007.
- [38] P. Van de Ven, N. Hegde, L. Massoulie, and T. Salonidis. Optimal Control of Residential Energy Storage Under Price Fluctuations. In *Energy*, 2011.
- [39] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy. Energy Storage in Datacenters: What, Where, and How Much? *SIGMETRICS*, 2012.
- [40] T. Zhu, A. Mishra, D. Irwin, N. Sharma, P. Shenoy, and D. Towsley. Efficient Renewable Energy Management for Smart Homes. In *BuildSys*, November 2011.

APPENDIX

Below are the additional constraints of the framework presented in section 4. Table-2 defines the notations. Constraint (13) limits the rate at which an energy storage's output power can increase, this is similar to acceleration of vehicles. Here $constant_k = \frac{r_k^{disch} * I}{R_k^{ramp}}$, and R_k^{ramp} is the power ramp-up time. As batteries lose some energy simply with passage of time, we model this battery self-discharge in constraint (14). constraint (15) bounds the lifetime of the storage. As the lifetime is primarily determined by the number of charge-discharge cycles and the depth of discharge, (15) bounds the number of times an energy storage can be discharged to its allowed depth of discharge in the given time horizon. Finally, we restrict the maximum volume for storage deployment that might be available at node u in (16) and (17).

$$D_{k,u,t} - D_{k,u,t-1} \leq C_{k,u} * Constant_k, \forall k, u, t \geq 2 \quad (13)$$

$$E_{k,u,t} = (1 - \mu_k) * E_{k,u,t-1} + e_k * S_{k,u,t-1} * I - D_{k,u,t-1} * I, \quad \forall k, u, t \geq 2 \quad (14)$$

$$\frac{I * \sum_{t=1}^T D_{k,u,t}}{DoD_k^{max} * C_{k,u}} \leq NumChDischCycles_k, \forall k, u \quad (15)$$

$$\sum_{k=1}^K \frac{C_{k,u}}{\phi_k^{energy}} \leq V_u^{max}, \forall u \quad (16)$$

$$\sum_{k=1}^K \frac{C_{k,u} * r_k^{disch}}{\phi_k^{power}} \leq V_u^{max}, \forall u \quad (17)$$