

## Optimizing software defined battery systems for transformer protection

Sonia Martin <sup>a,1,\*</sup> , Obidike Nnorom Jr. <sup>b,1</sup> , Philip Levis <sup>c</sup> , Ram Rajagopal <sup>b,d</sup>

<sup>a</sup> Department of Mechanical Engineering, Stanford University, 440 Escondido Mall, Stanford, 94305, CA, USA

<sup>b</sup> Department of Electrical Engineering, Stanford University, 350 Jane Stanford Way, Stanford, 94305, CA, USA

<sup>c</sup> Department of Computer Science, Stanford University, 353 Jane Stanford Way, Stanford, 94305, CA, USA

<sup>d</sup> Department of Civil & Environmental Engineering, Stanford University, 472 Via Ortega, Stanford, 94305, CA, USA

### HIGHLIGHTS

- Proposes protecting residential transformers by aggregating home battery systems.
- Joint optimization of batteries can reduce transformer aging by 48 %.
- Joint optimization also reduces costs compared to individual optimization.
- Battery virtualization achieves these gains with 75 % battery sharing.

### ARTICLE INFO

**Keywords:**

Battery energy storage systems  
Virtualization  
Transformer aging  
Optimization  
Model predictive control

### ABSTRACT

Residential electric vehicle charging causes large spikes in electricity demand that risk violating neighborhood transformer power limits. Battery energy storage systems reduce these transformer limit violations, but operating them individually is not cost-optimal. Instead of individual optimization, aggregating, or sharing, these batteries leads to cost-optimal performance, but homeowners must relinquish battery control. This paper leverages virtualization to propose battery sharing optimization schemes to reduce electricity costs, extend the lifetime of a residential transformer, and maintain homeowner control over the battery. A case study with simulated home loads, solar generation, and electric vehicle charging profiles demonstrates that joint, or shared, optimization reduces consumer bills by 56 % and transformer aging by 48 % compared to individual optimization. Hybrid and dynamic optimization schemes that provide owners with autonomy have similar transformer aging reduction but are slightly less cost-effective. These results suggest that controlling shared batteries with virtualization is an effective way to delay transformer upgrades in the face of growing residential electric vehicle charging penetration.

### 1. Introduction

Increasing residential electric vehicle (EV) penetration can harm distribution transformers [1]. Since utility companies chose many transformer sizes before widespread installation of residential EV chargers, high peak power from charging can lead to capacity violations and consequently, transformer aging [2–4]. This rapid aging forces expensive and time-consuming transformer upgrades [5]. As of 2023, distribution transformer prices in the U.S. have more than doubled, and the installation wait time exceeds one year [6]. As more EV chargers are installed in neighborhoods with limited transformer capacities, it

is in the utilities' best interest to implement techniques to prolong transformer lifetimes instead of upgrading these assets [7].

While controlling EV charging can help reduce peak loads and protect transformers [1,5,8–13], drivers may not respond perfectly to charging control signals, which risks unplanned transformer limit violations [14]. Instead, controlling stationary distributed energy resources, such as stationary battery energy storage systems (BESSs), offers a solution to protect transformers that is decoupled from drivers' charging behavior [15,16]. For example, a homeowner with a BESS can take advantage of excess solar photovoltaic energy generated to charge the BESS and then discharge it during an EV charging session to offset peak

\* Corresponding author.

Email addresses: [soniamartin@stanford.edu](mailto:soniamartin@stanford.edu) (S. Martin), [obdk@stanford.edu](mailto:obdk@stanford.edu) (O. Nnorom Jr.).

<sup>1</sup> Equal contribution.

demand. With this type of peak shaving, BESSs not only reduce peak power costs but also mitigate the large power spikes from EV charging that can harm transformers.

Many studies on BESSs co-located with EVs focus only on controlling a single BESS, with control objectives ranging from transformer protection [17,18] and cost minimization [19] to voltage control [20]. However, a single centrally controlled BESS does not account for individual homeowner needs if implemented in a residential neighborhood or distribution network [21]. In regions with frequent power outages or high peak prices, a homeowner buying into a BESS system may want to retain autonomy over their BESS to maintain reserve power.

Especially in communities, combining, or aggregating, multiple BESS units creates shared infrastructure from which entire neighborhoods can benefit [22–25]. Aggregation can increase both BESS utilization and decrease consumer costs: correctly sized shared systems that take into account solar and EV system sizes [26,27] have been shown to increase utilization by over 35 % and decrease costs by over 10 % [28–31].

However, similar to single BESS control, these aggregated systems are often static. Once a set of batteries is shared, it cannot easily be reverted back to individual control, and vice versa. With unpredictable outages, loads that change seasonally, and variable EV charging patterns, a static aggregation scheme may be suboptimal. In such a system where adaptability is crucial, there is no consensus on how exactly BESSs should be split among homeowners and how this split should change over time.

Battery virtualization is a new approach that simplifies operation and allows for more flexibility in the battery setup in the face of variable grid conditions. Virtualization provides an abstraction of a virtual resource that is equivalent to an underlying physical resource [25,32–34], hiding complexities of the physical resource and making it easier to interact with it. Battery virtualization involves treating each physical battery as an abstract entity that can be partitioned into separate parts or aggregated with other abstract, or software defined, batteries. For example, a group of homes in a neighborhood can use virtualization to combine all of their physical batteries and treat them as a single aggregate battery. Virtualization also enables a hybrid aggregation approach, in which an owner can maintain control of one battery partition while contributing another partition to a shared aggregate. If constituent batteries are added or drop offline [25,32], virtualization recalibrates without interruption. Lastly, virtualization helps the system adapt to seasonality changes and uncertainty in forecasts when running a model predictive control (MPC) scheme for BESS control [20,35–38].

Prior work has explored BESS aggregation [22–25] and cost minimization techniques [19] separately. Combining these two approaches can help protect transformers in the face of increased EV adoption. Flexible and dynamic BESS sharing as well as aggregation schemes allowing homeowners to retain autonomy over part of their BESS have not yet been investigated for this purpose.

This paper explores how BESS operators can leverage virtualization software [32] to seamlessly control a set of aggregated residential BESSs to both minimize costs and mitigate transformer loss of life while maintaining homeowner autonomy. In contrast to existing work that utilizes static, aggregated BESSs for transformer protection, we compare the performance of multiple static and dynamic aggregation schemes in a case study with simulated load and EV data. Lastly, we test the sensitivity of the algorithm to imperfect forecasting by implementing an MPC scheme and analyzing results across seasons.

We find that introducing flexibly aggregated BESSs reduces expensive transformer upgrades and decreases operational costs. Our proposed algorithms are useful for homeowners, energy aggregators, and utilities aiming to avoid these transformer upgrades, decrease electricity costs, and increase BESS utilization within a residential neighborhood. These cost reductions enable more widespread BESS and EV charger installations, paving the way for cleaner light-duty electric transportation.

This work makes the following contributions:

- Optimization schemes enabled by virtualization for a group of BESSs to reduce costs and transformer aging by sharing the batteries,
- A “hybrid” optimization scheme that flexibly combines individual and joint (shared) control using virtualization,
- A “dynamic” optimization scheme that leverages virtualization to dynamically balance individual and joint control,
- Identification of the best battery sharing fraction to simultaneously reduce costs and maintain homeowner autonomy, and
- Analysis of each scheme under imperfect forecasts, with variable pricing schemes, across seasons, with varying EV penetration, and considering battery aging.

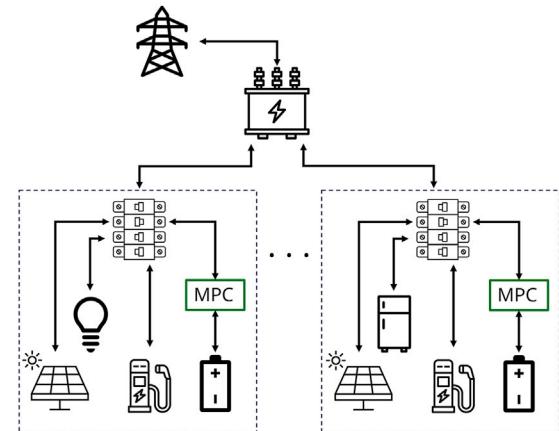
The paper is organized as follows: Section 2 explains the optimization schemes, Section 3 describes the transformer model, and Sections 4 and 5 present simulation results. Section 6 discusses the implications of these results while Section 7 outlines future work.

## 2. Control algorithms

This section discusses the optimization strategies, controller, system sizing, forecasting, and electricity cost structure. We consider a setting where a collection of homes in a neighborhood has solar systems, EV chargers, residential BESSs, and various home loads. Fig. 1 displays these components for two homes, connected to the distribution grid via a transformer. This work analyzes the homes downstream of one residential transformer. We assume that solar generation and all loads are fixed (i.e., no flexible demand) and that each home's battery is controllable.

We compare five different battery optimization schemes to demonstrate the effects of virtualization on cost and transformer aging:

- **Individual optimization** assumes that each home in a neighborhood has its own battery that it controls to minimize its individual cost,
- **Individual optimization with uneven transformer allocation** is similar to the first scheme but allows homes to have an uneven limit allocation instead of distributing the limit evenly,
- **Joint optimization** controls all of the batteries in a neighborhood as if they were a single aggregate battery,
- **Hybrid optimization** blends the individual and joint schemes, and



**Fig. 1.** Diagram of a residential power system. A set of homes, connected to the grid under a single transformer, contains home loads, solar panels, a battery, and an EV charger. Each of these loads is connected to the transformer via a circuit breaker. The model predictive controller determines the battery's output.

- **Dynamic partitioning optimization** builds on the hybrid scheme by adding dynamic partition values between retained and shared portions of the battery.

We describe the schemes below, with variables listed in [Table 1](#).

**Table 1**  
Algorithm variables.

Symbol	Variable name
$n$	Number of homes
$t$	Timestep
$T$	Total number of timesteps
$\lambda$	Transformer violation penalty weight
$K$	Nominal transformer limit [kW]
$\alpha$	(Dis)charging penalty weight
$\eta$	One-way battery efficiency
$\Delta t$	Timestep duration [h]
$E_{\max}$	Maximum battery capacity [kWh]
$h$	Hours between change in partition weights
$T^{\text{MPC}}$	Number of timesteps per MPC iteration
$B^{\text{dischg}}, B^{\text{chg}} \in \mathbb{R}^+$	Maximum battery (dis)charge power [kW]
$\mathbf{B}^{\text{dischg}}, \mathbf{B}^{\text{chg}} \in \mathbb{R}^{T \times n}$	Battery (Dis)charge power [kW]
$\mathbf{M} \in \mathbb{R}^{T \times n}$	Home meter reading [kW]
$\mathbf{C} \in \mathbb{R}^T$	Electricity cost [ $\frac{\text{S}}{\text{kWh}}$ ]
$\mathbf{S} \in \mathbb{R}^{T \times n}$	Home solar generation [kW]
$\mathbf{L} \in \mathbb{R}^{T \times n}$	Home load demand [kW]
$\mathbf{E} \in \mathbb{R}^{(T+1) \times n}$	Available battery capacity [kWh]
$\mathbf{K}^U \in \mathbb{R}^n$	Uneven transformer allocation [kW]
$\mathbf{B}^{\text{R}, \text{dischg}} \in \mathbb{R}^{T \times n}$	Retained discharge power [kW]
$\mathbf{B}^{\text{S}, \text{dischg}} \in \mathbb{R}^{T \times n}$	Shared discharge power [kW]
$\mathbf{B}^{\text{R}, \text{chg}} \in \mathbb{R}^{T \times n}$	Retained charge power [kW]
$\mathbf{B}^{\text{S}, \text{chg}} \in \mathbb{R}^{T \times n}$	Shared charge power [kW]
$\mathbf{B}^{\text{R}} \in \mathbb{R}^{T \times n}$	Retained net power [kW]
$\mathbf{B}^{\text{S}} \in \mathbb{R}^{T \times n}$	Shared net power [kW]
$\mathbf{W} \in \mathbb{R}^{T \times n}$	Retained battery fraction
$\mathbf{E}^{\text{R}} \in \mathbb{R}^{(T+1) \times n}$	Retained battery capacity [kWh]
$\mathbf{E}^{\text{S}} \in \mathbb{R}^{(T+1) \times n}$	Shared battery capacity [kWh]
$\mathbf{Z} \in \mathbb{R}^{T \times \frac{T \Delta t}{h}}$	Dynamic partitioning transformation matrix
$\mathbf{W}^{\text{C}} \in \mathbb{R}^{\frac{T \Delta t}{h} \times n}$	Dynamic partitioning constant weights
$C^{\text{batt}}$	Battery capital cost
$N^{\text{cyc}}$	Maximum cycles under warranty

### 2.1. Individual optimization

With individual optimization, each homeowner operates their own battery. The neighborhood transformer limit  $K$  is split evenly amongst each of the  $n$  homeowners, incentivizing each home's power consumption to remain below  $\frac{K}{n}$ . Each homeowner controls their battery to minimize costs (see tariff information in [Section 2.10](#)) while ensuring their power consumption does not exceed the individual transformer limit. The following convex optimization problem solves for the optimal charge and discharge battery output for each home independently:

$$\min_{\mathbf{B}^{\text{dischg}}, \mathbf{B}^{\text{chg}}} \sum_{i=0}^n \left( \mathbf{C}^T \max(\mathbf{M}_i, 0) \cdot \Delta t \right) \quad (1)$$

$$+ \sum_{i=0}^T \left( \lambda \max \left( \mathbf{M}_{t,i} - \frac{K}{n}, 0 \right)^2 + \alpha \left( \mathbf{B}_{t,i}^{\text{chg}} + \mathbf{B}_{t,i}^{\text{dischg}} \right) \right) \quad (2)$$

$$\text{s.t.} \quad \mathbf{M}_t = \mathbf{L}_t - \mathbf{S}_t + \mathbf{B}_t^{\text{chg}} - \mathbf{B}_t^{\text{dischg}} \quad \forall t \in [0, T] \quad (3)$$

$$\mathbf{E}_{t+1} = \mathbf{E}_t - \Delta t \left( \frac{\mathbf{B}_t^{\text{dischg}}}{\eta} - \eta \mathbf{B}_t^{\text{chg}} \right) \quad \forall t \in [0, T] \quad (4)$$

$$0 \leq \mathbf{E}_t \leq E_{\max} \quad \forall t \in [0, T+1] \quad (5)$$

$$0 \leq \mathbf{B}_t^{\text{chg}} \leq B_{\max}^{\text{chg}} \quad \forall t \in [0, T] \quad (6)$$

$$0 \leq \mathbf{B}_t^{\text{dischg}} \leq B_{\max}^{\text{dischg}} \quad \forall t \in [0, T] \quad (7)$$

The objective, [Eq. \(2\)](#), has three terms: the total electricity cost for the homeowners in the community, the transformer violation penalty, and

a penalty to discourage the solver from finding solutions in which the battery charges and discharges simultaneously (since this is impossible in practice). There is no net metering, so the positive values of the meter are multiplied by the timestep length and the cost at each timestep in units of  $\frac{\text{S}}{\text{kWh}}$ .

[Eq. \(3\)](#) constrains the utility meter reading to equal the sum of the homeowner's loads and battery output minus their solar generation. [Eqs. \(4\)–\(7\)](#) govern the battery state of charge evolution, capacity minimum and maximum, and charge and discharge limits, respectively.

### 2.2. Individual optimization with uneven transformer allocation

The individual optimization scheme penalizes each homeowner for exceeding their equal allocation of the total transformer limit. In practice, not all homes have the same maximum power draw [\[39\]](#). This uneven allocation scheme optimizes the fraction of the transformer limit allocated to each home to account for this variability in maximum home power draw.

The scheme has a different objective function and two additional constraints not found in individual optimization, as listed in the equations below:

$$\min_{\mathbf{B}^{\text{dischg}}, \mathbf{B}^{\text{chg}}} \sum_{i=0}^n \left( \mathbf{C}^T \max(\mathbf{M}_i, 0) \Delta t \right. \\ \left. + \sum_{i=0}^T \left( \lambda \max(\mathbf{M}_{t,i} - \mathbf{K}_i^U, 0)^2 + \alpha \left( \mathbf{B}_{t,i}^{\text{chg}} + \mathbf{B}_{t,i}^{\text{dischg}} \right) \right) \right) \quad (8)$$

$$\text{s.t.} \quad \mathbf{K}^U \geq 0 \quad (9)$$

$$\sum_{i=0}^n \mathbf{K}_i^U = K \quad (10)$$

+ remaining constraints from individual optimization

[Eq. \(8\)](#) shows the updated objective function which replaces  $K$  in the objective with  $\mathbf{K}^U$ , a vector that represents the transformer limit for each home in kW. Next, [Eqs. \(9\)](#) and [\(10\)](#) are added constraints to ensure each home has a nonnegative limit and that the sum of these limits over all the homes is equal to the nameplate limit  $K$ . These changes allow the solver to simultaneously find the minimum cost solution and optimal transformer limit allocation.

### 2.3. Joint optimization

With joint optimization, the homeowners aggregate their batteries together using virtualization [\[32\]](#). The controller manages the aggregated battery to prevent the homeowners' total power consumption from exceeding the transformer limit whenever possible.

The joint scheme allows homeowners to collaborate to protect the transformer. Depending on the exact home energy consumption profiles, the joint scheme can be advantageous compared to the individual scheme: if any given home's battery is underutilized in the individual scheme, another home can use it. The problem is as follows:

$$\min_{\mathbf{B}^{\text{dischg}}, \mathbf{B}^{\text{chg}}} \mathbf{C}^T \max(\mathbf{M}^1, 0) \cdot \Delta t + \sum_{i=0}^T \left( \lambda \max((\mathbf{M}^1)_i - K, 0)^2 \right. \\ \left. + \sum_{i=0}^n \alpha \left( \mathbf{B}_{t,i}^{\text{chg}} + \mathbf{B}_{t,i}^{\text{dischg}} \right) \right) \quad (11)$$

$$\text{s.t.} \quad \mathbf{M}_t = \mathbf{L}_t - \mathbf{S}_t + \mathbf{B}^{\text{chg}} - \mathbf{B}^{\text{dischg}} \quad \forall t \in [0, T] \quad (12)$$

$$\mathbf{E}_{t+1} = \mathbf{E}_t - \Delta t \left( \frac{\mathbf{B}_t^{\text{dischg}}}{\eta} - \eta \mathbf{B}_t^{\text{chg}} \right) \quad \forall t \in [0, T] \quad (13)$$

$$0 \leq \mathbf{E}_t \leq E_{\max} \quad \forall t \in [0, T+1] \quad (14)$$

$$0 \leq \mathbf{B}_t^{\text{chg}} \leq B_{\max}^{\text{chg}} \quad \forall t \in [0, T] \quad (15)$$

$$0 \leq \mathbf{B}_t^{\text{dischg}} \leq B_{\max}^{\text{dischg}} \quad \forall t \in [0, T] \quad (16)$$

The aggregate battery aims to minimize the total homeowners' cost. A given customer might be restricted below their fair share of power at certain times but can save money by participating in the aggregate. The utility bills the aggregate system as a whole: the first term in the objective, Eq. (11), multiplies the sum of all the home meters and the aggregate battery output by the TOU cost vector. Similar to the individual optimization scheme, the second term in Eq. (11) penalizes transformer limit violations, but here the system is penalized when the sum of the homes' net power and aggregate battery output exceeds the limit, instead of when an individual home's power peak exceeds its allotted limit. A penalty is also included in the objective to discourage the solver from choosing to charge and discharge the battery variables simultaneously. Eq. (12) constrains each home meter to equal load demand minus solar generation. Eqs. (13)–(16) govern the battery SOC evolution, capacity minimum and maximum, and charge and discharge limits, respectively.

#### 2.4. Hybrid optimization

In the hybrid scheme, each homeowner keeps a certain fraction of their battery and gives up the remaining fraction to be aggregated into a larger battery. This allows the homeowner to participate in joint optimization while maintaining control of some of their battery, e.g., for blackouts. These fractions are denoted by  $\mathbf{E}^R$  and  $\mathbf{E}^S$  for the retained and shared partitions, respectively, and weighted by an input  $\mathbf{W}$  between 0 and 1. In this scheme, these partitions are fixed for the whole time horizon. Virtualization enables this combination of individual optimization and joint optimization by treating the shared partitions as one virtual battery and each retained partition as additional virtual batteries.

The objective of the hybrid scheme includes terms for the individual, joint, and transformer violation objectives. This balances the homeowner's goal to act within their financial interest while also sharing part of their battery in the interest of the entire community. The problem is as follows:

$$\begin{aligned} \min_{\mathbf{B}^R, \mathbf{B}^S} \quad & \sum_{i=0}^n \mathbf{C}^T \max(\mathbf{M}_i + \mathbf{B}_i^R, 0) \cdot \Delta t \\ & + \sum_{t=0}^T \left( \mathbf{C}^T \max((\mathbf{M}^1 + \mathbf{B}^S \mathbf{1})_t, 0) \right)^2 \cdot \Delta t \\ & + \lambda \max((\mathbf{M}^1 + \mathbf{B}^S \mathbf{1})_t - K, 0)^2 \\ & + \sum_{i=0}^n \alpha \left( \mathbf{B}_{t,i}^R + \mathbf{B}_{t,i}^S \right) \end{aligned} \quad (17)$$

$$\text{s.t.} \quad \mathbf{M}_t = \mathbf{L}_t - \mathbf{S}_t \quad \forall t \in [0, T] \quad (18)$$

$$\mathbf{B}^R = \mathbf{B}^{R,\text{chg}} - \mathbf{B}^{R,\text{dischg}} \quad (19)$$

$$\mathbf{B}^S = \mathbf{B}^{S,\text{chg}} - \mathbf{B}^{S,\text{dischg}} \quad (20)$$

$$\mathbf{E}_{t+1}^R = \mathbf{E}_t^R - \Delta t \left( \eta \mathbf{B}_t^{R,\text{chg}} - \frac{\mathbf{B}_t^{R,\text{dischg}}}{\eta} \right) \quad \forall t \in [0, T] \quad (21)$$

$$\mathbf{E}_{t+1}^S = \mathbf{E}_t^S - \Delta t \left( \eta \mathbf{B}_t^{S,\text{chg}} - \frac{\mathbf{B}_t^{S,\text{dischg}}}{\eta} \right) \quad \forall t \in [0, T] \quad (22)$$

$$0 \leq \mathbf{E}_t^R \leq \mathbf{W}_t E_{\max} \quad \forall t \in [0, T+1] \quad (23)$$

$$0 \leq \mathbf{E}_t^S \leq (1 - \mathbf{W}_t) E_{\max} \quad \forall t \in [0, T+1] \quad (24)$$

$$0 \leq \mathbf{B}_t^{R,\text{chg}} \leq \mathbf{W}_t B_{\max}^{\text{chg}} \quad \forall t \in [0, T] \quad (25)$$

$$0 \leq \mathbf{B}_t^{S,\text{chg}} \leq (1 - \mathbf{W}_t) B_{\max}^{\text{chg}} \quad \forall t \in [0, T] \quad (26)$$

$$0 \leq \mathbf{B}_t^{R,\text{dischg}} \leq \mathbf{W}_t B_{\max}^{\text{dischg}} \quad \forall t \in [0, T] \quad (27)$$

$$0 \leq \mathbf{B}_t^{S,\text{dischg}} \leq (1 - \mathbf{W}_t) B_{\max}^{\text{dischg}} \quad \forall t \in [0, T] \quad (28)$$

The first term in Eq. (17) mirrors the objective of the individual optimization scheme, while the second term is similar to the objective for joint optimization. The transformer penalty only occurs if the

home loads, solar, and shared battery exceed the limit, i.e., the retained portion does not factor into the penalty, giving homeowners complete autonomy over it. The constraints here, Eqs. (18) and (21)–(28), are also similar to both prior schemes, but they are duplicated for each of the retained and shared partitions. The charge/discharge and energy constraints ensure each partition is only allotted  $\mathbf{W}$  or  $1 - \mathbf{W}$  of the power and energy limits. Eqs. (19) and (20) set the net battery power equal to the charge minus the discharge power for each partition.

#### 2.5. Dynamic partitioning

While the hybrid scheme takes a static value  $\mathbf{W}$ , virtualization allows for this weight to change over time. For example, if the occupants of one home go on vacation for a week then return, it is suboptimal if their retained battery fraction remains constant over a whole month. In traditional, static optimization, this fraction is fixed over time. However, with virtualization, this fraction can be dynamic as the virtualization software handles the re-partitioning of the batteries while respecting battery power constraints. We test a case in which  $\mathbf{W}$  is allowed to vary every two hours. The optimization problem is the same as in the hybrid case, with the following additional constraints:

$$0 \leq \mathbf{W} \leq 1 \quad (29)$$

$$\mathbf{W} = \mathbf{Z} \mathbf{W}^C \quad (30)$$

Eq. (29) restricts  $\mathbf{W}$  to between 0 and 1, while Eq. (30) (using  $h = 2$ ) ensures that  $\mathbf{W}$  can only change once every two hours.

#### 2.6. Model predictive control

While a convex optimization solver can easily solve the above problems, the battery output is only optimal given perfect foresight of the solar, load, and EV charging profiles. In real time operation, these profiles are not known. We use moving horizon MPC that updates initial values if the system deviates from its forecast. At each MPC iteration, the controller solves for the next day given solar and load forecasts and outputs the desired battery trajectory. Then, the controller selects the initial battery power command and sends this value to the battery. After one timestep, the next iteration begins with a new initial battery SOC that is read from the battery. This process repeats until the end of the overall optimization horizon. With this moving window, the controller corrects for changes in the system, either due to deviations between forecasts and actual values or experimental factors such as sensor bias or error. In the results section, we compare the MPC with a perfect foresight controller, which is a one-shot optimization over the entire time horizon.

#### 2.7. System sizing

In simulation, we model the batteries as the Tesla Powerwall 2. We assume each homeowner owns a single Tesla Powerwall 2 with a maximum capacity of 13.5 kWh and a maximum (dis)charge output of 5 kW [40]. We use a fixed transformer size of 25 kVA as it represents a transformer size commonly found in residential communities [2].

We select the number of homes located downstream of the transformer based on the home loads. Since older transformers were typically sized before high prevalence of EV chargers, we adopt a heuristic to select a realistic number of homes under each transformer. The heuristic proceeds as follows. First, we randomly sample homes from the dataset. Next, we add these sampled homes, one at a time, until the maximum sum of their loads (excluding EV charging) over a month exceeds the nominal transformer limit; this is called a "violation". This set of homes forms the neighborhood behind the transformer for a particular trial.

#### 2.8. Battery aging considerations

Although a complex battery aging model is out of the scope for this work, we examine the sensitivity of the total system costs to battery aging. Specifically, we add a term to the objective function that represents

a cost in \$ per roundtrip charge and discharge cycle. This cost per cycle is derived from the cost of purchasing a new Powerwall,  $C^{\text{batt}}$  [41] divided by the number of total cycles allowed under warranty [42]. This represents the effective cost per cycle for the Powerwall. The objective term is as follows in Eq. (31):

$$\frac{C^{\text{batt}}}{N^{\text{cyc}}} \frac{\left( \sum_{t=0}^T \mathbf{B}_t^{\text{chg}} + \sum_{t=0}^T \mathbf{B}_t^{\text{dischg}} \right) \Delta t}{2E^{\text{max}}} \quad (31)$$

## 2.9. Forecasting

We use a naive forecaster that provides the solar generation and home load demand (excluding EVs) at each home for the upcoming day by averaging values from the previous four days. The naive forecaster is inaccurate for EV loads as charging times for EVs can be short and somewhat sporadic. As a result, we assume that the EV charging profiles are known in advance. This assumes that homeowners will know one day in advance (the MPC time horizon) when they plan to charge their vehicle.

## 2.10. Costs

For the main simulations, we adopt the Pacific Gas & Electric's (PG&E) Home Charging EV2-A time-of-use (TOU) rate plan as a tariff structure for each home [43]. The pricing scheme, effective as of April 2024, is representative of the costs for homeowners located in San Jose with Level 2 EV charging. We also run a simulation that uses the EV-B and TOU-D PG&E rate structures [43]. For all schemes, we analyze the total costs of a community of homes using the same tariff structure. We do not determine the individual cost distribution for each home as the exact methods used for distributing the costs are outside the scope of this paper. This cost distribution would factor in homes' baseline loads, EV charger utilization, time of peak loads, to name a few inputs. Lastly, this work only considers operational costs of the system and does not factor in capital costs of purchasing a BESS, a solar array, or an EV charger.

## 3. Transformer model

We model the lifetime and aging of an oil-type, ONAN (natural cooling) residential distribution transformer using equations from the IEEE C57.96 standard [44]. We focus on oil-type transformers rather than dry-type transformers as they are common in residential settings [1]. We obtain the model parameters for a 25 kVA low voltage transformer, such as power loss at 100 % load and weight of the coils, from a datasheet [45]. As power flow analysis is outside the scope of this work, we assume a power factor of 1 and thus treat the nominal 25 kVA limit as 25 kW. Table 2 lists the variables in the transformer aging equations.

**Table 2**  
Transformer aging variables.

Symbol	Variable name
$\Delta\Theta^{TO,U}$	Ultimate top oil temperature rise [°C]
$\Delta\Theta^{TO,R}$	Rated top oil temperature rise [°C]
$L^{PU}$	Load [per unit]
$R$	Ratio of rated load loss to no-load loss [unitless]
$n$	Empirical constant
$\tau^{TO}$	Oil time constant [s]
$\tau^{TO,R}$	Oil time constant at rated load [s]
$m$	Empirical constant
$\Delta\Theta^{TO}$	Top oil temperature rise [°C]
$\Delta\Theta^H$	Hottest spot temperature rise [°C]
$\Delta\Theta^{H,R}$	Hottest spot temperature rise at rated load [°C]
$\Theta^H$	Hottest spot temperature [°C]
$\Theta^A$	Ambient temperature [°C]
$F^{AA}$	Rate of accelerated aging [unitless]
$LT_{\text{normal}}$	Normal transformer lifetime [h]
$F^{\text{EQA}}$	Equivalent aging factor [unitless]
$\%LOL$	Percent loss of life

A transformer's lifetime is correlated with its hottest spot temperature (HST),  $\Theta^H$ , or the point in the coils with the highest temperature. Exceeding a threshold temperature for long periods of time causes increased aging, so the nominal power limit corresponds to a hottest spot temperature that can be held without abnormal detrimental aging effects. We measure the transformer's aging via percent loss of life over a certain elapsed time period. Calculating the percent loss of life requires tracking the hottest spot temperature evolution over time using these equations:

$$\Delta\Theta_{t+1}^{TO,U} = \Delta\Theta^{TO,R} \left( \frac{\left( L_{t+1}^{PU} \right)^2 R + 1}{R + 1} \right)^n \quad (32)$$

$$\tau_{t+1}^{TO} = \tau^{TO,R} \frac{\left( \frac{\Delta\Theta_{t+1}^{TO,U}}{\Delta\Theta^{TO,R}} \right) - \left( \frac{\Delta\Theta_t^{TO}}{\Delta\Theta^{TO,R}} \right)}{\left( \frac{\Delta\Theta_{t+1}^{TO,U}}{\Delta\Theta^{TO,R}} \right)^{\frac{1}{n}} - \left( \frac{\Delta\Theta_t^{TO}}{\Delta\Theta^{TO,R}} \right)^{\frac{1}{n}}} \quad (33)$$

$$\Delta\Theta_{t+1}^{TO} = \left( \Delta\Theta_{t+1}^{TO,U} - \Delta\Theta_t^{TO} \right) \left( 1 - \exp \left( \frac{-t}{\tau_{t+1}^{TO}} \right) \right) + \Delta\Theta_t^{TO} \quad (34)$$

$$\Delta\Theta_{t+1}^H = \Delta\Theta^{H,R} \left( L_{t+1}^{PU} \right)^{2m} \quad (35)$$

$$\Theta_t^H = \Theta_t^A + \Delta\Theta_t^{TO} + \Delta\Theta_t^H \quad (36)$$

Eqs. (32)–(34) govern the change in top oil temperature due to transformer loading, and Eq. (35) calculates the hottest spot temperature rise due to loading. Eq. (36) yields the hottest spot temperature at each time step by summing the ambient temperature, top oil rise, and hottest spot temperature rise.

From the hottest spot temperature, we then calculate the rate of accelerated aging, equivalent aging factor, and the percent loss of life:

$$F_t^{AA} = \exp \left( \frac{15,000}{383} - \frac{15,000}{\Theta_t^H + 273} \right) \quad (37)$$

$$F^{\text{EQA}} = \frac{\sum_{t=1}^T \Delta t F_t^{AA}}{\sum_{t=1}^T \Delta t} \quad (38)$$

$$\%LOL = \frac{100}{LT_{\text{normal}}} F^{\text{EQA}} T \quad (39)$$

The rate of accelerated aging  $F^{AA}$  is a function of hottest spot temperature Eq. (37). Integrating  $F^{AA}$  over the simulation period and dividing by total simulation time yields the equivalent aging factor  $F^{\text{EQA}}$  shown in Eq. (38). Lastly, the percent loss of life shown in Eq. (39) relates the equivalent aging factor with the rated lifetime,  $LT_{\text{normal}}$ .

## 4. Case study

To simulate each optimization scheme, we analyze a dataset with 48 unique homes in San Jose, California in 2018. We use simulated home load data from NREL ResStock [46]; each home has a solar PV array and home loads. We generate Level 2 home EV charging data to assign to each home from the NREL OCHRE residential energy model [47]. We assume a perfect power factor and that the battery is able to execute commands exactly (i.e., there is no error or bias in the battery's setpoint and sensor readings).

To calculate the cost of electricity for consumers, we construct a cost vector  $\mathbf{C}$  based on PG&E's EV2-A TOU rate plan [43]. The total simulation time horizon is four weeks, and we separately simulate four weeks in January and July 2018 beginning on a Monday (January 8 and July 2, respectively). The MPC horizon is one day. These and other simulation parameter values are listed in Table 3. In this table, the battery capacity and power limit parameters indicate those of each individual battery, not the aggregate. The transformer penalty weight,  $\lambda$ , is set to  $\lambda = 100$  to balance the system electricity costs with the transformer penalty value. The final electricity costs reported do not include this penalty value.

**Table 3**  
Parameter values.

Parameter	Value	Parameter	Value
$T_{MPC}$	24	$\Delta\Theta^{TO,R}$	65 °C
$T$	1344	$\tau^{TO,R}$	8.738 s
$\lambda$	100	R	3.625
$\alpha$	0.01	$\Theta^A$	25 °C
$\Delta t$	0.5 h	$\Delta\Theta^{H,R}$	15 °C
$\eta$	0.9487	$LT^{\text{normal}}$	180,000 h
$E^{\max}$	13.5 kWh	$C^{\text{batt}}$	\$5550
$B_{\text{dischg}}^{\max}$	5 kW	$N^{\text{cyc}}$	1400
$B_{\text{chg}}^{\max}$	5 kW		
$B_{\text{max}}$			

For each simulation month, we run 50 trials, sampled as described in Section 2.7. This sampling yields a different number of homes in each trial for January and July. Based on this random selection of homes, we observe between 6 and 22 homes located downstream of the transformer. In particular, in January, each trial has an average of 12.6 homes with a standard deviation of 2.7 homes, and in July, each trial has an average of 16.1 homes with a standard deviation of 3.4 homes.

## 5. Results

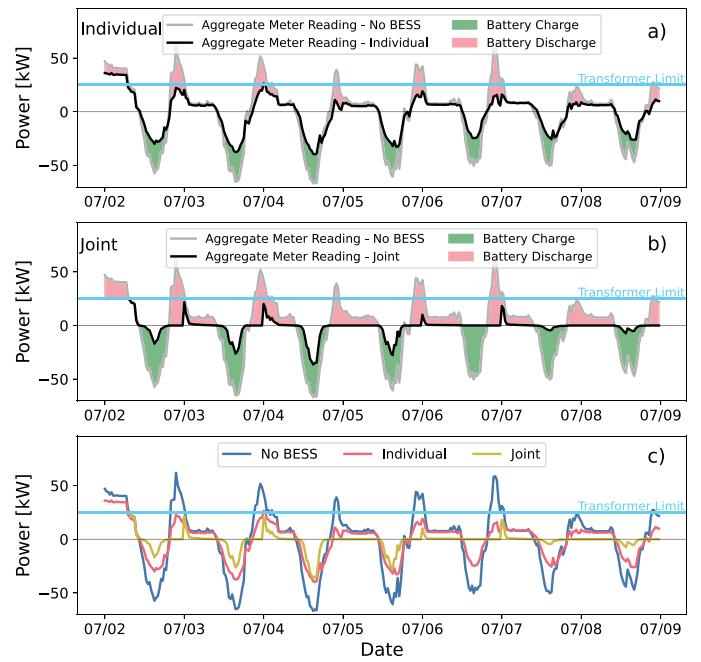
To highlight the impact of virtualization on battery control, we compare the individual, joint, hybrid, and dynamic schemes to a system without a battery. Analyzing both transformer protection and cost reduction metrics reveals differences in each scheme's performance. We also examine the effect of changing battery partition sizes in the hybrid case. Next, we analyze the sensitivity of the schemes to different pricing plans, imperfect forecasts, different seasons, varying EV penetration levels, and battery aging costs. Finally, we present experimental results from a small-scale physical testbed.

### 5.1. Meter and transformer impacts of battery aggregation

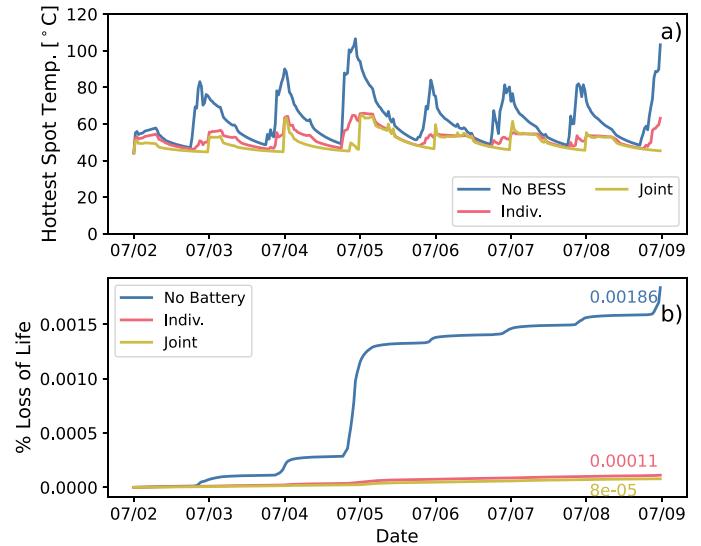
While the individual scheme, in which each homeowner controls their own BESS, shows a drastic improvement over a system with no BESSs, it still does not fully utilize the full BESS potential. Fig. 2 compares the aggregate meter values for the system in the first week of July 2018 for a single trial. On each of the plots (individual scheme in Fig. 2(a) and joint scheme in Fig. 2(b)), the black curve shows the aggregate (summed) net meter for the collection of homes in the trial. The gray curve shows the net meter with no BESSs present in the system. The shaded areas indicate the BESS behavior that makes up the difference between the black and gray curves (green shading is charging while red shading is discharging). As shown in the green shaded areas, the BESS takes advantage of excess solar to charge. It uses this to discharge in the evening to offset high evening loads (shaded red).

The joint scheme maintains the aggregate meter under the transformer limit for the entire simulation period, but the individual scheme violates the limit in the first few hours of the week, as shown by the black curve exceeding the blue transformer limit line. This result is highlighted in the direct comparison in Fig. 2(c). The joint scheme utilizes the full aggregate BESS capacity to minimize excess solar generation that would otherwise be curtailed. The individual scheme cannot capture this excess solar. For example, on 7/6, the joint scheme BESS uses all of the excess solar, which the individual scheme cannot achieve.

Fig. 3 displays the HST and percent loss of life during the first week of July 2018 for the same single trial. While there is a cumulative low loss of life without a BESS, the HST often spikes more than 20 °C above the individual and joint schemes with a BESS, whose HST remains below 70 °C during this week. The individual and joint schemes have similar losses of life, underscoring that just adding BESSs to a system without aggregation still has a substantial effect on transformer protection. Lastly, this plot demonstrates that loss of life is monotonically increasing: no loading pattern can decrease loss of life.



**Fig. 2.** Aggregate meter comparison between (a) individual, (b) joint, and (c) both schemes for the first week in July 2018. The gray curve indicates the aggregate meter (sum of net meters of all homes) before adding the BESSs, and the black curve indicates the meter after adding the BESSs. (c) Overlays the no BESS, individual, and joint optimization, highlighting that the joint scheme best utilizes available solar generation. By sharing the BESSs, the joint scheme can maintain the meter under the transformer limit for the whole week, while the individual scheme cannot.



**Fig. 3.** (a) HST with no BESS compared to individual and joint optimization. For this week of simulation in July 2018, having no BESS substantially raises HST values during the load peaks. (b) Percent loss of life for no BESS, individual, and joint optimization. The HST spikes correlate directly to larger increases in loss of life.

### 5.2. Transformer loss of life and cost reduction metrics

Given similarities in transformer protection, the main difference between the individual and joint schemes lies in the total electricity costs. Fig. 4 shows the loss of life (a) and total costs (b) for each scheme for four weeks of July 2018. Each scheme is effective in protecting the

transformer, but aggregating the batteries with the joint, hybrid, and dynamic schemes introduces cost savings of between 27 % and 56 % over the individual scheme.

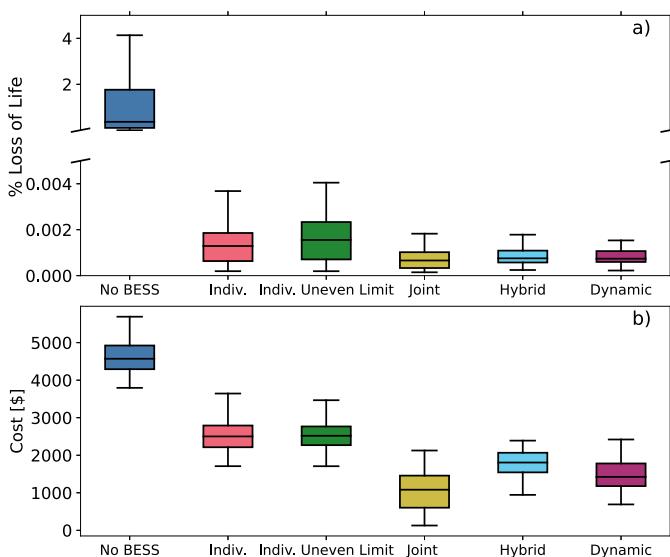
Although the median value of loss of life with no BESS is low at 0.36 %, almost 25 % of the trials in the no BESS case exceed 2 % loss of life in July. Continuing at this rate of degradation, these transformers will reach their end of life in just over four years, well below the expected lifetime of 20 years [44]. For the schemes with a battery, the loss of life incurred in July is negligible, since the battery allows the aggregate meter to remain at or below the transformer nameplate capacity.

The differences in costs for each scheme are more distinct. The joint scheme offers a 76 % median cost decrease compared to the no BESS system cost. The median costs slightly rise across the joint, hybrid, and dynamic cases from \$1082 to \$1119 and \$1158, respectively. The joint case has the lowest total cost because it utilizes the battery most efficiently with complete sharing.

However, the hybrid and dynamic schemes retain a portion of the battery that is not shared. This provides a crucial component for homeowner peace of mind: the ability to autonomously control their battery. For example, in a blackout scenario, a homeowner might want to leave some charge in reserve. Or, they might want the ability to participate in lucrative demand response programs. Given similar transformer aging and costs, the hybrid and dynamic schemes are almost as effective as the joint scheme but with the added benefit of a retained battery partition.

In the individual scheme with uneven transformer limits, each home has unique but static transformer limit allocation. Despite the flexibility, this scheme does not perform better than with an equal limit: transformer loss of life is slightly worse. Since homes better utilize their transformer limit, the aggregate meter reads closer to the transformer nameplate limit, elevating the HST and accelerating aging. As costs are primarily driven by battery availability, a higher transformer allocation does not translate into cost savings in the absence of battery aggregation.

Finally, the dynamic scheme has slightly lower costs than the hybrid scheme due to the time-varying partition weights. With these dynamic weights, underutilized retained partitions are converted into shared partitions to reduce costs.



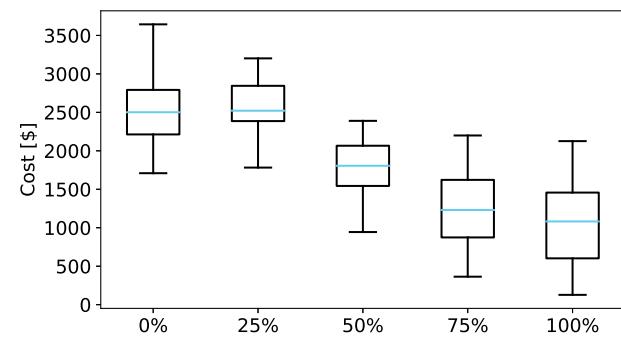
**Fig. 4.** (a) Percent loss of life over the month of July 2018 comparing each scheme. (b) Total system electricity costs for July 2018. The boxplots indicate results from 50 trials. Including a BESS decreases transformer aging substantially, and introducing aggregation through virtualization has a strong effect on cost reduction. Joint optimization, or complete sharing, performs the best in both transformer and cost metrics.

### 5.3. Sensitivity of hybrid scheme to partition fractions

In Fig. 4, the hybrid scheme is shown with fixed, 50 % partitions for the shared and retained battery partitions. This means each homeowner retains half and shares half of their battery. To examine the effect of changing this fixed partition value, Fig. 5 varies the fraction of shared battery from 0 % (individual scheme) to 100 % (joint scheme).

Sharing only 75 % of the battery yields similar total costs to 100 % sharing. The median cost for 75 % sharing is only 14 % worse than the median for 100 % sharing. For many homeowners, this difference may be irrelevant when considering the benefits of complete control over 25 % of their battery. For example, the homeowners could recoup this 14 % cost differential by using their 25 % retained partition to participate in lucrative demand response programs.

Understanding this critical point of 75 % sharing is important in implementing this algorithm in a real system. These results reveal that there needs to be more than 50 % sharing to achieve the cost benefits of joint optimization, but 100 % sharing is not necessary to reach these cost minimums.



**Fig. 5.** Sensitivity to different partition fractions for the hybrid scheme for July 2018 over 50 trials. While there is a large spread between trials, sharing only 25 % does not decrease costs, but sharing 50 % or more has substantial cost reductions. Complete (100 %) battery sharing yields the lowest costs.

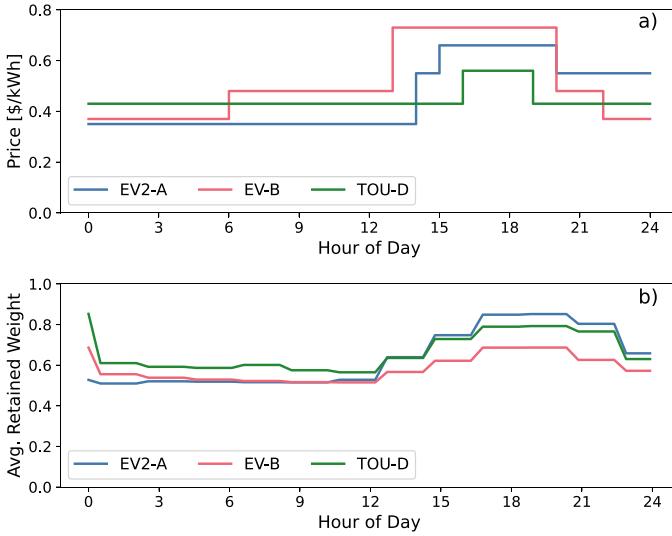
### 5.4. Dynamic partitioning with varying electricity pricing plans

When every home is subscribed to the same TOU rate structure, the dynamic scheme improves transformer loss of life but increases costs slightly compared to the hybrid scheme. In practice, homes in a neighborhood may be subscribed to different rate structures depending on their load usage patterns. When peak price periods are not consistent among homes, dynamic partitioning offers flexibility to share battery partitions when other homes need them and retain partitions for homes experiencing peak pricing.

To test this, we run a modified version of the individual, hybrid, and dynamic schemes in which each home is assigned a random price structure among the EV2-A, EV-B, and TOU-D structures. The individual term calculating the cost of the individual home meter (the first term in Eq. (17)) is subject to this randomly chosen price structure while the pricing for the joint term remains the EV2-A structure.

Fig. 6(a) compares the three price structures, which vary by both peak price value and timing. Fig. 6(b) plots the average retained partition weight (i.e., the partition of the battery kept by the homeowners) for the dynamic scheme, separating each group of homeowners by their pricing structure. All three groups' weights step up or down each time a pricing structure changes. This reflects the dynamic scheme's ability to re-adjust the partitions so the operation is cost-optimal. The retained weight is higher during peak times when the prices are highest, which is expected as homes need a greater fraction of the battery to offset home loads during peak periods.

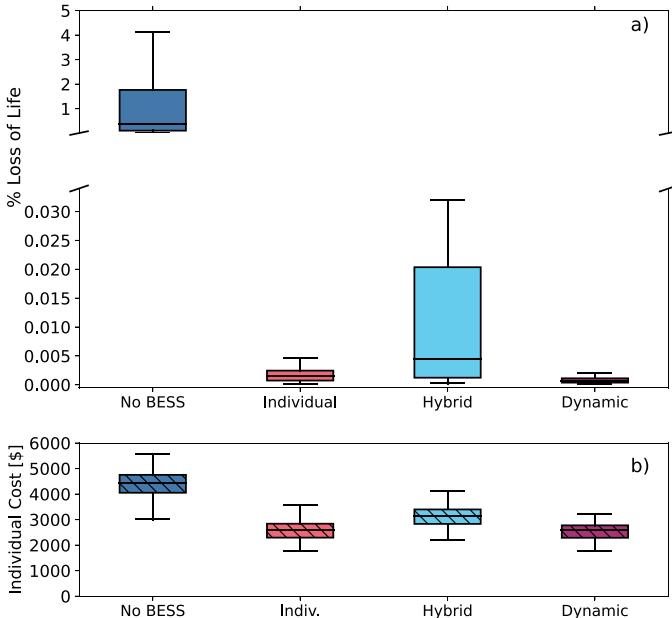
Fig. 7(a) shows the transformer loss of life for each scheme compared to the no BESS scenario. The dynamic scheme yields the lowest



**Fig. 6.** (a) Prices for one weekday for PG&E EV2-A, EV-B, and TOU-D rate structures. Both the magnitude and timing of peak pricing vary between the rate structures. (b) Average retained partition weight for each timestep for the dynamic scheme. The weights change corresponding to changes in the rate structures.

transformer loss of life. Fig. 7(b) examines the individual partition costs for each scheme, indicated with hatching to differentiate from the total costs shown in Fig. 4. For the hybrid and dynamic scheme, this is calculated by the optimal value of the first term in Eq. (17).

It is expected that the hybrid scheme has higher costs than the individual scheme, since only 50 % of each home's battery is represented in this cost calculation while the individual scheme utilizes 100 %. However, although the dynamic scheme shares only 38 % and retains 62 % of the battery on average across all homes and trials, it still achieves



**Fig. 7.** (a) Loss of life for no BESS, hybrid with 50 % weights, and dynamic schemes for 50 trials on July 18 with homes on varying rate structures. (b) Sum of individual meter costs for each scheme, representing the first term in the objective function for hybrid and dynamic schemes. The dynamic scheme has similar costs to the individual scheme due to its temporal flexibility.

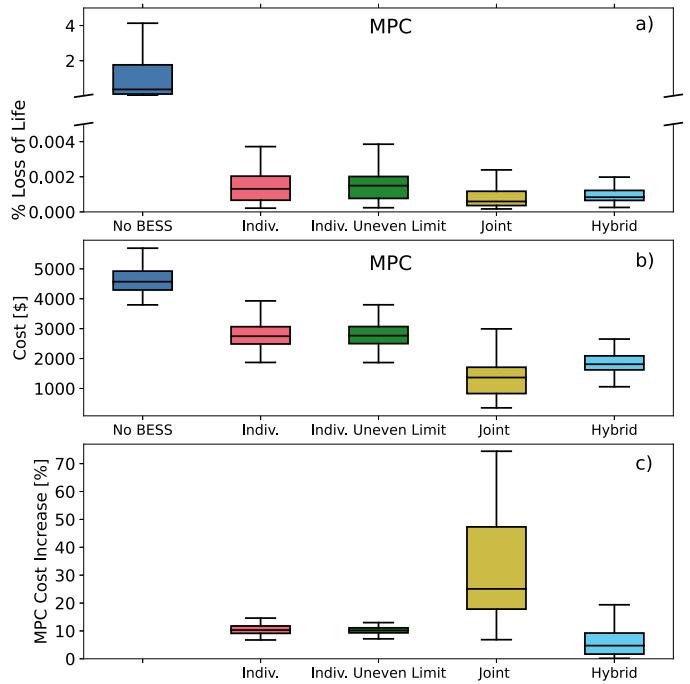
similar costs as the individual scheme. This result demonstrates the value of virtualization's temporal flexibility. By changing partitions over time, the dynamic scheme can retain only the minimum necessary battery fraction and share the remainder, leading to similar costs as the individual scheme but better transformer protection.

### 5.5. Sensitivity to imperfect forecasts

All the simulation results above represent optimization algorithms with perfect foresight, meaning the problem is solved with knowledge of exact solar and load profiles over the simulation horizon. In practice, this is not the case as forecasts are often imperfect. MPC adjusts for this by re-solving the optimization problem at each timestep with a short forecast horizon and updating the SOC based on actual solar and load values once they occur.

Fig. 8 shows the transformer aging and cost for the schemes using MPC instead of perfect foresight. We do not test MPC on the dynamic scheme as its performance depends on perfect load and solar foresight over the entire time horizon. Overall, transformer aging and costs follow the same patterns as seen with perfect foresight control. The individual scheme is not as cost-effective as the joint scheme. However, the individual scheme is still preferable to the no BESS system, even with the slightly worse performance of MPC.

Fig. 8(c) shows the percent cost increase for each scheme using MPC compared to the cost using perfect foresight. Cost increases for individual and individual with uneven limit are around 10 %. The cost increase spread grows with the schemes that incorporate battery sharing, but the median differences are still less than 25 %. This discrepancy is because with a shared battery, an inaccurate forecast for one home may affect the entire shared battery, but with the individual scheme, the inaccurate forecast would be confined to just a single home. This phenomenon explains the larger cost increase spread for the joint and hybrid schemes. The deviations for transformer aging are also small; the greatest aging



**Fig. 8.** (a) Percent loss of life over the month of July 2018 comparing each scheme solved with MPC. (b) Total system electricity costs for July 2018 with MPC. (c) Percent cost increases for each scheme using MPC compared to perfect foresight. The boxplots indicate results of 50 trials. Using MPC has minimal impact on transformer aging and a small impact on costs, with a 25 % median increase for joint optimization.

increase for July 2018 is 5 % for joint optimization. We also find that these results hold across seasons such that battery sharing with MPC is still cost-optimal compared to the No BESS case.

The MPC cost increases correspond to a range between a \$91.96 increase for the hybrid scheme to \$269.31 for the individual scheme. Although the joint scheme has the greatest percentage median cost increase, in unnormalized dollars, the individual scheme has the greatest increase. The results show that even with a cost increase, MPC implementation is still cheaper than with no BESS. Better forecasting can mitigate these increases even further. Ultimately, the cost increases only amount to between \$5.75 and \$18.50 per home compared to perfect foresight, indicating that deploying these algorithms in a real system is economically viable.

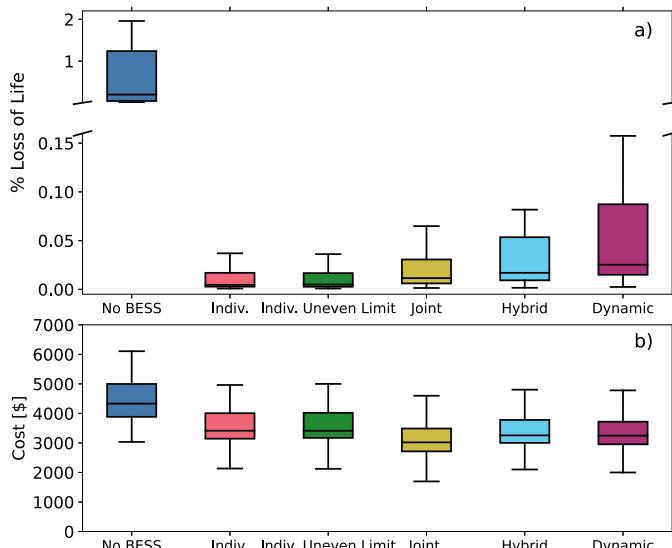
### 5.6. Seasonality effects

The July simulation results are not necessarily representative of system behavior during the rest of the year. To ensure the optimization yields similar transformer protection and cost savings, Fig. 9 presents results for each optimization scheme for January 2018.

In January, owning a BESS is still effective, but the cost savings are not as drastic. Although transformer loss of life is an order of magnitude worse than July, the aging still remains under 0.08 % for all schemes. On the cost side, individual and joint optimization only offer a 21 % and a 30 % cost reduction, respectively, compared to no BESS. Although January loads are lower because there is no need for air conditioning, lower solar generation is the main reason for the discrepancy between January and July. Less excess solar means there are fewer opportunities for the BESSs to charge for free from solar and discharge to offset power at expensive peak times. Regardless, these results prove that BESS optimization is worthwhile across the whole calendar year, though its impact (cost reduction or transformer protection) varies over the seasons.

### 5.7. Sensitivity to EV penetration and battery aging

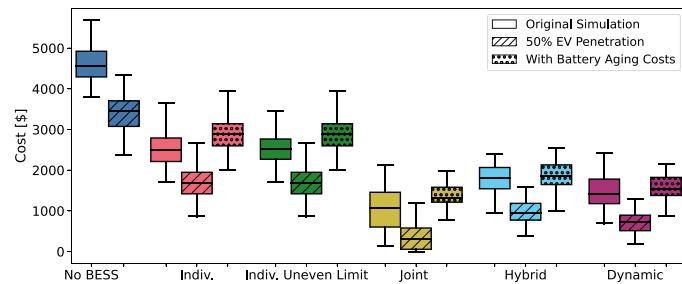
The results from Sections 5.1 to 5.6 assume both 100 % EV penetration (i.e., every home has one EV charger) and no economic effects from battery aging. However, in real neighborhoods, some homes may not have an EV and the cost of battery degradation may be substantial.



**Fig. 9.** (a) Percent loss of life over the month of January 2018 comparing each scheme. (b) Total system electricity costs for January 2018. The boxplots indicate results of 50 trials. Due to limited solar, loss of life is worse and cost reduction is not as effective compared to July performance.

We test the algorithm's sensitivity to both EV penetration and a cost on battery aging. Fig. 10 compares the costs for each scheme between the original simulation, a simulation with 50 % EV penetration, and a simulation that includes a battery degradation cost per cycle.

As expected, the costs are lower in each scheme with 50 % EV penetration. Joint optimization still yields the lowest cost as it best utilizes available solar generation. However, the battery aging costs do not follow a consistent pattern across all schemes. In the individual scheme, adding a cost per cycle increases the median cost by 15.1 %, but the costs only increase by 2.5 % and 8.2 % in the hybrid and dynamic schemes, respectively. This result suggests that the impact of battery aging costs is relatively smaller in these hybrid battery sharing schemes.



**Fig. 10.** Comparison of July 2018 costs for each scheme between original, 50 % EV penetration, and included battery aging cost simulations. The boxplots indicate results of 50 trials. With 50 % EV penetration, the median costs are consistently lower than with 100 % EV penetration across all schemes. Adding a battery aging cost has a greater effect when the batteries are not shared (i.e., in the individual schemes).

### 5.8. Small-scale physical system demonstration

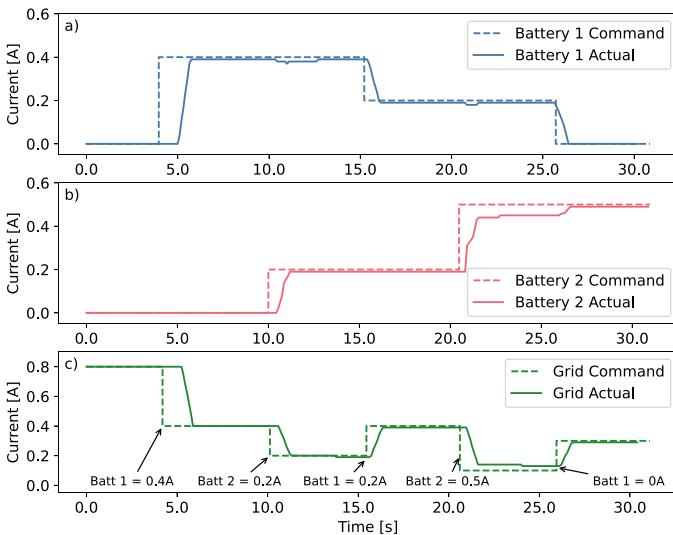
The results presented in Sections 5.1 to 5.7 are based on simulated home load, solar, and EV data. In real world implementation, factors such as communication delays and setpoint errors can affect the efficacy of the battery commands. To quantify these delays and errors, we design a small DC testbed study using a 20 Ah LiFePo4 battery and a 1.6 Ah lithium-ion battery connected to a power supply that emulates the grid and an electronic load that emulates home loads in a community. This study highlights how battery injection reduces the power pulled from the grid through the transformer and shows battery response time to setpoint commands.

To illustrate setpoint errors and battery response times, we run a 30 second experiment with varying battery setpoint commands under a constant load of 0.8 A. Fig. 11 shows the system response to battery injection. The current being pulled through the transformer, shown in Fig. 11(c), highlights the grid's response to changing battery setpoint commands. On average, the grid takes 1.00 s to respond to the setpoint commands while the batteries take 0.92 and 1.44 s, respectively. The setpoint errors are relatively minor for both batteries and the grid with RMSE values of 0.08 A, 0.06 A, and 0.24 A, respectively. In any real world electrical system, perturbations around the setpoint are anticipated. Overall, the communication delays are negligible under the time intervals simulated in this study.

## 6. Discussion

The results show that all of the optimization schemes reduce transformer aging substantially over the course of the month-long simulation periods. Joint, hybrid, and dynamic schemes highlight the impact of battery sharing on both reducing costs and mitigating transformer aging.

Virtualization is crucial to achieving maximum cost reduction in joint optimization. However, with only 50 % sharing, the hybrid scheme reduces the total community cost incurred during the month of July 2018 by 57 %. Sharing 75 % of their batteries allows homeowners to be within



**Fig. 11.** (a) Battery 1 command and actual current. (b) Battery 2 command and actual current. (c) Command and actual current provided by the grid, which is equal to 0.8 A minus the current from batteries 1 and 2. The commands sent to batteries 1 and 2 are annotated. While there is a settling time delay, it is only 1 s on average.

14 % of the optimal lowest cost while still maintaining control over a modest fraction of their batteries.

The effect of variable solar generation across seasons is substantial. Reduced solar generation in January leads to diminished cost savings, but the role of batteries is still prominent. These seasonality effects may also change in different geographical locations depending on solar and other variable energy generation penetration.

There are slight differences in transformer protection between optimization schemes. The patterns vary seasonally: in the summer, battery sharing leads to slightly better transformer protection, while in the winter battery sharing is slightly worse for the transformer. This phenomenon results from varying battery utilization among the schemes. Operating at the transformer limit results in greater aging than operating well below the limit, even if the system never exceeds the limit in either case. The difference between the individual scheme and individual with uneven transformer limit allocation scheme also occurs because of this utilization pattern. By optimizing each home's limit allocation, each home is more likely to operate closer to its limit, thus pushing the aggregate meter toward the nameplate limit.

Overall, complete sharing of batteries (joint optimization) is the most effective scheme for minimizing community costs and mitigating transformer aging. The results show a 76 % and 30 % decrease in costs compared to the no BESS case in July and January 2018, respectively. In the worst case of joint optimization in January 2018, the 0.07 % loss of life over the month would still yield a transformer lifetime of 119 years.

While the hybrid and dynamic schemes have comparable transformer protection to pure joint optimization, they are less cost-effective. Partial sharing of the batteries mitigates the worst of the peak loads to effectively reduce aging, but without complete sharing, there is still some underutilized storage accounting for the higher total costs.

The schemes are robust to both varying EV penetration and a battery aging cost. Similar trends hold when testing sensitivity to those parameters: joint optimization is still the most cost-effective, but by leveraging virtualization, hybrid and dynamic optimization are preferable to the individual cases while still allowing homeowner autonomy.

### 6.1. Implementation considerations

Each of the schemes provides value for transformer aging when compared to systems without batteries, but the flexibility of the hybrid and dynamic schemes can better meet various goals and requirements of

community members. To implement these schemes in practice, there are many additional considerations including imperfect load and solar forecasting, cost structure, battery aging, and algorithm scalability.

First, the MPC algorithm is susceptible to inaccurate solar or load forecasts, making the batteries charge at inopportune times. This can occur in two scenarios: when actual loads are greater than forecasted loads or when actual solar generation is less than forecasted generation. This inaccurate forecasting can lead to suboptimal battery behavior that momentarily violates the transformer limit. Ultimately, minimizing transformer aging and costs with MPC is sensitive to forecasting accuracy, and further work will explore MPC performance with a more accurate forecast.

Deploying the joint, hybrid, or dynamic schemes also requires a new electricity cost structure. This raises questions about cost fairness, battery partition sizes, and responsibility for system upgrades, operation, and maintenance costs. The utility must decide how to bill shared grid imports and exports and divide costs fairly among homeowners. To implement a new cost structure, each home needs its own meter to accurately measure load, solar, and battery output. A central aggregator could serve as a global controller and communicate with the utilities and homeowners about pricing, but the exact cost structure is dependent on existing local distribution grid infrastructure.

Battery aging is another concern in real world implementation. In each scheme, we find that the battery generally completes one full discharge and charge cycle per 24 h. This usage in the individual scheme is covered under Tesla's warranty policy: ten years of coverage with unlimited battery cycling for solar backup and load shifting or a 1400 charge and discharge cycle limit for other uses [42]. If homeowners use their battery partition for other programs such as grid frequency response or ancillary services, then battery usage should be restricted to using 140 full cycles per year, or 38 % of the battery capacity per day.

When adding a cost per cycle to the simulation, the total electricity cost increases for each scheme, from 2.5 % to 21.2 % depending on the scheme. This reflects that when adding a cost per cycle, it is financially advantageous for the system to complete fewer cycles to prolong the battery lifetime, even if doing so increases electricity costs. However, this cost increase associated with battery degradation still does not exceed the costs for the system with no batteries at all, indicating that batteries are still economically preferred.

While we test the algorithm on a small, 25 kVA residential transformer, the optimization schemes can be applied to systems with larger transformers, which may add more batteries and more decision variables in the optimization problem. Solving for a longer time horizon or a smaller time granularity also increases the number of decision variables. Since the optimization problem is convex, it is scalable to a larger problem and does not pose time constraints when solving at fast frequency with MPC. Virtualization software can also be scaled to large systems; in a simulation of 600 virtual batteries, latency from the virtualization software accounts for only half of the total system latency [32]. Lastly, incorporating a transformer penalty that only relies on the transformer nameplate limit and not on device-specific parameters also facilitates deployment in a variety of physical systems.

In real world implementation, virtualization is imperative to distribute battery charge and discharge instructions to each constituent physical battery. Virtualization serves multiple roles in this distribution process, such as allowing the system to re-calibrate if one battery drops offline and enabling the system to change battery partition weights frequently, as in the dynamic partitioning scheme. Without virtualization, homeowners could either control an individual battery and risk underutilizing an expensive asset, or they could buy into a centralized, shared battery with no individual autonomy. Virtualization makes this combination possible while ensuring the system is robust in the face of system outages, additions, or changes.

Each of these implementation challenges is critical for home and EV owners, policymakers, utilities, and aggregators to consider. While utilities and aggregators may be concerned with the scalability and flexibility

of the algorithm to deploy it in heterogeneous distribution systems, homeowners may be more focused on their direct costs of buying a battery and buying into a shared battery system. On the other hand, policymakers need to understand implementation challenges and associated costs to help design rebates or incentive programs that aim to ensure equitable deployment. All of these stakeholders must collaborate to fairly and successfully deploy this shared battery system.

## 6.2. Limitations

Battery sharing algorithms have limitations. This work does not focus on cost allocation fairness, forecasting accuracy and prediction models, and battery control execution.

The question of how to fairly allocate the benefits of battery sharing still remains. This paper only provides an assessment of total system cost and does not report a breakdown of costs among homeowners. In an aggregation scheme, splitting the cost savings equally among all homeowners would not be fair to homeowners who are already able to manage their loads individually. Two other potential options are scaling the savings based on a home's peak power or total energy throughput. Future work will explore the differences in cost allocation schemes, including which type of homeowner benefits from each scheme.

Our simulation also has the limitation of a simplified perfect foresight prediction model for EV charging. In reality, EV drivers may not know their exact charging needs one day in advance. A deviation from this prediction has the potential to exacerbate transformer aging with a high load, especially if multiple EVs in a neighborhood charge simultaneously. If hardware and computational capabilities allow, a short MPC timestep (i.e., quickly correcting forecasting errors) could rectify this problem. Alternatively, a smart charger or other peak load management system could serve as a backup to curtail load in case of incorrect EV demand forecasting.

Lastly, we only examine the effect of long latencies or battery set-point control errors in a small DC testbed. In a grid-scale system, the residential batteries may not be perfectly controlled. While the MPC corrects for setpoint errors, long latencies could reduce the algorithm's performance if they are beyond the seconds scale as observed in the DC testbed. To mitigate these effects, engineers should perform extensive physical testing with the algorithm to understand any latencies across different battery management systems.

## 7. Conclusion

This work presents a comparison between different battery sharing optimization schemes enabled by battery virtualization. The schemes aim to reduce consumer electricity bills and mitigate residential transformer aging. We find that sharing the batteries results in both lower costs and similar transformer protection than a scheme where owners optimize their batteries individually. While the joint scheme offers the highest battery utilization, the hybrid and dynamic schemes, or combinations of the individual and joint schemes, yield similar costs and transformer protection while still allowing homeowners to keep part of the battery for themselves. By employing virtualization, the dynamic scheme improves individual costs further by changing the retained and shared partition to maximize utilization. Lastly, the proposed model predictive controller allows for real time implementation that is robust to forecast error.

Future work will explore higher fidelity forecasting methods, simulations on different configurations of batteries and transformers, and experiments on physical battery systems in homes. To test and implement our algorithms in a neighborhood, a central controller or aggregator must also decide how to fairly allocate cost savings. Both fair pricing and best practices for collaboration between aggregators, utilities, and homeowners are other areas of future research.

Ultimately, residential battery virtualization empowers homeowners to optimize their batteries to protect shared neighborhood transformers and reduce costs for the entire community. As EV sales continue to

grow and more homeowners install Level 2 EV chargers, residential distribution transformer overload will become a more pressing problem. Our BESS optimization schemes offer the flexibility for homeowners to install chargers without compromising neighborhood transformer lifetimes.

## CRediT authorship contribution statement

**Sonia Martin:** Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Obidike Nnorom, Jr.:** Writing – review & editing, Writing – original draft, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Philip Levis:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization. **Ram Rajagopal:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization.

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

Sonia Martin reports financial support was provided by US Department of Energy. Obidike Nnorom, Jr. reports financial support was provided by Building Technologies Office. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgments

This work is supported by the U.S. Department of Energy, Office of Electricity under Award Number DEOE0000919 and the Building Technologies Office IBUILD Graduate Research Fellowship. Some of the computing for this project was performed on the Sherlock computing cluster, supported by Stanford University and the Stanford Research Computing Center.

## Data availability

Data will be made available on request.

## References

- [1] Hilshey AD, Hines PDH, Rezaei P, Dowds JR. Estimating the impact of electric vehicle smart charging on distribution transformer aging. *IEEE Trans Smart Grid* 2013;4(2):905–13. <https://doi.org/10.1109/TSG.2012.2217385>
- [2] Muratori M. Impact of uncoordinated plug-in electric vehicle charging on residential power demand. *Nat Energy* 2018;3(3):193–201. <https://doi.org/10.1038/s41560-017-0074-z>
- [3] Gong Q, Midlam-Mohler S, Marano V, Rizzoni G. Study of PEV charging on residential distribution transformer life. *IEEE Trans Smart Grid* 2012;3(1):404–12. <https://doi.org/10.1109/TSG.2011.2163650>
- [4] Soleimani M, Kezunovic M. Economic analysis of transformer loss of life mitigation using energy storage and PV generation, In: 2020 IEEE/PES transmission and distribution conference and exposition (T&D); 2020. p. 1–5. <https://doi.org/10.1109/TD39804.2020.9299895>
- [5] Sarker MR, Olsen DJ, Ortega-Vazquez MA. Co-optimization of distribution transformer aging and energy arbitrage using electric vehicles. *IEEE Trans Smart Grid* 2017;8(6):2712–22. <https://doi.org/10.1109/TSG.2016.2535354>
- [6] Kaufman A. The U.S. is running low on the machines needed to avoid blackouts; May 2023. [https://www.huffpost.com/entry/transformer-shortage\\_n\\_64645004e4b0005c6055c541](https://www.huffpost.com/entry/transformer-shortage_n_64645004e4b0005c6055c541)
- [7] Li Y, Jenn A. Impact of electric vehicle charging demand on power distribution grid congestion. *Proc Natl Acad Sci USA* 2024;121(18):e2317599121. <https://doi.org/10.1073/pnas.2317599121>
- [8] Powell S, Kara EC, Sevlian R, Cezar GV, Kılıçcote S, Rajagopal R. Controlled workplace charging of electric vehicles: the impact of rate schedules on transformer aging. *Appl Energy* 2020;276:115352. <https://doi.org/10.1016/j.apenergy.2020.115352>
- [9] Li S, Hu W, Cao D, Zhang Z, Huang Q, Chen Z, et al. EV charging strategy considering transformer lifetime via evolutionary curriculum learning-based multiagent deep reinforcement learning. *IEEE Trans Smart Grid* 2022;13(4):2774–87. <https://doi.org/10.1109/TSG.2022.3167021>
- [10] Affonso CDM, Kezunovic M. Technical and economic impact of PV-BESS charging station on transformer life: a case study. *IEEE Trans Smart Grid* 2019;10(4):4683–92. <https://doi.org/10.1109/TSG.2018.2866938>

[11] Visakh A, Thomas B, Selvan MP. Controlled charging of electric vehicles to reduce the aging of distribution transformers, In: 2022 22nd national power systems conference (NPSC); 2022. p. 24–9. <https://doi.org/10.1109/NPSC57038.2022.10069054>

[12] Rossi F, Diaz-Londono C, Li Y, Zou C, Gruosso G. Smart electric vehicle charging algorithm to reduce the impact on power grids: a reinforcement learning based methodology. *IEEE Open J Veh Technol* 2025;6:1072–84. <https://doi.org/10.1109/OJVT.2025.3559237>

[13] Sarmokadam S, Suresh M, Mathew R. Power flow control strategy for prosumer based EV charging scheme to minimize charging impact on distribution network. *Energy Rep* 2025;13:3794–809. <https://doi.org/10.1016/j.egyr.2025.03.032>

[14] Latinopoulos C, Sivakumar A, Polak JW. Response of electric vehicle drivers to dynamic pricing of parking and charging services: risky choice in early reservations. *Transp Res Part C Emerg Technol* 2017;80:175–89. <https://doi.org/10.1016/j.trc.2017.04.008>

[15] Navidi T, El Gamal A, Rajagopal R. Coordinating distributed energy resources for reliability can significantly reduce future distribution grid upgrades and peak load. *Joule* 2023;7(8):1769–92. <https://doi.org/10.1016/j.joule.2023.06.015>

[16] Endara ID, Macas DM, Valverde MP, García JM. Design of load management methods using distributed storage and their impact on the aging level of distribution transformers. In: Peng Z, editor. *Technological advancements and future directions in green energy: selected papers from ICGET 2024*. Cham: Springer Nature Switzerland; 2025. p. 19–31.

[17] Datta U, Kalam A, Shi J. Smart control of BESS in PV integrated EV charging station for reducing transformer overloading and providing battery-to-grid service. *J Energy Storage* 2020;28:101224. <https://doi.org/10.1016/j.est.2020.101224>

[18] Hong S-K, Lee SG, Kim M. Assessment and mitigation of electric vehicle charging demand impact to transformer aging for an apartment complex. *Energy* 2020;13(10):2571. <https://doi.org/10.3390/en13102571>

[19] Li J, Wu Z, Zhou S, Fu H, Zhang X-P. Aggregator service for PV and battery energy storage systems of residential building. *CSEE J Power Energy Syst* 2015;1(4):3–11. <https://doi.org/10.17775/CSEJPES.2015.00042>

[20] Moradi Amani A, Sajjadi SS, Somaweera WA, Jalili M, Yu X. Data-driven model predictive control of community batteries for voltage regulation in power grids subject to EV charging. *Energy Rep* 2023;9:236–44. <https://doi.org/10.1016/j.egyr.2022.12.089>

[21] Raghuvir RM, Bhalja BR, Agarwal P. Real-time Energy management system for an active distribution network with multiple EV charging stations considering transformer's aging and reactive power dispatch. *IEEE Trans Ind Appl* 2025;1–13. <https://doi.org/10.1109/TIA.2025.3552367>

[22] Ntube N, Li H. Stochastic multi-objective optimal sizing of battery energy storage system for a residential home. *J Energy Storage* 2023;59:106403. <https://doi.org/10.1016/j.est.2022.106403>

[23] Hafiz F, Rodrigo de Queiroz A, Fajri P, Husain I. Energy management and optimal storage sizing for a shared community: a multi-stage stochastic programming approach. *Appl Energy* 2019;236:42–54. <https://doi.org/10.1016/j.apenergy.2018.11.080>

[24] Henni S, Staudt P, Weinhardt C. A sharing economy for residential communities with PV-coupled battery storage: benefits, pricing and participant matching. *Appl Energy* 2021;301:117351. <https://doi.org/10.1016/j.apenergy.2021.117351>

[25] Lee S, Shenoy P, Ramamritham K, Irwin D. AutoShare: virtual community solar and storage for energy sharing. *Energy Inform* 2021;4(1):10. <https://doi.org/10.1186/s42162-021-00144-w>

[26] Khanal S, Khezri R, Mahmoudi A, Kahourzadeh S. Optimal capacity of solar photovoltaic and battery storage for grid-tied houses based on energy sharing. *IET Gener Transm Distrib* 2023;17(8):1707–22. <https://doi.org/10.1049/gtd2.12824>

[27] Merrington S, Khezri R, Mahmoudi A. Optimal sizing of grid-connected rooftop photovoltaic and battery energy storage for houses with electric vehicle. *IET Smart Grid* 2023;6(3):297–311. <https://doi.org/10.1049/stg2.12099>

[28] Walker A, Kwon S. Analysis on impact of shared energy storage in residential community: individual versus shared energy storage. *Appl Energy* 2021;282:116172. <https://doi.org/10.1016/j.apenergy.2020.116172>

[29] Kang H, Jung S, Kim H, Hong J, Jeoung J, Hong T. Multi-objective sizing and real-time scheduling of battery energy storage in energy-sharing community based on reinforcement learning. *Renew Sustain Energy Rev* 2023;185:113655. <https://doi.org/10.1016/j.rser.2023.113655>

[30] Barbour E, Parra D, Awwad Z, González MC. Community energy storage: a smart choice for the smart grid? *Appl Energy* 2018;212:489–97. <https://doi.org/10.1016/j.apenergy.2017.12.056>

[31] Roberts MB, Bruce A, MacGill I. Impact of shared battery energy storage systems on photovoltaic self-consumption and electricity bills in apartment buildings. *Appl Energy* 2019;245:78–95. <https://doi.org/10.1016/j.apenergy.2019.04.001>

[32] Martin S, Mosier N, Nnorom O, Ou Y, Patel L, Tribe O, et al. Software defined grid energy storage, In: *Proceedings of the 9th ACM international conference on systems for energy-efficient buildings, cities, and transportation, BuildSys '22*. New York, NY, USA: Association for Computing Machinery; 2022. p. 218–27. <https://doi.org/10.1145/3563357.3564082>

[33] Bashir N, Guo T, Hajiesmaili M, Irwin D, Shenoy P, Sitaraman R, et al. Enabling sustainable clouds: the case for virtualizing the energy system, In: *Proceedings of the ACM symposium on cloud computing*. Seattle WA USA: ACM; 2021. p. 350–8. <https://doi.org/10.1145/3472883.3487009>

[34] Souza A, Bashir N, Murillo J, Hanafy W, Liang Q, Irwin D, et al. Ecovisor: a virtual Energy system for carbon-efficient applications, In: *Proceedings of the 28th ACM international conference on architectural support for programming languages and operating systems; vol. 2*. Vancouver BC Canada: ACM; 2023. p. 252–65. <https://doi.org/10.1145/3575693.3575709>

[35] Kumar R, Wenzel MJ, Ellis MJ, ElBsat MN, Drees KH, Zavala VM. A stochastic model predictive control framework for stationary battery systems. *IEEE Trans Power Syst* 2018;33(4):4397–406. <https://doi.org/10.1109/TPWRS.2017.2789118>

[36] Clarke WC, Manzie C, Brear MJ. An economic MPC approach to microgrid control, In: *2016 Australian control conference (AuCC)*; 2016. p. 276–81. <https://doi.org/10.1109/AUCC.2016.7868202>

[37] Zamani V, Cortés A, Kleissl J, Martínez S. Integration of PV generation and storage on power distribution systems using MPC, In: *2015 IEEE power & energy society general meeting*; 2015. p. 1–5. <https://doi.org/10.1109/PESGM.2015.7286588>

[38] Nair UR, Sandelic M, Sangwongwanich A, Dragičević T, Costa-Castelló R, Blaabjerg F. An analysis of multi objective energy scheduling in PV-BESS system under prediction uncertainty. *IEEE Trans Energy Convers* 2021;36(3):2276–86. <https://doi.org/10.1109/TEC.2021.3055453>

[39] Pezeshki H, Wolfs PJ, Ledwich G. Impact of high PV penetration on distribution transformer insulation life. *IEEE Trans Power Del* 2014;29(3):1212–20. <https://doi.org/10.1109/TPWRD.2013.2287002>

[40] Tesla. *Tesla Powerwall 2 datasheet*. Technical Report. Jun. 2019.

[41] Tesla's Powerwall 2 and solar roof tiles: our review. <https://unboundsolar.com/solar-information/tesla-powerwall-for-solar>

[42] Tesla. *Tesla Powerwall Limited Warranty (USA)* Jan. 2025. <https://energylibrary.tesla.com/docs/Public/EnergyStorage/Powerwall/General/Warranty-us/Powerwall-Warranty-EN.pdf>

[43] PG&E. *Electric rates: current and historic electric rates* Apr. 2024. <https://www.pge.com/tariffs/en/rate-information/electric-rates.html#accordion-a84c67dc1-item-e10ee0cc5>

[44] IEEE guide for loading mineral-oil-immersed transformers and step-voltage regulators. IEEE std C57.91-2011 (revision of IEEE std C57.91-1995). 2012:1–123. <https://doi.org/10.1109/IEEESTD.2012.6166928>

[45] Larson Electronics. *25 KVA pad mount transformer - 12470Y/7200 grounded wye primary, 240/120V secondary - mineral oil fluid/ONAN*. <https://www.larsonelectronics.com/product/292915/25-kva-pad-mount-transformer-12470y-7200-grounded-wye-primary-240-120v-secondary-mineral-oil-fluid-onan>

[46] Wilson E, Parker A, Fontanini A, Present E, Reyna J, Adhikari R, et al. End-use load profiles for the U.S. building stock: methodology and results of model calibration, validation, and uncertainty quantification. Technical Report NREL/TP-5500-80889, 1854582, MainId: 78667 Mar. 2022 <https://doi.org/10.2172/1854582>

[47] Blonsky M, Maguire J, McKenna K, Cutler D, Balamurugan SP, Jin X. OCHRE: the object-oriented, controllable, high-resolution residential energy model for dynamic integration studies. *Appl Energy* 2021;290:116732. <https://doi.org/10.1016/j.apenergy.2021.116732>