

# Long Duration Battery Sizing, Siting, and Operation Under Wildfire Risk Using Progressive Hedging

Ryan Piansky,<sup>\*</sup> Georgia Stinchfield,<sup>†</sup> Alyssa Kody,<sup>‡</sup> Daniel K. Molzahn,<sup>\*</sup> and Jean-Paul Watson<sup>§</sup>

<sup>\*</sup> School of Electrical and Computer Engineering, Georgia Institute of Technology

Atlanta, Georgia USA, {rpiansky3, molzahn}@gatech.edu

<sup>†</sup>Department of Chemical Engineering, Carnegie Mellon University

Pittsburgh, Pennsylvania USA, gstinchf@andrew.cmu.edu

<sup>‡</sup> Energy Systems and Infrastructure Analysis, Argonne National Laboratory

Lemont, Illinois USA, akody@anl.gov

<sup>§</sup> Lawrence Livermore National Laboratory

Livermore, California USA, watson61@llnl.gov

**Abstract**—Battery sizing and siting problems are computationally challenging due to the need to make long-term planning decisions that are cognizant of short-term operational decisions. This paper considers sizing, siting, and operating batteries in a power grid to maximize their benefits, including price arbitrage and load shed mitigation, during both normal operations and periods with high wildfire ignition risk. We formulate a multi-scenario optimization problem for long duration battery storage while considering the possibility of load shedding during Public Safety Power Shutoff (PSPS) events that de-energize lines to mitigate severe wildfire ignition risk. To enable a computationally scalable solution of this problem with many scenarios of wildfire risk and power injection variability, we develop a customized temporal decomposition method based on a progressive hedging framework. Extending traditional progressive hedging techniques, we consider coupling in both placement variables across all scenarios and state-of-charge variables at temporal boundaries. This enforces consistency across scenarios while enabling parallel computations despite both spatial and temporal coupling. The proposed decomposition facilitates efficient and scalable modeling of a full year of hourly operational decisions to inform the sizing and siting of batteries. With this decomposition, we model a year of hourly operational decisions to inform optimal battery placement for a 240-bus WECC model in under 70 minutes of wall-clock time.

**Index Terms**—Batteries, Investments, Progressive Hedging, PSPS, Wildfires

## I. INTRODUCTION

Wildfires pose a growing threat due to accelerating climate change [1]. While power system infrastructure is only responsible for a small fraction of all wildfire ignitions, these often

lead to more severe and expansive fires compared to other ignition sources [2]. In parts of the Western United States during periods of elevated wildfire threat, system operators proactively de-energize power lines with the goal of reducing the risk of wildfire ignition. These de-energizations are referred to as “Public Safety Power Shutoff” (PSPS) events [3]. When a PSPS event is invoked, system operators identify power lines to de-energize that pose a high ignition risk. These lines are selected using information about the lines (e.g., condition, age, routing), environmental factors (e.g., humidity, wind, temperature), wildfire spread models, etc. [4].

Although PSPS events can mitigate acute ignition risk, they also can result in load shedding as de-energized portions of the grid cannot serve all load demands. Furthermore, power outages can have significant negative economic and societal impacts [5]. To aid in reducing the impact or extent of PSPS events, utilities are investing in various types of infrastructure such as undergrounded lines [6], microgrids [7], and covered conductors [8] as well as vegetation management [8]. Utility-scale battery energy storage systems [9], [10] are another investment that can help mitigate the impact of PSPS events on customers during wildfire season [11], [12]. However, unlike some of these other investment options, utility-scale batteries can also be beneficial during periods without wildfire ignition threats to improve system operations and renewable integration [13]. This leads to the following question: *How do we optimally size, site, and operate utility-scale batteries on a transmission network considering both high wildfire-risk periods during which PSPS events are enacted and also normal operations?*

During normal operations, we may aim to locate batteries in congested areas of the network or near large concentrations of renewables [14]. During high wildfire-risk conditions, the topology of the network will change due to PSPS de-energization events. Therefore, congested areas of the network may shift and some parts of the network may experience load shedding. These seasonal changes make it difficult to identify appropriate sizing and siting decisions without considering long, multi-month time horizons. Simultaneously, to properly

---

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 and was supported by both the LLNL LDRD Program under Project 22-SI-008 and the US Department of Energy Office of Electricity’s Advanced Grid Modeling program. D.K. Molzahn and R. Piansky acknowledge support from the NSF AI Institute for Advances in Optimization (AI4OPT), #2112533. The authors would also like to thank Gurobi for providing the academic license used to run tests and D.L. Woodruff at University California – Davis with assistance on implementation. G. Stinchfield and R. Piansky are co-first authors.

evaluate the benefits of battery systems, one must consider decisions at the operational time scale, such as hourly charging and discharging behavior. This paper proposes a method to take into account a *full year of hourly operation decisions* spanning both wildfire and non-wildfire threat time periods.

Finding optimal locations to place utility-scale batteries on a large power grid is a computationally challenging problem [13]–[15]; see [9], [10] for literature reviews. Given the need to consider detailed operational decisions across multiple scenarios, the tractability of this problem scales poorly in both the number of scenarios and grid size. Prior work has developed algorithms using smaller test systems, shorter time horizons, or larger operational time steps [13], [15]. Work has also been done to model battery behavior in conjunction with an AC power flow network representation, but again, on smaller test networks [14]. Stochastic programming techniques have been used for sizing and siting problems to consider variability in renewable forecasts [16]. Similarly, a decomposition method based on progressive hedging has been used for problems related to wildfire operational problems that optimize line switching decisions [17], but these problems did not consider batteries. The need to model both extreme wildfire conditions with PSPS events and normal operations requires a level of temporal granularity and length that further challenges the optimization of battery sizing and siting.

In this paper, we propose a method for optimally sizing, siting, and operating long duration utility-scale battery systems considering both PSPS events and normal operations. To enable computational tractability, we employ a progressive hedging (PH) algorithm, which is a scenario-based decomposition technique originally developed for stochastic programming [18]. Traditionally, to model uncertainties, PH algorithms consider many scenarios that are all coupled by one set of common variables. In contrast, since our setting instead considers a temporal decomposition which represents a year of hourly operations via week-long scenarios, our algorithm must consider both coupling with a common set of variables and coupling between scenarios at temporal boundaries. Specifically, these scenarios must agree on (1) the sizing and siting of utility-scale batteries, and (2) the state-of-charge values at temporal boundaries. This setting thus requires non-trivial modifications to standard progressive hedging algorithms to account for this additional temporal coupling.

For week-long scenarios representing periods during wildfire season, sufficiently high wildfire risks result in line de-energizations. We model the wildfire risk posed by each line using real wildfire risk data from United States Geological Survey (USGS). Our approach allows us to model long-duration energy storage while capturing the different wildfire and load conditions at various times over the year.

The optimization model for the placement problem considering multiple scenarios is defined in Section II, while our solution approach using a customized progressive hedging algorithm is detailed in Section III. Our experimental methodology is described in Section IV followed by discussions of specific case studies and results in Section V. We conclude

in Section VI with a summary of our contributions and a discussion of future work.

## II. OPTIMIZATION MODEL

We now describe our long duration battery sizing, siting, and operational optimization model. We set a 100MW per-unit (p.u.) basis.

### A. Network, Demand, and Generator Variables

We define index sets over a transmission network as follows:

- $\mathcal{N}$ , the set of buses
- $\mathcal{L}$ , the set of transmission lines
- $\mathcal{G}$ , the set of generators
- $\mathcal{T} = \{1, \dots, T\}$ , the (ordered) set of time periods

For each line  $\ell \in \mathcal{L}$ , we specify the following parameters:

- $b^\ell$ , the line susceptance in p.u.
- $\bar{p}^\ell$ , the power flow limit in p.u.
- $r_t^\ell$ , the unitless wildfire risk associated with line  $\ell$  being energized at time  $t$
- $n^{\ell, \text{fr}}$  and  $n^{\ell, \text{to}}$ , *to* and *from* buses, respectively, where positive power flows from the *from* bus to the *to* bus
- $\bar{\delta}^\ell$  and  $\underline{\delta}^\ell$ , upper and lower voltage angle difference limits in radians, respectively
- $\ell \in \mathcal{L}_t^{\text{off}} \subseteq \mathcal{L}$ , a subset of lines that are de-energized (switched off) during time  $t$ . Note: the line's energization state is dictated by the wildfire risk on a given day and is thus not constant for all  $t \in \mathcal{T}$

For each generator  $i \in \mathcal{G}$ , we specify the following parameters:

- $\bar{g}^i$  and  $\underline{g}^i$ , upper and lower generation limits, respectively, in p.u.
- $n^i$ , bus  $n$  at which generator  $i$  is located
- $c_j^i$ , the  $j^{\text{th}}$  order cost term for generator  $i$  with  $j \in \{1, 2, \dots, J\}$  for a provided  $J^{\text{th}}$  degree polynomial cost function

For each bus  $n \in \mathcal{N}$ , we specify the following parameters:

- $p_{l,t}^n$ , power demand at time  $t \in \mathcal{T}$  in p.u.
- $\mathcal{G}^n$ , the set of generators located at bus  $n$
- $\mathcal{L}^{n, \text{to}}$  and  $\mathcal{L}^{n, \text{fr}}$ , the subset of lines  $\ell \in \mathcal{L}$  with bus  $n$  as the designated *to* bus, and bus  $n$  as the designated *from* bus, respectively

We model the operational problem for each time period using a B- $\theta$  DC optimal power flow representation, necessitating the following variables:

- $p_{g,t}^i$ , output of generator  $i \in \mathcal{G}$  at time  $t \in \mathcal{T}$  in p.u.
- $g_{\text{slack},t}^n$ , slack generation at bus  $n \in \mathcal{N}$  at time  $t \in \mathcal{T}$  in p.u.
- $\theta_t^n$ , voltage angle, in radians, at bus  $n \in \mathcal{N}$  at time  $t \in \mathcal{T}$
- $p_{\text{ls},t}^n$ , load shedding at bus  $n \in \mathcal{N}$  at time  $t \in \mathcal{T}$  in p.u.
- $p_t^\ell$ , power flow from bus  $n^{\ell, \text{fr}}$  to bus  $n^{\ell, \text{to}}$  on line  $\ell \in \mathcal{L}$  at time  $t \in \mathcal{T}$  in p.u.

## B. Network Operational Bounds and Constraints

We use a B- $\theta$  DC OPF with the following constraints

$$\underline{g}^i \leq p_{g,t}^i \leq \bar{g}^i, \quad \forall i \in \mathcal{G}, \forall t \in \mathcal{T}, \quad (1a)$$

$$0 \leq p_{ls,t}^n \leq p_{l,t}^n, \quad \forall n \in \mathcal{N}, \forall t \in \mathcal{T}, \quad (1b)$$

$$0 \leq g_{slack,t}^n \leq 10000, \quad \forall n \in \mathcal{N}, \forall t \in \mathcal{T}, \quad (1c)$$

$$\underline{\delta}^\ell \leq \theta_t^{\ell,fr} - \theta_t^{\ell,to} \leq \bar{\delta}^\ell, \quad \forall \ell \in \mathcal{L} \setminus \mathcal{L}^{off}, \forall t \in \mathcal{T}, \quad (1d)$$

$$-b^\ell(\theta_t^{\ell,fr} - \theta_t^{\ell,to}) \leq p_t^\ell \leq -b^\ell(\theta_t^{\ell,fr} - \theta_t^{\ell,to}), \quad \forall \ell \in \mathcal{L} \setminus \mathcal{L}^{off}, \forall t \in \mathcal{T}, \quad (1e)$$

where (1a) bounds the output of each generator  $i \in \mathcal{G}$  based on its physical limits, (1b) constrains the allowable load shed at each bus based on the load demanded, (1d) bounds the voltage angle difference across each energized line, and (1e) models the power flow on each energized line  $\ell \in \mathcal{L}$  using the B- $\theta$  DC approximation. To help with feasibility in the decomposed model described in Sections III and IV, we introduce slack generation with bounds described in (1c).

Power flow on each line is bi-directional with the thermal limit set as a bound, unless the line is de-energized, in which case the power flow is set to 0:

$$p_t^\ell = 0, \quad \forall \ell \in \mathcal{L}_t^{off}, \forall t \in \mathcal{T}, \quad (2)$$

$$-\bar{p}^\ell \leq p_t^\ell \leq \bar{p}^\ell, \quad \forall \ell \in \mathcal{L} \setminus \mathcal{L}_t^{off}, \forall t \in \mathcal{T}. \quad (3)$$

Note that the line de-energizations are based solely on the time-varying wildfire risks and are not decision variables in the optimization formulation.

## C. Grid-Scale Battery Modeling

We allow batteries to be placed at a subset of buses  $\mathcal{N}^{batt} \subseteq \mathcal{N}$ . This allows for flexibility in large systems where a heuristic needs to be used to identify a subset of candidate battery locations to improve tractability. Unless otherwise specified,  $\mathcal{N}^{batt} = \mathcal{N}$ .

Battery parameters are defined as:

- $\bar{E}$  and  $\underline{E}$ , the upper and lower energy storage limits of the battery, respectively, in p.u.
- $E_0^n$ , the initial charge of batteries at bus  $n \in \mathcal{N}^{batt}$  in p.u.
- $\bar{p}_c$  and  $\underline{p}_c$ , the upper and lower charge/discharge rate limits for a single battery, respectively, in p.u. in a single time interval
- $e$ , the efficiency of charging and  $\frac{1}{e}$  the efficiency of discharging
- $h$ , the hourly carry-over rate to model self-discharge
- $X_{max}$ , the maximum number of batteries allowed at a single bus
- $X_{total}$ , the total number of batteries allowed across the entire network

For each bus  $n \in \mathcal{N}^{batt}$ , we introduce the following variables:

- $x^n \in \mathbb{R}^+$ , number of batteries placed at bus  $n$
- $p_{c,t}^n$ , charging rate at bus  $n$  at time  $t \in \mathcal{T}$  in p.u.
- $p_{d,t}^n$ , charging rate at bus  $n$  at time  $t \in \mathcal{T}$  in p.u.

TABLE I: Battery Parameters

	Parameter	Value
minimum storage limit	$\underline{E}$	0 p.u.
maximum storage limit	$\bar{E}$	1.0 p.u.
minimum charge rate	$\underline{p}_c$	0 p.u./hour
maximum charge rate	$\bar{p}_c$	1.0 p.u./hour
maximum batteries	$X_{max}$	4 batteries
efficiency	$e$	95%
hourly carryover	$h$	0.999958
total batteries	$X_{total}$	10 batteries

Limits are imposed on the number of batteries introduced to the network and the amount placed at each bus:

$$0 \leq x^n \leq X_{max}, \quad \forall n \in \mathcal{N}^{batt}, \quad (4)$$

$$\sum_{n \in \mathcal{N}^{batt}} x^n \leq X_{total}. \quad (5)$$

Let  $E_t^n(\cdot)$  be the total energy, or state-of-charge (SOC), stored in all batteries placed at bus  $n \in \mathcal{N}^{batt}$  at time  $t \in \mathcal{T}$ . The stored energy changes as the batteries charge and discharge:

$$E_{t+1}^n(x, p_c, p_d) = hE_t^n(x, p_c, p_d) + ep_{c,t}^n - \frac{1}{e}p_{d,t}^n \quad (6)$$

The energy stored in the set of batteries at bus  $n \in \mathcal{N}^{batt}$  must abide by lower and upper storage bounds:

$$x^n \underline{E} \leq E_{t+1}^n(x, p_c, p_d) \leq x^n \bar{E}, \quad \forall n \in \mathcal{N}^{batt}, \forall t \in \mathcal{T}. \quad (7)$$

Charging and discharging is enforced via lower and upper rate limits:

$$x^n \underline{p}_c \leq p_{c,t}^n \leq x^n \bar{p}_c, \quad \forall n \in \mathcal{N}^{batt}, \quad \forall t \in \mathcal{T}, \quad (8)$$

$$x^n \underline{p}_c \leq p_{d,t}^n \leq x^n \bar{p}_c, \quad \forall n \in \mathcal{N}^{batt}, \quad \forall t \in \mathcal{T}. \quad (9)$$

Consistent with the scale of ongoing battery installations [19], Table I summarizes the battery parameters used in our numerical tests. Batteries are modeled to have a capacity of 1.0 p.u. (100 MWh) in our problem formulation [20]. This is consistent with utility-scale lithium-ion battery installations [21]. A charge/discharge efficiency of 95% (roughly a 90% round-trip efficiency) is used and an hourly carryover rate is used to model a 0.1% daily self-discharge loss [22]. Generalizations can be made to introduce other models of losses, such as the formulation in [23], or multiple types of batteries with distinct characteristics.

With the inclusion of grid-scale batteries, power balance is enforced across all buses at all times:

$$\sum_{\ell \in \mathcal{L}^{n,fr}} p_t^\ell - \sum_{\ell \in \mathcal{L}^{n,to}} p_t^\ell = \sum_{i \in \mathcal{G}^n} p_{g,t}^i + g_{slack,t}^n - p_{l,t}^n + p_{ls,t}^n - p_{c,t}^n + p_{d,t}^n, \quad \forall n \in \mathcal{N}, \forall t \in \mathcal{T}. \quad (10)$$

## D. Objective Function

A specified number of batteries are placed in the system with no associated cost. This allows for the optimization to focus on minimizing operational costs over the time horizon,  $\mathcal{T}$ .

Modeled costs include production for power generation and penalties imposed for unserved load, as follows:

$$C_{gen} = \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{G}^n} \sum_{j=0}^{J-1} c_i^j \cdot (p_{g,t}^i)^j, \quad (11)$$

$$C_{loadshed} = \sum_{t \in \mathcal{T}} \sum_{n \in \mathcal{N}} K_{ls} \cdot p_{ls,t}^n. \quad (12)$$

$$C_{slack} = \sum_{t \in \mathcal{T}} \sum_{n \in \mathcal{N}} K_{slack} \cdot g_{slack,t}^n. \quad (13)$$

Here,  $K_{ls}$  represents the cost per p.u. of demand not served, in this model set to be \$20,000/p.u. and the cost of slack generation  $K_{slack}$  is set to  $50K_{ls}$ . The total system operating cost is then given by:

$$C_{tot} = C_{gen} + C_{loadshed} + C_{slack}. \quad (14)$$

The optimization problem we formulate for sizing, siting, and operating systems with long duration energy storage is thus:

$$\begin{aligned} \min_{p_g, \theta, p^e, p_{ls}, x, p_b, r_g, r_b} \quad & (14) \\ \text{s.t.} \quad & (1) - (3), (4) - (10). \end{aligned} \quad (15)$$

Note that this problem takes the form of a linear program (LP).

### E. Thresholded Line Switching and Investments

Unlike in previous work [24], line de-energization decisions are predetermined by a risk threshold instead of a binary decision variable embedded in the optimization problem. In a similar manner to [24], each transmission line in the network is assigned a daily risk of wildfire ignition. In this paper, lines above a set threshold on a given day are de-energized and thus included in  $\mathcal{L}_t^{\text{off}}$ . Accordingly, the line energization status on a given day is a *parameter* rather than a decision variable such as in, e.g., [7], [24]–[26]. Extensions to simultaneously consider long duration battery sizing, siting, and operation along with optimizing line de-energization decisions is a topic of ongoing work.

## III. DECOMPOSITION VIA PROGRESSIVE HEDGING

Stochastic programs are often employed to enable power grid optimization models to hedge against the prevalence of uncertainties such as load and renewables production, and ultimately make more informed and reliable decisions [27]. The block-angular structure in deterministic multi-scenario planning problems considering large numbers of time periods (including the model we propose in Section II) is identical to that of a stochastic program, and we can thus leverage mathematical machinery from stochastic programming to address our deterministic problem (15). In practice, such problems are typically too large to solve directly without tailored optimization approaches. The Progressive Hedging (PH) algorithm, first proposed in [28], was developed to handle this challenge. PH exploits the structure of stochastic programs by breaking the larger problem into multiple smaller, more computationally tractable sub-problems that can be solved in parallel.

The primary input to a stochastic program is a set of representative scenarios, which we denote  $\mathcal{S} = \{1, \dots, S\}$ . In




Fig. 1: Two-stage stochastic program representation

our battery optimization model, scenarios correspond to operational time periods. Scenario data in a stochastic program is partitioned into multiple stages,  $\mathcal{J} = \{2, \dots, J\}$ . Typically, stage indexes start at  $j = 2$  because constant, scenario-independent data is given at stage  $j = 1$ . Across all stages and scenarios are realizations of random variables, stored in  $\xi$ , drawn from the discrete probability space  $\Xi$ . A stage  $j \in \mathcal{J}$  represents a point in time across the scenarios where uncertain parameters *up to and including* stage  $j$ , which corresponds to the vector  $\xi^j(s)$ , become known and specified decisions must be made. A scenario  $s \in \mathcal{S}$  contains a full set of realized random variables  $\xi(s)$  for all stages  $j \in \mathcal{J}$  and is weighted by the probability of that realization occurring,  $\mathbb{P}(s)$ . In the context of our deterministic battery sizing and siting formulation, we could use  $\mathbb{P}(s)$  to weight each time period by its frequency of occurrence if we employed representative weeks to model longer periods. However, in our implementation, we model each week individually and thus weight each period the same.

Here, we consider a two-stage problem ( $\mathcal{J} = \{2\}$ ) in which first-stage decisions (storage sizing and siting) are made before the random variables (wildfire risk and load demand vectors) are realized or observable; all operational decisions are second-stage. While we only consider two-stage problems, we note that the modeling approach and solution approach based on PH can easily be extended to the multi-stage case. This could allow for modeling of uncertainties in load and risk profiles within each scenario (week) of the simulation, which is a subject of our future work.

We denote first- and second-stage variables generally as  $x \in \mathbb{R}^n$  and  $y \in \mathbb{R}^{S \times m}$ , respectively. In a solution to a two-stage stochastic program, first-stage variables are identical across scenarios, while second-stage variables are scenario-dependent. We will denote a scenario-specific set of second-stage variables as  $y_s \in \mathbb{R}^m$ . First-stage variables must be indexed over the stage at which they are made. Let  $x^j$  denote the decision vector for stage  $j$ ,  $\bar{x}^j$  hold all decision for stages up to and including stage  $j$ , and  $x$  refer to decisions made for all stages in  $\mathcal{J}$ . This first- and second-stage relationship is shown graphically in Figure 1.

The general formulation of a stochastic program is:

$$\begin{aligned} \min_{x, y_s} \sum_{s \in \mathcal{S}} \mathbb{P}(s) \left[ f_1(x^1) + \sum_{j=2}^J f_j(x^j; \bar{x}^{j-1}, \xi^j(s)) \right] \quad (16) \\ x \in X_\xi, \xi \in \Xi, \end{aligned}$$

where  $f_i$ ,  $i = 1, \dots, J$ , are generic functions associated with each stage. For a two-stage stochastic program,  $\mathcal{J} = \{2\}$  and  $x^2$  maps directly to  $y$ . The optimization objective in a two-stage stochastic program generally is a function of both first- and second-stage variables. However, PH-based solution methods apply in situations where dependence of cost for one class of variables may be absent.

The first term in (16) represents the costs associated *only* with the first stage, and thus  $f_1$  is a function of the first-stage decisions  $x^1$ . The second term in (16) represents the expected costs associated with subsequent stages  $j \in \mathcal{J}$ . For this term,  $f_j$  is a function of all  $j - 1$  stage decisions  $\bar{x}^{j-1}$  and sampled random realizations for all scenarios  $s \in \mathcal{S}$ . For a two-stage program, this is just a function of  $x$  and  $y_s$  for each scenario. A feasible solution requires all  $j$ -stage variables to be the same across all scenarios  $s \in \mathcal{S}$ , so we can only select first-stage variables  $x$  that are feasible for *all* random parameters considered. For a two-stage program, second-stage variables  $y_s$  must be feasible for their associated scenario  $s$  and the random variables realized for that scenario  $\xi(s)$ .

PH decomposes problem (16) by scenario, in contrast to stage-based decomposition approaches such as the L-shaped method [29]. In a two-stage program, first-stage variables are decoupled, which effectively creates independent sub-problems corresponding to the individual scenarios; we denote the resulting scenario-specific first-stage variables as  $x_s$ . This yields an alternative version of an extensive form of a stochastic program, as follows:

$$\min_{x, y_s} \sum_{s \in \mathcal{S}} \mathbb{P}(s) \left[ f_1(x_s^1) + \sum_{j=2}^J f_j(x_s^j; \bar{x}_s^{j-1}, \xi^j(s)) \right] \quad (17a)$$

$$x = x_s, \quad \forall s \in \mathcal{S}, x \in X_\xi, \xi \in \Xi. \quad (17b)$$

Recall that a valid solution to the overall problem is contingent on the first-stage variables agreeing across all scenarios  $s \in \mathcal{S}$ . To ensure the optimization does not make decisions based on knowledge of the future, all scenario-specific first-stage variables must be linked together via non-anticipativity constraints to ensure  $x = x_1 = \dots = x_s$ , as shown in (17b).

To solve this problem, PH iteratively reaches a solution that satisfies all non-anticipativity constraints and scenario-specific constraints. Let us refer to the iteration number as  $v$ . PH drops the non-anticipativity constraints (17b) and solves each scenario-based sub-problems independently. This step becomes trivially parallelizable since each sub-problem is separable from each other. The objective of each decomposed problem is appended with a proximal term, as shown in (18):

$$\frac{\rho}{2} \|x - \bar{x}^{(v)}\|^2. \quad (18)$$

This term penalizes the objectives based on the magnitude of disagreement between all first stage variable values and also generally prohibits unbounded sub-problems. Here,  $\rho$  is the quadratic penalty parameter.

The general PH algorithm for a two-stage stochastic program is outlined in Algorithm 1. PH has convergence guarantees in the context of linear programs with continuous

variables, and acts as a heuristic with gap-closing capabilities when discrete decision variables are involved. PH employs and updates scenario-specific weight vectors  $w_s^{(j,v)}$  at each iteration  $v$  of the algorithm. These vectors function similarly to dual weights, with the overall goal of pushing each sub-problem to optimality while *simultaneously* encouraging agreement of first-stage variable values across all scenarios  $s \in \mathcal{S}$ . After solving the decomposed sub-problems, the aggregated first-stage variables  $\bar{x}^{(j,v+1)}$  are calculated. These are used as a measure of consensus among all of the dis-aggregated first-stage variables  $x_s$ . Next, PH updates the price vector  $w_s^{(j,v+1)}$  for all scenarios  $s \in \mathcal{S}$ . This weighting factor encourages first-stage variables that are far from the average of all scenario first-stage variables  $\bar{x}$  to move closer to the average when solving the next iteration's sub-problems. There are certain properties of the price vectors that must be ensured to maintain algorithmic guarantees. PH maintains these requirements by the way price vectors are initialized and updated [28]. Sub-problems are solved again, and the iterative process continues until a specified convergence criteria is met. For example, this could be when the deviations between the dis-aggregated first-stage variable and their averages falls below a specified tolerance.

#### IV. APPLICATION OF PROGRESSIVE HEDGING TO BATTERY SIZING, SITING, AND OPERATING PROBLEMS

We use PH, as described in Section III, to solve the battery sizing, siting, and operation optimization problem introduced in Section II. We leverage the capabilities of the open-source Python package `mpi-sppy` [30], which is built on 1) the algebraic modeling language Pyomo [31] to specify optimization sub-problems and 2) MPI (Message Passing Interface) – specifically, `mpi4py` – to efficiently execute PH in parallel.

We adapt the scenario-based approach of PH to allow for decomposition by time periods  $\mathcal{T}$ , in addition to allowing

---

#### Algorithm 1 Progressive Hedging Algorithm for Optimizing Battery Sizing, Siting, and Operation

---

- 1: **Initialize:**  $v \leftarrow 0$  and  $w_s^{(j,v+1)} \leftarrow 0, \forall s \in \mathcal{S}, j \in \mathcal{J}$   
Compute  $\forall s \in \mathcal{S}$

$$x_s^{(v+1)} \in \underset{x_s}{\operatorname{argmin}} f_1(x_s^1) + \sum_{j=2}^J f_j(x_s^j; \bar{x}_s^{j-1}, \xi^j(s))$$

- 2: **Iteration Update:**  $v \leftarrow v + 1$

- 3: **Aggregation:**  $\bar{x}^{(v+1)} \leftarrow \sum_{s \in \mathcal{S}} \mathbb{P}(s) x_s^{(v)}$

- 4: **Price Update:**  $w_s^{(j,v+1)} \leftarrow w_s^{(j,v)} + \rho (x_s^{(j,v)} - \bar{x}^{(j,v)})$

- 5: **Decomposition:**  
Compute  $\forall s \in \mathcal{S}$

$$x_s^{(v+1)} \in \underset{x_s}{\operatorname{argmin}} f_1(x_s^1) + \sum_{j=2}^J f_j(x_s^j; \bar{x}_s^{j-1}, \xi^j(s))$$

$$+ \sum_{j=1}^{J-1} [w_s^{(j,v+1)\top} x_s^j + \frac{\rho}{2} \|x_s^j - \bar{x}^{(j,v+1)}\|^2]$$

- 6: **Termination:** If criterion met, Stop. Otherwise  $\rightarrow$  Step 2.
-




Fig. 2: Time-period decomposition structure of battery sizing, siting, and operation model.

for first-stage variables (relating to state-of-charge variables) that do not appear in every time period’s sub-problem. We emphasize that this is a key modification relative to traditional progressive hedging algorithms to account for the temporal decomposition necessary to model the batteries’ states-of-charge. We begin by decomposing the optimization model into a set of sub-problems  $\mathcal{P} = \{1, \dots, P\}$ , each representing a contiguous period of time from the full time horizon  $\mathcal{T}$ . In the WECC-240 case study described subsequently, we assign groups of contiguous days to form a total of  $P = 50$  time periods for roughly week-long subproblems. Each individual time  $t$  within a time period  $p$  represents a single hour.

The resulting decomposition structure is depicted in Figure 2, and is mathematically identical to a two-stage stochastic program decomposition by scenario. In the first stage, we have battery placement and sizing decisions,  $x^n$ ; these must be non-anticipative relative to all time periods in  $\mathcal{P}$ . The first-stage cost  $f_1$  is  $C_{gen}$ , per (15). All operational decisions (including load shedding, generator outputs, power flow, voltage angles, and states-of-charge) are second-stage variables that can vary by sub-problem.

The decomposition discussed so far addresses all aspects of our optimization model with the exception of battery state-of-charge variables that couple time periods of adjacent sub-problems. We must ensure the states-of-charge of every placed battery are consistent at the start and end of temporally adjacent sub-problems. These additional complicating constraints are absent in traditional stochastic programming contexts. To address this final aspect of our optimization model in the context of PH-based decomposition, we extend the index set for all state-of-charge variables  $E_t^n(\cdot)$  for each time period  $p \in \mathcal{P}$  by including the last unit of time in the immediately preceding time period  $p - 1$ . (No such modification is performed for the first time period.) This structure is depicted graphically in Figure 3. These shared variables between temporally adjacent sub-problems are treated as additional first-stage variables. However, these variables do not appear in all sub-problems, requiring introduction of corresponding variable-specific probabilities (i.e., scenario weights in the context of the battery sizing, siting, and operation problem) – specifically equal to 0.5 for sub-problems with the variable and 0 otherwise. Finally, we note that the state-of-charge variables are not referenced in the objective function (state-of-the-charge is not explicitly costed), such that determination of PH penalty parameters  $\rho$  for such variables is more challenging [18].




Fig. 3: Battery state-of-charge decomposition structure in the battery sizing, siting, and operation model.




Fig. 4: Geo-location of the WECC-240 system network.


## V. CASE STUDIES

To demonstrate the capabilities of our algorithm, we consider a synthetic test network geolocated in the western United States with 240 buses (WECC-240) [32], for which we associate temporally and spatially varying wildfire risk data from the USGS Wind-Enhanced Fire Potential Index [33]. Wildfire data is defined for the end of April through November. The network topology and electrical parameters are adopted from [34] augmented with data originating from [35] and [36]. To avoid infeasibility, power inputs represented as DC lines modeled as negative loads were deactivated for these simulations and slack generation is included at each bus. The geo-located structure of the WECC-240 network topology is provided in Figure 4. We next consider both direct solution of the optimization model described in Section II and solution using PH-based decomposition (per Sections III and IV).


### A. Solution of the Extensive Form

We first conduct tests on the WECC-240 case using the full extensive form model outlined in Section II, i.e., considering direct solution of (16) without decomposition as in Algorithm 1. The corresponding models were solved using Gurobi 10.0.0 [37]. To implement the optimization formulations, we used Julia 1.8.0 [38] with JuMP v1.11.1 [39] along with the data input functionality of PowerModels.jl v0.19.8 [40]. Computations were conducted on the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology [41]. Each case was run on one node with 8 cores with 24 GB of memory. Each node has Dual Intel Xeon Gold 6226 CPUs @ 2.7 GHz.

Figure 5 shows the resulting battery locations during nominal operation (for a month in April of 2021) and during a



(a) April 2021



(b) June 2021

Fig. 5: Optimal battery placements for the WECC network in April 2021 (top) and in June 2021 (bottom). Red circles are sized proportionally to the number of batteries placed at that bus with the largest circle representing 4 batteries.

period of high wildfire risk (for a month in June of 2021). From these tests, we observe that optimal placement location varies based on the time of year. Given the computational difficulties with long time horizons, the deterministic model cannot be used to find optimal placement values over the span of a full year. For the one-month extensive form in June 2021, the model contains over 1.2 million constraints, 1.9 million variables, and 5 million nonzero values. The one-month simulations took nearly 25 minutes of wall-clock time. A one-year simulation was terminated without a solution after a 24-hour build time. Instead, we turn to the PH algorithm described in Sections III and IV. Results from the temporally decomposed problem are discussed next.




Fig. 6: Placement decisions on the WECC-240 network from the full-year simulation decomposed into 120 time periods. Red circles are sized proportionally to the number of batteries placed at that bus with the largest circle representing 1.05 batteries.

### B. Decomposition Results

For PH, we decompose the full one-year time horizon into 120 time periods, most with 72 hours, and a handful of scenarios including an additional 24 hours, to model the full year. We report PH solutions that are fully non-anticipative, i.e., are valid incumbents for the extensive form optimization model. Following the methodology outlined in Section IV, we solved a decomposed, one-year horizon battery sizing, siting, and operating problem for the WECC-240 network. We used Gurobi v10.0.2 [37] to solve this problem in Python v3.9.12 and modeled the optimization formulation in Pyomo v6.6.1 [31]. We employed `mpi-sppy` v0.12 [30] to execute PH in parallel, as outlined in Section III. This was executed on the quartz HPC at Lawrence Livermore National Laboratory. We executed this problem using 60 nodes of the HPC. Each node has two 18-core Xeon E5-2695 processors (2.1 GHz) and 128 GB of memory. Furthermore, we decomposed the problem into 120 scenarios. We used 240 MPI ranks to solve the problem, with 120 dedicated to running the PH algorithm on each sub-problem and 120 dedicated to executing an asynchronous incumbent-finder. (A rank is a unique number that identifies each parallel process.) Each Gurobi solve was limited to 36 threads, and the  $\rho$  parameter was set to a constant value of 0.001.

For the decomposed problem, the reported solution was found after 50 iterations of the PH algorithm and incumbent finders. We reached a relative optimality gap of 0.023% and an absolute gap of 0.367 in under 70 minutes. Note, a total slack generation of 0.0106 p.u. was included in the solution throughout the year across all buses to ensure feasibility. The extensive form of a one-year run, as mentioned, did not *build*




Fig. 7: Load shedding and battery discharge for the WECC-240 network in June of 2021. Note that battery discharging decisions and load shed are both used to serve power in addition to generation (not shown).

within 24 hours.

Comparing the placement locations from the full-year decomposition, seen in Figure 6, to the the month-based placements in Figure 5, we observe that the final sites have some overlap with decisions from the extensive model of April but are largely different the placement decisions in June. Results from April and June are shown as a subset of the results of the full year. As can be seen in Figure 7, during the month of June, there is considerable load shedding during days with many lines de-energized. Battery operation at bus 6401 in the WECC-240 can be seen in Figure 8. It is interesting to note that there is very minimal difference in the operation of the batteries between April and June in these results. This might be explained by the selected battery shown largely responding to local demand and not being significantly influenced by line de-energizations. As seen in Figure 5, batteries are most helpful in the southwestern US for the June 2021 wildfire risk profile. When looking at the battery placements based only on April and June, we see drastically different results. The results from the full year decomposed run reflects this as battery installations are distributed throughout the network to aid in operations at all times of the year in different regions.

## VI. CONCLUSION

We have introduced a new approach to solving utility-scale battery sizing, siting, and operations optimization problems, considering the mitigation of wildfire risk and associated load shedding. Direct approaches to solving an extensive form of the optimization model fail to scale to significant time horizons, e.g., on the order of a year. Long time horizons are critical to capture operations in both nominal and wildfire seasons. We introduce a decomposition approach based on a progressive hedging (PH) algorithm. Our approach allows us to capture both a standard planning/operations decomposition structure, but also non-anticipative variables (state-of-charge in our model) that couple sub-problems representing temporally adjacent operational periods. Implemented in a scalable open-source stochastic programming library (`mpi - sppy`) and

executed using parallel computing clusters, our PH-based decomposition yields near-optimal battery placements across a year-long time horizon in tractable run times. The ability to consider long time horizons is shown to be critical, in that the resulting battery placements can differ from placements obtained using shorter time horizons. Our implementation of PH with temporal decomposition allows for benefits in solve time as well as considerable benefits in parallelizing the build time for large models.

In addition, these tools can be used for other forms of infrastructure upgrades such as line undergrounding. Similarly, longer time scales could be considered. Instead of making investment decisions based on one year, a rolling-budget schedule could be introduced that allows infrastructure to be placed over multiple years with changing load demand and varying wildfire risk. Given the nature of the progressive hedging algorithm, this formulation could also be extended to consider multiple wildfire risk profiles and other uncertainties within each time period.

## REFERENCES

- [1] S. Martinuzzi, A. J. Allstadt, A. M. Pidgeon, C. H. Flather, W. M. Jolly, and V. C. Radeloff, "Future changes in fire weather, spring droughts, and false springs across US National Forests and Grasslands," *Ecological Applications*, vol. 29, no. 5, p. e01904, 2019.
- [2] J. E. Keeley and A. D. Syphard, "Twenty-first century California, USA, wildfires: Fuel-dominated vs. wind-dominated fires," *Fire Ecology*, vol. 15, no. 1, p. 24, 2019.
- [3] PG&E, "Learn about public safety power shutoffs," 2021. [Online]. Available: [https://www.pge.com/en\\_US/residential/outages/public-safety-power-shutoff/learn-about-psps.page](https://www.pge.com/en_US/residential/outages/public-safety-power-shutoff/learn-about-psps.page)
- [4] M. Sotolongo, C. Bolon, and S. H. Baker, "California power shutoffs: Deficiencies in data and reporting," *Initiative for Energy Justice*, October 2020.
- [5] G. Wong-Parodi, "Support for public safety power shutoffs in California: Wildfire-related perceived exposure and negative outcomes, prior and current health, risk appraisal and worry," *Energy Research & Social Science*, vol. 88, p. 102495, 2022.
- [6] PG&E, "10,000-Mile Undergrounding Program," July 2023. [Online]. Available: [https://www.pge.com/pge\\_global/common/pdfs/customer-service/other-services/electric-undergrounding-program/PGE-10K-UG-Program-Local-Work-Plan-Maps.pdf](https://www.pge.com/pge_global/common/pdfs/customer-service/other-services/electric-undergrounding-program/PGE-10K-UG-Program-Local-Work-Plan-Maps.pdf)
- [7] S. Taylor, G. Setyawan, B. Cui, A. Zamzam, and L. A. Roald, "Managing wildfire risk and promoting equity through optimal configuration of networked microgrids," in *Proceedings of the 14th ACM International Conference on Future Energy Systems*, 2023, pp. 189–199.
- [8] PG&E, "Hardening the electric system," 2023. [Online]. Available: [https://www.pge.com/en\\_US/safety/emergency-preparedness/natural-disaster/wildfires/cwsp-system-hardening.page?WT.mc\\_id=Vanity\\_systemhardening](https://www.pge.com/en_US/safety/emergency-preparedness/natural-disaster/wildfires/cwsp-system-hardening.page?WT.mc_id=Vanity_systemhardening)
- [9] Y. Yang, S. Bremner, C. Menictas, and M. Kay, "Battery energy storage system size determination in renewable energy systems: A review," *Renewable and Sustainable Energy Reviews*, vol. 91, pp. 109–125, 2018.
- [10] K. Marnell, M. Obi, and R. Bass, "Transmission-scale battery energy storage systems: A systematic literature review," *Energies*, vol. 12, no. 23, p. 4603, 2019.
- [11] S. Singer, "PG&E, Energy Vault plan largest US utility-scale battery, green hydrogen long-duration storage project," *Utility Dive*, Jan 2023. [Online]. Available: <https://www.utilitydive.com/news/california-wildfires-battery-storage-fuel-cells-hydrogen/639771>
- [12] C. Murray, "CAISO: 1GW dispatch during July wildfire a 'fabulous moment' for battery storage amidst 12-fold growth," *Energy Storage News*, March 2022. [Online]. Available: <https://www.energy-storage.news/caiso-1gw-dispatch-during-july-wildfire-a-fabulous-moment-for-battery-storage-amidst-12-fold-growth/>






Fig. 8: State-of-Charge of batteries at bus 6401 on the WECC-240 network in a week of April 2021 (left) and a week of June 2021 (right) with red bars showing the number of de-energized transmission lines on each day.

- [13] S. F. Santos, M. Gough, D. Z. Fitiwi, A. F. P. Silva, M. Shafie-Khah, and J. P. S. Catalão, "Influence of battery energy storage systems on transmission grid operation with a significant share of variable renewable energy sources," *IEEE Systems Journal*, vol. 16, no. 1, March 2022.
- [14] L. Fiorini, G. A. Pagani, P. Pelacchi, D. Poli, and M. Aiello, "Sizing and siting of large-scale batteries in transmission grids to optimize the use of renewables," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 7, no. 2, pp. 285–294, 2017.
- [15] J. Aguado, S. de la Torre, and A. Trivino, "Battery energy storage systems in transmission network expansion planning," *Electric Power Systems Research*, vol. 145, 2017.
- [16] D. Chen, Z. Jing, and H. Tan, "Optimal siting and sizing of used battery energy storage based on accelerating benders decomposition," *IEEE Access*, vol. 7, pp. 42 993–43 003, 2019.
- [17] Y. Zhou, K. Sundar, D. Deka, and H. Zhu, "Optimal power system topology control under uncertain wildfire risk," *arXiv:2303.07558*, 2023.
- [18] J.-P. Watson and D. L. Woodruff, "Progressive hedging innovations for a class of stochastic mixed-integer resource allocation problems," *Computational Management Science*, vol. 8, no. 4, p. 355, 2011.
- [19] A. Colthorpe, "US installed grid-scale battery storage capacity reached 9GW/25GWh in 'record-breaking' 2022," February 2023. [Online]. Available: <https://www.energy-storage.news/us-installed-grid-scale-battery-storage-capacity-reached-9gw-25gwh-in-record-breaking-2022/>
- [20] A. Akhil, G. Huff, A. Currier, B. Kaun, D. Rastler, S. Chen, A. Cotter, D. Bradshaw, and W. Gauntlett, "DOE/EPRI 2013 Electricity Storage Handbook in Collaboration with NRECA," Sandia National Laboratories, Tech. Rep., July 2013.
- [21] International Renewable Energy Agency (IRENA), "Innovation landscape brief: Utility-scale batteries," Tech. Rep., 2019. [Online]. Available: [https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Sep/IRENA\\_Utility-scale-batteries\\_2019.pdf](https://www.irena.org/-/media/Files/IRENA/Agency/Publication/2019/Sep/IRENA_Utility-scale-batteries_2019.pdf)
- [22] X. Luo, J. Wang, M. Dooner, and J. Clarke, "Overview of current development in electrical energy storage technologies and the application potential in power system operation," *Applied Energy*, vol. 137, pp. 511–536, 2015.
- [23] N. Nazir and M. Almassalkhi, "Guaranteeing a physically realizable battery dispatch without charge-discharge complementarity constraints," *IEEE Transactions on Smart Grid*, vol. 14, no. 3, pp. 2473–2476, 2023.
- [24] A. Kody, R. Piansky, and D. K. Molzahn, "Optimizing transmission infrastructure investments to support line de-energization for mitigating wildfire ignition risk," in *IREP Symposium on Bulk Power System Dynamics and Control-XI*, July 2022.
- [25] N. Rhodes, L. Ntamo, and L. Roald, "Balancing wildfire risk and power outages through optimized power shut-offs," *IEEE Transactions on Power Systems*, vol. 36, no. 4, pp. 3118–3128, 2020.
- [26] A. Kody, A. West, and D. K. Molzahn, "Sharing the load: Considering fairness in de-energization scheduling to mitigate wildfire ignition risk using rolling optimization," in *IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 5705–5712.
- [27] L. A. Roald, D. Pozo, A. Papavasiliou, D. K. Molzahn, J. Kazempour, and A. Conejo, "Power systems optimization under uncertainty: A review of methods and applications," *Electric Power Systems Research*, vol. 214, p. 108725, 2023, presented at the *22nd Power Systems Computation Conference (PSCC 2022)*.
- [28] R. T. Rockafellar and R. J.-B. Wets, "Scenarios and policy aggregation in optimization under uncertainty," *Mathematics of Operations Research*, vol. 16, no. 1, pp. 119–147, 1991.
- [29] R. Van Slyke and R.-B. Wets, "L-shaped linear programs with applications to optimal control and stochastic programming," *SIAM Journal on Applied Mathematics*, vol. 17, pp. 638–663, 1969.
- [30] B. Knueven, D. Mildebrath, C. Muir, J. D. Sirola, J.-P. Watson, and D. L. Woodruff, "A parallel hub-and-spoke system for large-scale scenario-based optimization under uncertainty," *Mathematical Programming Computation*, vol. 15, no. 4, pp. 591–619, 2023.
- [31] M. L. Bynum, G. A. Hackebeil, W. E. Hart, C. D. Laird, B. L. Nicholson, J. D. Sirola, J.-P. Watson, D. L. Woodruff *et al.*, *Pyomo-Optimization Modeling in Python*. Springer, 2021, vol. 67.
- [32] A Library of IEEE PES Power Grid Benchmarks: `pglib_opf_case240_pserc`. Accessed: 2022-02-18. [Online]. Available: [https://github.com/power-grid-lib/pglib-opf/blob/master/pglib\\_opf\\_case240\\_pserc.m](https://github.com/power-grid-lib/pglib-opf/blob/master/pglib_opf_case240_pserc.m)
- [33] US Geological Survey, "Wind-Enhanced Fire Potential Index," <https://www.usgs.gov/fire-danger-forecast/wildland-fire-potential-index-wfpi>, accessed: 2022-02-18.
- [34] IEEE PES Task Force on Benchmarks for Validation of Emerging Power System Algorithms, "The Power Grid Library for benchmarking AC optimal power flow algorithms," *arXiv:1908.02788*, Aug. 2019.
- [35] C. Barrows, A. Bloom, A. Ehlen, J. Ikäheimo, J. Jorgenson, D. Krishnamurthy, J. Lau, B. McBenett, M. O'Connell, E. Preston *et al.*, "The IEEE Reliability Test System: A Proposed 2019 Update," *IEEE Transactions on Power Systems*, vol. 35, no. 1, pp. 119–127, 2019.
- [36] J. E. Price and J. Goodin, "Reduced network modeling of WECC as a market design prototype," in *IEEE Power and Energy Society General Meeting*, 2011.
- [37] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: <https://www.gurobi.com>
- [38] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, "Julia: A fresh approach to numerical computing," *SIAM Review*, vol. 59, no. 1, pp. 65–98, 2017. [Online]. Available: <https://doi.org/10.1137/141000671>
- [39] I. Dunning, J. Huchette, and M. Lubin, "Jump: A modeling language for mathematical optimization," *SIAM Review*, vol. 59, no. 2, pp. 295–320, 2017.
- [40] C. Coffrin, R. Bent, K. Sundar, Y. Ng, and M. Lubin, "PowerModels.jl: An Open-Source Framework for Exploring Power Flow Formulations," in *19th Power Systems Computation Conference (PSCC 2018)*, 2018.
- [41] PACE, *Partnership for an Advanced Computing Environment (PACE)*, 2017. [Online]. Available: <http://www.pace.gatech.edu>