

Operations-Based Planning for Placement and Sizing of Energy Storage in a Grid With a High Penetration of Renewables

Krishnamurthy Dvijotham⁽¹⁾, Scott Backhaus⁽²⁾ and Misha Chertkov⁽³⁾

Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195⁽¹⁾

Material Physics and Applications Division⁽²⁾ and

Center for Nonlinear Studies and Theoretical Division,⁽³⁾

LANL, Los Alamos, NM 87545, USA

Abstract

As the penetration level of transmission-scale time-intermittent renewable generation resources increases, control of flexible resources will become important to mitigating the fluctuations due to these new renewable resources. Flexible resources may include new or existing synchronous generators as well as new energy storage devices. The addition of energy storage, if needed, should be done optimally to minimize the integration cost of renewable resources, however, optimal placement and sizing of energy storage is a difficult optimization problem. The fidelity of such results may be questionable because optimal planning procedures typically do not consider the effect of the time dynamics of operations and controls. Here, we use an optimal energy storage control algorithm to develop a heuristic procedure for energy storage placement and sizing. We generate many instances of intermittent generation time profiles and allow the control algorithm access to unlimited amounts of storage, both energy and power, at all nodes. Based on the activity of the storage at each node, we restrict the number of storage node in a staged procedure seeking the minimum number of storage nodes and total network storage that can still mitigate the effects of renewable fluctuations on network constraints. The quality of the heuristic is explored by comparing our results to seemingly “intuitive” placements of storage.

1 Introduction

Electrical grid planning has traditionally taken two different forms; operational planning and expansion or upgrade planning. The first is concerned with the relatively short time horizon of day-ahead unit commitment or hour-ahead or five-minute economic dispatch. The focus is on controlling assets that are already present on

the system to serve loads at minimum cost while operating the system securely. The second typically looks out many years or decades and is focused on optimal addition of new assets, again with a focus on minimizing the cost of electricity. When a system consists entirely of controllable generation and well-forecasted loads, the network power flows do not deviate significantly or rapidly from well-predicted patterns. In this case, expansion planning can be reasonably well separated from operational planning. In the simplest approach, expansion plans may be optimized against two extreme cases, e.g. the system’s maximum and minimum load configurations.

As the penetration of time-intermittent renewables increases, expansion and operational planning will necessarily become more coupled. For an electrical grid with large spatial extent, renewable generation fluctuations (here, we focus on wind generation) at well-separated sites will be uncorrelated on short time scales[1, 2], and the intermittency of this new non-controllable generation will cause the patterns of power flow to change on much faster time scales than before, and in unpredictable ways. New equipment (e.g. combustion turbines or energy storage) and control systems may have to be installed to mitigate the network effects of renewable generation fluctuations to maintain generation-load balance. It is at this point where operations planning must interface with expansion planning. The optimal placement and sizing of the new equipment, if required, depends on how the rest of the network and its controllable components respond to the fluctuations of the renewable generation. Overall, we desire to install a minimum of new equipment by placing it at network nodes that afford us a high degree of controllability, i.e. nodes where controlled power injection and/or consumption have a significant impact on the network congestion introduced by the renewable fluctuations. From the outset, it is not clear which nodes provide the best controllability. Im-

plicity, placing a minimum of new equipment implies that it will experience a high duty, thus avoiding outcomes where equipment is only used for a small fraction of time.

Before discussing our initial approach at integrating operational planning and expansion planning, we first summarize a few methods for mitigating the intermittency of renewable generation. When renewable penetration is relatively low and the additional net-load fluctuations are comparable to existing load fluctuations, a power system may continue to operate “as usual” with primary and secondary regulation reserves[3] being controlled via a combination of distributed local control, i.e. frequency droop, and centralized control, i.e. automatic generation control (AGC). In this case, *planning* for renewables may simply entail increasing the level of reserves to guard against the largest expected fluctuation in *aggregate renewable output*.

As the penetration level grows, system operators may simply continue to increase the level of reserves, however, if this fast-moving generation comprises natural gas combustion turbines, this simple planning will generally result in increased renewable integration costs[4] which are usually spread over the rate base. Alternatively, operational planning can be improved by incorporating long-range and short-range renewable generation forecasting to better schedule the controllable generation (energy and reserves) to meet net load and operate reliably[4, 5, 6]. Long-range forecasts, typically day-ahead or farther, are used in a unit commitment optimization to ensure adequate generation and reserves will be online to meet the expected net load and its uncertainty. Short-range forecasts, typically hour ahead or shorter, are used in an economic dispatch optimization that sets actual generation and reserve levels of the committed generation. Some have investigated rolling unit commitment[7, 4] where updated wind forecasts are used modify the unit commitment more frequently. Simulations have resulted in lower overall renewable integration costs.

Both unit commitment and economic dispatch seek minimize the cost of electricity, however, they must also respect system constraints including generation limits, transmission line thermal limits, voltage limits, system stability constraints, and N-1 contingencies. Previous works[4, 5, 7, 6] have generally looked at the effects of stochastic generation on the economics and adequacy of *aggregate* reserves while not considering such network constraints. These constraints may be respected for a dispatch based on a *mean* renewable forecast, however, if the number of renewable generation sites and their contribution to the overall generation is significant, verifying the system security of all probable renewable fluc-

tuations (and the response of the rest of the system) via enumeration will be exponentially complex in the number of sites resulting in a computationally intractable problem.

The approaches summarized above do not consider network constraints or the behavior of the system on time scales shorter than the time between economic dispatches (one hour in the case of [4]). In particular, they do not model how fast changes in renewable generation and the compensating response of regulation reserves interact with network constraints. In our initial study, we augment the approaches summarized above by focusing on the behavior of the electrical network at a finer time resolution and investigate how the control of energy storage affects its placement and sizing. We presume that the unit commitment problem has been solved, and at the start of a time period, we dispatch the controllable generation by solving a DC optimal power flow (DCOPF) based on the *site-specific mean forecast* for wind generation. In the time before the next DCOPF is executed, the wind generation fluctuates, and we control a combination of existing generators (using a simplified description of frequency droop control) and energy storage to avoid violations of network constraints and generation limits. For each level of wind penetration, we generate many different realizations of wind fluctuations, and we gather statistics on the activity of the energy storage at each node. The statistics from simulated system operations are then coupled to the expansion planning process by developing a heuristic to guide the optimal placement and sizing of storage throughout the network—a result that cannot be achieved with the previous approaches described above.

The rest of this manuscript is organized as follows. Section 2 lays out the system model including the development of our heuristic for placement and sizing of energy storage. Section 3 describes the simulations carried out on a slightly modified version of RTS-96[8] and the application of our design heuristic. Section 4 gives some interpretations of the results, and Section 5 wraps up with some conclusions and directions for future work.

2 Mathematical Formulation

We briefly describe the mathematical formulation of the optimal control problem and the algorithm used to solve it. We begin with a quick review of the DC power flow model and the proportional control scheme used to emulate frequency droop and AGC[9].

2.1 Background and Notation

Let $\mathcal{G} = (\mathcal{E}, \mathcal{V})$ denote the undirected graph underlying the power grid, and n denote the number of nodes in the grid. Nodes are classified into three types:

- 1 Loads \mathcal{G}_l (assumed fixed over time scale of the study relative to the renewable generation).
- 2 Traditional Generators \mathcal{G}_g (whose output can be controlled).
- 3 Renewable generators \mathcal{G}_r (generators based on renewable sources whose output fluctuates over time).

Given the vector \mathbf{p} of power generated/consumed at every node in the grid, the power flow equations determine how much power \mathbf{f}_{ij} flows through each line in the network from node i to node j . For any $S \subset \mathcal{V}$, we denote the power generation at nodes in S by $\mathbf{p}_S = \{\mathbf{p}_i : i \in S\}$ and use similar notation for other quantities (like energy). The DC power flow equations are a linearization of the exact AC power flow equations, and have been widely used to compute optimal economic dispatch of generators in the grid. If $\text{Neb}(j)$ denotes the set of neighbors of node j in the grid and \mathbf{x}_{ij} denote the inductance of the line from i to j , the DC power flow equations are given by:

$$\begin{aligned} \mathbf{p}_i &= \sum_{j \in \text{Neb}(i)} \frac{\theta_j - \theta_i}{\mathbf{x}_{ij}} \\ \mathbf{f}_{ij} &= \mathbf{x}_{ij}(\theta_j - \theta_i), \end{aligned}$$

where θ_i is the phase of the complex voltage at node i . Written in matrix form, the first set of equations become $\mathbf{p} = -\mathbf{L}_x \boldsymbol{\theta}$ where \mathbf{L}_x is the Laplacian of the graph \mathcal{G} with edge weights $1/\mathbf{x}_{ij}$. The Laplacian is not invertible, but if we are given a balanced power configuration, i.e., $\sum_i \mathbf{p}_i = 0$, we can always find a unique solution for the $\boldsymbol{\theta}$. This mapping can be constructed from the eigenvalue decomposition of $\mathbf{L}_x = \sum_i w_i \mathbf{e}_i \mathbf{e}_i^T$ by inverting the non-zero eigenvalues to get $\mathbf{M}_x = \sum_{i:w_i>0} \frac{1}{w_i} \mathbf{e}_i \mathbf{e}_i^T$ and $\boldsymbol{\theta} = \mathbf{M}_x \mathbf{p}$. Letting \mathbf{M}_x^i denote the i -th row of \mathbf{M}_x , we can summarize the power flow equations as:

$$\begin{aligned} \mathbf{f} &= \mathbf{M}\mathbf{p} \\ \mathbf{M}^{ij} &= \mathbf{x}_{ij}(\mathbf{M}_x^i - \mathbf{M}_x^j) \end{aligned} \quad (1)$$

The aim of our energy storage control scheme is to dispatch power from storage such that, in spite of fluctuations in the renewable generation, the following con-

straints are maintained all the time:

$$\begin{aligned} -\bar{\mathbf{f}} &\leq \mathbf{M}\mathbf{p} \leq \bar{\mathbf{f}} && \text{(Transmission Capacities)} \\ 0 &\leq \mathbf{p}_{\mathcal{G}_g} \leq \bar{\mathbf{p}}_{\mathcal{G}_g} && \text{(Generation Capacities)} \\ \sum_i \mathbf{p}_i &= 0 && \text{(Power Balance)} \end{aligned} \quad (2)$$

where the sum over i in the third constraint is a sum over three types of nodes.

2.2 Power Generation and Proportional Control

We will design a control scheme that operates in the duration T between two successive DCOFs. At the beginning of each such time horizon, we assume that a subset of active generators \mathcal{G}_g are committed (based on the predetermined unit-commitment). We set their generation levels according to a DCOF:

$$\min_{\mathbf{p}_{\mathcal{G}_g}} \sum_{i \in \mathcal{G}_g} w_i \mathbf{p}_i$$

Subject to Eqs. 1 and 2

where w_i are generation costs for each generator. Let \mathbf{p}^0 be the complete vector of initial generation at every node based on the DCOF. The output $\mathbf{p}_{\mathcal{G}_r}$ of renewable generation \mathcal{G}_r is taken to be the mean of the renewable generation $\mathbf{p}_{\mathcal{G}_r}^0$ over the time between DCOFs. We assume we know this value exactly, which implies we have perfect predictions of the mean output of the renewables.

In reality, the renewable generation $\mathbf{p}_{\mathcal{G}_r}$ fluctuates during this time causing a mismatch between total generation and total load. We assume that the generators \mathcal{G}_g respond by changing their output by dividing the mismatch in some predetermined proportion α (usually chosen based on generation capacities)[9]. For notational convenience, we take α to be a vector of length n , with entries $\alpha_i = 0$ for $i \notin \mathcal{G}_g$. Let $\mathbf{p}^r = \mathbf{p}_{\mathcal{G}_g} - \mathbf{p}_{\mathcal{G}_r}^0$ denote the vector of renewable generation change. We assume for simplicity that loads are fixed, which is reasonable since in many cases fluctuations in renewables will be much larger than fluctuations in loads. The generators \mathcal{G}_g will then respond as:

$$\mathbf{p}_{\mathcal{G}_g} = \mathbf{p}_{\mathcal{G}_g}^0 + \alpha_{\mathcal{G}_g} (-\sum \mathbf{p}^r),$$

where non-zero components of $\alpha_{\mathcal{G}_g}$ associated with generators are chosen such that the net balance of power is maintained, $\sum \mathbf{p}^r = \sum_{i \in \mathcal{G}_g} \mathbf{p}_i$.

2.3 Controlling Storage and Constraints on Storage

We use \mathbf{s} to denote the vector of energies stored at each node. However, the quantity we control directly is \mathbf{p}^s ,

the power drawn from each storage node at any given time. In our storage placement and sizing heuristic, we allow \mathbf{p}^s and \mathbf{s} to be as high as needed to ensure that the constraints (2) are satisfied. We also restrict the set of nodes that have storage to a set \mathcal{G}_s . However, again for notational convenience, \mathbf{p}^s, \mathbf{s} are still assumed to be vectors of length n and we set $\mathbf{p}_j^s = 0, \mathbf{s}_j = 0$ for $j \notin \mathcal{G}_s$. In the paper, we use the optimal energy storage control as a heuristic to decide placement and sizing of storage. This decision only assumes that we know what the renewable generation is going to be over the next window of time of length T until the next DCOPT. Thus, we would like this decision to be robust to the behavior of wind beyond T and that our control over the next time window can work independently of what the control was in the current time window. Hence we enforce the following invariance condition: No net energy is exchanged between the grid and the storage during the window of time T . Mathematically, this means that we require:

$$\mathbf{s}(0) = \mathbf{s}(T) \quad (3)$$

We also assume that the storage is not accounted for in the DCOPT, i.e. $\mathbf{p}^s(0) = 0$.

2.4 Overall Power generation

If storage is being dispatched at a node, it adds to the pre-existing generation \mathbf{p}^{ns} such that total generation is $\mathbf{p} = \mathbf{p}^{ns} + \mathbf{p}^s$, where \mathbf{p}^{ns} is the generation without including the local storage dispatch. For \mathcal{G}_l or \mathcal{G}_r , \mathbf{p}^{ns} is just the fixed load value or the time-variable renewable generation. The storage dispatch simply adds to these outputs.

For generators, i.e. \mathcal{G}_g , the situation is slightly more complicated. The dispatch of storage modifies the instantaneous load-generation imbalance so that the proportional control we proposed for \mathcal{G}_g is slightly modified. The controlled generation at a node responds to the sum of storage dispatch at all nodes, including its own, i.e.

$$\mathbf{p}_{\mathcal{G}_g}^{ns} = \mathbf{p}_{\mathcal{G}_g}^0 + \alpha_{\mathcal{G}_g} \left(- \sum_{i \in \mathcal{G}_r} \mathbf{p}_i^r - \sum_i \mathbf{p}_i^s \right).$$

Thus, the overall power generation from a \mathcal{G}_g node is

$$\mathbf{p} = \mathbf{p}^0 + \mathbf{p}^s + \mathbf{p}^r + \alpha \left(- \sum_{i \in \mathcal{G}_r} \mathbf{p}_i^r - \sum_i \mathbf{p}_i^s \right) \quad (4)$$

2.5 Optimal Control with perfect Forecasts

In this initial work on operations-based placement and sizing of energy storage, we implement an optimal con-




Figure 1: Penalty Functions

trol algorithm that assumes perfect forecasts of renewable generation, i.e. the output of each renewable generation node is known exactly over the time horizon T . The control algorithm uses this information to decide how much power to draw from or inject into each storage node in order to keep the grid within the constraints (2). We note that, at the same time, the generation \mathcal{G}_g is responding according to (4). The control algorithm computes this optimal dispatch by minimizing a cost function consisting of three terms. The cost function tries to enforce the constraints (2) in a soft manner: We do this because in some situations, it might be infeasible to maintain all the constraints strictly. In those cases, we allow violations to occur but the controller suffers a cost that blows up cubically with the magnitude of the constraint violation. Thus, the controller will do its best to keep the system within its constraints, but when it is infeasible to do so, it will allow the violations to occur. To make the notation compact, we define two auxiliary functions f and h , which are plotted in Figure 1 and defined below.

$$f(x, a, b) = \begin{cases} 0 & \text{if } a \leq x \leq b \\ (\kappa_f * (x - b))^3 & \text{if } x > b \\ (\kappa_f * (x - a))^3 & \text{if } x < a \end{cases}$$

$$h(x) = \log(\cosh(\kappa_h x))$$

where $\kappa_f = 50, \kappa_h = .001$.

The three costs in the penalty function are:

- 1 cli(\mathbf{p}): penalizes overloading of transmission lines:
cli(\mathbf{p}) = $\sum_{(i,j) \in \mathcal{E}} f(\mathbf{M}^{ij} \mathbf{p}, -\bar{\mathbf{f}}_{ij}, \bar{\mathbf{f}}_{ij})$
 - 2 cg(\mathbf{p}^{ns}): penalizes violation of capacity constraints of controllable generators:
cg(\mathbf{p}^{ns}) = $\sum_{i \in \mathcal{G}_g} f(\mathbf{p}_i^{ns}, 0, \bar{\mathbf{p}}_i)$
 - 3 csp(\mathbf{p}^s): penalizes the absolute value of power drawn from storage (increases very slowly):
csp(\mathbf{p}^s) = $\sum_i h(\mathbf{p}_i^s)$
- The fact that h increases very slowly ensures that

$\text{cli}(\mathbf{p}), \text{cg}(\mathbf{p}^{ns})$ dominate this cost by far, so that the controller will never choose to allow violations to occur in order to save on power capacity.

We define the total cost as $\text{cost}(\mathbf{p}, \mathbf{p}^{ns}, \mathbf{p}^s) = \text{cli}(\mathbf{p}) + \text{cg}(\mathbf{p}^{ns}) + \text{csp}(\mathbf{p}^s)$. The best control strategy minimizes this cost function integrated over time. In order to leverage practical numerical techniques, we discretize the time axis uniformly with time step Δ to get $T_f = \frac{T}{\Delta}$ time steps and assume that our controls are constant over each time period Δ . We seek

$$\begin{aligned} & \min_{\mathbf{p}_{\mathcal{G}_s}^s(0), \dots, \mathbf{p}_{\mathcal{G}_s}^s(T_f-1)} \sum_{t=0}^{T_f} \text{cli}(\mathbf{p}) + \text{cg}(\mathbf{p}^{ns}) + \text{csp}(\mathbf{p}^s) \\ & \text{subject to} \\ & \mathbf{s}(t) = \sum_{\tau=0}^{t-1} \mathbf{p}^s(\tau)\Delta \\ & \mathbf{s}(T_f) = \mathbf{s}(0) \\ & \mathbf{p}_i^s(t) = 0 \quad i \notin \mathcal{G}_s \\ & \mathbf{p}(t) = \mathbf{p}^0 + \mathbf{p}^s(t) + \mathbf{p}^r(t) + \alpha \left(- \sum_{i \in \mathcal{G}_r} \mathbf{p}_i^r(t) - \sum_i \mathbf{p}_i^s(t) \right) \\ & \mathbf{p}^{ns}(t) = \mathbf{p}(t) - \mathbf{p}^s(t) \end{aligned}$$

Our problem falls within the well-studied class of deterministic, discrete-time optimal control problems and can be solved using standard nonlinear programming techniques [10]. Note that \mathbf{p}^0 and \mathbf{p}^r are known from the DCOPF and the perfect renewable forecasts, respectively. Thus, we can express $\mathbf{p}(t), \mathbf{p}^s(t), \mathbf{p}^{ns}(t)$ directly in terms of $\mathbf{p}_{\mathcal{G}_s}^s(t)$ using the second, third and fourth constraints above to get $\mathbf{p}(\mathbf{p}_{\mathcal{G}_s}^s), \mathbf{p}^s(\mathbf{p}_{\mathcal{G}_s}^s), \mathbf{p}^{ns}(\mathbf{p}_{\mathcal{G}_s}^s)$. Further we can write $\mathbf{p}_{\mathcal{G}_s}^s(t) = \frac{\mathbf{s}_{\mathcal{G}_s}(t+1) - \mathbf{s}_{\mathcal{G}_s}(t)}{\Delta}$, with the convention that $\mathbf{s}_{\mathcal{G}_s}(T_f) = \mathbf{s}_{\mathcal{G}_s}(0)$ (in order to satisfy the constraint (3)). This leaves us with the following unconstrained optimization problem:

$$\begin{aligned} & \min_{\mathbf{s}_{\mathcal{G}_s}(0), \dots, \mathbf{s}_{\mathcal{G}_s}(T_f-1)} \sum_{t=0}^{T_f} \text{cli} \left(\mathbf{p} \left(\frac{\mathbf{s}_{\mathcal{G}_s}(t+1) - \mathbf{s}_{\mathcal{G}_s}(t)}{\Delta} \right) \right) + \\ & \quad \text{cg} \left(\mathbf{p}^{ns} \left(\frac{\mathbf{s}_{\mathcal{G}_s}(t+1) - \mathbf{s}_{\mathcal{G}_s}(t)}{\Delta} \right) \right) + \\ & \quad \text{csp} \left(\left(\frac{\mathbf{s}_{\mathcal{G}_s}(t+1) - \mathbf{s}_{\mathcal{G}_s}(t)}{\Delta} \right) \right) \quad (5) \end{aligned}$$

It can be shown that this is a convex optimization problem [11] that can be solved efficiently using Newton's method (since the Hessian has a sparse block structure arising from the pairwise interaction). In practice, a

Levenberg-Marquardt correction [12] is required to ensure convergence as the Hessian can become numerically singular. We use software from [13] to solve the problem.

2.6 Optimal Sizing and Placement of Storage

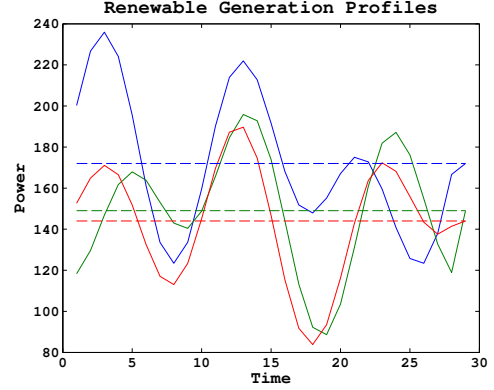


Figure 2: Three typical renewable generation profiles used in our study. The dashed lines are the time averages of the profiles.

We seek to develop heuristics to decide how to place storage and size its energy and power capacity. The high-level pseudocode given in algorithm 1, which uses the optimal control algorithm given above, is our initial attempt at this heuristic. Let \mathcal{G}_s denote the set of nodes with non-zero storage and $\text{perf}(\mathcal{G}_s)$ denote the inverse of the total storage power capacity used by the optimal control with access to unlimited storage only at nodes in \mathcal{G}_s . At every iteration of the outer loop, the set of active storage nodes \mathcal{G}_s is chosen. For each inner loop iteration, we create a trial where a random profile for renewable generation is created at each renewable node (details of how this is done are in Section 3 and typical examples of three profiles are shown in Fig. 2). For each trial of the renewable profile, the control is allowed access to an unlimited amount of storage power and energy at all the nodes in \mathcal{G}_s (and 0 storage at nodes not in \mathcal{G}_s), and we solve the control problem for the optimal storage power dispatch $\mathbf{p}_{\mathcal{G}_s}^{s*}(t)$. For each trial k and each node in \mathcal{G}_s , we find A^k , the maximum of the absolute value of the storage power at that node, which is a measure of the size of the storage power required at that node by the optimal storage dispatch for a given trial k . The values of A^k are averaged over many trials.

Next, the set of storage nodes \mathcal{G}_s is reduced by picking only those nodes $i \in \mathcal{G}_s$ such that $A_i \geq \gamma \max(A)$, i.e.

Algorithm 1 Optimal Sizing and Placement Algorithm

Choose thresholds $\varepsilon, \varepsilon'$
 $\mathcal{G}_s \leftarrow 1 : n$
repeat
 for $k = 1 \rightarrow N$ **do**
 Generate Random Time Series Profiles for the Renewables
 Solve (5) with storage only at nodes \mathcal{G}_s for the given profiles to get optimal dispatch $\mathbf{p}^{s*}(t)$
 $A^k \leftarrow \max_t |\mathbf{p}^{s*}(t)|$
 end for
 $A \leftarrow \frac{1}{N} \sum_{k=1}^N A^k$
 $\gamma \leftarrow \max\{\gamma : \{\text{perf}(\{i \in \mathcal{G}_s : A_i \geq \gamma \max(A)\}) < \text{perf}(\mathcal{G}_s) + \varepsilon\}\}$
 $\mathcal{G}_s \leftarrow \{i \in \mathcal{G}_s : A_i \geq \gamma \max(A)\}$
until $\gamma \leq \varepsilon'$

only the nodes that require, on average, a storage power capability within a fraction γ of the node with the maximum storage power capability. In essence, we are using our optimal control scheme to identify the nodes that exercise the greatest control over cost $(\mathbf{p}, \mathbf{p}^{ns}, \mathbf{p}^s)$. The parameter γ determines the severity of the “cut”. As γ is increased towards one, fewer and fewer nodes remain in \mathcal{G}_s . As our controllability is reduced, we might expect the level of constraint violations to increase. However, the violations seem to remain almost constant and we instead observe that at a critical value of γ , the total amount of storage power on the network, as required by the optimal control, suddenly increases. We set γ just below this threshold thereby selecting the reduced set of nodes that performs nearly as well as allowing the optimal control access to all of nodes.

Using the reduced set of nodes from above as our new starting point, the algorithm is repeated (i.e. we regenerate A and find the new maximal reduction in the number of nodes). This outer loop repeats until we are not able to reduce the number of nodes any farther without observing a significant increase in the total amount of storage power.

3 Simulations

We tested our optimal control and heuristic for storage placement and sizing on a modified versions of RTS-96[8]. The grid is shown in Fig. 3. Our modification includes the addition of three renewable generation nodes shown in blue. The capacities of the new lines connecting the renewables to their immediate neighbors are set higher than the capacity of the added renewable generation, otherwise, these lines would be overloaded in

Grid Structure




Figure 3: Our modified version of RTS-96. The added renewables are blue, loads are yellow and controllable generators are green.

nearly every trial.

3.1 Simulation Setup

We performed three sets of simulations (corresponding to outer iterations of algorithm 1), and for each setting, we perform $N = 2000$ trials. We define the renewable penetration value of a trial to be

$$\frac{\sum_{t=0}^{T_f} \sum_{i \in \mathcal{G}_r} \mathbf{p}_i^{ns}(t)}{\sum_{t=0}^{T_f} \sum_{i \in \mathcal{G}_l} \mathbf{p}_i^{ns}(t)}$$

The penetration is the fraction of the load that is served by the renewable generation. For each trial, a random value uniformly distributed between 0 and 50% is selected for the penetration level, and the mean values of the renewables are scaled so as to achieve these values. A zero-mean fluctuation is generated around the base value for each renewable:

$$\mathbf{p}_i^r(t) = \sum_k \frac{1}{k} \frac{\int_0^\Delta \sin(\omega_k(t\Delta + \tau) + \phi_k) d\tau}{\Delta}$$

where the ϕ_k are randomly chosen phases and ω_k are chosen such that the above sine waves are harmonics of the basic wave with time period T . Since, we discretize the time axis, we use $\frac{\int_0^\Delta \sin(\omega_k(t\Delta + \tau) + \phi_k) d\tau}{\Delta}$, the average power in the interval $t\Delta$ to $(t+1)\Delta$. The final profile is generated by scaling the disturbances so they are comparable to the mean generation (which is typical for turbulent fluctuations of wind) and adding it to

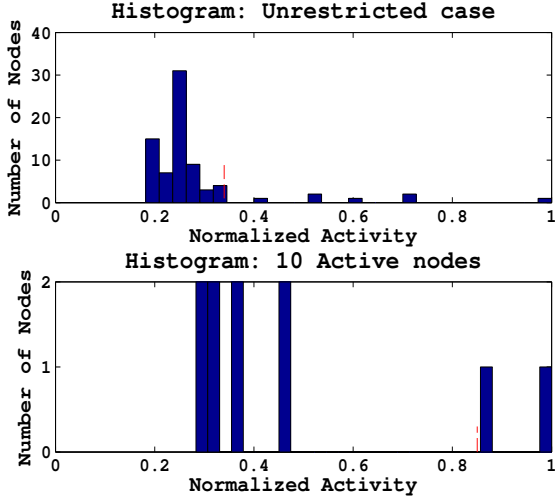


Figure 4: Activity Histograms: Red lines mark thresholds used for the reduction in the storage node set.

the mean: $\mathbf{p}_i^{ns}(t) = \mathbf{p}_i^0(t) * (1 + \mathbf{p}_i^r(t)), i \in \mathcal{G}_r$. Each instance of the renewable generation profiles we create is statistically independent and uncorrelated with every other instance. We do not intend for a sequence of these instances to approximate an actual time series of wind generation.

We start with the unrestricted case (where unlimited storage is available at all nodes) and compute the average activity of each node over all $N = 2000$ trials. The histogram of activities is shown in Fig. 4. The optimal γ for the first outer iteration of algorithm 1 is 0.34 and the resulting reduced set contains 10 nodes. When we tried to remove more nodes from the set, the average level of total network power and energy storage capacity showed a sudden increase indicating that further reductions based on this first stage of planning actually decreased the system performance.

We repeat this process, but this time starting with the 10-node set identified from the first iteration. Simulating another $N = 2000$ trials yields the second histogram in figure 4. When we restrict our optimal control to just these 10 nodes, their rank order changes relative to the rank order in the unrestricted case. It is this property that allows our heuristic to make further reductions in the number of nodes without significant impacts to overall performance.

The optimal γ for this second iteration is 0.85 leaving only two nodes for the third iteration. The algorithm terminates here as the controllable set cannot be shrunk any further without significance performance degradation. The 3 cuts (all nodes, 10 nodes, 2 nodes) obtained

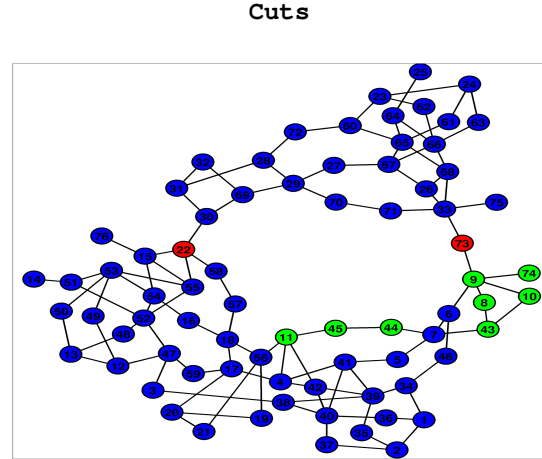


Figure 5: Our modified RTS-96 grid showing the three sets of nodes identified by our heuristic. The minimal set (two nodes) is shown in red. The ten node set includes the two red nodes and the eight additional green nodes. The maximal set includes all of the nodes.

by the algorithm are depicted on the grid in Fig. 5.

To study how the energy and power capacity requirements of storage depends on the level of renewable penetration, we first must quantify these measures in a consistent fashion. For a given renewable profile $\mathbf{p}^r(t)$, the optimal storage power dispatch is $\mathbf{p}^*(t)$, and $\mathbf{s}^*(t)$ is the resulting energy in storage. We define the energy in the renewable fluctuations to be $\mathbf{s}^r(t) = \sum_{\tau=0}^{t-1} \mathbf{p}^r(\tau)\Delta$, i.e. $\mathbf{s}^r(t)$ is the energy stored in a (hypothetical) battery that is connected directly to a renewable node and eliminates all fluctuations about the mean renewable genera-

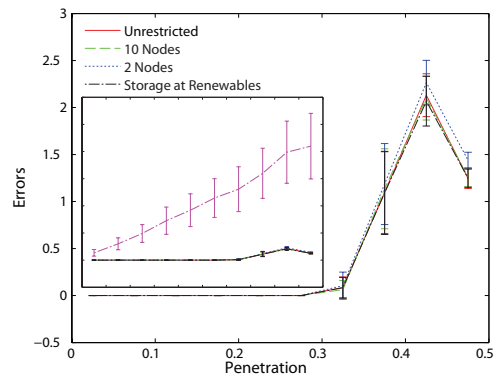


Figure 6: The level of constraint violations vs the penetration of renewable generation.

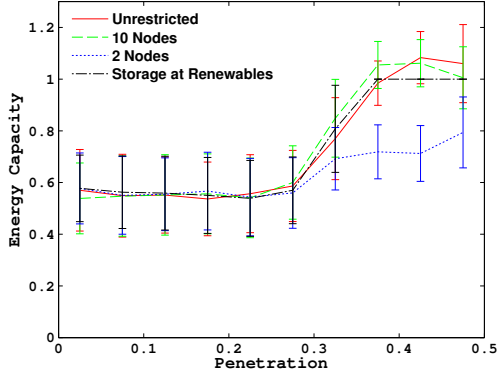


Figure 7: The normalized energy capacity of storage in the entire network vs the penetration of renewable generation.

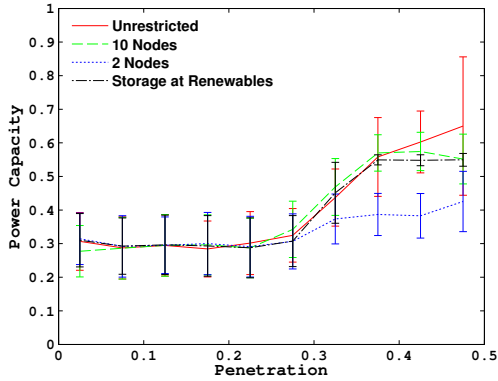


Figure 8: The normalized power capacity of storage in the entire network vs the penetration of renewable generation.

tion. We will only look at changes in energy stored over time, so $s^*(0)$ and $s^r(0)$ can be arbitrarily chosen. With these brief preliminaries, the parameters relevant to our study are:

Normalized Power Capacity: This quantifies the total power capacity of the storage relative to the sum of maximal power fluctuations over the renewables:

$$\frac{\sum_{j \in \mathcal{S}I} \max_t |\mathbf{p}_j^*(t)|}{\sum_{i \in \mathcal{G}_r} (\max_t \mathbf{p}_i^r(t) - \min_t \mathbf{p}_i^r(t))}$$

Normalized Energy Capacity: This quantifies the total energy capacity of the storage relative to the sum of maximal energy fluctuations over the renewables:

$$\frac{\sum_{j \in \mathcal{S}I} (\max_t s_j^*(t) - \min_t s_j^*(t))}{\sum_{i \in \mathcal{G}_r} (\max_t s_i^r(t) - \min_t s_i^r(t))}$$

Constraint Violations: This is defined as the sum of all constraint violations (as defined in (2)) averaged over time.

These quantities are all defined for a given renewable configuration. In order to investigate their typical behavior as a function of renewable penetration, we create a 10 equally-spaced intervals between 0 and .5. For each bin I , we compute statistics (mean and standard deviation) of all three quantities over trials with a penetration value in the bin I . The means and standard deviations for the three quantities are shown as functions of penetration in Figs. 8,7,6.

4 Interpretations

First, we note that Fig. 6 demonstrates the effectiveness of node-specific control for mitigating the impact of renewable generation fluctuations on the network. The upper curve in the inset in Fig.6 shows the level of constraint violations for a system that maintains generation-load balance purely via proportional control, i.e. only the controlled generators responding to a generation-load imbalance in proportion to their capacity. Without the node-specific control in our model of energy storage, the response of the generators is inflexible, and system begins to experience constraint violations for small values of renewable penetration. With the node-specific control allowed by our storage model, Figure 6 shows that the control can eliminate all constraint violations up to about 30% renewable penetration for all the node configurations we explored. We do not attribute this behavior to the storage itself, rather, this appears to be a property of the flexibility of node-specific control. Although we did not explore this possibility, we believe that if we would have incorporated the controlled generators \mathcal{G}_g into our control algorithm, we would have seen similar performance up to a renewable penetration of 30% for a system without any storage. It is interesting to note that more detailed studies have found similar thresholds for renewable penetration[4].

Figure 5 shows the location of the nodes with storage in the reduced-node sets identified by our heuristic in algorithm 1. The two-node set is perhaps the most illustrative of how optimal placement and sizing of storage in a network does yield itself to simple rules of thumb or intuition. Intuition may have led to the conclusion that storage should be placed at the site of the fluctuating renewables because it can effectively mitigate the fluctuating power flows without these flows ever being injected into the network. It seems plausible that this should lead to a minimal amount of storage and quality

control, however, the results in Fig. 5 and 6 show otherwise. The nodes in the two-node set do not include either of renewable generation sites. In fact, even when expanded to the ten-node set, only one of the three renewable sites is included. These results demonstrate that the nodes that have a high degree of control over congestion caused by renewable fluctuations are not necessarily the renewable nodes themselves.

To further demonstrate this point, we performed an additional simulation where the reduced-note set included just the three renewable sites. Figure 6 shows that the quality of control over the constraint violations is the same for this handpicked case as in the two-, ten-, and unrestricted-node sets. Figures 7 and 8 help to distinguish between these node sets. For renewable penetrations below about 30%, the network total power capacity and energy capacity of the storage is nearly the same for each case. Beyond 30% where we begin to see a few constraint violations in Fig. 6, the two-node set shows significantly lower total network power and energy capacity compared to the ten or unrestricted-node cases. The handpicked case is comparable to the ten-node case. A lower network total capacity implies that the nodes identified in the two-node case have the highest degree of controllability over the network congestion caused by the fluctuating renewables.

A possible reason for the two-node set showing better controllability is evident from Fig. 5. These two nodes sit at the end or the middle of crucial transmission lines that link two major regions in RTS-96. In reference to Fig. 5, we name these regions “upper” and “lower”. We conjecture that by controlling the power injections at these nodes, our control scheme can effectively control to which region the renewable generation is directed. Therefore, our control can direct the renewable fluctuations to the region that, at that point in time, has the greatest ability to assist in mitigation. In addition to simply absorbing the fluctuations locally, the storage may be used to redirect the renewable fluctuations throughout the network to leverage other resources to assist in the control.

5 Conclusions and Future Work

In this initial study, we have developed and demonstrated a method for coupling operations with system expansion planning for the optimal placement and sizing of storage in a grid with a significant penetration of time-intermittent renewable generation. Operations are incorporated into the planning via simulations of an optimal control scheme that uses perfect renewable generation forecasts to dispatch energy storage to eliminate

violation of network and generation constraints. By simulating the system on a relatively short time scale, we build up significant statistics on the storage activity at all network nodes. We have developed a heuristic that uses these statistics to reduce the number of storage-enabled network nodes while maintaining or improving system performance.

Somewhat unexpectedly, our method does not preferentially place energy storage at the nodes with renewable generation. Instead, our method apparently favors nodes at critical junctions between major subcomponents of the network. We have proposed that these nodes provide for enhanced controllability because, in addition to simply buffering the fluctuations of the renewables, controlled power injections at these nodes can modify overall network flows and direct fluctuating power flows to regions that are better positioned to mitigate them.

There is much follow on work needed to expand the concept presented in this manuscript and to verify some of its conjectures. A few examples include:

- Additional networks should be considered, including different configurations of renewable generation on the present network, to investigate whether our node-reduction heuristic is robust.
- The explanation of the placement of storage proposed in this work should be verified by correlating the response of the storage with nearby power flows.
- The operational simulation should be made more realistic by incorporating ramping constraints on the controllable generation.
- The method should be extended to consider alternatives to proportional control for changing the outputs of the controllable generators.

6 Acknowledgments

We are thankful to the participants of the Optimization and Control for Smart Grids LARD DR project at Los Alamos and Smart Grid Seminar Series at CNLS/LANL for multiple fruitful discussions. The work at LANL was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396.

References

- [1] Madeleine Gibescu, Arno J. Brand, and Wil L. Kling. Estimation of variability and predictability of large-scale

- wind energy in the netherlands. *Wind Energy*, 12(3):241–260, 2009.
- [2] Andrew Mills and Ryan Wiser. Implications of wide-area geographic diversity for short-term variability of solar power. Technical report, LBNL-3884E, 2010.
- [3] Eric Hirst and Brendan Kirby. Separating and measuring the regulation and load-following ancillary services. *Utilities Policy*, 8(2):75–81, June 1999.
- [4] P. Meibom, R. Barth, B. Hasche, H. Brand, C. Weber, and M. O’Malley. Stochastic optimization model to study the operational impacts of high wind penetrations in ireland. *Power Systems, IEEE Transactions on*, PP(99):1–12, 2010.
- [5] F. Bouffard and F.D. Galiana. Stochastic security for operations planning with significant wind power generation. In *Power and Energy Society General Meeting - Conversion and Delivery of Electrical Energy in the 21st Century, 2008 IEEE*, pages 1–11, july 2008.
- [6] E. Hirst. Integrating wind energy with the bpa power system: Preliminary study. Technical report, Consulting in Electric-Industry Restructuring, 2002.
- [7] A. Tuohy, E. Denny, and M. O’Malley. Rolling unit commitment for systems with significant installed wind capacity. In *Power Tech, 2007 IEEE Lausanne*, pages 1380–1385, july 2007.
- [8] C. Grigg, P. Wong, P. Albrecht, R. Allan, M. Bhavaraju, R. Billinton, Q. Chen, C. Fong, S. Haddad, S. Kuruganty, W. Li, R. Mukerji, D. Patton, N. Rau, D. Reppen, A. Schneider, M. Shahidehpour, and C. Singh. The iee reliability test system-1996. a report prepared by the reliability test system task force of the application of probability methods subcommittee. *Power Systems, IEEE Transactions on*, 14(3):1010–1020, aug 1999.
- [9] M. Chertkov, M. Stepanov, F. Pan, and R. Baldick. Exact and efficient algorithm to discover stochastic contingencies in wind generation over transmission power grids. 2011. Accepted for invited session on Smart Grid Integration of Renewable Energy: Failure analysis, Microgrids, and Estimation at CDC/ECC 2011. [www.arxiv:1104.0183](http://www.arxiv.org/abs/1104.0183).
- [10] J.T. Betts. *Practical methods for optimal control using nonlinear programming*. Number 3. Society for Industrial Mathematics, 2001.
- [11] S.P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge Univ Pr, 2004.
- [12] J. Nocedal and S.J. Wright. *Numerical optimization*. Springer verlag, 1999.
- [13] M. Schmidt. Matlab *minfunc*.