Sizing Solar Panels and Storage for Multiple Roofs

Anonymous Author(s)

ABSTRACT

Choosing the number of solar panels and the amount of storage needed to meet a certain fraction of the load in a microgrid setting is a difficult problem that needs to balance the competing objectives of efficiency, robustness, and cost. Prior work in this area makes the unrealistic assumption that solar panels are to be installed on a single roof that is capable of supporting all the panels required. In reality, we may need to deploy solar panels on several roof segments, each of limited size, and each with its own tilt, orientation and installation cost. This paper presents an algorithm for sizing solar panels and storage in this context. We evaluate the robustness of our approach using traces derived from the Pecan Street Dataport dataset and demonstrate the value of our approach by using it to size a hypothetical installation on the British Antarctic Survey's research base in Antarctica.

ACM Reference Format:
Anonymous Author(s). 2021. Sizing Solar Panels and Storage for Multiple Roofs. In *Proceedings of* . ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 INTRODUCTION

With the rapid decline in the cost of both solar photovoltaic (PV) generation and storage, solar and storage systems are increasingly adopted to provide carbon-free renewable energy throughout the day [11]. However, PV systems and storage are still quite expensive in absolute terms, so it is necessary to find the smallest possible size of the system that meets electricity load needs [14]. Specifically, given a *solar profile*, the typical hourly solar generation from a unit solar panel, and the *load profile*, typically characterized as hourly demand for a year or more, we would like to choose the number of solar panels and the amount of storage (a *sizing*) so that the load profile is met with a certain quality of service (QoS). We call a sizing that meets the desired QoS criterion a *feasible* sizing.

Choosing a feasible sizing is complex because it needs to balance multiple objectives, including:

- *Pareto efficiency*: It should not be possible to simultaneously reduce both the number of solar panels and the size of the storage without violating the QoS criterion.
- *Robustness*: Feasibility should not be violated due to small variations in solar energy and load in the future. Note that robustness and efficiency cannot be simultaneously satisfied.

, ,

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM…$15.00
https://doi.org/10.1145/nnnnnnn.nnnnnnn

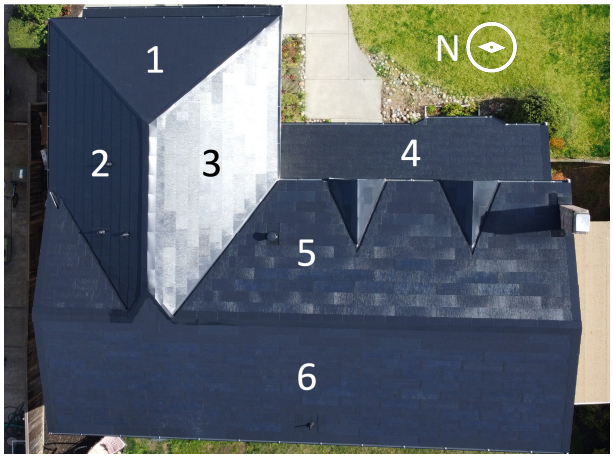

Figure 1: A rooftop with six roof segments.

- *Cost*: The cost should be as low as possible, taking into account the cost of panels and storage, the installation cost, and the need for robustness.

This important and practical problem has been studied in the literature using approaches ranging from optimization and stochastic network calculus to simulation [14]. However, prior work makes the simplifying assumption that solar panels are to be installed on a single site, typically a rooftop, and that the rooftop is capable of supporting as many solar panels as desired. Moreover, all portions of the roof are assumed to have the same tilt and orientation. In reality, we may wish to deploy solar PV on several roof segments[1], as in Figure 1, each of limited size, and each with its own tilt, orientation and installation cost. While this may appear to be a trivial change, the problem of sizing solar and storage, even for a system with only two roof segments, is *far* more complex.

Intuitively, the reason for this complexity is that some roof segments may have a higher potential for solar energy production or may better match the load profile, but also have higher installation costs or smaller panel capacity. In contrast, others may have lower installation costs and larger capacity, but have a poor solar production potential. Choosing the number of panels to place on each roof segment to produce sufficient energy to meet the QoS criterion at the least overall cost is a complex problem and the focus of our work. Specifically, our contributions are:

- We formally state the problem of sizing solar PV and storage on a set of roof segments with different orientations and installation costs and present an algorithm to solve it.
- We implement the algorithm and evaluate its correctness on realistic solar and load traces.
- We use our approach to size a hypothetical solar PV and storage for the British Antarctic Survey's research base in Antarctica, based on real data.

[1]For uniformity of notation, we use 'roof segment' to mean either the entire roof or a part of it, on one building or several.

2 RELATED WORK

A person looking to install a rooftop solar PV and battery system wants to know the least-cost sizing that meets their system performance target, e.g. meeting 50% of their load. The optimal sizing depends on the location, tilt, and orientation of each roof segment, as well as the building's load profile and the fixed and marginal costs of system components and installation. In practice, the sizing of solar PV and battery systems has often been done using a "rule-of-thumb"; for example, Tesla [2, 3] takes into account a location, roof layout, and a proxy for load, such as the monthly electricity bill, to calculate a system sizing. Such methods are imprecise, since they rely on aggregate PV generation and load metrics.

In the literature, the sizing of PV and storage systems has been studied in several contexts, including micro-grids [7, 8, 10] and building-scale systems [14, 18, 19]. At a high level, the strategy used by existing work is to assume the availability of load and PV generation measurements for one or more years, which in recent years have become more readily available for consumers [4, 16], and use them to compute a sizing that would be optimal over the given data. Recent work takes into account the stochastic nature of PV generation and load profiles, and proposes a method based for computing the cheapest *robust* PV panel and battery sizes that meet target performance with a specified level of statistical confidence; this work relies on the assumption that solar panels are installed on a single roof segment, and cannot be applied to multi-roof problems [13, 14]. In this paper, we look at computing a robust sizing that considers multiple roof segments.

3 PROBLEM FORMULATION

This section presents our model for the sizing problem. We assume that a user, such as a prospective system owner, can obtain (a) a set of solar traces corresponding to the generation from a single panel placed on each roof segment and (b) a representative load trace. We also assume that the solar panels are connected to the same storage unit and that there is no loss of power on these connections. Our goal is to find a feasible, robust sizing that is Pareto efficient, and therefore has the least cost. Table 1 shows the notation used in the remainder of the paper. Using this notation, we can make this goal more precise, as follows.

Goal: Given $N, T, \varepsilon, c_i^f, c_i^m, c^b, a_i^{max}, b^{max}, N$ solar traces $\mathcal{S}_i$, and the load trace $\mathcal{L}$, find $\mathcal{A}$ and b such that

- It is feasible, i.e. for all time t, with probability greater than 1-γ

$$\frac{u(\mathcal{A}, b, t, t+T)}{|\mathcal{L}[t, t+T]|} \leq \varepsilon \tag{1}$$

- It minimizes the cost function

$$C(\mathcal{A}, b) = \sum_{i|a_i \neq 0} (c_i^f + c_i^m * a_i) + (c^b * b) \tag{2}$$

where the bound is *robust*, that is, holds for future loads that are statistically the same as in the past. Note that, for any non-zero load profile, a sizing with a sufficiently small number of solar panels is always infeasible.

Symbol	Meaning (units)
N	Number of roof segments
i	Index of roof segment
T	Time period over which the which QoS criterion is computed (days)
γ	confidence level for robustness
ε	the target upper bound on unmet energy as a fraction of overall load $\|\mathcal{L}\|$
c_i^f	Cost of installing a panel on the i^{th} roof segment ($)
c_i^m	Marginal cost of installing a panel on the i^{th} roof segment ($)
c^b	Marginal cost of storage ($/kWh)
a_i^{max}	Largest possible number of panels on the i^{th} roof segment
b^{max}	Largest possible battery size (kWh)
n	Total number of hours provided in each solar/load trace
$\mathcal{S}_i$	An hourly solar power generation trace for one panel on the i^{th} roof segment; $\|\mathcal{S}_i\| = n$
$\mathcal{S}$	A vector of hourly solar power generation trace for panels on each roof segment; $\mathcal{S} = \{\mathcal{S}_1, \ldots, \mathcal{S}_N\}$
$\mathcal{S}_i[t_1, t_2]$	A subset of the ith solar trace from time step t_1 to time step t_2 (kW)
$\|\mathcal{S}_i[t_1, t_2]\|$	Total energy produced by panels allocated in the i^{th} segment from time step t_1 to time step t_2 (kWh)
$\mathcal{L}$	An hourly load trace (kW); $\|\mathcal{L}\| = n$
$\mathcal{L}[t_1, t_2]$	A subset of the load trace from time step t_1 to time step t_2 (kW)
$\|\mathcal{L}[t_1, t_2]\|$	Total energy used in the load trace from time step t_1 to time step t_2 (kWh)
a_i	Number of panels, as computed by the algorithm, allocated to i^{th} roof segment; $0 \leq a_i \leq \mathcal{A}_i^{max}$
$\mathcal{A}$	The sizing, a vector of allocations $a_1, a_2, ..., a_N$
b	Size of battery, as computed by the algorithm; $0 \leq b \leq b^{max}$ (kWh)
u	Unmet energy for a given allocation vector, battery size, load trace, and solar trace (kWh)

Table 1: Table of notation. Note that $\mathcal{A}, b$, and u are outputs of the algorithm and depend on the trace pair $(\mathcal{S}_i, \mathcal{L})$, but, for clarity of notation, this dependency is not explicitly denoted.

4 SOLUTION APPROACH

Our solution extends prior work by Kazhamiaka *et. al.* [14]. The key idea in their work is to simulate the behaviour of an ensemble of historical load and solar traces to construct a *distribution* of feasible and efficient sizings; this distribution is analyzed to identify the sizing with the lowest cost that will meet the QoS with a given confidence.

Specifically, their sizing algorithm first computes the *feasible Pareto frontier* of solar and storage (a, b) sizing tuples corresponding to each solar and load trace pair. By definition, decreasing the number of panels in a sizing on this frontier necessarily increases the storage capacity and *vice versa*. To compute the first point on

the frontier, they start at the maximum number of panels and use simulations to compute the amount of unmet load as they progressively increase the battery size from 0. This finds the minimum battery size needed to ensure that the QoS target is met (so that the sizing is both feasible and efficient). They then reduce the amount of solar generation by removing one panel and recompute the minimum battery size to find the next point on the frontier.

By following this process, they compute the minimum amount of storage required for each choice of number of panels to meet the QoS target, for each solar/load trace pair. Given a set of such curves, they compute a Chebyshev bound on the set of Pareto frontiers. The least-cost point on this bound is a feasible, robust sizing that also meets the QoS target (it is not Pareto efficient, but no Pareto efficient sizing is robust, by definition).

4.1 Complications

We cannot directly use the algorithm from Kazhamiaka *et. al.* [14] because at each step of their sizing algorithm, they compute the next step in the frontier after reducing the number of panels by one. However, when there are N roof segments, it is not obvious which roof segment to select for reduction, since the removal of a panel from different roof segments impacts the net cost differently.

To fix ideas, consider a site with two roof segments, labelled 1 and 2. Each segment has its own fixed and marginal costs for installing a panel, c_i^f and c_i^m. Moreover, each has its own energy generation capacity because of differences in tilt and orientation. Removing a panel on one of the roof segments decreases the marginal installation cost, but potentially leads to an infeasible sizing, depending on the shape of the load profile, complicating the choice.

4.1.1 A greedy approach is sub-optimal. Can we use a greedy approach, that is reducing the number of panels by one on the roof segment that maximally reduces the overall cost? Unfortunately, the following counterexample proves that a greedy approach is sub-optimal.

Consider two roof segments A and B with identical solar production. Let the installation cost of p panels on A be 0, if p=0, and 1+2p otherwise and on B be 0 if p=0, and 5+p otherwise. At initialization time, the sizing algorithm allocates the maximum number of panels to each roof segment, i.e. 3 panels on A and 2 on B (denoted (3,2)) for total cost of 1+2*3 + 5+2*1 = 14 units (Table 2). The choices taken by a greedy algorithm, as shown, would result in a least cost allocation of 10. However, if we were to reduce B by two panels, with an allocation of (3, 0), the cost would only be 7 units, so the greedy approach is not optimal.

We can avoid having to consider both fixed and marginal costs by amortizing the fixed cost to define c_i^a the amortized cost of installing one panel on roof segment i as

$$c_i^a \triangleq c_i^m + \frac{c_i^f}{s_i} \tag{3}$$

thus reducing fixed costs to zero. However, as we show next, this does not solve the problem.

4.1.2 The impact of unmet load. We now show, using a counterexample that, even with zero fixed costs, a greedy approach is not optimal.

	Greedy		Optimal	
	Allocation	Cost	Allocation	Cost
Initialize	(3,2)	14	(3,2)	14
Step 1	(2,2)	12	(3,1)	13
Step 2	(1,2)	10	(3,0)	7

Table 2: Counterexample showing that a greedy choice is suboptimal.

Consider, again, a system with two roof segments. Given an arbitrary load profile $\mathcal{L}$, let the solar generation from one panel on roof segment 1 be $\mathcal{L} + 0.25 + \epsilon$ and from one panel on roof segment 2 be $\mathcal{L} + 0.50 + \epsilon$. Also, assume that each roof segment only supports 2 panels. Finally, let the amortized cost per panel for the first roof segment be 10 units and the amortized cost for the second roof segment be 100 units. Thus, if we wanted to maximally reduce installation costs, we would choose to reduce the number of panels on the second roof segment.

When we start with the maximum allocation, which is two panels on both roof segments, denoted (2,2), the total production is $\mathcal{L} + 1.5 + 4\epsilon$, so there is no unmet load and no need for a battery, so the only costs are the PV costs, in total 2*10 + 2*100 = \$220.

Now, to reduce the costs the most, we should reduce the number of panels on the second roof segment by 1, leading to the allocation (2,1). Again this has a total solar production at time t of $\mathcal{L} + 1 + 3\epsilon$, so there would be no unmet load, and no battery, and the total cost would be 2*10 + 1*100 = \$120. If we now wish to reduce the number of panels by one again, the way to reduce installation costs would be to remove the second panel on the second roof segment, leading to the allocation (2,0). But this will lead to an unmet load of approximately 0.5 units during each time step, i.e., infeasibility. Instead, we are better off with the allocation (0,2), that is, removing both panels from the first roof segment and using two on the second, since the cost here is \$200 and there is no unmet load. But this allocation (0,2) cannot be obtained from the prior greedy allocation (2,1). Hence, in this example, the greedy approach doesn't work.

Taken together, the two counterexamples show that a greedy choice in reducing the number of panels is sub-optimal. However, there is no other obvious way to make this choice. Thus, as described in the next section, we use a stochastic gradient descent approach to adjust the allocation of roof segments.

5 SIZING ALGORITHM

This section provides a detailed description, pseudocode, and visualization of the algorithm, including the generation of solar and load traces, the system simulation and stochastic gradient descent process, the statistical bound, and how to extract a robust sizing.

5.1 Algorithm Overview

Before diving into the technical details of our algorithm, we give an overview of our solution with reference to later subsections. Additional notation used in this section can be found in Table 3.

At a high level, the process of computing a sizing consists of the following three steps:

(1) Acquire PV generation traces for each roof segment and load traces covering a span of at least one year and with

Symbol	Meaning (units)
η	Number of subintervals sampled, one of each a least-cost sizing is found. Calculated at the beginning of the algorithm as detailed in Section 5.5.
ζ	AdaDelta hyperparameter. Fudge factor for numerical stability. Default 0.1. Refer to Algorithm 2.
ρ	AdaDelta hyperparameter. Decay rate used in calculating exponentially moving average for gradients. Default 0.9. Refer to Algorithm 2.
ξ	A PV-battery sizing pair $(\mathcal{A}, b)$.
μ_η	The empirical mean of all η least-cost sizings.
Σ_η	The empirical covariance of all η least-cost sizings.
Ξ	A set of Chebyshev upper-bounded sizings, found in Section 5.4.
ξ^*	The output of the algorithm, the minimum-cost sizing on boundary Ξ.
Λ^2	The "practicality sizing" factor. Determines the distance between ξ^* and μ_η.
β	Hyperparameter to strike a balance between η and Λ^2. Default 0.1. For details see Section 5.5.

Table 3: Table of additional notation for the sizing algorithm.

a granularity of at most one hour. These can be from on-site measurements or synthetically generated. Sample η intervals from these traces to form an ensemble of (possibly overlapping) PV generation and load *trace tuples* $(\mathcal{S}_i\ , \mathcal{L})$.

(2) For each trace tuple, find the sizing that meets the target performance criterion; this results in a set of sizings.

(3) Compute a statistical bound on the set of sizings, and select a sizing from those found along the bound.

The number of trace tuples, η, represents a trade-off between algorithm runtime and the tightness of the statistical bound on the sizing. The runtime scales linearly with η, and the bound gets asymptotically tighter with more samples; the precise relationship is described in Section 5.5.

Our approach to finding the least-cost sizing in the search space is to adapt a stochastic gradient descent algorithm. We focus our search on the Pareto-optimal sizing subspace, which consists of sizings that have two properties: 1) the sizing is feasible, i.e., satisfies the unmet load target, and 2) no sizing that is smaller in every dimension is feasible. The unmet load of a sizing is calculated by simulating the operation of the system over the given data. The cost function for computing the gradient to guide the direction of the search is the total system cost. The simulation and stochastic gradient descent process are detailed in Section 5.3.

We find the least-cost sizing for every data sub-sample, creating a set of such sizings. We compute the statistical bound on the distribution of sizings represented by the set. As in Reference [13], an empirical multivariate Chebyshev bound [20] is used to calculate a hyper-ellipsoid bound that is centered at the mean the sizings set and scaled according to the desired confidence γ. Finally, we generate a set of sizings that lie on the upper portion of the bound and select the lowest-cost sizing from this set as our output. The process of finding the Chebyshev bound is detailed in Section 5.4.

5.2 Obtaining Solar and Load Traces

The algorithm relies on hourly PV generation (one per roof segment) and load traces $(\mathcal{S}_i, \mathcal{L})$. The user may have access to load data if they own a smart electricity meter, although it is possible to obtain both load and per-roof-segment PV generation traces using proxy techniques, which we describe next.

5.2.1 Solar Traces. In the United States, hourly solar traces can be generated through tools such as NREL's PVWatts calculator [6]. The calculator requires parameters that are more easily obtained, including the geographical location of the roof segments, and, for each roof segment, its tilt from level, orientation from true north, and an estimated performance loss percentage due to soiling/snow, wiring and connection losses, and inverter inefficiency, typically 14% of the total solar radiation received. It then calculates the solar traces based on previously stored solar radiation data measured near that location in a typical year.

Alternatively, if hourly horizontal solar irradiance is measured through a pyranometer, then in-plane irradiance can be calculated using the equations in [9] using a roof segment's location, tilt, orientation, and the estimated performance loss percentage.

5.2.2 Load Traces. Generating synthetic load traces is more challenging than generating solar traces as electricity usage pattern, unlike solar activity, depends on less-predictable human behaviours. Recent work has shown that an ARMA model for generating load traces, when trained on load patterns from neighbouring houses where data is available, generates traces that can be used for sizing [21]. Using a load profile database like the one provided by EERE [1] its often possible to find a load dataset that closely matches the monthly aggregate load values of the target site, which are typically available.

5.3 Finding Minimum-Cost Sizings

Given the PV generation and load traces for several years, we sample η sub-intervals from each trace using a sliding window approach to ensure equal representation in all years and seasons. This creates an ensemble of shorter PV generation and load traces, with the length of each trace T as specified by the user. For each PV-load pair in the ensemble, we want to compute a sizing that meets the target performance at a minimum cost. A naive way to do this is to exhaustively try all sizing combinations and simulate the operation of each system to check if it meets the target performance for the given trace pair. As in [13], system operation can be simulated to obtain the unmet load, denoted as $u(\mathcal{A}, b, t_1, t_2)$ over the time interval $[t_1, t_2]$, given the corresponding solar and load traces and solar/load allocation $(\mathcal{A}, b)$. However, the space of all possible sizings is very large and makes an exhaustive search impractical.

A more refined search algorithm, such as stochastic gradient descent, could be used to quickly find the minimum-cost sizing if the search space is convex and differentiable, which does not hold at the edges where the $\mathcal{A}_i$ goes from 0 to 1 due to fixed per-roof-segment installation costs. To get around this problem, we split the search space into several convex search spaces where fixed costs are ignored: one for each combination of roof segments. For example, if the search space has three roof segments A, B, C, then we have seven sub-spaces: A, B, C, AB, AC, BC, ABC. We then use

a stochastic gradient descent algorithm to efficiently search each space and find least-cost sizing across all of search spaces taking into account fixed costs *post hoc*. This mechanism, along with selecting the sub-intervals, are demostrated later in Algorithm 4 under Section 5.6.

Stochastic gradient descent requires a starting point, i.e., a sizing on the Pareto frontier, and a cost function. To find an initial Pareto-optimal sizing, we maximize the number of panels on each roof segment and calculate the minimum storage size b^* that leads to a feasible sizing through binary search. The cost function is simply the cost of the system: $C(\mathcal{A}^*) = C(\mathcal{A}^*, b^*)$ via Equation (2). The pseudocode for generating a cost for such a solar allocation $\mathcal{A}^*$ is specified in Algorithm 1.

Algorithm 1 Cost Function for a Solar Allocation

function C($\mathcal{A}$)
  $b^* \leftarrow min_{0\le b\le b^{max}}\ s.t.\ \frac{u(\mathcal{A},b,t_1,t_2)}{|\mathcal{L}[t_1,t_2]|} \le \varepsilon$, via binary search
  if b^* does not exist **then**
    return ∞
  end if
  return $C(\mathcal{A}, b^*) = \sum_{i|a_i\neq 0}(c_i^f + c_i^m * a_i) + (c^b * b^*)$
end function

Given the cost function for a solar allocation, we then adjust $\mathcal{A}^*$ via an iteration of a stochastic gradient descent algorithm, with cost function C and a finite difference approximation of gradient, i.e.

$$\frac{\partial C}{\partial a_i}\mathcal{A}^* \approx C(\{a_1, \ldots, a_i + 1, \ldots, a_n\}) - C(\mathcal{A}^*) \tag{4}$$

In our implementation, we experimented with several different stochastic gradient descent algorithms and settled on AdaDelta [23]. Its exponentially decaying gradient mechanism eliminates the need of setting an initial learning rate and we empirically found that it works consistently well with different orders of magnitudes of solar/load traces and search spaces. Other algorithms such as RMSProp [22] and AdaGrad [12] work equally well but require manual adjustment of learning rate to be effective. In our implementation, we set AdaDelta's two hyperparameters, ρ and ζ, to 0.9 and 0.1 respectively. The first hyperparameter represents the decay rate used to calculate an exponentially decaying running average for gradients and the objective function, and the second is used in a division to maintain numerical stability. In each iteration adjusting $\mathcal{A}^*$, we add a multivariate normal random perturbation of zero mean and one unit of variance to avoid search paths being trapped in local minima or saddle points. The stopping condition of the algorithm is when the cost C($\mathcal{A}^*$) exceeds m, the decaying average of the cost function for previous iterations. The pseudocode for this process is described in Algorithm 2.

Figure 2 shows a typical search path of AdaDelta in two dimensions of PVs. The background color gradient shows the cost at each PV sizing computed via grid search. The cost value starts high near $\mathcal{A}^{max}$ and gradually decreases as allocations from both roof segments decrease, and the gradient stays relatively constant. However, at pv1=3 the cost starts to increase with a unit decrease of pv1, and the search path changes direction to trade-off fewer allocations of pv2 and replace it with more allocations of pv1, keeping

Algorithm 2 Find the least-cost sizing through system simulation and stochastic gradient descent

function FIND_SIZING(t_1, t_2, subsets)
  $\mathcal{A}^* \leftarrow \{\mathcal{A}_i^{max} : i \in \text{subsets}\}$
  search_path $\leftarrow$ empty list
  $S \leftarrow \vec{0}$ ▷ S, Δ are AdaDelta intermediate variables
  $\Delta \leftarrow \vec{0}$
  $m \leftarrow$ C($\mathcal{A}^*$)
  while C($\mathcal{A}^*$) $\le m$ **do**
    search_path.add($\mathcal{A}^*$)
    $\nabla C(\mathcal{A}^*) \leftarrow \{\frac{\partial C}{\partial a_i}\mathcal{A}^* : i \in \text{subsets}\}$
      ▷ Approximation by equation (4)
    $S \leftarrow \rho S + (1-\rho)\nabla C(\mathcal{A}^*)^2$
    $G \leftarrow \frac{\sqrt{\Delta}+\zeta}{\sqrt{S}+\zeta} \circ \nabla C(\mathcal{A}^*)$ ▷ Element-wise product
    $\mathcal{A}^* \leftarrow \mathcal{A}^* - G + N(0,1)$
    $\Delta \leftarrow \rho\Delta + (1-\rho)G^2$
    $m \leftarrow \rho m + (1-\rho)$ C($\mathcal{A}^*$)
  end while
  return $arg\ min_{\mathcal{A}\in\text{search\_path}}$ C($\mathcal{A}$) $\cup \{0 : i \notin \text{subsets}\}$
end function

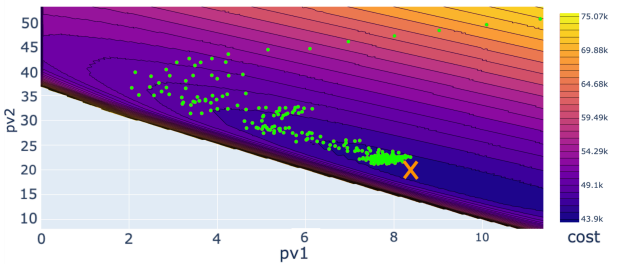

Figure 2: Search path for a typical AdaDelta run. The X marks the optimal (min-cost) sizing.

in mind cost efficiency. Ultimately the search terminates near the true minimum in the search space, denoted by the X mark.

5.4 Chebyshev Bound

As in Reference [13], we rely on a multivariate concentration bound, based on Theorem 1 in Reference [20] to find a robust sizing by treating the set of points on the Pareto frontier as i.i.d. samples from an unknown distribution. The bound is parameterized by the covariance matrix of the set of sizings, the mean of each dimension, and the desired confidence γ.

We denote each of the η sizing pairs $\mathcal{A}^{[t,t+T]}, b^{[t,t+T]}$ as an $N+1$-dimensional sizing $\xi^{(k)} \triangleq (a_1^{[t,t+T]}, a_2^{[t,t+T]}, \ldots, a_N^{[t,t+T]}, b^{[t,t+T]})$. Then, the *unbiased empirical covariance* Σ_η is defined by

$$\Sigma_\eta \triangleq \frac{1}{\eta-1}\sum_{k=1}^{\eta}(\xi^{(k)} - \mu_\eta)(\xi^{(k)} - \mu_\eta)^{\mathrm{T}} \tag{5}$$

where

$$\mu_\eta \triangleq \frac{1}{\eta}\sum_{k=1}^{\eta}\xi^{(k)} \tag{6}$$

is the empirical mean. Note that Σ_η is required to be non-singular (invertible) for subsequent computations. When it is not, it implies that there is no variability in one or more dimension i for all least-cost sizings simulated. In practice, this can happen when a roof segment is not used in any of the min-cost sizings, hence $\xi_i^{(k)} = 0\,\forall k$. Alternatively, it can be that the number of PV panels on roof segment i are always maximized, i.e., $\xi_i^{(k)} = a_i^{max}\,\forall k$. In either case, the matrix becomes invertible if we remove the i^{th} dimension from subsequent calculation (the i^{th} row and column of Σ_η) and lock in $\mu_{\eta,i}$ as final sizing for that dimension. We can repeat this operation until Σ_η is non-singular.

The multivariate Chebyshev bound is expressed as an $(N+1)$ dimensional hyper-ellipsoid of sizings Ξ that bounds the γ probability density mass of the empirical distribution. Since we are interested in the upper-bound of each dimension, we also specify that any point in Ξ must be non-dominated by other points inside the hyper-ellipsoid, i.e. if $\xi \in \Xi$, then no sizings inside the hyper-ellipsoid (which may not satisfy our robustness requirement) can be strictly larger than ξ. We denote this in equation (8).

To define Ξ, we first denote

$$L(\xi) \triangleq (\xi - \mu_\eta)^T \Sigma_\eta^{-1} (\xi - \mu_\eta) \tag{7}$$

for any sizing ξ. L represents the distance between ξ and μ_η. Then we have

$$\Xi = \{\xi : L(\xi) = \Lambda^2; \nexists \xi' > \xi, L(\xi') < \Lambda^2.\} \tag{8}$$

where Λ^2 satisfies the following equation:

$$\frac{(N+1)(\eta^2 - 1 + \eta\Lambda^2)}{\eta^2\Lambda^2} = (1-\gamma). \tag{9}$$

Finally, the output of the algorithm (ξ^*) is the cheapest sizing in Ξ,

$$\xi^* = arg\,min_{(\mathcal{A},b)=\xi\in\Xi} C(\mathcal{A}, b) \tag{10}$$

Given the equation L of the hyper-ellipsoid, we want to find the boundary through Tabu search, using function L to determine if a point is inside or outside the boundary. We shall start at sizing μ_η and for each dimension i, using binary search to determine the least-cost point outside the boundary as the starting point of the search, as implemented on and after line 10.

We then perform a queued breadth-first Tabu search starting with these sizings. For each sizing in the search queue, we examine its $(N+1)$ neighbor points - higher-by-one neighbors for points inside and lower-by-one neighbors for points outside. We admit into the Chebyshev bound Ξ points that are outside or on the boundary with its lower-by-one neighbors inside of the boundary. This guarantees that for each sizing admitted, there cannot be a strictly larger sizing inside the hyper-ellipsoid, given its convexity. Algorithm 3 describes the computation of the bound on lines 2-17, and the search for the cheapest system on the bound on lines 18-33.

In practice, the algorithm is implemented in addition to a hashmap to cache the previously searched sizings to improve performance, and boundary limits to prune unrealistic sizings. See Figure 3 for an example of a three-dimensional (two roof segments and one storage), non-dominating partial hyper-ellipsoid Ξ that represents a particular sizing set.

Algorithm 3 Chebyshev Bound-Finding Algorithm

1: **function** FIND_BOUND($\Lambda^2, \xi^{(1)}, \ldots, \xi^{(k)}$)
2: $\mu_\eta \leftarrow \frac{1}{\eta}\sum_{k=1}^{\eta} \xi^{(k)}$
3: $\Sigma_\eta \leftarrow \frac{1}{\eta-1}\sum_{k=1}^{\eta}(\xi^{(k)} - \mu_\eta)(\xi^{(k)} - \mu_\eta)^{\mathrm{T}}$
4: singular_dimensions ← empty list
5: **while** Σ_η is singular **do**
6: Find i such that $\Sigma_{\eta_i} = 0$
7: singular_dimensions.add(i, μ_{η_i})
8: Remove i^{th} row and column for Σ_η, remove μ_{η_i}
9: **end while**
10: search_queue ← empty queue
11: **for** $i \in \{1, \ldots,\}\setminus$ singular_dimensions **do**
12: $\xi^* \leftarrow \mu_\eta$
13: $\xi^*{}_i \leftarrow arg\,min_{\mu_{\eta_i} \le \xi_i \le a_i^{max}} L(\{\mu_{\eta_1}, \ldots, \xi_i, \ldots, \mu_{\eta_N}\})$
14: **if** $\xi^*{}_i$ exists **then**
15: search_queue.enqueue(ξ^*)
16: **end if**
17: **end for**
18: $\Xi \leftarrow$ empty set
19: **while** search_queue is not empty **do**
20: $\xi^* \leftarrow$ search_queue.dequeue()
21: dir $\leftarrow L(\xi^*) \ge \Lambda^2\ ?\ 1 : -1$
22: all_inside? ← true
23: **for** $i \in \{1, \ldots,\}\setminus$ singular_dimensions **do**
24: $\xi_{\text{neighbor}} \leftarrow \{\xi_1^*, \ldots, \xi_i^* + \text{dir}, \ldots, \xi_N^*\}$
25: **if** $L(\xi_{\text{neighbor}}) \ge \Lambda^2$ **then**
26: all_inside? ← false
27: **end if**
28: search_queue.enqueue(ξ_{neighbor})
29: **end for**
30: **if** $L(\xi^*) \ge \Lambda^2$ and all_inside? **then**
31: $\Xi \leftarrow \Xi \cup \xi^*$
32: **end if**
33: **end while**
34: **return** $\{\xi \cup \{\mu_{\eta_i} : i \in$ singular_dimensions$\} : \xi \in \Xi\}$
35: **end function**

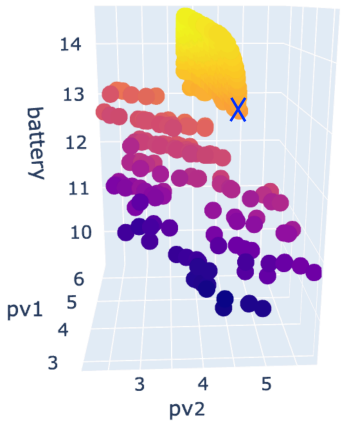

Figure 3: A typical three-dimensional Chebyshev bound Ξ as shown on the upper yellow "dome". The points below are 132 least-cost feasible sizings computed for each data subsample. Warmer colors represent higher total sizing cost. The X marks the least-cost sizing found on the Chebyshev bound.

5.5 Computing η, the Number of Samples

We now discuss how to compute η, the number of data sub-samples. Recall from equation (9) that given fixed N, γ, and η, we compute Λ^2 which determines how close the Chebyshev bound is to μ_η. η and Λ^2 are inversely related; the larger η is, the smaller Λ^2. In other words, using more samples can get us a cheaper system with the same robustness guarantee. From Theorem 2 in Reference [20], when $\eta \to \infty$, Λ^2 approaches

$$\Lambda^{2^*} = \frac{N+1}{(1-\gamma)} \tag{11}$$

from above. This is the lower limit of the gap between the center of the ellipse (at μ_η) and the bound.

Since we can obtain many samples from a large enough PV generation and load dataset, and the bound only gets asymptotically tighter with more samples, we need to decide how many samples is enough. We want to find an η that balances the runtime and system cost. The trade-off can be controlled by introducing a parameter $\beta > 0$ and setting $\Lambda^2 = (1+\beta)\Lambda^{2^*}$. For example, if $\beta = 0.1$, we get a Λ^2 that is a factor 1.1 greater than the minimum Λ^{2^*} achieved when $\eta \to \infty$. From equation (9), this reduces to

$$\eta^2\beta - \eta\Lambda^2 + 1 = 0, \tag{12}$$

which forms a quadratic equation that has one real solution greater than 1, giving us the number of samples corresponding to β:

$$\eta = \frac{\Lambda^2 + \sqrt{\Lambda^4 - 4\beta}}{2\beta}. \tag{13}$$

In practice, we find that setting $\beta = 0.1$ gives a reasonable trade-off, requiring 220 simulations for 85% confidence level and 660 for 95% confidence level over two roof segments, and performing within 1% cost of the sizing result for a lower β value such as 0.01, which requires 10 times more computation.

5.6 Putting it Together

Recall that the algorithm consists of three parts: sampling η sub-intervals from a PV generation and load dataset, running simulations and gradient descent to find the minimum cost sizing for each data sample $\xi^{(1)}, \ldots, \xi^{(\eta)}$ over each separated convex search space, and computing the Chebyshev bound Ξ. Algorithm 4 provides pseudocode that composes the algorithms described in this section to computer a min-cost robust sizing.

Algorithm 4 Robust Sizing Algorithm

function SIZING($N, n, T, \varepsilon, c_i^f, c_i^m, c^b, a_i^{max}, b^{max}, \mathcal{S}_i, \mathcal{L}$)
  $\Lambda^2 \leftarrow (1+\beta)\frac{N+1}{(1-\gamma)}$
  $\eta \leftarrow \frac{\Lambda^2 + \sqrt{\Lambda^4 - 4\beta}}{2\beta}$
  $\Xi^* \leftarrow$ empty set
  for subsets $\in P(\{1, \ldots, N\}) \setminus \varnothing$ **do**
    ▷ Use power set to separate convex search spaces
    for $i \in 1..\eta$ **do**
      $t \leftarrow i\lfloor\frac{n-T}{\eta}\rfloor$ ▷ sample with sliding window
      $\xi^{(k)} \leftarrow$ FIND_SIZING($t, t+T$, subsets)
    end for
    $\Xi \leftarrow$ FIND_BOUND($\Lambda^2, \xi^{(1)}, \ldots, \xi^{(\eta)}$)
    $\xi^* \leftarrow arg\,min_{(\mathcal{A},b)\in\Xi} C(\mathcal{A}, b)$
    $\Xi^* \leftarrow \Xi^* \cup \xi^*$
  end for
  return $arg\,min_{(\mathcal{A},b)\in\Xi^*} C(\mathcal{A}, b)$
end function

6 EVALUATION

We evaluate our multi-roof sizing algorithm using two datasets. The first dataset is extracted from the Pecan Street Dataport [5] and has four years of PV and residential load data. Here, we combine data from multiple homes to synthesize a multi-roof sizing problem. The second is a pyranometer and load dataset from the British Antarctic Survey's Rothera Antarctic research station, where solar PV can be deployed on up to five roof segments.

6.1 Evaluating Robustness on Residential Load

We evaluate the robustness of our algorithm using four years of residential load and PV generation data measured at 49 homes in Austin, Texas [5]. We run leave-one-year-out experiments, where 3 years of data are used as input to the sizing algorithm, and the final year is used as a validation year to check whether the sizing met the unmet load and confidence targets. Given a confidence γ, we expect that the computed sizing will have at most 1-γ fraction of the tests exceed the unmet load target ε.

To create a multi-roof sizing scenario, we use PV generation data from two homes to represent two roof segments, dubbed pv1 and pv2, that have noticeably different PV generation profiles, one peaking in the morning hours and one in the afternoon, as shown in Figure 4. In addition, we found that each home in the dataset has a distinct load pattern. Hence, to compare the sizings according to the shape of each home's load profile rather than its magnitude, we rescale the data so that on average, across four years, each solar panel generates 0.2 kW and the mean load is 2kW. We assume that each home can install up to 60 panels on each roof segment, and up to 120 kWh of storage. We have also assumed a fixed cost of \$2000 for each roof segment, a variable cost of \$500 per panel [17], and a battery cost of \$500 per kWh, similar to the current cost of a Tesla Powerwall [2].

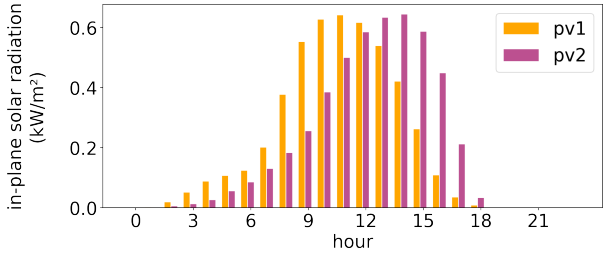

Figure 4: Average hourly solar generation for two roof segments in the Pecan Street dataset. Note that pv1 peaks in the morning hours and pv2 peaks in the afternoon hours.

We evaluate our sizing method for two QoS targets, $\varepsilon = 0.1$ and $\varepsilon = 0.5$. The first target represents a scenario where the homeowner wishes to achieve near-total grid independence, while the second represents a more fiscally-prudent scenario where solar PV primarily meets loads by day and the grid is used to meet load at night. We also evaluate two confidence levels, $\gamma = 0.85$ and $\gamma = 0.95$, with the former level expected to produce a cheaper albeit less-robust sizing.

Figure 5 shows the aggregate results of 196 tests consisting of 4 leave-one-year-out experiments across 49 homes. Note that as the target validation loss increases from 10% to 50%, the distribution visibly shifts to the right, as expected. Moreover, when the target loss is 10% ($\varepsilon = 0.1$), only 4.1% of the tests exceeded the loss at 85% confidence level, and 1% of tests at 95% confidence level. Similar results are seen with a loss target of 50% ($\varepsilon = 0.5$), with 2.5% of tests exceeding the loss target at 85% confidence level, and 2% at 95% confidence level. These results empirically demonstrate that the sizings computed by our algorithm are feasible and robust. They are also reasonably tight, as seen by the increase in the density of the loss distribution left of the loss target indicated by the red line in each figure, though the sizing is somewhat conservative for $\epsilon = 0.5, \gamma = 0.85$.

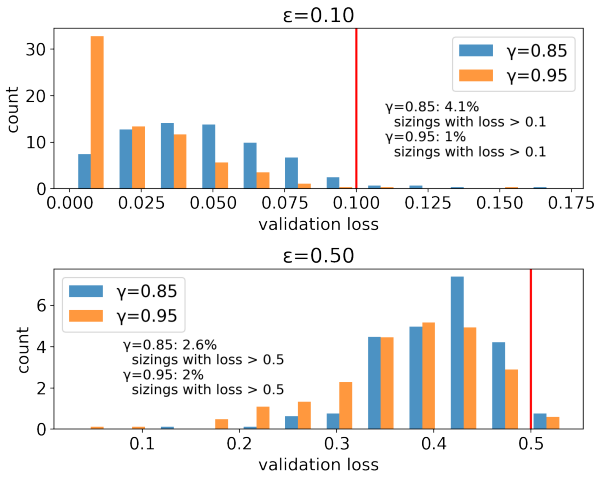

Figure 5: Aggregated leave-one-year-out test results on 49 houses. Percentage of tests that land within the unmet load target for $\varepsilon = 0.1$ and 0.5, with the red vertical line indicating target unmet load levels.

Figure 6 shows the average sizing results across all homes and years given different unmet load target and confidence level combinations. Notably, pv2 is slightly more favored than pv1. This can be explained by observation that the majority of houses have load peaks in the evening, and storage costs are reduced when more panels are allocated to pv2. Also, there is a 4-5× increase in battery and 2-3× increase in PV allocation when ε goes from 0.5 to 0.1. This is because meeting the more stringent loss target requires the system to generate and store PV generation by day for use at night.

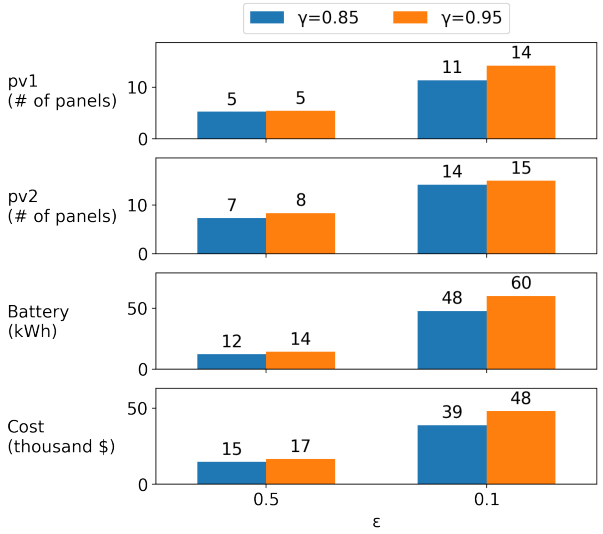

Figure 6: Average sizing results across all houses and years, including PV allocation for the two roof segments, battery amount and total cost of the system.

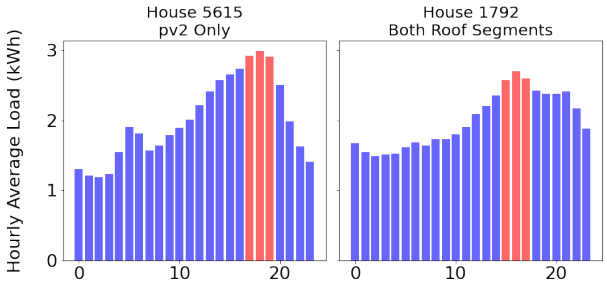

Figure 7: Two typical residential load patterns. House 5615 shows load more concentrated in the evening and its sizing contains only pv2. House 1792's load is more evenly distributed and its sizing contains both roof segments.

We also found that the sizing for the majority of the homes includes panels on both pv1 and pv2, despite each having enough space to hold all the allocated panels. We hypothesize that the differing peaking times for the two roof segments cause PV generation to be more spread out over the course of the day, which reduces the need for storage. To confirm this, we studied, in more detail, the two load patterns shown in Figure 7. Note that the first house has a more pronounced evening load peak, which better matches the generation profile from pv2. Indeed, for this home, our algorithm suggests a sizing that uses only pv2. In contrast, the sizing for the second home uses both roof segments. This confirms our intuition that the optimal sizing attempts to match PV generation profiles to load profiles.

To sum up, our experiments using the Pecan Street Dataport dataset confirm that our algorithm produces feasible sizings that are robust to variations in the solar and load profiles.

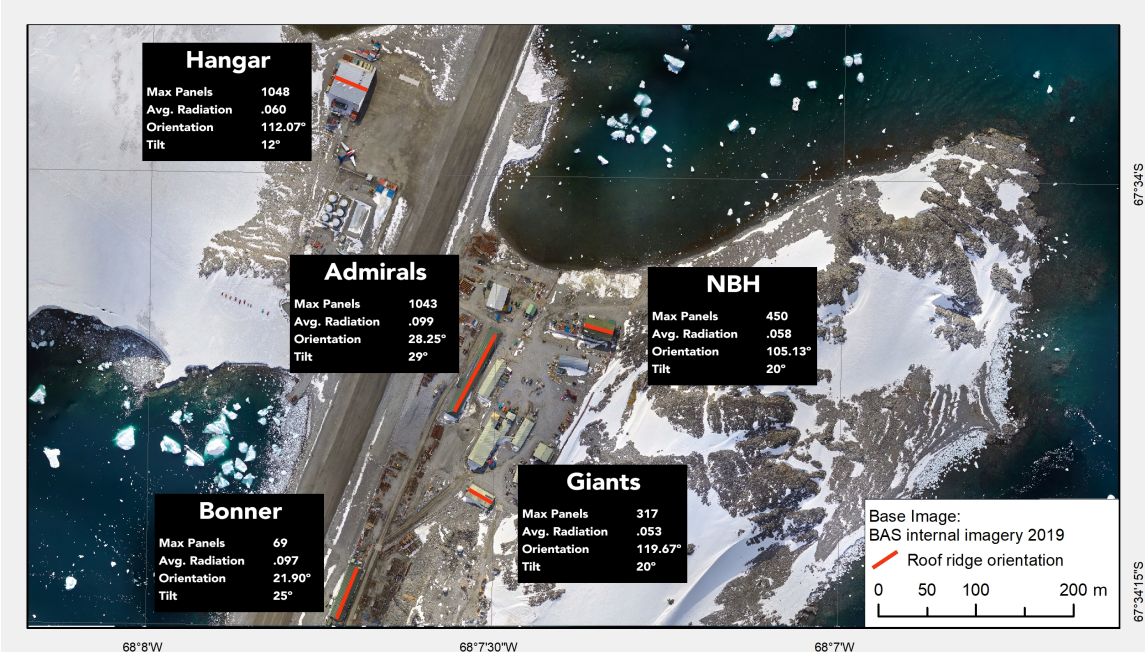

Figure 8: Satellite picture of the Rothera research station. For each building with a usable roof segment, we annotate the building name, the number of 1 m^2 panels that it can support, the average PV generation (kW/panel), the orientation angle from true north, and the tilt angle to the ground plane.

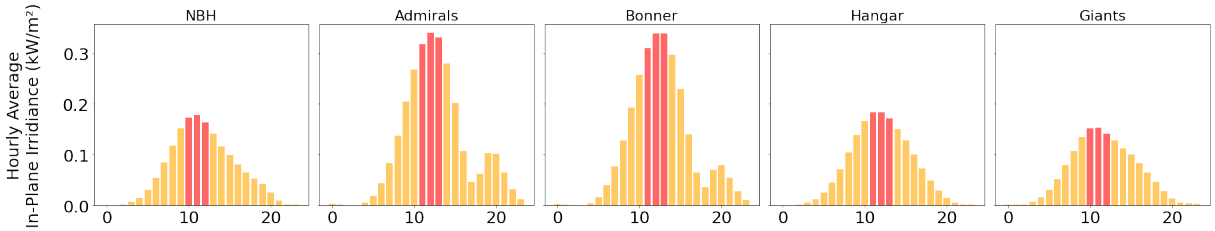

Figure 9: Solar patterns averaged by hour for the five roof segments in the Rothera Station, with the three peak hours emphasized in red. Note that the Admirals and Bonner buildings receive more sunlight, on average, compared to the others.

6.2 Decarbonizing the BAS Rothera Station in Antarctica

We have applied our sizing algorithm to the use case of partly decarbonizing the British Antarctic Survey's (BAS) Rothera research station, based on real traces. The average load of the station is 95 kW, with a peak of 130 kW. So far, the station has been powered by diesel power generators that use 60 m^3 of fuel per month on average, which is carbon-intensive and expensive to deliver to such a remote location. Through its current modernization program, the BAS aims to decarbonize the station by 2030. We study how the station might hypothetically install a solar+storage system to contribute to decarbonization.

The BAS provided us with six years of hourly horizontal solar irradiance as measured through a pyranometer, and one year of representative hourly load data from 2015, which was before the initiation of the modernization program. At present, five buildings in the station have been identified as suitable for solar PV installation, as shown in Figure 8. We therefore use the irradiance trace to compute five separate PV generation traces, according to the tilt and orientation of each potentially-suitable roof segment as suggested in Reference [9]; the resulting set of daily average PV generation profiles is shown in Figure 9. Note that the Admirals and Bonner have higher solar generation potential than the other three and are therefore the best candidates for PV panels. However, the roof segment on Bonner is relatively small. Three other roof segments are less desirable, with Hanger having a slightly larger average radiation and the largest available area. Collectively, the five roof segments can support up to 450 kWp of PV generation.

We also trained an ARMA model on the load data to create synthetic load traces spanning six years, using the process described in Reference [21].

The costs used in this study are a £5000 fixed cost for equipping each roof, £1.18/Wp marginal PV cost (equivalent to £188 per m^2

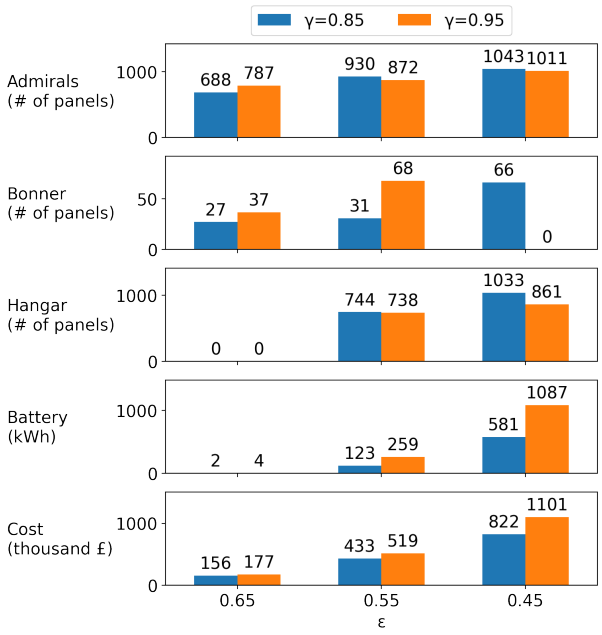

Figure 10: Results of sizing, including number of solar panels, kWh of batteries, and cost for different confidence levels and target unmet load. Note that allocation for less-efficient roof segments NBH and Giants are 0 across all tested loss/confidence combinations.

panel which is assumed to be 160 Wp), and £670/kWh for batteries. This data was used to compute the sizing for several system performance targets, with ε ranging from 0.35 (meeting 65% of load) to 0.65 (meeting 35% of load), and confidence level γ at 0.85 and 0.95.

The results of the sizing are summarized in Figure 10, which includes the per-roof segment panel allocation, battery size, and total cost. No feasible sizing existed for $\varepsilon = 0.35$, due to the fact that very little sunlight is seen in Antarctica between May and August. Other ε targets resulted in varied systems sizings: with $\varepsilon = 0.65$, the sizing requires very little battery capacity and most of the solar panels are on a single roof segment, while for $\varepsilon = 0.45$ the sizing was split more equally across two of the roof segments and required roughly 500-1000 kWh of storage to meet night-time load, depending on the desired confidence level. Notably, for $\varepsilon = 0.45$ and $\gamma = 0.95$, the preferred roof segments are not always those which receive the most radiation (Admirals and Bonner), but a combination of those that spread out PV generation over the course of a full day (Admirals and Hangar). This is in contrast to a naive greedy approach that would have placed panels on the most efficient available roof segment, as discussed in Section 4.1.

Note that because we only had access to only one year of load data, we were unable to carry out a leave-one-out analysis to evaluate the robustness of our sizing.

7 DISCUSSION AND CONCLUSION

In this work, we present an algorithm for choosing the number of solar panels and the amount of storage needed to meet a certain fraction of the load in a microgrid setting. Unlike prior work, which assumes that the desired number of panels can always be accommodated on a rooftop, we take into account the pragmatic issue that roofs typically incorporate multiple roof segments, each with its own panel capacity, tilt, and orientation. This unexpectedly leads to a much more complex sizing problem. Our solution, which is based on stochastic gradient descent, allows us to compute sizings despite the non-linear nature of the problem. We demonstrate the robustness of our approach using a leave-one-out analysis and the Pecan Street dataset. We also use our approach to compute the sizing needed for different levels of decarbonization of the British Antarctic Survey station in Rothera.

Our approach deals with non-linearity by using stochastic gradient descent to find least-cost, feasible sizings over repeated trials and the computation of a Chebyshev bound over the results of these trials. Unfortunately, this approach has some limitations, as discussed next.

First, when using solar and load traces from multiple years, one of the years may have atypically low PV generation or high load. A sizing computed using an atypical year tends to be more conservative, as is a Chebyshev bound that includes this sizing. Unfortunately, there does not appear any systematic approach to identify anomalous traces in a fairly small set of traces.

Second, to make our search more efficient, we partition the non-linear search space into a set of convex sub-spaces. This is reasonable for a small number of roof segments (for example, in Rothera, we used this approach with 5 roof segments). However, this approach does not scale well with the number of roof segments due to a combinatorial explosion.

Third, we compute the Chebyshev bound using a multidimensional breath-first Tabu search. This turns out to be memory intensive, especially when searching a large N-dimensional space where each point has N neighbours.

Fourth, a Chebyshev bound sometimes leads to unintuitive results, an example being the recommendation of a handful of panels on the Bonner building in Rothera. This is because the Bonner building is part of a least-cost sizing for some runs of the stochastic gradient descent algorithm and not others. When computing a Chebyshev bound, however, this results in a small number of panels being allocated to this building, a counter-intuitive result.

Finally, on an unrelated matter, we note that our sizing approach assumes a microgrid setting, that is, with no payments for over-generation from the grid using net-metering or feed-in-tariff. If such payments are introduced, then the sizing problem becomes far more complex, since the storage operation policy depends in detail upon the nature of the grid payment scheme [15]. We defer analysis of this more complex scenario, as well as to overcome the limitations listed above, to future work.

8 ACKNOWLEDGEMENT

We would like to thank the Pecan Street Dataport for access to their solar and load traces. Thanks also to the British Antarctic Survey for sharing their load and solar data sets, as well as the building orientation and size parameters, from Rothera. Finally, thanks to Jackson Taylor for his roof.

REFERENCES

[1] Commercial and residential hourly load profiles for all tmy3 locations in the united states. https://openei.org/doe-opendata/dataset/commercial-and-residential-hourly-load-profiles-for-all-tmy3-locations-in-the-united-states. Accessed: 2021-02-15.

[2] Tesla powerwall design. https://www.tesla.com/powerwall/design. Accessed: 2021-02-15.

[3] Tesla solar panel sizing and design. https://www.tesla.com/support/energy/solar-panels/going-solar/sizing-and-design. Accessed: 2021-02-15.

[4] U.S. Energy Information Administration, Smart Meter Installations. https://www.eia.gov/tools/faqs/faq.php?id=108&t=3. Accessed: 2021-02-15.

[5] Pecan street dataport. https://dataport.cloud/, 2021.

[6] Pvwatts calculator. https://pvwatts.nrel.gov/, 2021. Accessed: 2021-02-15.

[7] Clairand, J.-M., Arriaga, M., Cañizares, C. A., and Álvarez-Bel, C. Power generation planning of galapagos' microgrid considering electric vehicles and induction stoves. *IEEE Transactions on Sustainable Energy 10*, 4 (2018), 1916–1926.

[8] Das, I., and Cañizares, C. A. Renewable energy integration in diesel-based microgrids at the canadian arctic. *Proceedings of the IEEE 107*, 9 (2019), 1838–1856.

[9] Gao, X. Should you clean your solar panels now? Master's thesis, 2014. Accessed: 2021-02-15.

[10] Hafez, O., and Bhattacharya, K. Optimal planning and design of a renewable energy based supply system for microgrids. *Renewable Energy 45* (2012), 7–15.

[11] Hidalgo-León, R., Siguenza, D., Sanchez, C., León, J., Jácome-Ruiz, P., Wu, J., and Ortiz, D. A survey of battery energy storage system (bess), applications and environmental impacts in power systems. In *2017 ieee second ecuador technical chapters meeting (etcm)* (2017), IEEE, pp. 1–6.

[12] J. Duchi, E. H., and Singer, Y. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research* (2011), 2121–2159.

[13] Kazhamiaka, F., Ghiassi-Farrokhfal, Y., Keshav, S., and Rosenberg, C. Robust and practical approaches for solar pv and storage sizing. In *Proceedings of the Ninth International Conference on Future Energy Systems* (2018), pp. 146–156.

[14] Kazhamiaka, F., Ghiassi-Farrokhfal, Y., Keshav, S., and Rosenberg, C. Comparison of different approaches for solar pv and storage sizing. *IEEE Transactions on Sustainable Computing* (2019), 1–1.

[15] Kazhamiaka, F., Rosenberg, C., and Keshav, S. Practical strategies for storage operation in energy systems: design and evaluation. *IEEE Transactions on Sustainable Energy 7*, 4 (2016), 1602–1610.

[16] Krah, K. Behind-the-meter solar+ storage modeling tool comparison. Tech. rep., National Renewable Energy Lab.(NREL), Golden, CO (United States), 2019.

[17] Matasci, S. How much do solar panels cost? 2021 guide: Energysage. https://news.energysage.com/how-much-does-the-average-solar-panel-installation-cost-in-the-u-s/, Feb 2021.

[18] Orlando, M., Bottaccioli, L., Patti, E., Macii, E., Vinco, S., and Poncino, M. Optimal configuration and placement of pv systems in building roofs with cost analysis. In *2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC)* (2020), IEEE, pp. 1411–1416.

[19] Ru, Y., Kleissl, J., and Martinez, S. Storage size determination for grid-connected photovoltaic systems. *IEEE Transactions on Sustainable Energy 4*, 1 (Jan 2013), 68 – 81.

[20] Stellato, B., Van Parys, B. P., and Goulart, P. J. Multivariate Chebyshev inequality with estimated mean and variance. *The American Statistician 71*, 2 (2017), 123–127.

[21] Sun, S., Kazhamiaka, F., Keshav, S., and Rosenberg, C. Using synthetic traces for robust energy system sizing. In *Proceedings of the Tenth ACM International Conference on Future Energy Systems* (New York, NY, USA, 2019), e-Energy '19, Association for Computing Machinery, p. 251–262.

[22] Tieleman, T., and Hinton, G. Rmsprop: Divide the gradient by a running average of its recent magnitude. *Neural Networks for Machine Learning* (2012).

[23] Zeiler, M. D. Adadelta: An adaptive learning rate method. *CoRR abs/1212.5701* (2012).