Using Storage to Minimize Carbon Footprint of Diesel Generators for Unreliable Grids

Sahil Singla, Yashar Ghiassi-Farrokhfal, and Srinivasan Keshav
David R. Cheriton School of Computer Science
University of Waterloo

Technical Report CS-2013-16

Abstract

Although modern society is critically reliant on power grids, even modern power grids are subject to unavoidable outages. The situation in developing countries is even worse, with frequent load shedding lasting several hours a day due to a large power supply- demand gap. The standard solution for residences is, therefore, to backup grid power with local generation from a diesel generator (genset).

When carbon emission matters, a hybrid battery-genset is preferable to a genset only system. Designing such a hybrid system is entangled with the tradeoff between cost and carbon emission. Towards the analysis of such a hybrid system, we first compute the minimum required battery size for eliminating the use of genset. Such a battery must guarantee a target *loss of power* probability for an unreliable grid. We then compute the minimum required battery for a given genset and a target allowable carbon footprint.

Drawing on recent results, we model both problems as buffer sizing problems that can be addressed using stochastic network calculus. Specifically, a numerical study shows that, for a neighbourhood of 100 homes, we are able to estimate the carbon footprint reduction, compared to an exact numerical analysis, within a factor of 1.7.

Index Terms

Smart grids, Batteries, Diesel engines, Performance analysis, Power system reliability, Power demand.

This work was done at the David R. Cheriton School of Computer Science, University of Waterloo.

I. INTRODUCTION

The power grid underlies most modern societies: power failures can affect critical institutions such as hospitals, water treatment facilities, aircraft control towers, and Internet data centres. Despite this great reliance on electrical power, as the aftermath of super-storm Sandy vividly demonstrated, even modern power grids are subject to unavoidable outages due to storms, lightning strikes, and equipment failures. The situation in developing countries is worse, with daily load shedding lasting two-to-four hours due to demand spikes and unreliable generation [26].

In the face of this inherent unreliability, a common solution is for critical facilities, and even some individual homes, to augment grid power with local generation, typically from a diesel generator (genset). This, however, increases the carbon footprint of the load [23]. One of the best solutions to this problem is using a hybrid system that combines gensets with storage batteries (along with two-way inverters to convert between AC and DC power). We study the use of such a hybrid system to allow a set of homes in a single residential neighbourhood to avoid power outages. Because storage battery is expensive, the design of a hybrid system requires a trade-off between carbon emission and cost. In particular, to completely eliminate the carbon footprint, we should choose the minimum battery size that can guarantee a target loss of power probability in the total absence of a genset.

The design of such a hybrid system for unreliable grids is a technically complex problem. In this work, for the first time, we analytically approach this problem. We first consider the case where reducing residential neighbourhood carbon emission to zero is desirable and we size a battery in the absence of a genset. Thereafter, we consider the case where both constraints exist and we compute the minimum battery size for a given maximum allowable genset carbon emission. We note that these two problems are technically different as they have different constraints; loss of power probability is the constraint in the absence of genset and total carbon footprint is the constraint in the presence of genset (the presence of genset ensures demand is always met).

The complexity of the problem comes from choosing an appropriate model for the stochastic nature of household loads [16]. Indeed, it has been shown to be isomorphic to the complex (but well-known) problem of choosing a buffer large enough to smooth the data being generated by a variable-bit-rate traffic source [2], [30]. Therefore, drawing on recent results, we approach the solutions to our problem using the powerful techniques of stochastic network calculus.

We have numerically evaluated the accuracy of our algorithms using real traces of electrical loads collected over 12 months from 4500 homes in Ireland[1] [9]. Analysis shows that we are able to estimate the carbon footprint reduction, compared to an exact numerical analysis, within a factor of 1.7. Moreover, given a target requirement of a power outage being upper bounded by one day in ten years (a standard risk target in power engineering), our approach computes a battery size that is only about 10% more than the minimum battery required had the future load been exactly known.

The key contributions of our work are:

1) We use a stochastic network calculus approach to characterize the stochastic electrical load and find the genset carbon emission as a function of the battery size, thus allowing us to compute the battery size needed to limit carbon emissions of a residential neighbourhood.
2) Given a load characterization, we analytically compute the smallest battery size necessary to eliminate the use of genset and meet a target loss of power probability for a load connected to a highly *unreliable* power grid.
3) We use measured electricity consumption data from 4500 Irish homes to compare our carbon emission and battery size bounds to those computed (a) empirically, (b) from bounds obtained using teletraffic theory, and (c) using a multi-state fluid Markov model. We find that our bounds are quite close to the empirical optimal and far better than the two other approaches.

The rest of the paper is laid out as follows. Section II starts with a motivation for the problem along with the prior work. Section III presents the background needed from teletraffic networks and stochastic network calculus, and discusses the prior known bounds. The model used to analyze battery sizing problem along with our assumptions is explained in Section IV. The bounds on minimum battery size, both in presence and absence of genset, are computed in Section V. These bounds are heavily dependent on the accuracy of statistical sample path envelope on effective demand. Hence, in Section VI, we explain how to obtain a tight envelope for any given load dataset. We show the tightness of our bounds numerically on the Irish dataset in Section VII by comparing them to the bounds obtained had the future load been exactly known

[1] Although loads in Ireland are not the same as loads in developing countries, our methodology can be applied to traces collected from any country.

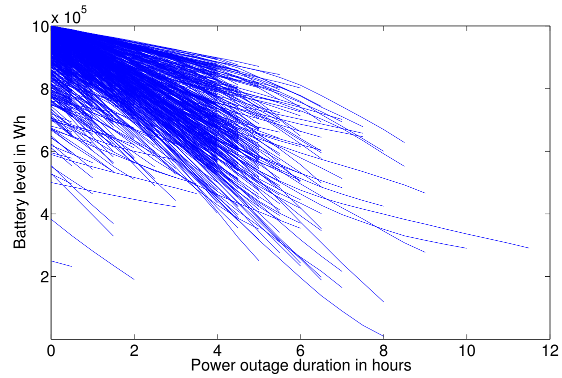

Fig. 1: Trajectories of battery level during power outage periods for $B = 1MWh$ and $C = 10^5W$.

and the bounds from prior work. Finally, Section VIII concludes the paper with the limitations and future work.

II. Motivation and Related Work

The difficulty of choosing a battery size sufficient to meet a stochastic load when confronted with stochastic power outages is illustrated in Figure 1. In this figure, the X-axis shows the duration of a power outage, that is, a time during which the grid does not supply power, and the Y-axis shows the state of charge of a battery of size 1 MWh that charges itself when the grid is available and discharges during an outage. Each line (trajectory) in the figure represents a particular power outage incident; the Y-intercept of the line is the initial state of charge when the outage started and the line ends when the power outage ends. Each trajectory was computed using real measured loads from the Irish dataset (details in Section VII).

We define *loss of power* when the battery level drops to 0, a situation that we would very much like to avoid. In absence of a genset, this would imply that the demand cannot be met by the system, and in presence of a genset, loss of power would imply that the genset is emitting carbon to meet the demand. In this numerical evaluation, the loss of power state happens only once. Note that had the battery size been smaller, say 400 KWh, we would have had a much larger number of loss of power events.

It is clear that the probability of a loss of power depends on many factors, including the battery size, its charging rate, a characterization of the load during a loss of power event, the distribution of the inter-outage intervals, and the distribution of outage durations. Moreover, what we seek to compute is the tail probability of the state-of-charge distribution of the battery. This is an inherently complex problem. However, it has been recently shown that this storage sizing problem is very similar to the buffer-sizing problem in a telecommunications network. We are inspired by new theoretical results in [30], [3], which adopt a queuing-theoretic buffer-sizing analysis to size batteries, to use a similar approach based on stochastic network calculus, to study the problem of battery sizing to limit genset carbon emission for unreliable grids.

Specifically, we consider two scenarios. In the first scenario, we study a system that completely avoids genset by only using storage battery. Here, our goal is to find the minimum required battery size that meets a given target loss of power probability, such as one day loss of power in every ten years. In the second scenario, we study the battery-genset hybrid system where the demand can be always met as the genset steps in when the battery is fully discharged. However, we still seek to size the battery so that the carbon emission from the diesel generator must not exceed a target threshold.

Prior work in this area relies primarily on empirical numerical analysis rather than analytical modeling [4], [14], [27]. Carbon emission due to a battery-genset hybrid system has been studied mainly through the notion of genset *efficiency*, which is the diesel consumption per unit energy production [4], [25]. We note that some of the prior analytical works (e.g. [3]) assume load stationarity. In contrast, our approach does not need to explicitly assume stationarity.

A probabilistic loss of power formulation is presented in [30] for an intermittent power resource (e.g., wind or solar power) serving a stochastic demand with the aid of batteries. The framework in [30] considers a battery only system, and genset existence or the trade-off between storage size and genset carbon emission is not part of that framework. Moreover, even for the battery only system, that framework cannot be directly used for an unreliable system due to a major system model difference; In an unreliable gird scenario, if the grid is available then it is assumed to be large enough to serve both the instantaneous demand and charge the battery at its maximum charging rate C. This is not the case for renewable energies as it can be available in a time slot yet not large enough to meet the demand and charge the battery. In other words, in an unreliable gird scenario, the intermittent source can be considered as an infinite source with vacation time, whereas renewable energies are limited stochastic sources.

III. Background and Preliminaries

The problem of battery sizing in power distribution systems can be mapped to the problem of buffer sizing in the teletraffic network. This analogy has been used in some recent papers for battery sizing, borrowing state-of-the-art analytical results on probabilistic buffer sizing from teletraffic theory [2], [30], [32]. We briefly discuss the essence of this analogy in this section.

A. Deterministic loss of power vs. deterministic loss of packet

Suppose that an arrival process A enters a buffer (queue) of size B, which can serve traffic at rate C, and let A' be the corresponding departure process. We assume discrete time model, where the events can only happen at discrete time instants, i.e., $t = 0, 1, \ldots$. We denote the total arrival from process A in time interval $[0, t]$ by $A(t)$ and we use $A(s, t)$ to mean $A(t) - A(s)$. The *backlog* $b(t)$ at any time t is defined to be the buffer content at that time and is given by the following recursive equation:

$$b(t) = \min(B, [b(t-1) + A(t-1,t) - C]_+) \ , \tag{1}$$

where $[x]_+ = \max(0, x)$ for any value of x. Eq. (1) is equivalent to the following non-recursive expression [13]:

$$\begin{aligned} b(t) = \min_{0 \le u \le t} \Big(\max_{u \le s \le t} \Big(& A(s,t) - C.(t-s), \\ & A(u,t) - C.(t-u) + B \Big) \Big) \ . \end{aligned} \tag{2}$$

The *loss of packet* due to buffer overflow at any time t is

$$l(t) = [A(t-1,t) - C + b(t-1) - B]_+ \ . \tag{3}$$

Eqs. (2-3) can be combined to extract the following loss characterization [13]:

$$\begin{aligned} l(t) = \min_{0 \le u \le t-1} \Big(\max_{u \le s \le t-1} \Big(& [A(s,t) - C.(t-s) - k(t)]_+, \\ & [A(u,t) - C.(t-u) + k(u) - k(t)]_+ \Big) \Big) \ , \end{aligned} \tag{4}$$

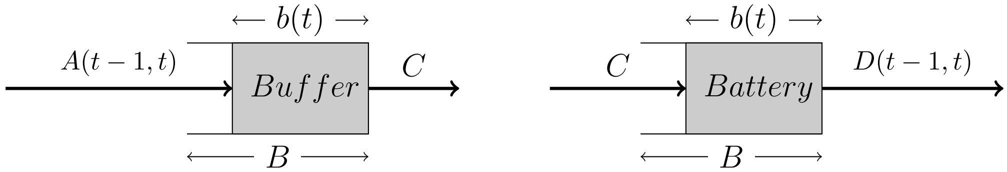

Fig. 2: Loss of packet due to buffer overflow vs. loss of power when a demand finds the battery empty.

where

$$k(t) = \begin{cases} B & t > 0 \\ 0 & t = 0 \end{cases} \tag{5}$$

There is an analogous problem in the power system as follows. A constant power source with rate C is feeding a battery with size B, and the battery is used to serve an intermittent demand D (see Figure 2). The *deficit battery charge* b^d is defined as the amount of energy needed to fully charge the battery and is given by

$$b^d(t) = \min(B, [b^d(t-1) + D(t-1,t) - C]_+) \ , \tag{6}$$

where $D(t-1,t)$ is the demand at time slot t. A comparison between Eq. (1) and Eq. (6) suggests that the deficit battery charge is mapped to the backlog status, if the demand D is mapped to the arrival traffic A, the power supply is mapped to the capacity (service rate) of the link, and the battery size is mapped to the buffer size.

The *loss of power* event is defined as the event that a demand finds the battery empty or, equivalently, finds the deficit charge of the battery full. The loss of power at time t is given by $[b^d(t-1) - C + D(t-1,t) - B]_+$, which is comparable to Eq. (3). Thus, the non-recursive equations for the backlog and loss of packet in a finite buffer system in Eqs. (2) and (4) can, respectively, be used to compute the deficit state of the charge and the loss of power in energy systems [2], [31].

B. Probabilistic loss formulations

A probabilistic loss analysis requires an accurate characterization of the statistical properties of the arrival (demand) and the link capacity (power source). The probabilistic loss formulation has been studied extensively in the asymptotic regime when the number of independent arrivals is large (*many sources asymptotic regime*) [15]. Since there are not *many independent* sources of energy or demand sharing the same battery in a distribution system, the many sources asymptotic results cannot be used in this context. However, there are alternative approaches:

1) Kesidis bound: A traffic arrival process A is called a *peak-rate constraint leaky-bucket arrival* if it satisfies

$$\forall s,t: \quad A(s,t) \leq \min(\pi(t-s), \sigma + \rho(t-s)) \tag{7}$$

for some π, σ and ρ such that $\rho \leq \pi$. If a peak-rate constrained leaky-bucket arrival process which is also *stationary* is fed to a link with total capacity C, where $\rho \leq C \leq \pi$, then the *stationary* backlog status b satisfies the following [20]:

$$\Pr\{b > B\} \leq \frac{\sigma - \frac{\pi-\rho}{\pi-C}B}{\frac{C}{\rho}\sigma - B} \ . \tag{8}$$

This formulation is used in [3] to compute transformer/storage sizing in the distribution networks.

2) Using Network Calculus for a measurement trace: Network Calculus allows probabilistic performance analysis including loss probability for a large class of arrivals. This theory uses upper bounds on the traffic arrivals and lower bounds on the available service on any time scales to compute performance bounds. Interested readers can refer to [5], [10], [18] for a complete tutorial on Network Calculus.

There are several probabilistic upper bounds on the arrivals proposed in the literature (see [24] for a review of existing bounds). Here we use a concept called the *statistical sample path envelope* [8]. A non-decreasing function $\mathcal{G}$ is a statistical sample path envelope for an arrival process A with bounding function ε if it satisfies the following at any time $t \geq 0$ and for any $\sigma \geq 0$

$$\Pr\left\{\max_{s \leq t}\{A(s,t) - \mathcal{G}(t-s)\} > \sigma\right\} \leq \varepsilon(\sigma) \ , \tag{9}$$

where $\varepsilon(\sigma)$ is non-increasing in σ.

Due to the complexity of Eq. (4), the following upper bound on Eq. (4) is used in [13] to extend the traffic loss formulations to the probabilistic settings:

$$l(t) \leq \max_{0\leq s\leq t}([A(s,t) - C.(t-s) - B]_+) \ . \tag{10}$$

This probabilistic upper bound formulation in [13] is used to obtain an upper bound on the loss of power probability in [30], employing the underlying mapping between the loss of power and the loss of traffic:

Theorem 1 (Loss of power probability [30])**.** *Suppose that $\mathcal{G}$ is the statistical sample path envelope for a demand process D in the sense of Eq. (9) with bounding functions ε_g. Then, the loss of power probability satisfies the following:*

$$\Pr\{l(t) > 0\} \leq \varepsilon_g\left(B - \max_{0\leq\tau\leq t}(\mathcal{G}(\tau) - C\tau)\right) \ . \tag{11}$$

If a measurement trace of a process A is given, the statistical sample path envelope $\mathcal{G}$ can be computed by (a) constructing a set Y, consisting of the sample values for the event $\max_{s\leq t}\{A(s,t) - \mathcal{G}(t-s)\}$ for each trajectory and at any time t, and (b) using the complementary cumulative distribution function (CCDF) of the sample set Y as a bounding function for Eq. (9) [11].

3) Using Network Calculus for Markovian arrivals: Sometimes an arrival model is given instead of a measurement set. One of the most general and widely-used models is the multi-state Markovian (MSM) fluid flow process, which is a Markov chain with finite states, with states representing the rate at which traffic is generated at a certain time. Theorem 1 can be used to obtain a loss probability for MSM processes since there is a statistical sample path envelope for this type of processes, which can be computed as follows:

An M-state fluid flow Markov chain with transition matrix Q and traffic rate at state i being r_i satisfies [19]

$$\forall t: \quad E\left[e^{\beta A(t)}\right] \leq e^{\rho(\beta)t} \ , \tag{12}$$

where $\rho(\beta)$ is the largest eigenvalue of the matrix $\frac{1}{\beta}Q + R$ and $R = diag(r_i)$. Combining Chernoff bound and Eq. (12) yields the following at any time t and s $(\leq t)$ and any σ (≥ 0)

$$\Pr\{A(s,t) > \rho(t-s) + \sigma\} \leq e^{-\beta\sigma} \tag{13}$$

which shows that a multi-state Markovian process is a special case of the large class of exponentially bounded burstiness traffic sources (EBB) [33]. For such a traffic, $\mathcal{G}$ from the following is a statistical sample path envelope in the sense of Eq. (9) with bounding function ε for any $\gamma > 0$ [12]

$$\mathcal{G}(t) = (\rho(\beta) + \gamma)t; \qquad \varepsilon(\sigma) = \frac{e^{-\beta\sigma}}{1 - e^{-\beta\gamma}} \tag{14}$$

One can insert the statistical sample path envelope for a multi-state Markovian model from Eq. (14) to Theorem 1 to obtain a loss of power probability for such a demand.

IV. System Model

The system model considered in this paper is illustrated in Figure 3. The grid and the battery are used to serve 'most' of the demand. The grid is available irregularly. If the grid is available, then it is used to serve the demand and charge the battery. When the grid is not available (power outage), the charge of the battery is used to serve the demand. In the presence of a genset (not shown in Figure 3), the genset starts running only when the battery is empty during an outage.

Denote by $d(t)$ the energy demand at time slot t, and by $D(t)$ the cumulative energy demand in time interval $[0, t]$[2]. To simplify notation, we define $D(s, t) = D(t) - D(s)$. The charging rate of the battery is represented by C and the battery size by B. We assume a discrete time model, where $t = 0, 1, \ldots$. Let $x(t)$ be a binary random variable representing the availability of the utility grid at time slot t, *i.e.*,

$$x(t) = \begin{cases} 0 & \text{If grid is unavailable (power outage) at slot } t \\ 1 & \text{If grid is available at time slot } t \end{cases} \tag{15}$$

We define $x^c(t) = 1 - x(t)$ as the complement of $x(t)$. If the grid is available at time slot t (i.e., $x(t) = 1$), the energy demand is served by the grid and the battery will be charged by as much as C energy unit in that time slot. On the other hand, if the grid is not available (i.e., $x(t) = 0$), the energy demand must be served by the energy stored in the battery.

[2]The energy demand varies widely over the course of a year showing marked seasonality. Our analysis is agnostic to the time interval over which the demand is modeled. In practice, however, similar to the concept of busy-hour sizing in a telecommunication network, we advocate the sizing of a battery keeping in mind the underlying non-stationarity of the demand process [3].

Name	*Description*
Power outage	Electricity from the grid is unavailable
Loss of power	An outage period with battery being empty
B	Storage battery capacity
C	Battery charging rate
ϵ^*	Target loss of power probability
$x(t)$	Grid availability at time t
$x^c(t)$	Grid unavailability at time t
$d(t)$	Power load at time t
$d^e(t)$	Effective power load to battery at time t
$b(t)$	Battery charge level at time t
$b^d(t)$	Battery deficit charge at time t
$l(t)$	Amount of loss of power at time t
$\mathcal{G}$	Statistical sample path envelope
ε_g	Bounding function for sample path envelope

TABLE I: Notation

We follow two objective functions for battery sizing in this paper. First, in the absence of genset, we size the battery B such that given some statistical properties of the energy demand process, the probability of loss of power is kept below a target threshold ϵ^*. Second, in the presence of genset, we size the battery such that the total carbon footprint is kept below a certain threshold.

We have the following assumptions in our formulations

1) Battery charging rate is upper bounded by C but there is no constraint on the discharge rate. [3]
2) Genset is large enough to meet the maximum aggregate loads.

The above assumptions typically hold in practice since battery charge rate depends on the technology of the battery, and for technologies such as lead-acid battery, the discharge rate is multiple times higher than the charging rate [6]. Moreover, the marginal cost of increasing genset size is negligible compared to the marginal cost of a battery [28], [29].

[3]Note that a constant battery charge-discharge energy loss factor can be incorporated in C as it would just reduce the original battery charging rate by this factor.

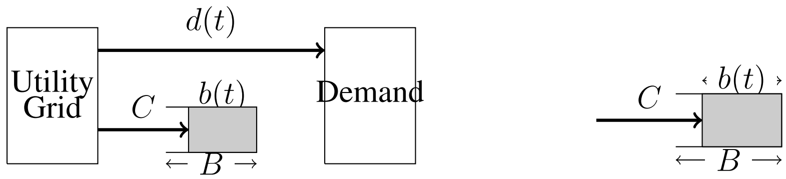

(a) Grid available

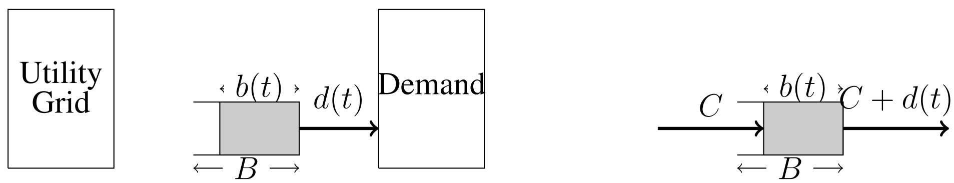

(b) Power Outage

Fig. 3: Storage battery model before and after transformation

V. Battery Sizing Formulation

The unreliable grid described in our problem statement can be converted to a reliable *compound* power source i.e., one that is augmented by batteries. In this section we compute the required battery size such that a target loss of power probability can be guaranteed in the absence of a genset and a target carbon emission threshold can be met in the presence of a genset. The size of the battery is a function of the stochastic nature of demand and grid unavailability.

A. *Loss of power formulation*

As discussed in Section II, due to major differences in system models we cannot apply the loss of power formulation in [30] to our problem. However, the following trick converts our system model to their system model, allowing us to use their loss of power formulation.

Consider the input and output processes to the battery separately for the cases where the grid is available ($x(t) = 1$) and unavailable ($x(t) = 0$). If $x(t) = 1$, the arrival energy process to the battery at any time instant t is a constant, C, and the departure energy process is zero. If $x(t) = 0$, the arrival energy process is zero and the departure energy process is $d(t)$. The battery state-of-charge does not change if we assume that the real arrivals and departures to the battery are both shifted by the same constant at any time. Therefore, we can assume C and $d(t) + C$, respectively, as the arrival and departure processes when $x(t) = 0$ (instead of 0 and $d(t)$) and have the same battery state of charge (see Figure 3). Combining the two cases $x(t) = 0$ and $x(t) = 1$ with the above substitution, we can assume that the battery is always charged with the rate C and discharged by the *effective demand* d^e defined as:

$$d^e(t) = [d(t) + C](1 - x(t)) = [d(t) + C]x^c(t) \tag{16}$$

which is the portion of the demand that the battery must serve. Using the above transformation and Eq. (6), we have:

$$b^d(t) = \min(B, [b^d(t-1) + d^e(t) - C]_+) \ . \tag{17}$$

With the above mapping, the loss formulation in the previously mentioned finite buffer system is given by $l(t) = [b^d(t-1) + d^e(t) - C - B]_+$. Hence, we can use Eq. (4) to describe the following loss process as follows:

$$\begin{aligned} l(t) = \min_{0 \le u \le t-1} \Bigg(& \max_{u \le s \le t-1} ([D^e(s,t) - C.(t-s) - k(t)]_+, \\ & [D^e(u,t) - C.(t-u) + k(u) - k(t))]_+ \Bigg) \ , \end{aligned} \tag{18}$$

where D^e is the cumulative version of the effective demand (Eq. (16)) and k is as expressed in Eq. (5).

B. Battery sizing for eliminating genset

The exact loss description from Eq. (18) is difficult to use in practice. Instead, we use the following upper bound (from [17]) to derive an upper bound on the loss probability:

$$\begin{aligned} l(t) \le \min \big(& [d^e(t) - C]_+, \\ & \max_{0 \le s \le t-1} ([D^e(s,t) - C(t-s) - B]_+) \Big) \end{aligned} \tag{19}$$

$$\begin{aligned} = \min \big(& d(t)x^c(t), \\ & \max_{0 \le s \le t-1} ([D^e(s,t) - C(t-s) - B]_+) \Big) \ , \end{aligned} \tag{20}$$

where in Eq. (19) two specific values for u in the minimization of Eq. (18) are chosen: $u = t-1$ and $u = 0$. Eq. (20) uses the definition of $d^e(t)$ and that $d(t) > 0$. The above inequality can be used to compute a probabilistic upper bound on the loss probability for our problem. See Appendix for a proof of the following lemma.

Lemma 1 (Amendment to Theorem 1 [30])**.** *Suppose that an unreliable grid uses a battery of size B with charging rate C to serve a demand. Suppose also that $x^c(t)$ represents the grid unavailability at time slot t ($x^c(t) = 1$ if the grid is unavailable and $x^c(t) = 0$, otherwise) and $\mathcal{G}$ is a statistical sample envelope for process D^e with bounding functions ε_g in the sense of Eq. (9). Then, the loss of power probability satisfies the following*

$$\Pr\{l(t) > 0\} \leq \min\Bigg(\Pr\{x^c(t) > 0\}, \varepsilon_g \left(B - \max_{\tau \geq 0}(\mathcal{G}(\tau) - C\tau) \right) \Bigg) . \tag{21}$$

We can use Lemma 1 to compute the minimum battery size satisfying a target loss of power probability ϵ^* by bounding $\Pr\{l(t) > 0\}$ with ϵ^*. We observe that the first term, $\Pr\{x^c(t) > 0\}$, in the minimum expression of Eq. (21) is independent of the battery size. We can therefore set battery size to be zero whenever the first term forms the minima. Intuitively this means that there is no need of a battery if the probability of power outage is less than ϵ^*.

From Lemma 1, If $\mathcal{G}$ is a statistical sample path envelope on the effective demand in the sense of Eq. (9) with bounding function ε_g, then using Eq. (21), we get

$$\begin{aligned} &\min\left(Pr\{x^c(t) > 0\}, \varepsilon_g \left(B - \max_{\tau \geq 0}(\mathcal{G}(\tau) - C\tau) \right) \right) \leq \epsilon^* \\ &\implies B \geq \left(\max_{\tau \geq 0}(\mathcal{G}(\tau) - C\tau) + {\varepsilon_g}^{-1}(\epsilon^*) \right) I_{(\Pr\{x^c(t)=1\} > \epsilon^*)} \end{aligned} \tag{22}$$

where I_{expr} is the indicator function, which is 1 if $expr$ is true and is 0, otherwise.

C. Battery sizing for limiting genset carbon footprint

For the battery-genset hybrid system, at any instant that the grid is unavailable, the demand can be either served by the battery or by the genset. It is desirable to reduce carbon emission from the genset by using the battery to store ‘greener’ energy produced by the grid and consuming it when an outage occurs. The optimum scheduling algorithm between the battery and the genset must minimize the carbon emission from the genset while keeping the loss of power below an acceptable threshold. Scheduling is trivial if the genset size is larger than the maximum (worst-case) demand load (i.e., $\max_t d(t)$). This is because the demand can be always met by using the genset and carbon emission is minimized by always scheduling energy from the battery whenever it is not empty. The scheduling, however, is non-trivial if the genset size is smaller than the maximum demand load. This is illustrated by the following example.

Suppose the aggregate load is fixed to 100kW for 5 successive hours when power from the grid is unavailable. Assume the battery is fully charged to its capacity of 100kWh at the beginning of the first hour and genset serves demand up to the rate of 80kW. If the battery is used merely to meet the load in the first hour, there will be a loss of power in the remaining 4 hours as genset cannot meet load > 80kW and the battery is already exhausted. On the other hand, if genset is being used to its capacity for all the 5 hours with battery supporting the remaining 20kW every hour, there will be no loss of power! Consequently, as this example shows, if genset size is less than the maximum load ($\max_t d(t)$) exhausting the battery before using the genset is not the optimal strategy since that may lead to a larger loss of power.

For simplicity, therefore, we can assume that the genset capacity is large enough to meet the maximum aggregate load (Assumption 2 in Section IV). This is reasonable because the marginal cost of increasing genset capacity is small compared to the storage battery.

The objective in battery sizing in the presence of genset is to keep the carbon emission below a certain threshold. The carbon emission is proportional to the cumulative demand that cannot be served by the battery when the grid is unavailable, i.e.,

$$\text{carbon emission} \sim \sum_t l(t) \ . \tag{23}$$

Using the analogy between queueing theory and the distribution power system, this quantity corresponds to the total loss in a finite-buffer queue. In spite of extensive efforts, the problem is still open for non-Poisson arrivals. The complexity of the problem arises from the fact that the total loss is a function of the number and length of the busy periods, which occurs in a total time interval. Liu and Cruz [22] show that a probabilistic upper bound on the total loss must account for the numbers and lengths of the busy periods leads to cumbersome formulations, which cannot be used in practice. Here we compute an upper bound on the expected value of the total loss (carbon emission) in a time interval of size T using Eq. (20) as follows:

$$
\begin{aligned}
& E\left[\sum_{t=1}^{T} l(t)\right] = \sum_{t=1}^{T} E[l(t)] \\
& \leq \sum_{t=1}^{T} E\Bigg[\min\Bigg(d(t)x^c(t), \\
& \qquad\qquad \max_{0\leq s<t}\left([D^e(s,t) - C(t-s) - B]_+\right)\Bigg)\Bigg] \qquad (24) \\
& \leq \sum_{t=1}^{T} \min\Bigg(E\left[d(t)x^c(t)\right], \\
& \qquad E\left[d(t)x^c(t) I_{\max([D^e(s,t)-C(t-s)-B]_+)>0}\right]\Bigg) \qquad (25) \\
& \leq \sum_{t=1}^{T} \min\left(E\left[d(t)x^c(t)\right], \right. \\
& \qquad\qquad \left. E\left[d(t) I_{\max([D^e(s,t)-C(t-s)-B]_+)>0}\right]\right) \qquad (26) \\
& \approx \min\Bigg(\sum_{t=1}^{T} E\left[d(t)x^c(t)\right], \\
& \Pr\{\max\left([D^e(s,t) - C(t-s) - B]_+\right) > 0\}. \sum_{t=1}^{T} E[d(t)]\Bigg), \qquad (27)
\end{aligned}
$$

where we use Eq. (20) to obtain Eq. (24). The first term in Eq. (25) is trivial as it is the first term in the minima of the previous line. The second term in Eq. (25) only accounts for the sign of the second term in the minima of Eq. (24) and sets the whole expression to zero if that term is zero, otherwise it returns the first term in Eq. (24). Next line uses the fact that $0 \leq x^c(t) \leq 1$. We assume that the two processes in the second term in Eq. (26) are independent to derive Eq. (27). This assumption holds for statistically independent increments processes, which is widely assumed in the literature (e.g., Kelly [19]). In addition, we numerically find in Section VII-B that the inaccuracy due to this approximation step is small. Numerically, we also find that the upper bound used to derive Eq. (25) from Eq. (24) is quite tight. More precisely, we observe that Eq. (20) evaluates to its first term if the second term is positive.

Suppose $\mathcal{G}$ is a statistical sample path envelope on D^e with bounding function ε_g. From Eq. (27), the following is an upper bound on the expected total fuel consumption.

$$\min\Big(\sum_{t=1}^{T} E\left[d(t)x^c(t)\right], \\ \varepsilon_g\left(B - \max_{\tau\geq 0}(\mathcal{G}(\tau) - C\tau)\right) . \sum_{t=1}^{T} E[d(t)]\Big) \tag{28}$$

This formulation can be used for battery sizing when the objective is not to exceed a certain average fuel consumption.

VI. Constructing an envelope for the effective demand

To evaluate our derivations, we use half hourly electricity consumption data from more than 4500 Irish homes produced as part of CER Smart Metering Project [9]. We assume that all homes belong to the same region. The effective demand from Eq. (16) is derived from randomly selected 100 homes in this dataset. Since we could not find an available trace for grid availability, we use a simple model to represent that process. We discuss that in the following section.

A. Modeling grid availability process

We model the grid availability process by an ON-OFF Markov model, where the grid is available in the 'ON' state and is unavailable in the 'OFF' state. The transition rates between $ON \rightarrow OFF$ and $OFF \rightarrow ON$ are, respectively, λ and μ. With these parameters, on average, the grid spends $\frac{\lambda}{\lambda+\mu}$ and $\frac{\mu}{\lambda+\mu}$ fraction of time, respectively, in the OFF and ON states. The stability condition in our problem enforces the long-term average rate of the demand during an average-length OFF state to be less than the average charging that the battery receives from the grid during an average-length ON state. That is, if L represents the long-term average rate of the demand (i.e., $L = \frac{\sum_{t=1}^{T} d(t)}{T}$) and $\rho' = \frac{\lambda}{\mu}$, then

$$\frac{\lambda}{\lambda+\mu} L \leq \frac{\mu}{\lambda+\mu} C \implies \rho' L \leq C \ . \tag{29}$$

In our numerical studies we consider one of two cases: (1) We assume that we have a measurement trace of the effective demand (2) We assume that we have a multi-state Markov model for the demand.

Recall that the effective demand is the product of the demand (for which we have a measurement trace) added to the battery charging rate and the grid unavailability process. For the former case, where we need a measurement trace for the effective demand, we construct sample trajectories using the ON-OFF grid availability process. We construct a separate grid availability trajectory for each trajectory of the demand. To have a large enough dataset, we simulate 100 different sample paths of effective demand of a housing complex of 100 homes.

For the second case, where we assume a multi-state Markovian (MSM) process for the effective demand, we use the approach from [1]. As effective demand is zero whenever the grid is available ($x^c(t) = 0$), we first classify only the demands at all time slots with power outages into M Markovian states using the k-means clustering algorithm (we choose $M = 5$). Then, we add a new state representing availability of the grid. The emission rate of this state is 0 and we are at this state whenever the grid is available. We add C to the old values of the emission rates of the other states. The transition rate, q_{ij}, from state i to state j for such an $M + 1$ multi-state Markov chain can be calculated from the dataset using

$$q_{ij} = \frac{\text{Number of transitions from state } i \text{ to state } j}{\text{Total time spent in state } i}$$

B. Choosing parameters for $\mathcal{G}$ and ε_g

We use a leaky-bucket envelope $\mathcal{G}(t) = \sigma + \rho t$ with design parameters ρ, σ (≥ 0) as a statistical sample path envelope for the effective demand in the sense of Eq. (9) with bounding function ε_g. The design parameters in this modeling are the values of σ and ρ and the choice of distribution for ε_g. Empirically, we try different values of the leaky bucket parameters to find the minimum upper bound on the optimal battery size using Eq. (22), i.e.

$$B \geq \min_{\rho \leq C; \sigma} \left(\sigma + {\varepsilon_g}^{-1}\left(\epsilon^*\right)\right) I_{(\Pr\{x^c(t)=1\} > \epsilon^*)} \tag{30}$$

For the bounding function ε_g, we recall that this is the CCDF on the event $\max_{0 \leq s \leq t}([D^e(s,t) - \mathcal{G}(t-s)]_+)$. We observe that there is a large fraction of the elements evaluated as zero in that event. This happens with probability p_0 and consists of the cases where the arrival process D^e does not exceed the leaky bucket envelope. Let $\delta_0(x)$ be the delta function, which is 1 when $x = 0$, and is 0, otherwise. We can now write ε_g using the delta function as:

$$\varepsilon_g(x) = p_0 \delta_0(x) + (1 - p_0)\varepsilon_\delta(x)$$

$$\text{or,} \qquad \forall x > 0: \qquad \varepsilon_g(x) = (1-p_0)\varepsilon_\delta(x) \tag{31}$$

where ε_δ is a CCDF function with $\varepsilon_\delta(0) = 1$ [4]. Eq. (30) can now be rewritten as:

$$B \geq \min_{\rho \leq C;\sigma} \left(\sigma + {\varepsilon_\delta}^{-1} \left(\frac{\epsilon^*}{1-p_0} \right) \right) I_{(\Pr\{x^c(t)=1\} > \epsilon^*)} \tag{32}$$

We first tried fitting an exponential distribution to the function ε_δ, but the distribution failed to pass the Kolmogorov-Smirnov (KS) fitting test (we always use the default MATLAB parameters, i.e. hypothesis rejected at 0.05 significance level). Therefore, we use either of the following more complicated distributions:

1) **Weilbull distribution:** For parameters a and b (> 0), Weibull distribution is given by

$$\varepsilon_\delta(x) = abx^{b-1}e^{-ax^b}$$

We use the default MATLAB function *wblfit* to find the parameters a, b, and find that the distribution passes the KS fitting test.

2) **Hyper-exponential distribution:** Although Weibull distribution is a good fit for the tail bound distribution, because of its heavy-tailed property, it is less preferable than hyper-exponential distribution. We therefore try fitting a two-phase hyper-exponential distribution with three parameters, p, β_1, β_2, i.e.

$$\varepsilon_\delta(x) = pe^{-\beta_1 x} + (1-p)e^{-\beta_2 x}$$

To fit hyper-exponential distribution to dataset X, we use the standard approach (pg. 143 [21]) of empirically trying different values of parameter p and finding the corresponding values of parameters β_1 and β_2 by matching the first and the second order moments, i.e.

$$E[X] = \frac{p}{\beta_1} + \frac{1-p}{\beta_2}; \qquad E[X^2] = \frac{2p}{{\beta_1}^2} + \frac{2(1-p)}{{\beta_2}^2} \tag{33}$$

We perform KS fitting test to check the validity of the obtained parameters.

The approach used to model D^e in this section can also be employed to model the demand thus facilitating a generic performance analysis for intermittent demand energies.

[4]For convenience, we overload the term bounding function to also refer to ε_δ

VII. Numerical Examples

In the following section, we compare our battery sizing approach with the optimal benchmark for the same dataset as in Section VI. We use the Markov-modulated ON-OFF process as described in Section VI-A to generate grid unavailability trace. Unless otherwise stated, we set the parameters to $\mu = 1\ hr^{-1}$ and $\lambda = \frac{1}{11}\ hr^{-1}$, which correspond to an average of two power outages in a day with average duration of a power outage being one hour, which is common for developing countries such as India. Moreover, unless otherwise stated, the target violation probability in examples is assumed to be one day loss of power in ten years, i.e. $\epsilon^* = 2.7 * 10^{-4}$ and battery charging rate $C = 100$kW.

A. Battery sizing for eliminating genset

In this section we evaluate our analytical results on the required battery size to eliminate the use of genset. The accuracy of our analysis depends on the accuracy of loss of power formulation from Section V-B and our effective demand modeling from Section VI. We use the following curves in our evaluations:

- *Dataset quantile*: This is the battery sizing obtained by using the effective demand measurement trace as the input to the exact recursive loss of power description in Eq. (3). The battery size obtained by this method is the least value that satisfies the target loss of power probability for exactly known future demand.
- *Our bound*: Here, we use the parameter fitting techniques as explained in Section VI-B. We use leaky-bucket as the statistical sample path envelope on D^e and fit a hyper-exponential (or Weibull) distribution on ε_δ. Using Lemma 1 for loss of power formulation, we find the least battery size that satisfies the target loss of power probability (as given in Eq. (22)).
- *Ideal D^e model*: Here, we apply the effective demand measurement trace to Eq. (35) for loss of power formulation and find the least battery size that satisfies the target loss of power probability. By avoiding to use Lemma 1 in this approach, we remove the inaccuracy imposed by D^e modeling (i.e. from Eq. (35) to Eq. (36)). This allows us to distinguish between the inaccuracy that is induced by D^e modeling and that from Lemma 1.
- *Kesidis bound*: Here, we first compute a deterministic peak-rate constrained leaky-bucket description for the effective demand process in the sense of Eq. (7). We then use the Kesidis

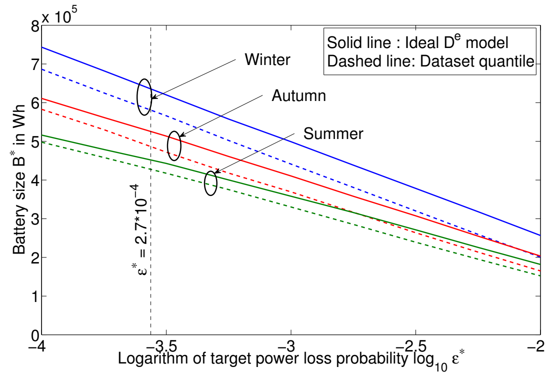

Fig. 4: **Accuracy of loss of power formulation:** *Dataset quantile* and *Ideal D^e model* battery sizes as a function of target violation probability in three seasons

battery overflow probability from Eq. (8) for loss of power formulation and find the least battery size that satisfies the target loss of power probability.

- *MSM bound*: We obtain the statistical sample path envelope for the arrival process D^e by first modeling it as a multi-state Markovian process (Section VI-A). Then we insert the statistical sample path envelope for Markovian processes from Section III-B3 in the loss of power formulation in Lemma 1 to obtain bounds on the battery size.

To study the accuracy of our analysis and possible sources of inaccuracy we must consider two major issues (1) how accurately we could model the effective demand using envelopes, (2) how accurately our formulation can compute the loss of power or carbon emission values given that there is no inaccuracy in modeling the effective demand. We study each issue in detail as they are also important for our second problem of battery sizing while ensuring that the genset carbon footprint is below a threshold.

1) Accuracy of loss of power formulation: As we only consider two points in the minimization from Eq. (18) to Eq. (19), we examine the accuracy of the loss of power formulation. This evaluation is necessary as this upper bound has never been evaluated before, even for the looser bound used in [30].

The accuracy of our loss of power formulation can be examined by the comparison of 'dataset quantile' that uses the dataset trace and exact recursive loss equation with 'Ideal D^e model' (to remove the inaccuracy of demand modelling), which also uses the dataset trace but applies the dataset to our upper bound loss formulation in Eq. (35). Figure 4 illustrates this comparison as a function of the target loss of power probability ϵ^*. To model accurately the fluctuations of the demand we consider battery sizing for different weather seasons by classifying the demand dataset into three seasons: Winter (December-March), Summer (April-July), and Autumn (August-November). We compute the battery size for each season separately. Note that a battery size that satisfies the loss of power requirements throughout the year would be the maximum of the battery sizes among all seasons.

We first observe from Figure 4 that for one day in ten years target loss of power probability (i.e., $2.7 * 10^{-4}$), 'Ideal D^e model' is within 10% of 'dataset quantile' implying that our loss of power formulation is reasonably tight. We also notice that different seasons can have significantly different battery size requirements (up to 35%) probably due to power-hungry heating appliances used in cold weather.

Our loss of power formulation in Eq. (21) consists of the minimum of two terms. The first term $\Pr\{x^c(t)\} > 0$ states that the loss of power formulation at any time instant cannot be larger than the power outage probability. This is a trivial bound, which becomes the dominant term in the minimization when the battery size is not large (to see the effect of this term one can consider the extreme case of battery-less scenario). The second term in Eq. (21) is quite accurate in identifying loss of power events, i.e. $\Pr\{l(t) > 0\}$ as can be seen in Figure 4 .

Finally, we note that while hyper-exponential distribution is sufficient to characterize seasonal demand as it passes KS-test, it seems to be an invalid distribution to describe the annual demand. This is due to seasonal changes of home-load, such as average demand, especially due to heating and cooling elements, which exhibit non-exponential, possibly heavy-tailed behaviour.

2) Accuracy of D^e modeling: We evaluate the performance of the fitting technique in modeling D^e by comparing the battery sizes obtained by this model and those from 'dataset quantile' and 'Ideal D^e model'. In fact, the difference between the battery sizing by envelope fitting with that of 'Ideal D^e model' indicates the accuracy of fitting since the 'Ideal D^e model' shows the battery sizing scheme when the D^e modeling inaccuracy is eliminated. We study both Weibull and hyper-exponential as tail bounds in the examples in this section. We also conduct a thorough sensitivity analysis of the above methods.

Figure 5 compares the battery size computed using the methods as a function of the violation probability. The similar slopes observed from this figure for different methods implies that the ratio of the battery sizes from each pair of methods is almost fixed even as the violation

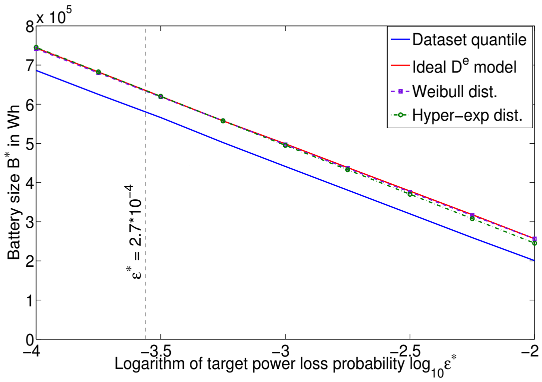

Fig. 5: **Accuracy of D^e modeling:** Minimum battery size B^* as a function of violation probability ϵ^*.

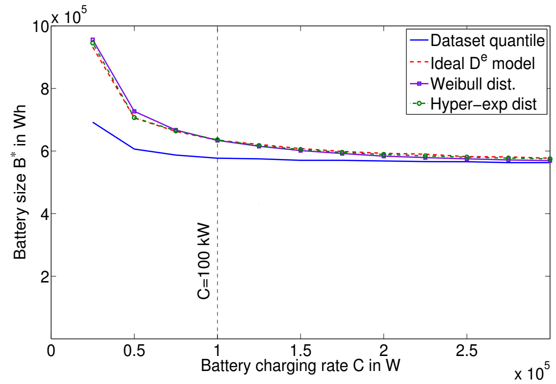

Fig. 6: **Accuracy of D^e modeling:** Minimum battery size B^* as a function of battery charging rate C for a fixed target loss of power probability $\epsilon^* = 2.7 * 10^{-4}$

probability is varying.

Figure 6 illustrates the battery sizing as a function of battery charging rate C with a target violation probability $\epsilon^* = 2.7 * 10^{-4}$. As the battery charging rate increases, all curves converge to the battery size required by an ideal battery, which can be charged instantaneously. The battery required for this ideal case must be large enough to serve the demand during the time when the grid is unavailable.

Finally, Figure 7 shows the battery sizing from different methods as a function of $\rho' = \frac{\lambda}{\mu}$ and for $\epsilon^* = 2.7*10^{-4}$. We set the average power outage duration to be one hour, i.e. $\mu = 1\ hr^{-1}$, and only vary λ. Due to limited space, we don't show error bars for 'Ideal D^e model' and 'Weibull distribution'. We observe that even for very high power outage periods, like $\rho' = 0.33$ (25% power outage), the bounds are within 15% of the 'Dataset quantile'. This sensitivity analysis indicates that our D^e modeling and loss of power formulations are quite tight.

3) Comparison with Kesidis and MSM: In this section, we evaluate the loss of power formulation from Lemma 1 with those obtained from other existing techniques: 'Kesidis bound', and multi-state Markov chain 'MSM'. We use 'Dataset quantile', and 'our bound' using hyper-exponential distribution as benchmarks. Figure 8 compares the battery sizes computed by the above methods as a function of the violation probability. We find that our approach outperforms the competing approaches. It can also be observed that the Kesidis bound is almost insensitive to the violation probability for the range of study. In addition, the Kesidis bound is comparably

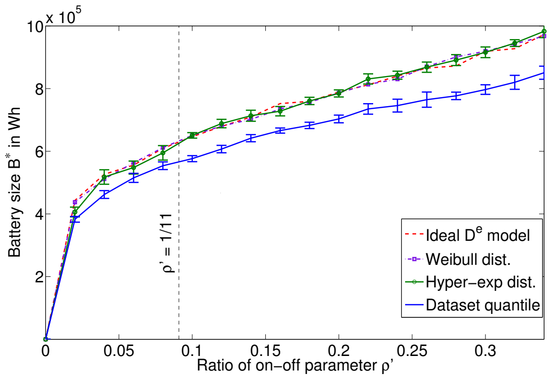

Fig. 7: **Accuracy of D^e modeling:** Minimum battery size B^* as a function of ρ' for a fixed target loss of power probability $\epsilon^* = 2.7 * 10^{-4}$. The bars show 95% confidence interval.

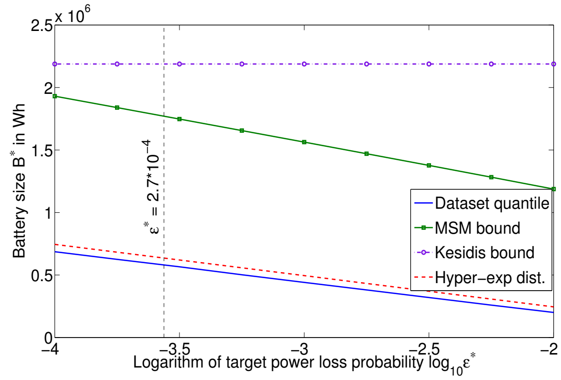

Fig. 8: **Comparison with other bounds:** Minimum battery size B^* as a function of violation probability ϵ^*.

loose since it is based on the assumption that the effective demand is regulated, which requires a deterministic peak-rate constrained leaky-bucket envelope on the effective demand. The tightness of Kesidis bound is highly affected by the tightness of deterministic envelope in describing the regulated traffic. The more bursty the traffic, the looser the bound becomes. Moreover, the MSM bound is also not as tight as the envelope fitting and this is due to the inaccuracy induced by employing the union bound to compute a statistical sample path envelope for the exponentially bounded burstiness processes (including MSM processes) as also observed in [7].

B. Battery sizing for limiting genset carbon footprint

In this section, we use all approaches to find the minimum the battery size that keeps the carbon footprint of a given genset below a certain threshold. Based on the results from Section VII-A2, we know that our effective demand modeling using envelopes is very accurate. Hence we do not repeat evaluating the accuracy of D^e modeling. The plots of this section show the following curves:

- *Dataset quantile*: Similar to the 'Dataset quantile' in the absence of genset, we apply the dataset trace to Eq. (3) and compute the exact carbon emission for our dataset.
- *Our bound*: This is the upper bound on carbon emission from Eq. (28) using envelope fitting as explained in Section VI-B to model D^e with hyper-exponential distribution on ε_δ.
- *Ideal D^e model*: We remove the inaccuracy induced by independence assumption in Eq. (27)

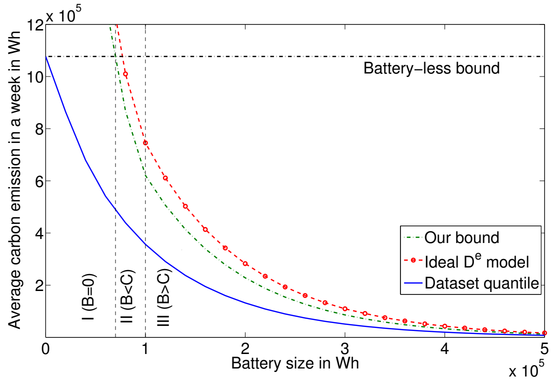

Fig. 9: **The impact of battery on carbon emission:** Carbon emission from genset as a function of battery size B

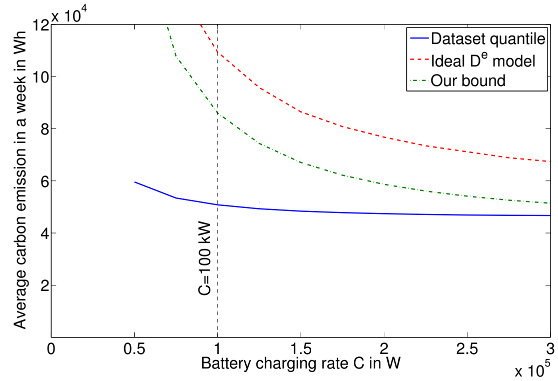

Fig. 10: **The impact of charging rate on carbon emission:** Carbon emission from genset as a function of battery charging rate C, and constant battery size $B = 3 * 10^5 Wh$

by applying the dataset trace directly to Eq. (26) to compute upper bounds on carbon emission.

We compute the carbon emission as a function of the battery size for charging rate of $C = 100$kW in Figure 9. For any given carbon emission threshold, one can use this plot to compute the required battery size. The plot is divided into three regions I, II, and III. Region I corresponds to the case where Eq. (27) is evaluated to $\sum_{t=1}^{T} E[d(t)x^c(t)]$, which is the battery-less scenario. If the carbon emission target threshold is as large as any value in this region we can remove the battery. Region II corresponds to the case when $C > B$. If the target carbon emission threshold falls into this region, one might choose the battery size to be equal to the energy generated in one time slot with the charging rate, i.e., $B = C$. Finally, Region III corresponds to the case of $C < B$ for which the curves in that region must be used to size the battery.

From Figure 10 we can see that our bound becomes more accurate for larger charging rates. Even for a relatively small charging rate like 100kW, we find that our bound is within a factor of 1.7 from the 'dataset quantile'. We can observe from Figures 9-10 that the carbon emission using Eq. (28) are slightly below the 'Ideal D^e model'. This is because of the independence assumption made for 'our bound' in Eq. (27).

An upper bound on expected total carbon footprint can be also obtained by using Kesidis or multi-state Markov chain to bound $\Pr\{\max([D^e(s,t) - C(t-s) - B]_+) > 0\}$ in Eq. (27). Such an approach, however, is extremely inaccurate (more than a factor of 10) as carbon footprint is directly proportional to any slackness in the loss of power probability. This is different to battery sizing that intuitively depends on the logarithm of slackness in loss of power probability (consider hyper-exponential bounding function in Eq. (32)). Hence we omit presenting those figures.

VIII. CONCLUSIONS

Motivated by the need of reducing carbon footprint of diesel generator operation to mitigate power outages, we present an analytical technique based on the stochastic network calculus for a housing complex to choose a battery size to trade-off carbon. We solve the problem in two parts. First, we study the problem of eliminating the use of genset and find the smallest battery size needed to ensure a given target loss of power probability. Numerical evaluations show that the sizing using our methodology is within 10% of the minimum battery size required had the future load been exactly known. In contrast, the battery size computed using classical methods is far larger than necessary.

Secondly, we study the trade-off between the size of battery and genset carbon emission. For a given battery size, our computation of the carbon emission is within a small factor (1.7) of the value obtained through numerical evaluation. This allows us to find the battery size needed to limit the carbon footprint of a diesel generator. Given that prior work in the area of teletraffic analysis has had limited success in computing upper bounds on the total loss of buffer for non-Poisson arrivals, we believe that our work is of general interest, even in the area of teletraffic analysis.

Our results are necessarily limited by the lack of demand and outage data from developing countries. We model the demand from a neighbourhood of homes in a developed country by the demand of 100 randomly selected Irish homes and we assume that outages are modeled by a two-state Markov model. These limit the strength of our numerical results. Nevertheless, our general approach can be used to study real datasets when they are available.

We are also interested in eliminating the need to assume statistically independent increments when studying the total loss of power. Finally, an interesting related problem is to work with more complex models of backup generation to analytically study how a battery can improve its efficiency.

APPENDIX

PROOF OF LEMMA 1

The loss probability at any time t satisfies the following

$$
\begin{aligned}
&\Pr\{l(t) > 0\} \\
&\leq \Pr\Bigg\{\min\Bigg(d(t)x^c(t), \\
&\quad \max_{u \leq s \leq t-1}\Big([D^e(s,t) - C(t-s) - B]_+\Big)\Bigg) > 0\Bigg\} && (34) \\
&\leq \min\Bigg(\Pr\left\{d(t)x^c(t) > 0\right\}, \\
&\quad \Pr\left\{\max_{0 \leq s \leq t-1}(D^e(s,t) - C(t-s) - B) > 0\right\}\Bigg) && (35) \\
&\leq \min\Bigg(\Pr\{x^c(t) > 0\}, \varepsilon_g\left(B - \max_{\tau \geq 0}(\mathcal{G}(\tau) - C\tau)\right)\Bigg) && (36)
\end{aligned}
$$

where we use Eq. (20) to derive Eq. (34). Eq. (35) is an upper bound on Eq. (34) using the fact that $P(X \cap Y) \leq \min(P(X), P(Y))$ for any events X and Y. Eq. (36) uses the assumption that $\mathcal{G}$ is a statistical sample path envelope for the process D^e.

REFERENCES

[1] O. Ardakanian, S. Keshav, and C. Rosenberg. Markovian models for home electricity consumption. In *Proceedings of ACM SIGCOMM Green Networking Workshop*, 2011.

[2] O. Ardakanian, S. Keshav, and C. Rosenberg. On the use of teletraffic theory in power distribution systems. In *Proceedings of ACM e-Energy*, 2012.

[3] O. Ardakanian, C. Rosenberg, and S. Keshav. On the Impact of Storage in Residential Power Distribution Systems. In *Proceedings of ACM SIGMETRICS Greenmetrics Workshop*, 2012.

[4] P. Arun, R. Banerjee, and S. Bandyopadhyay. Optimum sizing of battery-integrated diesel generator for remote electrification through design-space approach. *Energy*, 33(7):1155–1168, 2008.

[5] J. Boudec and P. Thiran. *A Theory of Deterministic Queuing Systems for the Internet Lecture Notes in Computer Science*. Springer, 2001.

[6] Isidor Buchmann. Charging the lead-acid battery. *Cadex Elec. Inc. https://www.classle.net/book/lead-acid-batteries*, Oct 2012.

[7] A. Burchard, J. Liebeherr, and F. Ciucu. On superlinear scaling of network delays. *IEEE/ACM Transactions on Networking (TON)*, 19(4):1043–1056, 2011.

[8] A. Burchard, J. Liebeherr, and S. D. Patek. A Min-Plus Calculus for End-to-End Statistical Service Guarantees. *IEEE Transactions on Information Theory*, 52(9):4105 – 4114, 2006.

[9] CER. Irish Social Science Data Archive (ISSDA). *http://www.ucd.ie/issda/data/commissionforenergyregulation/*, Feb 2013.

[10] C.S. Chang. *Performance guarantees in communication networks*. Springer Verlag, 2000.

[11] F. Ciucu and J. Schmitt. Perspectives on network calculus: no free lunch, but still good value. In *Proceedings of ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 311–322, 2012.

[12] Florin Ciucu, Almut Burchard, and Jrg Liebeherr. A network service curve approach for the stochastic analysis of networks. In *Proceedings of ACM SIGMETRICS*, pages 279–290, 2005.

[13] RL Cruz and H.N. Liu. Single Server Queues with Loss: A Formulation. *In Proceedings of CISS 93, Johns Hopkins University*.

[14] MA Elhadidy and SM Shaahid. Optimal sizing of battery storage for hybrid (wind+ diesel) power systems. *Renewable Energy*, 18(1):77–86, 1999.

[15] A. Elwalid, D. Heyman, TV Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers with applications to video teleconferencing. *IEEE Journal on Selected Areas in Communications*, 13(6), 1995.

[16] F. Galiana, E. Handschin, and A. Fiechter. Identification of stochastic electric load models from physical data. *Automatic Control, IEEE Transactions on*, 19(6):887–893, 1974.

[17] Y. Ghiassi-Farrokhfal and F. Ciucu. On the impact of finite buffers on per-flow delays in FIFO queues. In *Proceedings of ITC*, pages 1 – 8, September 2012.

[18] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer Verlag, 2008.

[19] F. P. Kelly. Notes on Effective Bandwidths. In *Stochastic Networks: Theory and Applications.Royal Statistical Society Lecture Notes Series, 4*, pages 141 – 168. Oxford University Press, 1996.

[20] G. Kesidis and T. Konstantopoulos. Extremal shape-controlled traffic patterns in high-speed networks. *IEEE Transactions on Communications*, 48(5):813 – 819, 2000.

[21] L. Kleinrock. Queueing Systems. Volume 1: Theory. 1975.

[22] H. Liu. *Buffer Size and Packet Loss in a Tandem Queueing Network*. PhD thesis, Unversity Of California, San Diego, 1993.

[23] A.C. Lloyd and T.A. Cackette. Diesel engines: environmental impact and control. *Journal of the Air & Waste Management Association*, 51(6):809–847, 2001.

[24] S. Mao and S.S. Panwar. A survey of envelope processes and their applications in quality of service provisioning. *IEEE Communications Surveys & Tutorials*, 8(3):2–20, 2006.

[25] Ramya Winstead Mitch Peterson. Definition of effective metrics for evaluating the efficiency of hybrid power systems at telecom cell sites: A system-level approach for determining "tower energy efficiency". *Cummins Power*, Dec 2008.

[26] H.S. Ramakrishna Rahul Tongia, Mohd. Saquib. India Power Supply Position 2010. *Centre for Study of Science, Technology, and Policy CSTEP*, Aug 2010.

[27] A.L. Schmid and C.A.A. Hoffmann. Replacing diesel by solar in the Amazon: short-term economic feasibility of PV-diesel hybrid systems. *Energy Policy*, 32(7):881–898, 2004.

[28] Rahul Seth. personal communication, 2012. Sudhir Powering The Nation.

[29] T.H. Sloane, F. Mann, and H. Kaveh. Powering the last mile: an alternative to powering FITL. In *Proceedings of IEEE Telecommunications Energy Conference, 1997. INTELEC 97., 19th International*, pages 536–543, 1997.

[30] K. Wang, F. Ciucu, C. Lin, and S. H. Low. A Stochastic Power Network Calculus for Integrating Renewable Energy Sources into the Power Grid. *IEEE Journal on Selected Areas in Communications*, 30(6):1037 – 1048, July 2012.

[31] K. Wang, S. Low, and L. Chuang. How stochastic network calculus concepts help green the power grid . In *Proceedings of IEEE SmartGridComm*, pages 55 – 60, October 2011.

[32] K. Wu, Y. Jiang, and D. Marinakis. A stochastic calculus for network systems with renewable energy sources. In *Proceedings of IEEE Conference on Computer Communications Workshops (INFOCOM)*, pages 109–114, 2012.

[33] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Transactions on Networking*, 1(3):372 – 385, June 1993.