

Sizing Storage for Reliable Renewable Integration

Vivek Deulkar

Department of Electrical Engineering
Indian Institute of Technology Bombay,
Mumbai, India

Email: vivekdeulkar@iitb.ac.in

Jayakrishnan Nair

Department of Electrical Engineering
Indian Institute of Technology Bombay,
Mumbai, India

Email: jayakrishnan.nair@ee.iitb.ac.in

Ankur A. Kulkarni

Systems and Control Engineering,
Indian Institute of Technology Bombay,
Mumbai, India

Email: kulkarni.ankur@iitb.ac.in

Abstract—The inherent intermittency of wind and solar generation presents a significant challenge as we seek to increase the penetration of renewable generation in the power grid. Increasingly, energy storage is being deployed alongside renewable generation to counter this intermittency. However, a formal characterization of the reliability of renewable generators bundled with storage is lacking in the literature. The present paper seeks to fill this gap. We use a Markov modulated fluid queue to model the loss of load probability (LOLP) associated with a renewable generator bundled with a battery, serving an uncertain demand process. Further, we characterize the asymptotic behavior of the LOLP as the battery size scales to infinity. Our results shed light on the fundamental limits of reliability achievable, and also guide the sizing of the storage required in order to meet a given reliability target. Finally, we present a case study using real-world wind power data to demonstrate the applicability of our results in practice.

I. INTRODUCTION

Electric supply is an indispensable part of modern life and is thus required to meet extremely stringent requirements of reliability. Classically, loss of load has been caused due to operational reasons, such as a generator undergoing maintenance, grid conditions, such as the overdrawing of power, or due to extraneous circumstances, such as natural calamities. With increasing penetration of renewable generation, the natural variability of the output of these generators adds a new, supply-side cause for the loss of load. Fortunately, with the growing capacity of renewable generation, we are also witnessing a tremendous fall in storage prices. Thanks to this, an increasing number of renewable generators are countering their variability, not with conventional, fast-ramping generation, but rather with storage [1], [2]. Thus, we believe that the renewable generator of the future will not be a standalone renewable generator, but rather a renewable generator *bundled* with a battery.

Keeping in mind reliability as one of the central concerns of the electricity infrastructure, the introduction of a battery-renewable generator bundle raises some basic questions. To begin, how does one account for this bundle in calculations for system reliability? How does this reliability change with increasing variability of the renewable source (wind or solar)? How does this change with increasing capacity of the battery? If one targets a certain level of reliability, how much battery storage is required to attain this level? And finally, are there fundamental limitations on the performance of a bundle, in

the sense that are there levels of performance that are simply unattainable?

A moment's thought reveals that answers to these questions cannot be obtained by only considering one snapshot in time. To understand this, consider the hypothetical scenario where there is no battery and only a renewable generator attached to a constant load. Then the loss of load probability (LOLP) would be simply the probability that the instantaneous output of the generator drops below the load, which one could potentially calculate via meteorological data. However, introducing storage changes the picture dramatically. Even while the instantaneous output of the generator may drop to a low level, there may well be charge left in the battery to meet the load requirements, and thus, using the bundle, the load could still be met. But the battery is charged by the excess output of the renewable generator, whereby the charge in the battery at any time depends on the history of generation (and load) realized until that time. It is easy to see that characterization of the LOLP in this case is a nontrivial matter.

This paper develops an analytical framework for characterizing the LOLP of a battery-renewable generator bundle. Our framework yields crisp answers to the sizing questions raised above. For a target level of reliability, it provides order-optimal estimates of the minimum battery size one requires to meet that reliability level, in terms of the statistical properties of the renewable source and the load. It also reveals that there are hard impossibilities: for certain ranges of these statistical parameters, no amount of battery suffices to bring the LOLP to zero. These results could be applied in conjunction with a costing exercise to ascertain the right battery size to be bundled with a renewable generator. One could also potentially use our characterization of the steady-state LOLP within a larger calculation of network-level reliability.

We model the *net generation*, i.e., the renewable generation minus the demand, as a continuous time Markov chain evolving over a finite state space. The battery serves as a buffer of finite capacity that is charged at the available rate when the net generation is positive and is discharged at the deficit rate when the net generation is negative. The battery charging process is subject to 'boundary conditions': it cannot be charged above its capacity and cannot be discharged below zero. Any positive net generation produced when the battery is fully charged is unusable. The LOLP is then the long run fraction of time

the battery is discharged to zero. We find that when the *drift*, i.e., the steady-state average net generation, is negative, then a battery of any finite size results in an LOLP that remains bounded away from zero. In other words, the LOLP cannot be made arbitrarily small by choosing a large enough battery size when the drift is negative. However, when the drift is positive, the LOLP drops exponentially with battery size, allowing it to be made arbitrarily close to zero by choosing a suitably large battery size. The rate of decrease of LOLP with increase in battery size is dictated by a decay rate, which can be characterized as the smallest positive generalized eigenvalue of the rate matrix associated with the net generation. This decay rate characterization can in turn be used to estimate the battery size required to achieve a given target LOLP.

II. MODEL AND PRELIMINARIES

Consider a storage battery of capacity b_{\max} which is charged or discharged by a *net generation* process associated with rate $r(t) = g(t) - d(t)$, where $g(t)$ and $d(t)$ denote, respectively, the generation and demand at time t . The energy content of the battery, denoted by $b(t)$, evolves as a regulated process having upper cap b_{\max} and lower cap 0. Thus, $b(t)$ evolves as

$$\frac{d}{dt}b(t) = \begin{cases} 0 & \text{if } b(t) = 0 \text{ and } r(t) < 0, \\ 0 & \text{if } b(t) = b_{\max} \text{ and } r(t) > 0, \\ r(t) & \text{otherwise.} \end{cases} \quad (1)$$

Note that a fully charged battery cannot be charged further with a rate $r(t) > 0$. Similarly an empty battery cannot be discharged further with a rate $r(t) < 0$. Excluding these two boundary cases, the rate of change of the battery level is governed by the net generation rate $r(t)$. We assume that the rate $r(t)$ is dependent on the state of a background Markov process, which collectively captures supply (generation) side variability as well as demand side variability.

Let $\{X(t)\}$ denote the background Markov process. We assume that $\{X(t)\}_{t \geq 0}$ is an irreducible, time-reversible, continuous-time Markov chain (CTMC) over a finite state space S . For every state $i \in S$, we associate a *net generation rate* $r_i \in \mathbb{R} \setminus \{0\}$ with which the battery is charged or discharged. Thus, $r(t) := r_{X(t)}$, i.e., the rate of charging/discharging of the battery is a function of the state of the background CTMC $\{X(t)\}_{t \geq 0}$. It is easy to see now that $\{(b(t), X(t))\}$ is a Markov process that evolves over the state space $[0, b_{\max}] \times S$. Note that this model also captures charge/discharge rate constraints on the battery; these would simply be reflected in the range of values taken by the net generation rates r_i .

The above mathematical model, wherein the occupancy of a buffer (or battery) is modulated by a background Markov process, is referred to in the queueing literature as a Markov Modulated Fluid Queue (MMFQ); see [3], [4]. In this paper, we use a finite-buffer MMFQ model to analyse the reliability of a renewable generator bundled with a battery.

Next, we describe how to characterize the invariant distribution of the Markov process $\{(b(t), X(t))\}$, which then leads

to a characterization of the loss of load probability (LOLP). Note that we are assuming that the process $X(t)$ has no state i where the net generation rate is zero. This allows us to partition the state space S as follows: $S = S_+ \cup S_-$, where

$$S_+ = \{i \in S : r_i > 0\}, \quad S_- = \{i \in S : r_i < 0\}.$$

We assume that both S_+ and S_- are non-empty.¹

Let (b, X) denote the steady state of the Markov process $\{(b(t), X(t))\}$. We capture the invariant distribution of this process as follows:

$$F_i(x) = \mathbb{P}[b \leq x, X = i] \quad \forall i \in S, x \in [0, b_{\max}].$$

The invariant distribution is governed by the ODE

$$\frac{d}{dx}F(x) = R^{-1}Q^T F(x), \quad (2)$$

where $F(\cdot) = [F_1(\cdot), F_2(\cdot), \dots, F_{|S|}(\cdot)]^T$, Q denotes the transition rate matrix associated with the CTMC $\{X(t)\}$, and $R := \text{diag}(r_1, r_2, \dots, r_{|S|})$ (see [3], [4]).² The invariant distribution can now be computed using the following boundary conditions:

$$F_i(0) = 0 \quad \forall i \in S_+; \quad F_i(b_{\max}) = \pi_i \quad \forall i \in S_-, \quad (3)$$

where $\pi = (\pi_i, i \in S)$ denotes the invariant distribution of the CTMC $\{X(t)\}$.

The probability that the battery content is less than or equal to x in steady state is given by $\sum_{i \in S} F_i(x)$. This probability is of particular relevance for $x = 0$. Indeed, the quantity $\sum_{i \in S} F_i(0) = \sum_{i \in S_-} F_i(0)$ is the long run fraction of time the battery is empty, and is also the long run fraction of time that the demand remains unfulfilled. In other words, this is the loss of load probability (LOLP), i.e.,

$$\text{LOLP} = \sum_{i \in S_-} F_i(0).$$

The LOLP, which can only be expressed in closed form for very simple cases (see below), can be computed numerically by solving the ODE (2) using the boundary conditions (3). However, this computation does not provide insights into the structural dependence of the LOLP on the supplyside and demandside uncertainty (captured by the CTMC $\{X(t)\}$) and the capacity b_{\max} of the battery. In Section III, we analyse the large buffer asymptotics of the LOLP, which sheds light on the limits of reliability achievable in a given setting, as well as the storage capacity required to achieve a certain (small) LOLP target.

Finally, we define a quantity that plays a key role in the large buffer asymptotics, namely the *drift* associated with the supplyside and demandside uncertainty. The drift Δ is defined as the steady state average net generation, i.e.,

$$\Delta := \sum_{i \in S} \pi_i r_i.$$

¹Indeed, if either S_+ or S_- is empty, then the battery would forever remain completely charged or completely discharged.

²Since $r_i \neq 0$ for all $i \in S$, R^{-1} exists.

Note that $\Delta < 0$ (respectively, $\Delta > 0$) implies that the time-average generation is less than (respectively, greater than) the time-average demand.

We conclude this section by considering the special case where the background CTMC $\{X(t)\}$ has only two states. This simple scenario, which admits a closed form characterization of the LOLP, motivates the general large buffer asymptotics derived in Section III.

Two state example

Consider the special case $S = \{1, 2\}$, where the generation alternates between two values 0 and $g > 0$ while the demand takes a constant value $d \in (0, g)$. Specifically, we set $r_1 = -d$, $r_2 = g - d$. In this case,

$$Q = \begin{bmatrix} -a & a \\ b & -b \end{bmatrix} \quad R = \begin{bmatrix} -d & 0 \\ 0 & g - d \end{bmatrix},$$

where $a, b > 0$ are the state transition rates for the generation process.

In this case, the drift is given by $\Delta = \frac{ag-ad-bd}{a+b}$, and the LOLP can be shown to be

$$\text{LOLP} = \frac{-\frac{\Delta}{d}}{1 - \frac{ag-ad}{bd} \exp\left\{\frac{(a+b)\Delta}{(g-d)d} b_{\max}\right\}}.$$

It is easy to see that LOLP is a strictly decreasing function of b_{\max} . However, the limiting behavior of the LOLP as $b_{\max} \rightarrow \infty$ depends critically on whether the drift is positive or negative. When $\Delta < 0$, then

$$\frac{ad+bd-ag}{(a+b)d} \text{LOLP} \xrightarrow{b_{\max} \uparrow \infty} -\frac{\Delta}{d} > 0.$$

This means that the LOLP remains bounded away from zero for any finite b_{\max} . In other words, when the drift is negative, an LOLP less than $-\frac{\Delta}{d}$ is simply unattainable no matter how large the battery capacity. This is consistent with Theorem 1 in Section III, which establishes a positive lower bound on the LOLP for any battery size b_{\max} when the drift is negative.

On the other hand, when $\Delta > 0$,

$$\text{LOLP} \sim Ae^{-\lambda b_{\max}},$$

where $A = \frac{b}{a} \frac{(a+b)}{(g-d)} \Delta$ and $\lambda = \left(\frac{a}{d} - \frac{b}{g-d}\right) > 0$.³ This implies that when the drift is positive, the LOLP decays exponentially with the battery size, implying that an arbitrarily small LOLP target is achievable with a large enough battery. Moreover, we note that the decay rate λ is in fact the positive eigenvalue of $R^{-1}Q^T$. This is consistent with Theorem 2 in Section III, which establishes an exponentially decaying (in the battery size) upper bound on the LOLP when the drift is positive.

³We use $f(t) \sim g(t)$ to mean that $\lim_{t \rightarrow \infty} \frac{f(t)}{g(t)} = 1$.

III. LARGE BATTERY APPROXIMATIONS

In this section, we analyse the behavior of the LOLP as the battery size b_{\max} scales to infinity. Our results shed light on the feasibility of meeting reliability targets, and also guide the sizing of the battery required to meet a given reliability target.

As suggested by the two-state example in Section II, the asymptotic behavior of the LOLP as $b_{\max} \rightarrow \infty$ depends on whether the drift is positive or negative. Accordingly, we consider these cases separately. Due to space constraints, we omit the proofs of our results here; these can be found in the extended version of this paper [5].

A. Negative drift: Asymptotic LOLP lower bound

We now consider the case $\Delta < 0$, i.e., the time-average generation is less than the time-average demand. One would therefore expect that LOLP cannot be made arbitrarily small in this case. This is proved formally in Theorem 1, which also provides a lower bound on the LOLP that is achievable with any finite battery size.

Let $r_{\min} := \min\{r_i, i = 1, \dots, |S|\}$. Note that $r_{\min} < 0$, since we assume that S_- is non-empty; $|r_{\min}|$ is simply the maximum rate of discharge of the battery.

Theorem 1. *If $\Delta < 0$, then $\text{LOLP} > \frac{-\Delta}{-r_{\min}}$ for any value of b_{\max} . Moreover,*

$$\lim_{b_{\max} \rightarrow \infty} \text{LOLP} \geq \frac{-\Delta}{-r_{\min}},$$

with equality if $|S_-| = 1$.

Theorem 1 is a consequence of the law of large numbers for Markov chains. It states that when the steady state average demand exceeds the steady state average generation, then an LOLP less than or equal to $\frac{-\Delta}{-r_{\min}}$ is unattainable no matter how large a battery we deploy. Moreover, this bound is loose in general; it is tight when the background CTMC has only a single state of discharge.

Connection with the two state example: In the two-state example considered in Section II, note that $r_{\min} = -d$ and $|S_-| = 1$. In this example, when $\Delta < 0$, recall that indeed, $\text{LOLP} > \frac{-\Delta}{d}$, with $\lim_{b_{\max} \rightarrow \infty} \text{LOLP} = \frac{-\Delta}{d}$.

B. Positive drift: Asymptotic LOLP upper bound

We now consider the case $\Delta > 0$, i.e., the time-average generation exceeds the time-average demand. In this case, one might expect that it is possible, with a large enough battery, to store the excess generation when the instantaneous generation exceeds demand, and to use this stored energy to almost always fulfil the deficit when the instantaneous generation drops below the demand. Theorem 2 shows that this is indeed the case, and that the LOLP decays exponentially with the battery size (when the drift is positive).

Theorem 2. *If $\Delta > 0$, then*

$$\limsup_{b_{\max} \rightarrow \infty} \frac{\log \text{LOLP}}{b_{\max}} \leq -\lambda_c, \quad (4)$$

where λ_c is the smallest positive eigenvalue of $R^{-1}Q^T$.

Theorem 2 provides an upper bound on the LOLP that is exponentially decaying in the battery size when the drift is positive. Specifically, (4) implies that for any $\lambda \in (0, \lambda_c)$,

$$\text{LOLP} \leq e^{-\lambda b_{\max}}$$

for large enough values of b_{\max} .⁴ This ensures that any arbitrarily small LOLP target be achieved with a suitably large battery. Moreover, the above bound can be utilized to obtain a conservative estimate of the battery size required in order to meet a given (small) LOLP target; we address battery sizing in detail as part of our case study (see Section IV).

In fact, we conjecture that the bound in the statement of Theorem 2 is tight, i.e., we conjecture that

$$\lim_{b_{\max} \rightarrow \infty} \frac{\log \text{LOLP}}{b_{\max}} = -\lambda_c. \quad (5)$$

This is also consistent with the results of our simulation experiments (see Section IV).

Connection with the two state example: Recall that in the two state example considered in Section II, we saw that when $\Delta > 0$, $\text{LOLP} \sim Ae^{-\lambda_c b_{\max}}$, where λ_c is the only positive eigenvalue of $R^{-1}Q^T$. Note that the decay rate in this case matches the bound in Theorem 2 exactly.

IV. CASE STUDY

In this section, we demonstrate the applicability of the results presented in Section III in practice. We fit a Markov model to a real-world trace of wind power generation, allowing us to validate the predictions from our analytical results against empirical observations. Further, we address the question of battery sizing in order to meet a given reliability target.

A. Data collection

We collected time series data corresponding to three years of wind power generation (December 2014 to December 2017) within the jurisdiction of the Bonneville Power Administration (BPA) (see [6]). The data samples are five minutes apart, and ranges from 0 to 4500 MW.

As expected, the data is highly non-stationary in nature, exhibiting diurnal as well as seasonal variations. Since our Markov modeling is best suited to stationary data, we extracted the samples corresponding to the months of February and March from 9pm to 3am for fitting a Markov model; this restricted dataset is henceforth referred to as the ‘stationary wind data’. For comparison, we also fit a Markov model to the entire (highly non-stationary) time series.

B. Data processing and Markov modeling

We now describe how we fit a Markov model to the above wind data.⁵ We first quantize the data into $N = 20$ bins, the

⁴The existence of at least one positive eigenvalue of $R^{-1}Q^T$ is guaranteed when $\Delta > 0$; this is proved in [4].

⁵This has been attempted before by several authors, including [7]–[10]. However, these prior works evaluate the ‘fit’ quality of their Markov models using the mean and auto-correlation function. In contrast, we match the reliability implied by the Markov model against the empirical reliability, which is a more direct indicator of the usefulness of the model.

bin edges being (in MW): [0, 60, 120, 180, 240, 300, 450, 600, 900, 1200, 1500, 1800, 2100, 2400, 2700, 3000, 3300, 3600, 3900, 4200, 4500]. This non-uniform binning is done to ensure a roughly even distribution of samples across bins. The N bins constitute the state space for our Markov model.

Given this state space, we obtain the empirical transition probability matrix T as follows:

$$T[i, j] = \frac{\# \text{ transitions occurring from bin } i \text{ to bin } j}{\text{total } \# \text{ transitions occurring out of bin } i}$$

T is the maximum likelihood estimator of the transition probability matrix corresponding to a discrete-time Markov chain (DTMC) model for the wind power sampled at $\tau = 5$ min intervals. To obtain a continuous-time Markov chain (CTMC) description, we note that the transition rate matrix Q of the CTMC is related to T as follows: $T = e^{Q\tau}$. Using the first-order Taylor series approximation for small τ , we get $e^{Q\tau} \approx I + Q\tau$, where I is the identity matrix.⁶ Accordingly, we set $Q = (T - I)/\tau$. This Q matrix defines a CTMC description of the wind power data.

To define the net generation corresponding to each state, we assume a constant demand d over time. Thus, the net generation rate r_i corresponding to bin i equals $g_i - d$, where g_i denotes the bin-center corresponding to bin i . Note that we can control the drift Δ by varying d .

C. Evaluating the goodness of fit

We now evaluate the quality of our Markov models by comparing the LOLP implied by these models with the empirical LOLP implied by the data. This also allows us to demonstrate the applicability of the conclusions of Theorems 1 and 2 in practice. In Figures 1 and 2, we plot the LOLP as a function of the battery size b_{\max} setting $d = 1800$ MW ($\Delta < 0$) and $d = 1200$ MW ($\Delta > 0$), respectively. We do this for the ‘stationary wind data’ as well as the entire time series. Specifically, we plot the following quantities:

- *Simulated (cont. time) LOLP:* This is the LOLP computed by simulating the CTMC model for wind power generation obtained from the data.
- *Simulated (discrete time) LOLP:* This is the LOLP computed by simulating the DTMC model for wind power generation obtained from the data, taking the generation to be constant over 5 minute intervals.
- *Empirical LOLP:* This is the LOLP computed by simulating the battery evolution using the wind power generation trace, again assuming the generation to be constant over 5 minute intervals.

Note that in all the plots, the simulated LOLP from our CTMC model closely matches the simulated LOLP from the DTMC model. This essentially validates our first order Taylor approximation for fitting the transition rate matrix Q from the empirical transition probability matrix T . Moreover, we note that the simulated LOLP from the Markov models more

⁶This Taylor approximation is valid to long as τ is smaller than the typical transition times of the CTMC.

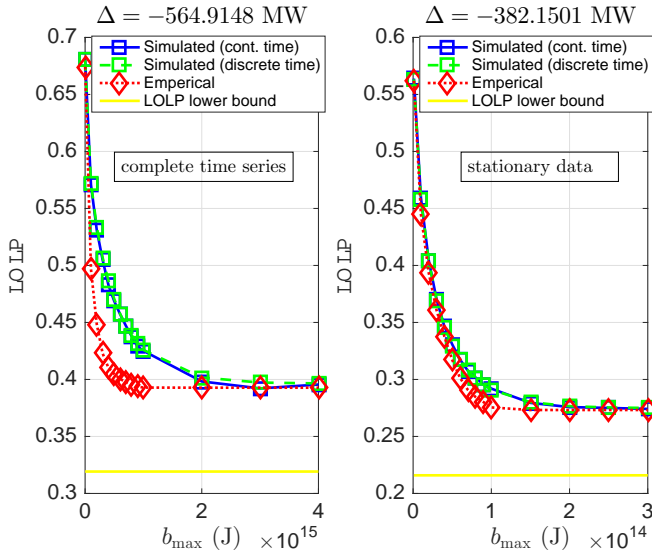


Fig. 1: LOLP vs battery size for $d = 1800$ MW ($\Delta < 0$)

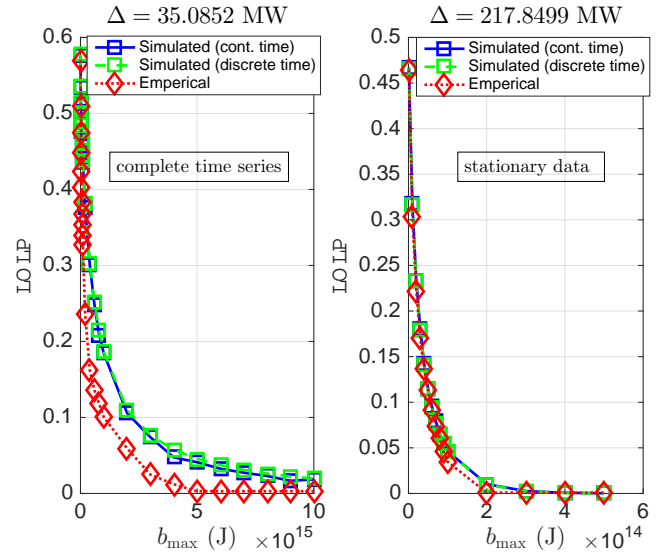


Fig. 2: LOLP vs battery size for $d = 1200$ MW ($\Delta > 0$)

closely matches the empirical LOLP for the stationary wind data than for the entire time series. This suggests that the Markov models are a better fit on the stationary data than on the complete, highly non-stationary time series. In practice, this means we should fit different Markov models to capture wind variability in different parts of the day in each season.

Focusing specifically on Figure 1, which corresponds to the negative drift scenario, we make the following observations.

- The empirical as well as simulated LOLP converges, as b_{\max} becomes large, to a value which is lower bounded by the bound specified in Theorem 1.
- The empirical LOLP is less than the LOLP implied by the Markov models. In other words, our models tend to overestimate the LOLP.
- The LOLP corresponding to a given battery size⁷ is greater for the entire time series as compared to the stationary data, suggesting that the former dataset is more ‘variable’ than the latter.

Focusing next on Figure 2, which corresponds to the positive drift scenario, we note that the LOLP decays to zero as b_{\max} becomes large, consistent with Theorem 2. Moreover, we see that the Markov models tend to overestimate the LOLP (as before). To illustrate the exponential decay of LOLP with battery size clearly, we plot the simulated LOLP from the CTMC model on a log-linear scale in Figure 3. Note that the plot looks asymptotically linear (establishing the exponential decay), with a slope that closely matches the decay rate bound from Theorem 2. This is consistent with our conjecture on the tightness of the bound of Theorem 2.

⁷We plot battery size in SI units (Joules). However, the engineering practice is to measure battery capacity in kiloWatt-hour (kWh), where $1 \text{ kWh} = 3.6 \times 10^6 \text{ J}$.

D. Battery sizing

The above results support our conjecture (see (5)) that when $\Delta > 0$, the LOLP decays exponentially with battery size with a decay rate equal to λ_c . This suggests when b_{\max} is large, the LOLP may be approximated as

$$\text{LOLP} \approx ce^{-\lambda_c b_{\max}}. \quad (6)$$

This further implies that the battery size required to maintain the LOLP at δ is given by

$$b_{\max} \approx \frac{\log(c)}{\lambda_c} + \frac{\log(1/\delta)}{\lambda_c}.$$

Since the pre-factor c in (6) is unknown here, a natural approximation would be to estimate the battery size required as

$$b_{\max} \approx \frac{\log(1/\delta)}{\lambda_c}. \quad (7)$$

Clearly, we would expect the above estimate to be accurate upto an additive offset. Moreover, we would expect that the error of our estimate would be small in relative terms for small δ .

To validate (7), consider the CTMC model for the stationary wind data, with $d = 1200$ MW. For this model, we compare the minimum storage size required to bring the simulated LOLP below δ with the estimate (7); see the left panel of Figure 4. Notice the constant offset between the two curves, as predicted. However, we note the (unknown) offset results in a roughly 40% error in battery size requirement when $\delta = 10^{-3}$. For lower values of δ , the relative error would of course be smaller. This means that for moderate values of reliability target δ , the estimate (7) can be used to make ballpark estimates of the storage size required.

However, (6) can also be used for *relative* storage sizing as follows: Note that (6) suggests that shrinking the LOLP be a factor of ϵ would require an increase in battery size

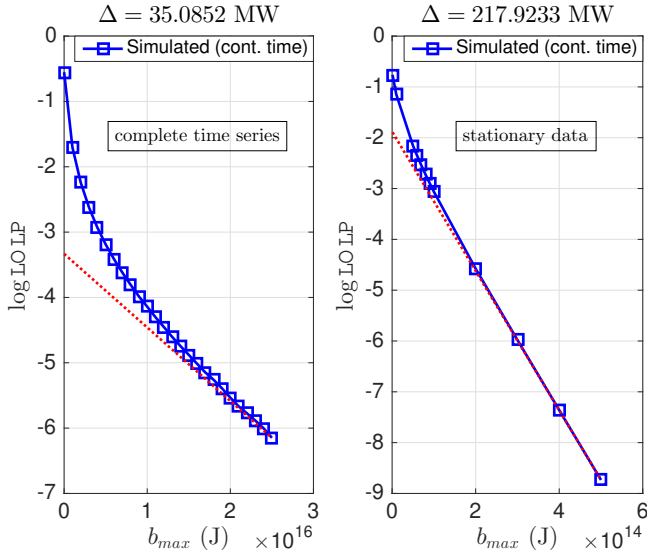


Fig. 3: Simulated LOLP vs battery size plot on log-linear scale for for $d = 1200$ MW (positive drift). The dotted red line has slope $-\lambda_c$.

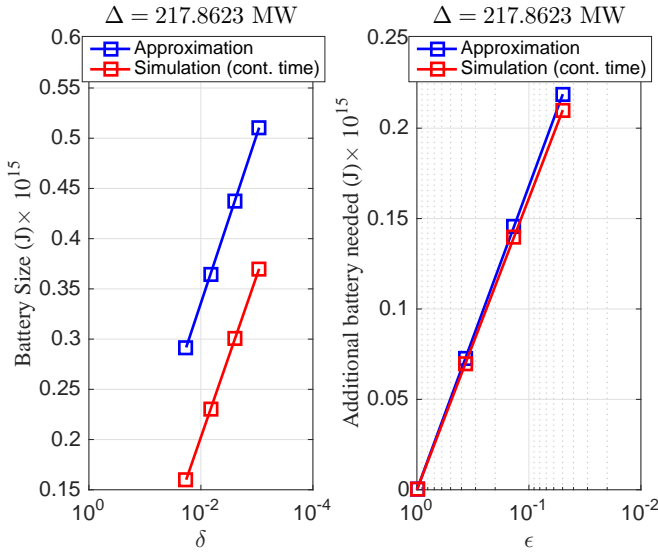


Fig. 4: Validation of battery sizing approximations

of $\frac{\log(1/\epsilon)}{\lambda_c}$. To validate this approximation, we consider the following baseline scenario. Setting $d = 1200$ MW with the stationary wind data, and $b_{\max} = 0.25 \times 10^{12}$ J, the simulated LOLP equals $L = 0.018$. In the right panel of Figure 4, we plot the *additional* battery size required to make the LOLP L/ϵ versus ϵ , using the above approximation, as well as by simulating the CTMC model. Note that the approximation is remarkably accurate, even for moderate values of LOLP.

This shows that (6) is an accurate description of the LOLP as b_{\max} becomes large, and can be used in practice to guide battery sizing decisions.

V. CONCLUDING REMARKS

In this paper, we developed an analytical framework for characterizing the reliability of a renewable generator bundled with a battery. We analysed how the reliability, captured by the LOLP, scales as the battery size scales to infinity. Our results highlight the achievable limits of reliability, and provide useful guidelines for sizing storage in practice.

This work motivates future research along several directions. A first step would be to refine our large-buffer asymptotics in the positive drift case, proving formally the conjecture (5). Indeed, characterizing the exact asymptotics of the LOLP, rather than logarithmic asymptotics, would result in even more accurate storage sizing prescriptions. At a higher level, we believe our formulations are a natural first step to analyse the economies of scale that would result from sharing of storage, between renewable generators or electricity prosumers. Finally, we note that our work motivates more sound stochastic modeling of renewable generation, to improve the real-world applicability of analytical reliability characterizations (as in the present paper).

ACKNOWLEDGEMENTS

The authors acknowledge support from the DST grant DST/CERI/MI/SG/2017/077, and the Bharti Centre for Communication at IIT Bombay.

REFERENCES

- [1] E. S. Association, "U.S. Energy Storage Project Pipeline Doubles in 2018, Nears 33 GW," <http://www.energystorage.org/news/esa-news/us-energy-storage-project-pipeline-doubles-2018-nears-33-gw/>, 2018, [Online; accessed 7-December-2018].
- [2] E. I. Administration, "U.S. Battery Storage Market Trends," <https://www.eia.gov/analysis/studies/electricity/batterystorage/>, 2018, [Online; accessed 7-December-2018].
- [3] D. Anick, D. Mitra, and M. M. Sondhi, "Stochastic theory of a data-handling system with multiple sources," *Bell System Technical Journal*, vol. 61, no. 8, pp. 1871–1894, 1982.
- [4] D. Mitra, "Stochastic theory of a fluid model of producers and consumers coupled by a buffer," *Advances in Applied Probability*, vol. 20, no. 3, pp. 646–676, 1988.
- [5] V. Deulkar, J. Nair, and A. A. Kulkarni, "Sizing Storage for Reliable Renewable Integration," 2019. [Online]. Available: www.ee.iitb.ac.in/jayakrishnan.nair/papers/PowerTech-TR.pdf
- [6] B. P. Administration, "Wind flow data," <https://transmission.bpa.gov/business/operations/wind/>, 2018, [Online; accessed 20-November-2018].
- [7] K. Brokish and J. Kirtley, "Pitfalls of modeling wind power using markov chains," in *Power Systems Conference and Exposition, 2009. PSCE'09. IEEE/PES*. IEEE, 2009, pp. 1–6.
- [8] H. Nfaoui, H. Essiarab, and A. Sayigh, "A stochastic markov chain model for simulating wind speed time series at tangiers, morocco," *Renewable Energy*, vol. 29, no. 8, pp. 1407–1418, 2004.
- [9] A. Shamshad, M. Bawadi, W. W. Hussin, T. Majid, and S. Sanusi, "First and second order markov chain models for synthetic generation of wind speed time series," *Energy*, vol. 30, no. 5, pp. 693–708, 2005.
- [10] G. Papaefthymiou and B. Klockl, "Mcmc for wind power simulation," *IEEE transactions on energy conversion*, vol. 23, no. 1, pp. 234–240, 2008.

APPENDIX

A. Proof of Theorem 1

The proof of Theorem 1 is based on energy conservation of the battery content together with the law of large numbers applied to the background Markov process. We use a following preliminary result in the proof.

Lemma 1. *Let $b(t)$ and $b_\infty(t)$ be finite capacity and infinite capacity battery content respectively at any instant t . Then, $b(t) \leq b_\infty(t) \forall t \geq 0$ on each sample path.*

The proof of this lemma is fairly simple which is based on the sample path argument and hence, we omit the proof.

Consider a time interval $[0, t]$. Let $\mathcal{O}_{cum}(t)$ be the total amount of energy due to battery overflow in the interval $[0, t]$ as the process $X(t)$ evolves. The overflow is because of the finite battery capacity. Similarly, let $\ell_{cum}(t)$ be the total amount of energy equivalent to lost load in the same interval. Load lost condition occurs when the battery is empty and any unserved load (or demand) which is present causes the discharging rate ($r_{X(t)} < 0$). Let $d_{cum}(t)$ be the total drained out battery content in the interval $[0, t]$ while $g_{cum}(t)$ be total injection into the battery in the same interval. We define

$$\ell_{avg} := \lim_{t \rightarrow \infty} \frac{\ell_{cum}(t)}{t} \quad \mathcal{O}_{avg} := \lim_{t \rightarrow \infty} \frac{\mathcal{O}_{cum}(t)}{t}.$$

As defined in Section III-A, r_{min} represents the maximum battery discharging rate. Analogously, let r_{max} be the maximum charging rate of battery i.e. $r_{max} := \max\{r_i | i = 1, \dots, |S|\}$. With these notations we have the following result.

Lemma 2. *If the average drift is negative i.e. $\Delta < 0$ then $\ell_{avg} = -\Delta + \mathcal{O}_{avg} \geq -\Delta$. Moreover, $\lim_{b_{max} \rightarrow \infty} \ell_{avg} = -\Delta$.*

Proof: Applying energy conservation, we get

$$\begin{aligned} g_{cum}(t) - \mathcal{O}_{cum}(t) &= d_{cum}(t) - \ell_{cum}(t) + b(t) \\ \Rightarrow \frac{\ell_{cum}}{t} &= -\frac{(g_{cum}(t) - d_{cum}(t))}{t} + \frac{\mathcal{O}_{cum}(t)}{t} + \frac{b(t)}{t} \\ \Rightarrow \frac{\ell_{cum}}{t} &= -\frac{1}{t} \int_0^t r_{X(s)} ds + \frac{\mathcal{O}_{cum}(t)}{t} + \frac{b(t)}{t} \\ \Rightarrow \lim_{t \rightarrow \infty} \frac{\ell_{cum}}{t} &= -\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t r_{X(s)} ds + \lim_{t \rightarrow \infty} \frac{\mathcal{O}_{cum}(t)}{t} \\ &\quad + \lim_{t \rightarrow \infty} \frac{b(t)}{t}. \end{aligned}$$

Since the battery capacity is finite, $\lim_{t \rightarrow \infty} \frac{b(t)}{t} = 0$ and law of large numbers for Markov chains implies that $\lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t r_{X(s)} ds = \Delta$. Therefore we get

$$\ell_{avg} = \lim_{t \rightarrow \infty} \frac{\ell_{cum}}{t} = -\Delta + \mathcal{O}_{avg}.$$

Also $\mathcal{O}_{cum}(t) \leq r_{max} \times \{\text{total time spent in } r_{max} \text{ state with full battery content in the interval } [0, t]\}$. Dividing by t and taking limit $t \rightarrow \infty$, we get,

$$\mathcal{O}_{avg} \leq r_{max} \mathbb{P}[b = b_{max}] \quad (8)$$

$$\leq r_{max} \mathbb{P}[b_\infty \geq b_{max}] \quad (9)$$

where b is the stationary battery content corresponding to finite capacity battery content $b(t)$ whereas b_∞ is the stationary battery content corresponding to the infinite capacity battery content $b_\infty(t)$. Both types of batteries are governed by the same underlying Markov process $X(t)$. Note that we have used Lemma 1 in (8) to get (9).

When the drift is negative, as $b_{max} \rightarrow \infty$, $\mathbb{P}[b_\infty \geq b_{max}] \sim k \exp\{-\lambda_c b_{max}\}$ where $-\lambda_c$ is the largest negative eigenvalue (having minimum magnitude) of $R^{-1}Q^T$ and k is some constant (see [4]). This implies $\mathcal{O}_{avg} \xrightarrow{b_{max} \uparrow \infty} 0$. In other words,

$$\lim_{b_{max} \rightarrow \infty} \ell_{avg} = -\Delta + \lim_{b_{max} \rightarrow \infty} \mathcal{O}_{avg} = -\Delta. \quad \blacksquare$$

With this result we now prove Theorem 1.

Proof: Total energy equivalent to lost load in the interval $[0, t]$ is upper bounded as $\ell_{cum}(t) \leq -r_{min} \times \{\text{total time spent in } r_{min} \text{ state with empty battery in the interval } [0, t]\}$. Dividing by t and taking limit $t \rightarrow \infty$, we get,

$$\ell_{avg} = \lim_{t \rightarrow \infty} \frac{\ell_{cum}(t)}{t} \leq -r_{min} \text{LOLP}.$$

From Lemma 2, $\ell_{avg} \geq -\Delta$. Therefore,

$$\text{LOLP} \geq \frac{-\Delta}{-r_{min}}.$$

When the background CTMC has only a single state of discharge i.e., $|S_-| = 1$, then $\ell_{avg} = -r_{min} \text{LOLP}$. From Lemma 2, $\lim_{b_{max} \rightarrow \infty} \ell_{avg} = -\Delta$ which gives us

$$\lim_{b_{max} \rightarrow \infty} \text{LOLP} = \frac{-\Delta}{-r_{min}}$$

(when background process has single state of discharge). \blacksquare

B. Proof of Theorem 2

Proof: If we consider the reversed process $X^r(t)$ of $X(t)$ in which $r_{X^r(t)} = -r_{X(t)}$, then a process $X(t)$ with $\Delta > 0$ will have reversed process $X^r(t)$ with $\Delta^r < 0$. Reversed process $X^r(t)$ will have $Q^r = Q$ and $R^r = -R$.

Also time spent at empty battery condition, $b(t) = 0$, with process $X(t)$ when the drift is positive is equivalent to time spent at full battery condition, $b^r(t) = b_{max}$, with reversed process $X_r(t)$ when the drift is negative.

Therefore we have,

$$\begin{aligned} \text{LOLP} &= \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mathbb{1}_{\{b(s)=0\}} ds \\ &= \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mathbb{1}_{\{b^r(s)=b_{max}\}} ds \\ &\leq \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t \mathbb{1}_{\{b_\infty^r(s) \geq b_{max}\}} ds \\ &= \mathbb{P}(b_\infty^r \geq b_{max}) \end{aligned}$$

where b_∞^r is the stationary battery content with infinite capacity.

To find $\mathbb{P}(b_\infty^r \geq b_{max})$, we use the same argument which is used in the proof of Lemma 2 i.e. when the drift is negative, as

$b_{\max}^r \rightarrow \infty$, $\mathbb{P}[b_{\infty}^r \geq b_{\max}] \sim k \exp\{-\lambda_c^r b_{\max}^r\}$ where $-\lambda_c^r$ is the largest negative eigenvalue (having minimum magnitude) of $R^r{}^{-1}Q^r{}^T$ and k is some constant (see [4]). Since $R^r = -R$ and $Q^r = Q$, λ_c is the smallest positive eigenvalue of $R^{-1}Q^T$. This gives us

$$\lim_{b_{\max} \rightarrow \infty} \frac{\log(\mathbb{P}(b_{\infty}^r \geq b_{\max}))}{b_{\max}} = -\lambda_c.$$

Therefore when the drift is positive, we get

$$\lim_{b_{\max} \rightarrow \infty} \frac{\log \text{LOLP}}{b_{\max}} \leq -\lambda_c.$$

■