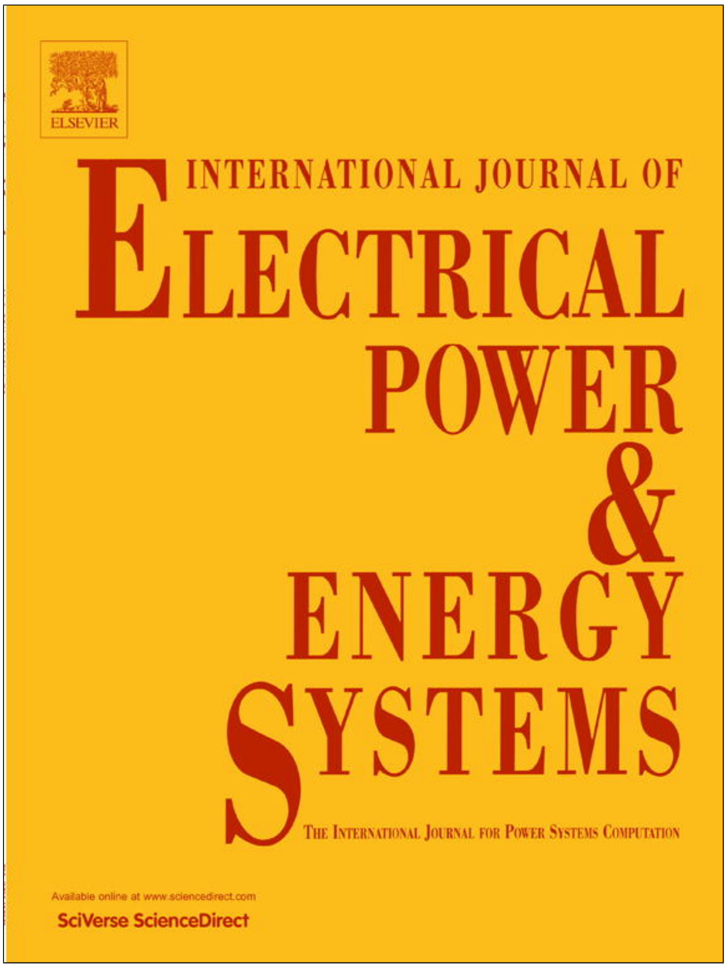
ELSEVIER
INTERNATIONAL JOURNAL OF
ELECTRICAL
POWER
&
ENERGY
SYSTEMS
THE INTERNATIONAL JOURNAL FOR POWER SYSTEMS COMPUTATION
Available online at www.sciencedirect.com
SciVerse ScienceDirect

Electrical Power and Energy Systems 44 (2013) 743–751

Contents lists available at SciVerse ScienceDirect

Electrical Power and Energy Systems

journal homepage: www.elsevier.com/locate/ijepes

Efficient simulation of blackout probabilities using splitting

John F. Shortle

Department of Systems Engineering & Operations Research, George Mason University, 4400 University Dr., MSN 4A6, Fairfax, VA 22030, United States

ARTICLE INFO

Article history:
Received 2 November 2011
Received in revised form 20 July 2012
Accepted 25 July 2012

Keywords:
Cascading blackouts
Rare-event simulation
Splitting

ABSTRACT

Standard Monte-Carlo simulation may be computationally intractable when the events of interest are extremely rare. This paper applies the rare-event simulation technique of splitting to the problem of estimating large-scale blackout probabilities. First, a stochastic model of cascading line failures is developed. Then, a simple network is presented and an analytical solution is derived for the simple network. Exploration of the analytical solution provides some guidance for setting splitting parameters in more complicated networks. In particular, geometrically increasing levels typically give an improvement over equally-spaced levels, due to the cascading nature of blackouts. The splitting methodology is applied to several different network topologies of varying complexity – a mesh network, a grid network, and the IEEE 118-bus network. Numerical results indicate that splitting has the potential to be effective on problems for which standard simulation may be infeasible.

1. Introduction

Monte-Carlo simulation is an important tool in investigating the reliability of power grids, since it can account for stochastic aspects of the system (e.g., [10]). However, standard Monte-Carlo simulation may be computationally intractable when the events of interest are extremely rare. This paper adapts a rare-event simulation methodology, called *splitting*, to achieve significant gains in computational efficiency to estimate large-scale blackout probabilities.

To illustrate the problem with standard Monte-Carlo simulation, consider an event that occurs with probability $\gamma = 10^{-9}$. On average, it takes 10^9 simulation runs to observe a single occurrence of the rare event. It takes many more runs to estimate the rare-event probability with a tight confidence interval. For example, to estimate the rare-event probability with a relative error[1] of 1%, 10^{13} replications are required. Supposing that 100,000 simulation replications can be conducted per second, the complete experiment would take about 3 years.

There are two main approaches that have been used to improve the efficiency of rare-event simulations: Importance sampling and splitting (e.g., [21]). The idea of importance sampling (e.g., [13,15]) is to change the sampling distribution so that rare events are more probable. For example, Bae and Thorp [2] use importance sampling to study hidden failures in a power grid. Another approach is splitting (see [16,17] for an overview). The basic idea of splitting is to create separate copies of the simulation whenever the simulation gets close to the rare event of interest. Effectively, this multiplies "promising" runs that are more likely to reach the rare event. One potential advantage of splitting is that (roughly speaking) it does not require a change to the model itself, while importance sampling requires going "inside" the model to change the sampling distributions. Thus, in the context of splitting, the cascading-blackout model can be developed separately and independently of the rare-event methodology.

Splitting has been used in a number of different application areas, such as queueing [12] and aviation safety [3], but it has not been applied to cascading blackouts. The objective of this paper is to (a) demonstrate that splitting can be effectively used on simple cascading-failure models and (b) gain insight on how splitting might be applied to more complex cascading-failure models.

The first part of this paper develops a stochastic model for cascading transmission-line failures. The model is similar to, and inspired by, a number of models in the literature – for example, Bae and Thorp [2], Carreras et al. [4], Chen et al. [7], Newman et al. [20], Wang et al. [26]. While the model is relatively simple, the spirit is to capture global blackout dynamics through representative behavior of key elements. Like other models in the literature, the model treats the intermediate failures of transmission lines as instantaneous events. (This is in contrast to time-based models, e.g., [1], in which line failures occur over time as a result of gradual changes in line temperature.) Following the failure of one line, or a set of lines, the power flows are recalculated, and the subsequent failures of the remaining lines are determined via a probabilistic model. The process repeats until there are no new failures, at which point the blackout ends. Although these models ignore many complexities, such as transient behavior (e.g., [19]) and the

E-mail address: jshortle@gmu.edu

[1] The relative error is the standard deviation of the estimator divided by its mean. The standard estimator is the number of observed rare events divided by the number n of simulation runs. The relative error for this estimator is $\sqrt{\gamma(1-\gamma)/n}/\gamma \approx 1/\sqrt{n\gamma}$, assuming at most one rare event can occur per simulation run.

0142-0615/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ijepes.2012.07.055

distribution network (e.g., [14]), the models are useful in understanding system-level effects. For example, a variety of these models (e.g., [7]) have explained the historically-observed power-law distribution of blackout size [25, chapter 7] and [5].

The next part of this paper applies the splitting methodology to several different network topologies of varying complexity. First, the methodology is applied to a very simple 2-node network connected by a set of identical links. For this simple network, the cascading blackout model can be solved analytically, which makes it possible to analytically derive the optimal parameters for the splitting methodology. While the network is extremely simple, results from the simple model give insight into application of splitting for more complex models.

The splitting methodology is then applied to more complex network structures – a mesh network, a grid network, and the IEEE 118-bus network. Numerical results indicate that splitting has the potential to be effective on problems for which standard simulation may be infeasible. The results also suggest how splitting might be applied to more complex models. Because of the cascading nature of blackouts in which line failures tend to accelerate, it is advantageous to choose levels within the splitting framework that increase in separation. In the numerical experiments conducted in this paper, geometrically increasing levels give an improvement over equally-spaced levels. The choice of the level function may be less critical than the locations of the levels. In most cases, a modified-allocation splitting scheme [22] gives an improvement over a standard equal-allocation splitting scheme.

2. Cascading failure model

This section describes a stochastic model for line failures in a power grid. The model is similar to, and inspired by, models found in Chen et al. [7] and Bae and Thorp [2]. One unique contribution of this model is the stochastic rule for a line failure, which depends not just on the present load, but also the largest previously observed load on the line (see discussion in Section 2.5).

Fig. 1 shows the basic logic of the cascading line-failure model. The simulation begins by randomly selecting a line to trip. Then, it cycles through a loop in which line failures in the previous iteration may lead to line failures in the next iteration. This cascading effect continues until there are no new line failures. The stochastic step in this loop is the probabilistic check for new line failures. The model is implemented in C++ and uses the random number generator from L'Ecuyer et al. [18]. The following sub-sections describe each part of the model in more detail.

2.1. Network connectivity

When a line fails, it is possible for the network to split into disconnected islands. This step determines the current number of islands and their associated network structures. Initially, there is only one island, the original network itself. Since disconnected islands are mathematically independent, subsequent steps in the loop must be repeated once for each island.

Island identification is accomplished via the following algorithm, which loops through the set of working lines. It takes the first working line and combines the two associated buses into a single island. Then it takes the next working line and combines the two associated buses into a single island. The algorithm continues through all working lines. Theodoro et al. [24] gives an alternate method of island identification, based on algebraic-graph methods.

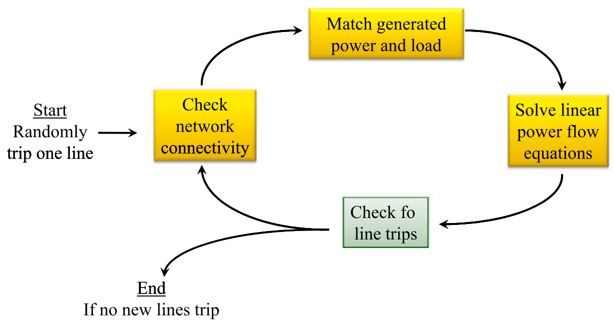

Fig. 1. Overall logic of cascading line-failure model.

2.2. Match generated power and load

This step matches total generated power and total load within each island of the network. In standard power-flow models, matching generation to load is accomplished via a "slack" bus. All other generators are assumed to have a fixed power output. In the context of cascading blackouts, the network frequently breaks up into separate islands, in which case a new slack bus must be designated for each island. Since it may not be obvious how to make this choice, it is simply assumed that each generator can freely and instantly adjust its generated power between zero and some node-dependent maximum.

There are two cases to consider when matching generation and load on an island. In the first case, when capacity exceeds load, the model sets the production at each generator. It does this by balancing production among the generators as uniformly as possible subject to capacity constraints. The following simple example illustrates the algorithm. Suppose the total load on an island is 18 MW and there are three generators with capacities of 4, 8, and 9 MW. To start, 4 MW of production are assigned to each generator, since 4 MW is the smallest capacity. This gives a total production of 12 MW. An additional 6 MW must be generated. This is divided evenly between the other two generators, so that generator 1 produces 4 MW; generators 2 and 3 produce 7 MW each.

When load exceeds capacity, the same idea is applied in reverse for load shedding. The following example illustrates. Suppose an island has four buses with loads of 2, 6, 10, and 12 MW (a total of 30 MW). Suppose the generation capacity of the island is 10 MW. Thus, 20 MW must be shed. It is not possible to shed load evenly among the buses (20/4 = 5 MW per bus), since the first bus has only 2 MW of load. Instead, 2 MW are initially shed from each bus, for a total of 8 MW. An additional 12 MW must be shed. This is divided evenly among the other three buses. The final results are: bus 1 sheds 2 MW; buses 2–4 shed 6 MW each. Some models in the literature (e.g., [4]) use a linear program to shed load.

2.3. Power-flow equations

The nonlinear power-flow equations are:

$$\begin{aligned} P_i &= \sum_{k=1}^{n} |V_i||V_k|[g_{ik}\cos(\theta_i - \theta_k) + b_{ik}\sin(\theta_i - \theta_k)], \\ Q_i &= \sum_{k=1}^{n} |V_i||V_k|[g_{ik}\sin(\theta_i - \theta_k) + b_{ik}\cos(\theta_i - \theta_k)], \end{aligned} \quad (1)$$

where P_i is the net real power injected into bus i, Q_i is the net reactive power injected into bus i, $|V_i|$ is the root-mean-square voltage at bus i, θ_i is the phase angle of the voltage at bus i, and g_{ik} and b_{ik} are the real and imaginary parts of element (i,k) of the network admittance matrix. The linear (or DC) power-flow equations are derived by assuming that θ_i is small, $|V_i| \approx 1$, and g_{ik} is small. Then (1) simplifies to:

$$P_i = \sum_{k=1}^{n} b_{ik}(\theta_i - \theta_k), \quad i = 1, \ldots, n. \tag{2}$$

Without loss of generality, we can assume that $\theta_1 \equiv 0$, since the power angles are relative. Typically, the n unknown variables in (2) are $\theta_2, \ldots, \theta_n$ and P_1. Node 1 is assumed to be a "slack" node, and P_1 is allowed to vary in order to achieve an exact balance between generated and consumed power (that is, to satisfy $\sum_i P_i = 0$). In our case, since the P_i have already been set to achieve this balance, we can remove one equation from (2) and then solve for $\theta_2, \ldots, \theta_n$. If the network consists of several disconnected islands, the power-flow equations are solved separately for each island. The power flow from bus i to j is

$$P_{ij} \equiv b_{ij}(\theta_i - \theta_j). \tag{3}$$

2.4. Line failure model

Line failures occur according to the following assumptions:

1. Each line j has a random capacity limit C_j chosen according to some probability distribution F_j.
2. Line j fails (with probability one) as soon as the load on the line equals or exceeds the capacity C_j. As long as the load remains below the capacity, the line remains in a working state (with probability one).

A stochastic model for line capacity (Assumption 1) can be motivated in several ways. First, line contact with vegetation is inherently stochastic: Vertical sag depends on stochastic factors such as wind and ambient temperature, and vertical vegetation growth can also be considered stochastic (e.g., [8]). Second, protective devices, which may operate correctly or incorrectly, can be a source of randomness in line capacity (e.g., [2,7], use hidden failures of protective relays as a basis for cascading blackout models). Finally, what matters is the *difference* between load and capacity. Here, capacity is treated as random, which means the difference is random. Thus, stochastic capacity might also be regarded as a surrogate for stochastic demand.

One way to implement the stochastic capacity model is to randomly generate all line capacities C_j at the start of each simulation replication. Following the initial trip of a line in the network, the simulation then evolves *deterministically*, because the line capacities C_j are already known. In other words, the stochastic elements are determined in the *initial* steps of the replication and the final outcome follows deterministically (through a somewhat complicated set of relationships in Fig. 1) from these initial conditions.

To make use of splitting techniques, we modify the implementation so that the simulation proceeds *stochastically* through each loop of Fig. 1. We can accomplish this by treating the *excess* capacity of a line as an undetermined random variable, conditioned on the line working in the previous iteration of the loop. That is, if x'_j is the largest previously observed power flow across line j and x_j is the present power flow across line j, then line j fails with probability $\Pr\left\{C_j \leqslant x_j | C_j > x'_j\right\}$.

2.5. Model summary

The following algorithm summarizes the model implementation:

1. *Initialization*:
   (a) Set all lines to the working state. Solve the power flow equations. Let x_j be the power flow (in magnitude) across line j.
   (b) Assume that $C_j \geqslant x_j$ for each j. (Line capacities are assumed sufficient to handle normal operating conditions when all lines are working.)
   (c) Choose a line at random to trip.

2. *Repeat*:
   (a) Set $x'_j = \max\left(x_j, x'_j\right)$, for each j. (The maximum function ensures that x'_j never decreases from one iteration to the next.)
   (b) Check network connectivity. For each island in the network:
      i. Match generated power and load.
      ii. Solve the power-flow equations.

   (c) Let x_j be the new power flow (in magnitude) across line j.
   (d) For each *working* line j such that $x_j > x'_j$, trip line j with probability

$$\Pr\left\{C_j \leqslant x_j | C_j > x'_j\right\} = \frac{F_j(x_j) - F_j\left(x'_j\right)}{1 - F_j\left(x'_j\right)}. \tag{4}$$

   (e) Exit loop if there are no new line trips (the blackout ends).

One distinctive feature of this model is that the failure probability of a line depends on both the present load x_j and the largest previously observed load x'_j. In particular, (4) implies that the overall survival probability of a line depends on its maximum load, but not on its intermediate sequence of loads. To illustrate, consider a line j that is working at a load of 10 MW and then undergoes a sudden increase in load to 20 MW. Applying (4), the probability that the line survives to a load of 20 MW is:

$$\Pr\{C_j > 20 | C_j > 10\} = 1 - \Pr\{C_j \leqslant 20 | C_j > 10\} = \frac{1 - F_j(20)}{1 - F_j(10)}.$$

Now suppose the line undergoes a different sequence of loads: 10 → 15 → 19 → 20 MW. The starting and ending loads are the same. The probability that the line survives to a load of 20 MW is the probability that the line survives each step. Multiplying (4) three times in succession gives:

$$\frac{1 - F_j(15)}{1 - F_j(10)} \cdot \frac{1 - F_j(19)}{1 - F_j(15)} \cdot \frac{1 - F_j(20)}{1 - F_j(19)} = \frac{1 - F_j(20)}{1 - F_j(10)}. \tag{5}$$

The overall survival probability is the same, independent of the intermediate steps. In contrast, a model in which the survival probability at each intermediate step depends only on the present load does not have this property.

In addition to all previously stated assumptions, we note the following assumptions and limitations: (1) Only line failures are considered; generator failures are not considered. (2) All generators can adjust output instantaneously within some specified minimum and maximum range. (3) Load shedding occurs in a reactive manner (when demand exceeds capacity on an island so the network is forced to shed load), but not in a proactive manner (e.g., to reduce the future potential for cascading failures). (4) Line failures occur instantaneously. (5) Loads are deterministic, and there are no time-of-day fluctuations. While the model is relatively simple, the objective is to see how to best use splitting for less complex models and then to generalize to more complex models in future work.

3. Splitting

This section gives a basic introduction to splitting in rare-event simulation. For a more extensive introduction, see Garvels [11], L'Ecuyer et al. [16], Rubino and Tuffin [21], Shortle and L'Ecuyer [23]. The explanation given here is in the context of power-grid blackouts, though the discussion can be generalized to other applications.

The basic idea of splitting is to generate independent copies of the simulation whenever it gets near the rare event – here, a major blackout. In this way, the simulation spends more time on runs that are more likely to reach the rare event. Let X_t represent the state of the power grid at time t. Depending on the model, X_t may include information like whether or not each line is working or failed, the load shed at each bus, and so forth. $\{X_t\}$ is assumed to be a Markov process, so the process can be simulated forward in time starting from X_t. (This assumption is not restrictive. If past history is needed, the necessary prior history can be imbedded within the state space.)

Let $h(\cdot)$ be a map from the state space to $\mathbb{R}^+$, denoting some measure of the size of the blackout. For example, throughout most of this paper, $h(X)$ is defined as the number of failed transmission lines when the system is in state X. $h(\cdot)$ is called the *importance function* or *level function*. A major blackout is defined as the set of states $\mathscr{R}$ whose level is at least as large as some constant $l > 0$. That is, $\mathscr{R} \equiv \{X : h(X) \geqslant l\}$. The simulation proceeds until either the rare-event set is reached or the blackout ends. When the blackout ends, the system is assumed to return to a state such that $h(X_t) = 0$. The probability to estimate is

$$\gamma \equiv \Pr\{\{X_t\} \text{ reaches } \mathscr{R} \text{ before } \{h(X_t)\} = 0\}.$$

The model in Section 2 starts with the random failure of one line. Thus, $h(X_0) = 1$. The simulation proceeds until $h(X_t) \geqslant l$ or until there are no new line failures, in which case it is assumed that the blackout ends and all lines are repaired, so $h(X_t) = 0$.

To implement splitting, define a sequence of M levels $0 < l_1 < \cdots < l_M$, where $l_M \equiv l$. Let D_j be the event that $\{X_t\}$ crosses level l_j before the blackout ends – that is $\{h(X_t)\} \geqslant l_j$ before $\{h(X_t)\} = 0$. Let $p_1 \equiv \Pr\{D_1\}$ and $p_j \equiv \Pr\{D_j|D_{j-1}\}$ for $j = 2, \ldots, M$. Roughly speaking, p_j is the probability of reaching level j starting from level $j-1$. Since $D_M \subset D_{M-1} \subset \cdots \subset D_1$,

$$\gamma = \Pr\{D_M\} = \Pr\{D_1\}\Pr\{D_2|D_1\}\cdots\Pr\{D_M|D_{M-1}\} = p_1 p_2 \cdots p_M.$$

The basic idea is to estimate each p_j separately, rather than estimating γ on its own.

There are a variety of ways to implement splitting. One approach, called *fixed-effort* splitting, pre-selects the number of runs at each level. Our implementation is an iterative variation of fixed-effort splitting. It works as follows (see also [22]). Let A_0 denote a set of possible starting states. For the model in Section 2, A_0 is a set corresponding to all possible single-line failures. Draw a state at random, with replacement, from the set A_0. From this starting point, simulate until either the system crosses level l_1 or the blackout ends. If the system crosses level l_1, add the terminating state to the set A_1. These states become the starting states for simulating from level l_1 to level l_2. More generally, a *stage-j* run consists of the following sequence: Draw a state at random, with replacement, from the set A_{j-1}; from this state, simulate until either the system crosses level l_j or the blackout ends; if the system crosses level l_j, add the terminating state to the set A_j. There may be duplicate copies of states in A_j.

The simulation proceeds by conducting one stage run at a time. Once a stage run is complete, the simulation must decide which stage j to simulate next. Two methods are used:

- *Equal allocation*. Each stage is given an (approximately) equal number of runs. The simulation conducts one stage-1 run, then one stage-2 run, and so forth, up to a stage-M run, and then it repeats. The simulation skips any stage j where A_{j-1} is empty, which is necessary in the early phases of the simulation, since initially A_0 is the only non-empty set.
- *Modified allocation* [22]. Following the completion of each stage run, determine the index j that maximizes:

$$n_j \hat{b}_j \left(1 + \frac{n_j \hat{p}_j}{1 - \hat{p}_j}\right), \quad j = 1, 2, \ldots, m, \tag{6}$$

where n_j is the number of stage-j runs conducted so far, $\hat{p}_j$ is the sample estimate of p_j and $\hat{b}_j$ is the sample estimate of the time to conduct a stage-j run, based on the simulation runs conducted so far. The maximization is conducted only over indices j for which A_j is non-empty.

The simulation stops as soon as the computing budget is used. An unbiased estimator for p_j is $\hat{p}_j \equiv |A_j|/n_j$, and an unbiased estimator for the rare-event probability is $\hat{\gamma} \equiv \hat{p}_1 \hat{p}_2 \cdots \hat{p}_M$.

The rule in (6) is an iterative implementation of an *optimal* allocation rule given in Shortle et al. [22]. The optimality of the rule was derived under the assumption that the probability of advancing from level $j-1$ to j does not depend on the starting state in A_{j-1}. This assumption does *not* hold for the models analyzed here. (To illustrate why, consider two different system states X_1 and X_2 such that $h(X_1) = h(X_2)$, where $h(X)$ is the number of failed lines in state X. Even though X_1 and X_2 are states with equal numbers of failed lines, the failed lines in X_1 may be more "critical" than those in X_2. Thus, the rare event may be more likely to result from X_1 than from X_2, violating the assumption.) Because (6) is not technically optimal for problems considered here, this paper refers to the rule as *modified-allocation* splitting rather than optimal splitting. Although the method is not technically optimal, it performs better than the equal-allocation rule for the examples in this paper (Sections 4 and 5). For other technical conditions and implementation issues, see Shortle et al. [22]. One slight modification here is that, as a warmup period, an equal-allocation scheme is used until at least one occurrence of the rare-event is observed, then the modified-allocation scheme is used. This helps to provide an initial base of samples with which to estimate p_j and b_j.

4. A simple network

This section considers a very simple network consisting of a generator and a load L connected by a set of N identical lines (Fig. 2). Although this example is extremely simple, the results provide insights for more complicated networks.

For this network, the model in Section 2 can be solved analytically as a Markov chain. The state of the chain is specified by two parameters (m,n), where n is the *present* number of failed lines and $m \leqslant n$ is the *previous* number of failed lines in the most recent pass through the loop in Fig. 1. Let y_n be the load on each working line when there are n failed lines – that is, $y_n = L/(N-n)$, which can also be obtained from (2) and (3).

To determine the transition probabilities, consider the system in state (m,n). Each working line is presently carrying a load $y_n = L/(N-n)$; each working line *was* carrying a load of $y_m = L/(N-m) \leqslant y_n$ in the previous iteration of the algorithm. Thus, the

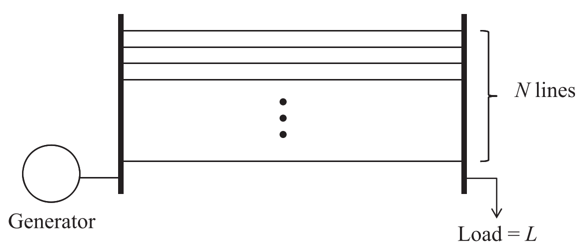

Fig. 2. Simple network model.

failure probability of any working line i is $\Pr\{C_i \leqslant y_n | C_i > y_m\} = [F(y_n) - F(y_m)]/[1 - F(y_m)]$; see (4). Let p_{mnj} be the single-step transition probability from state (m,n) to $(n,n+j)$, that is, the probability that exactly j lines fail in one step starting from state (m,n). The number of line failures during a single iteration is a binomial random variable, so

$$p_{mnj} = \binom{N-n}{j}\left(\frac{F(y_n)-F(y_m)}{1-F(y_m)}\right)^j\left(\frac{1-F(y_n)}{1-F(y_m)}\right)^{N-n-j}, \quad (7)$$

where $0 \leqslant m \leqslant n \leqslant N$ and $0 \leqslant j \leqslant N-n$. The blackout ends when there are no new line failures. In particular, (n,n) is a terminating state ($p_{nn0} = 1$ for each n).

To evaluate this Markov chain, let r_{mn} be the probability that the system ever reaches state (m,n). By assumption, $r_{01} = 1$, since the system is initialized with all lines working ($m = 0$) followed by the failure of a single line ($n = 1$). More generally, r_{mn} can be computed iteratively via the relation:

$$r_{mn} = \sum_{i=0}^{m-1} r_{im} p_{i,m,n-m}, \quad 0 < m \leqslant n \leqslant N. \quad (8)$$

The probability that the blackout terminates with exactly n failed lines is r_{nn}. The probability that n or more lines fail in a blackout is

$$R_n \equiv \sum_{j=n}^{N} r_{jj}. \quad (9)$$

In summary, the blackout-size distribution R_n is specified analytically via (7)–(9), though explicitly writing out the complete closed-form expression is cumbersome.

Consider a numerical example with $N = 100$, $L = 100$ MW, and exponential capacity distribution $F(x) = 1 - e^{-\lambda x}$ where $\lambda = .3296$. Each line carries 1 MW and the average line capacity is 1/.3296 $\approx$ 3.034 MW. (The parameter λ was chosen to yield a final rare-event probability of $\gamma \approx 10^{-10}$.) Fig. 3 shows R_n, calculated via (7)–(9). For small n, $\log R_n$ is roughly linear in n, indicating geometrically decaying probabilities. This means that line failures are roughly independent and that the failure of one line has a small impact on the potential failures of other lines. On the other hand, for large n, there are fewer working lines to carry the overall load, so the failure of one line is more likely to result in additional failures. This cascading effect accelerates as lines fail. For example, the failure of 80 lines is virtually guaranteed to result in the failure of all 100 lines, since Pr{100 lines fail | 80 or more lines fail} = $R_{100}/R_{80} \approx 1$.

To apply splitting to this problem, we first choose an importance function $h(\cdot)$. A natural choice is $h(m,n) = n$, that is, the number of presently failed lines. The rare-event set is then $\mathscr{R} \equiv \{(m,n) : h(m,n) \geqslant 100\}$. Next we choose the number M and locations l_j of the intermediate levels. For the sake of example, let $M = 10$. A natural choice for the locations of the levels is to distribute them evenly over [0,100]. That is,

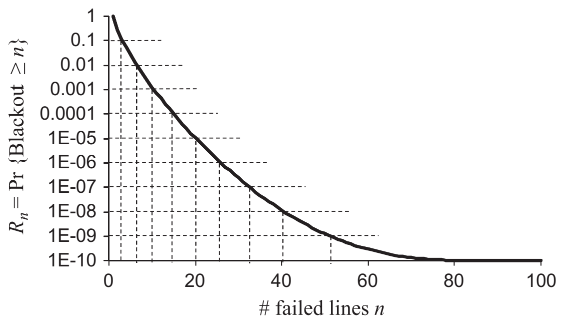

Fig. 3. Distribution of blackout sizes for simple network.

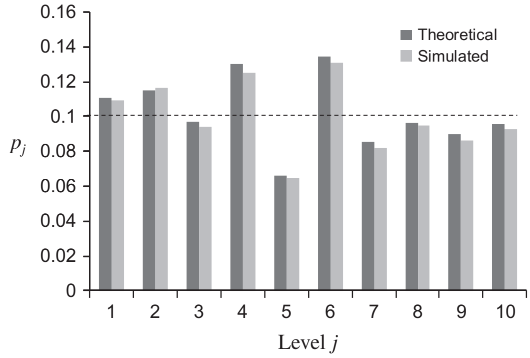

Fig. 4. Level probabilities p_j associated with (10).

levels $l_j = \{10, 20, 30, \ldots, 100\}$.

An alternate choice is to distribute the levels so that they are evenly spaced *in terms of probability*. Here, with $\gamma \approx 10^{-10}$ and 10 levels, the levels should be spaced apart *in probability* by a factor of about 10^{-1}. Fig. 3 shows this graphically. The levels are the values on the x-axis that achieve approximately equal spacing of the corresponding probabilities on the y-axis (on a log scale). For example, $n = 3$ corresponds to $R_n \approx 10^{-1}$, $n = 6$ corresponds to $R_n \approx 10^{-2}$, and so forth (dashed lines in the figure). The complete set of levels that achieves approximately equal spacing in terms of probability is

$$\text{levels } l_j = \{3, 6, 10, 14, 20, 25, 32, 40, 51, 100\}. \quad (10)$$

The spacing of the levels gets further and further apart. For example, the probability of advancing from 3 to 6 line failures is about the same as the probability of advancing from 51 to 100 line failures (roughly 10^{-1}). Fig. 4 shows the actual realized level probabilities. Since the levels must be integers, it is not possible to *exactly* achieve $p_j = 10^{-1}$. Sample simulated values are also shown, which simply helps to verify that the simulation experiment is implemented correctly.

Fig. 5 shows the results of several simulation experiments. Three different simulation techniques are used: standard simulation, equal-allocation splitting, and modified-allocation splitting (see Section 3). Two different level sets are used for splitting, as given previously. Each experiment represents 20 replications with a computing budget of 5 min per replication (100 min of simulation time per sample variance given in the figure). The variance is proportional to the time needed to achieve a given relative error.[2] Thus, a variance reduction by a factor of 10 equates to a speed improvement by a factor of 10.

The first observation from the figure is that standard simulation is inadequate for this problem. In the allotted simulation time, standard simulation did not generate *any* observation of the rare event. The second observation is that equally spaced levels *in terms of probability* (3,6,...) are better than equally spaced levels (10,20,...). This is particularly true when using equal-allocation splitting. The third observation is that modified-allocation splitting does better than equal-allocation splitting. In theory, if the level probabilities p_j are equal and the simulation-stage costs b_j are equal, then the modified-allocation rule in (6) reduces to the equal-allocation scheme. Here, in the case of levels {3,6,10,...}, the level probabilities are chosen to be *as equal as possible*, but they

[2] The variance of $\hat{\gamma}$ is inversely proportional to the computing time. This is true for standard simulation as well as fixed-effort splitting; e.g., see Shortle et al. [22], Section 5.1. Thus, the computing time to achieve a fixed relative error (standard deviation of the estimator $\hat{\gamma}$ divided by its mean γ) scales with the variance of the estimator.

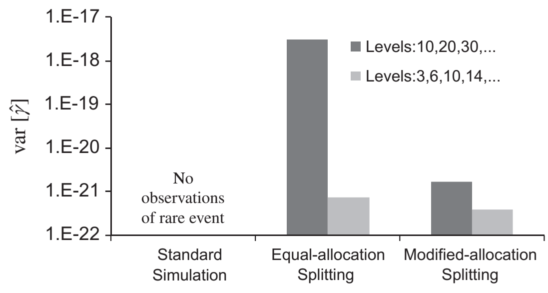

Fig. 5. Sample variances using splitting.

are not exactly equal (Fig. 4). This illustrates that the modified-allocation rule can still realize a significant benefit over an equal-allocation rule even when the level-probabilities are near equal.

We now consider an alternate choice for the level function $h(m,n)$. Choosing a good level function is generally regarded as a critical problem in splitting [17]. The level function should represent the "proximity" of the system to the rare event, meaning that the probability of reaching the rare event from any point on the same level contour should be roughly the same. The analysis so far has used $h(m,n) = n$, the number of presently failed lines. One problem with this choice is that the probability of reaching the rare event may vary significantly across a level contour. For example, the system is much more likely to reach the rare event from state (1, 60) than from (59, 60), even though both states correspond to 60 failed lines. This is because, in the first case, the system moves from 1 failed line to 60 failed lines in one step, so the load on the remaining 40 lines suddenly jumps from $100/99 \approx 1.01$ to $100/40 = 2.5$. In the second case, the load on the remaining 40 lines moves from 100/41 to 100/40, barely any change. Thus, the failure probability in (4) is much greater in the first case.

It is possible to analytically calculate the level function $h(m,n)$ that achieves (approximately) equal distance in probability to the rare event over each level contour. Let q_{mn} be the probability of reaching the rare event from state (m,n). Then q_{mn} can be calculated iteratively via:

$$q_{mn} = \sum_{k=n+1}^{N} p_{m,n,k-n} q_{nk},$$

where the boundary conditions are $q_{mN} = 1$ for $m \leqslant N$. Fig. 6 shows the resulting level function $h(m,n)$ that achieves approximately equal level probabilities, equally spaced in probability by a factor of 10^{-1}.

Repeating the same experiment using the modified levels in Fig. 6 gives the following results (Table 1): With equal-allocation splitting, the variance reduces by a factor of about 2.2; with modified-allocation splitting, the variance reduces by a factor of about 1.7. This shows that the new level function provides some variance reduction over the "obvious" choice of $h(m,n) = n$, though perhaps this difference is not as much as might be expected. There are two reasons that the choice of level function may not be so critical for this network: (1) The system tends to visit states near the diagonal. States like (1, 60) are virtually impossible to visit. Thus, the system acts more like a one-dimensional system in which the proximity to the rare event is effectively characterized by the second parameter. Thus, the dependence of the level function on the first parameter is not so important. (2) It is possible for the system to "skip over" a level contour as it up-crosses into a new level. This means that the entrance states into a given level do not all lie on the level contour. Thus, no matter what level function is used, it is impossible to create a scenario in which the probability of advancing from level $j-1$ to j is the same for every entrance state into level $j-1$. This may reduce the benefit of choosing the level function as above.

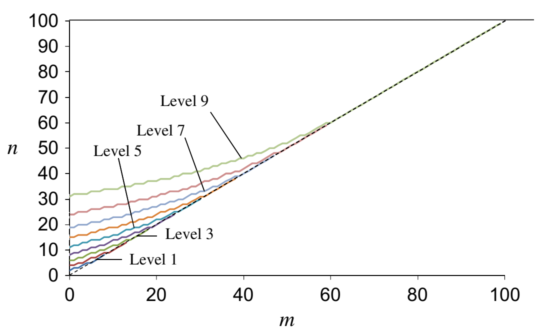

Fig. 6. Level function that yields approximately equal level probabilities within and between levels.

Table 1
Sample variance var[$\hat{\gamma}$] observed using different level functions.

Level function	Equal allocation	Modified allocation
$h(m,n) = n$	7.28×10^{-22}	3.89×10^{-22}
$h(m,n)$ from Fig. 6	3.37×10^{-22}	2.29×10^{-22}
Factor improvement	2.2	1.7

Finally, we note that the model given here is similar in spirit to an analytical model given in Dobson et al. [9]. In both models, the system state can be described with respect to the *number* of failed lines, without regard to the *location* of those lines within the network. A key difference is that, in Dobson et al. [9], a line failure results in the *same* load increase on all other lines, regardless of how many remaining lines are working. Here, the load increase is a function of the number of remaining working lines. Despite this difference, the blackout distributions (e.g., Fig. 3) of the two models have qualitatively similar shapes.

5. Test networks

This section considers a set of test networks of varying complexity. Examples include a mesh network, a grid network, and the IEEE 118-bus network.

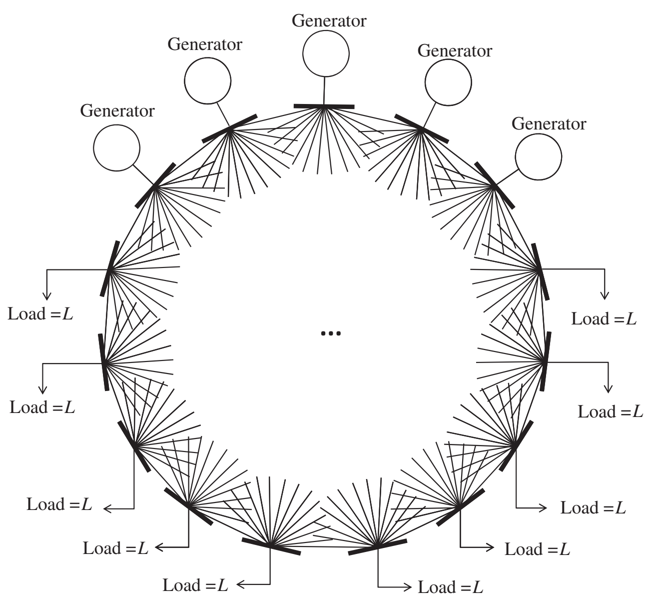

Fig. 7. 15-bus mesh network with 5 generators and 10 loads.

We start with a 15-bus mesh network in which each bus is connected to every other bus (Fig. 7), for a total of $(15^2 - 15)/2 = 105$ lines. The network has 5 identical generators and 10 identical loads as shown. The capacity of each line is assumed to follow a normal distribution with mean μ and standard deviation σ. The mean capacity μ is set equal to twice the largest power flow observed on any line when all lines are working. The standard deviation is set as $\sigma = 0.2\mu$. (The value of 0.2 is somewhat arbitrary. Roughly speaking, σ/μ is a ratio of "noise" to rated line capacity. Noise refers here to the variability in the *actual capacity* of a line, though it can also be thought of as a surrogate for variability in the *load* on a line. When $\sigma/\mu = 0.2$, there is roughly a 2.3% chance that line load exceeds actual line capacity, when the line load is 60% of its rated capacity, which is two standard deviations below the mean.)

Fig. 8 shows the simulated blackout-size distribution for this network. Similar to the behavior in Fig. 3, an accelerating or cascading effect is observed. A key difference is that Fig. 8 shows two ranges in which the blackout accelerates, with a brief "pause" in between. Roughly, what is happening is the following. Consider one of the load buses. It has 14 lines attached to it. If one of these 14 lines fails, the remaining 13 lines must carry the load of the failed line in order to meet the demand at that load bus. Thus, this set of 14 lines behaves somewhat like the simple network discussed in Section 4. That is, once a few of the 14 lines have failed, it is likely that all 14 lines will fail. Once this happens, the load bus is disconnected from the network. At this point, the cascade decelerates because the total load on the network drops. But as additional lines fail, the cascade accelerates again.

Fig. 9 shows the sample variance in estimating Pr{50 or more lines fail} $\approx 10^{-6}$. Three different simulation techniques are used as well as two different sets of 10 levels: Evenly spaced levels $\{5, 10, 15, \ldots, 50\}$ and geometrically spaced levels $\{2, 3, 4, 5, 8, 11, 16, 23, 34, 50\}$ ($l_j = \lceil 50^{j/10} \rceil$, where $\lceil \cdot \rceil$ is the ceiling function). Geometrically spaced levels are motivated as a qualitative approximation to the levels derived in (10) for the simple network. Geometrically spaced levels provide an improvement over evenly spaced levels, and modified-allocation splitting provides an improvement over equal-allocation splitting. Using both improves efficiency by a factor of about 25. Standard simulation is not well suited for this experiment. It only generates a few observations of the rare event. (By chance, no observations were observed in the geometric-level case.) The simulation budget is 8 min, and 60 replications are conducted for each simulation methodology (i.e., 8 h of computer time for each method).

The next example considered is the grid network in Fig. 10. Each bus has a load L. There are four generators near the middle of the grid. Carreras et al. [6] consider similar networks as well as a variety of other simple structures. In the example here, the line capacity parameter μ is set according to an "$N-1$" criterion. Specifically, μ is set equal to twice the largest power flow observed on any line, considering the set of all single-line failures. All lines have the same capacity distribution – a normal distribution with mean μ and standard deviation σ, where $\sigma = 0.2\mu$.

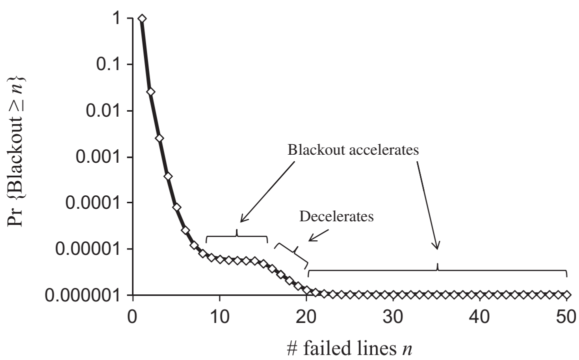

Fig. 8. Blackout distribution for 15-bus mesh network with 5 generators and 10 loads.

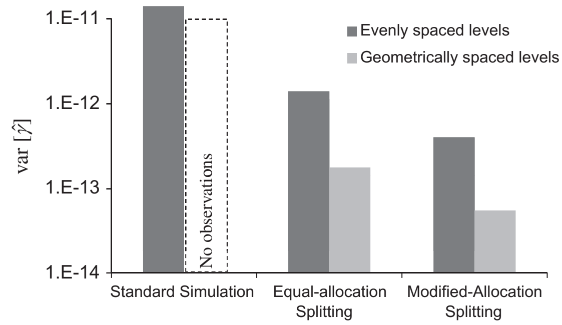

Fig. 9. Sample variances observed for 15-bus mesh network.

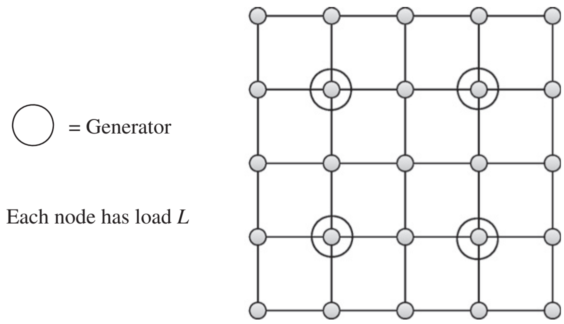

Fig. 10. Example 5 × 5 grid network.

Fig. 11 shows the simulated blackout-size distribution for the grid network. The left graph shows the distribution with respect to the number of failed lines. The right graph shows the distribution with respect to the number of buses without power (which is proportional to shed load, since the buses are assumed to have identical loads.) A loss of 21 buses represents a maximal system failure, since 4 buses are directly connected to generators which, by assumption, do not fail. The key observation from the figure is that the two blackout distributions are qualitatively very similar. This suggests that the number of failed lines may be a good surrogate for the amount of load shed.

One benefit of defining levels with respect to the number of failed lines is that it provides greater fidelity for splitting in the early stages of a blackout. If the level function is defined with respect to lost buses, the first possible event at which the simulation can be split is the loss of a single bus. This event occurs with probability 10^{-5} in the right graph (somewhat rare). There is no way to further break this event into intermediate less-rare events. On the other hand, if the level function depends on the number of failed lines, then it is possible to split the simulation earlier – say, after two lines have failed but perhaps prior to the loss of any load. Thus, there is some efficiency benefit in defining levels with respect to failed lines, though levels can also be defined with respect to shed load or a combination of failed lines and shed load.

Fig. 12 shows the sample variance in estimating Pr{32 or more lines fail} $\approx 10^{-10}$. Evenly spaced levels $\{4, 8, 12, \ldots, 32\}$ and geometrically spaced levels $\{2, 3, 4, 6, 9, 14, 21, 32\}$ ($l_n = \lceil 32^{n/8} \rceil$) are considered. The results are similar to those from the mesh network. Geometrically-spaced levels provide an improvement over evenly spaced levels (by roughly one order of magnitude), and modified-allocation splitting provides an improvement over equal-allocation splitting (by roughly one order of magnitude). In the case of stan-

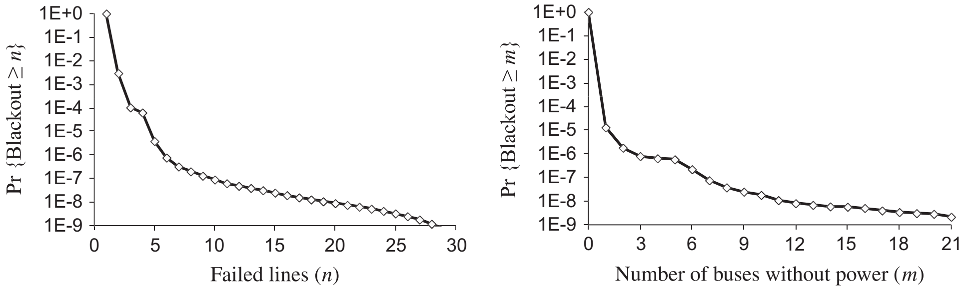

Fig. 11. Blackout distribution, by failed lines and number of buses without power.

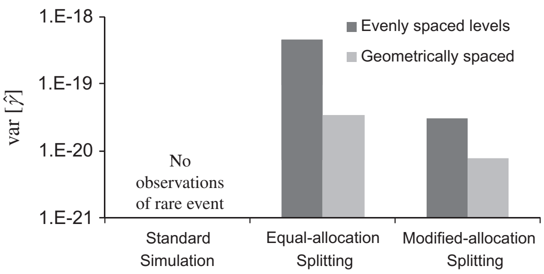

Fig. 12. Sample variances observed for 15-bus mesh network.

dard simulation, *zero* instances of the rare event are observed. This is because the rare-event probability is very low. The computing budget and replications are the same as before.

The final example is the IEEE 118-bus network. The network topology, bus loads, generator locations, and line parameters (conductance/susceptance) are obtained from http://www.ee.washington.edu/research/pstca/. To set line capacities, all line-capacity distributions are assumed to follow a normal distribution with mean μ and standard deviation σ, where $\sigma = 0.2\mu$. The average line capacity μ is set equal to (1/0.85) times the largest power flow observed on any line, considering the set of all single-line failures (i.e., in the worst-case single failure, no line is operating above 85% of its expected capacity.) Generator failures are not considered in this paper, so generator capacities are assumed to be sufficiently large to mimic infinite capacity.

Fig. 13 shows sample variances observed. In this experiment, the simulation budget is 10 min per replication, and 95 replications are conducted for each of the four cases in the figure. The rare event is defined as the failure of 16 or more lines (out of 179) with a probability on the order of 10^{-8}. Evenly spaced levels $\{4, 8, 12, 16\}$ and geometrically spaced levels $\{2, 4, 8, 16\}$ are considered. The results are similar to those from the grid network. In particular, geometrically-spaced levels provide an improvement over evenly spaced levels by roughly one order of magnitude.

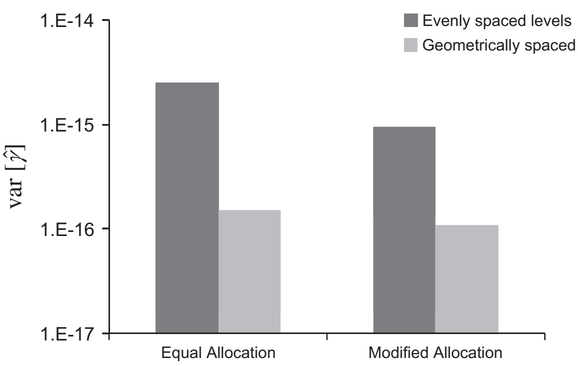

Fig. 13. Sample variances observed for IEEE network.

The modified-allocation scheme provides a slight benefit over the equal-allocation scheme, though the benefit is less compared with the previous examples. There are a couple of possible explanations. First, because of the larger network size, the warmup period for the modified-allocation scheme takes a relatively larger percentage of the computing budget (about 60% of the budget on average for this experiment). This reduces the difference between the two methods, since the warmup period for the modified-allocation scheme uses an equal-allocation scheme. Second, the modified-allocation scheme is not *guaranteed* to provide a reduction in variance. The scheme was derived under a simplifying assumption in which the probability of advancing from one level to the next does not depend on the starting state from the lower level [22]. This assumption does not hold for the problems in this paper. Network symmetry may help the modified-allocation scheme, because symmetry implies that some (but not all) states at a given level are equivalent in terms of the likelihood of reaching the rare event. The IEEE 118-bus network has no symmetry, which may reduce the effectiveness of the modified-allocation scheme compared with the earlier examples.

6. Conclusions

This paper applied the rare-event simulation technique of splitting to the problem of estimating large-scale blackout probabilities. First, we developed a stochastic model of cascading line failures, similar to a number of models in the literature. The model accounts for network topology, physical line characteristics (capacity, conductance, susceptance), demand at each bus, and generation capacities. Network characteristics outside of these inputs are ignored.

Second, we presented a very simple network and derived an analytical solution. Exploration of the analytical solution provided some guidance for setting splitting parameters in more complicated problems. Ideally, levels should be chosen so that the probabilities of advancing from one level to the next are equal. For the simple network, this corresponded to choosing levels l_j that grew farther and farther apart with increasing j. This principle translated well to more complex networks.

For the simple network, we also derived a level function $h(\cdot)$ with the property that each contour of the function contains states that are (approximately) equally likely to reach the rare event. This is ideal in the sense that equal emphasis is placed on paths that have the same probability of reaching the rare event. Without this property, splitting may inadvertently emphasize paths that are unlikely to reach the rare event (e.g., [12]). While this level function yielded an improvement in efficiency over the "obvious" level function, the improvement was not large. The spacing of the levels seemed to be more important than the choice of the level function. This may suggest that the choice of the level function is less critical for these types of blackout models compared with other rare-event applications. This is also a subject for future research.

Third, we applied splitting to a number of different network topologies – a mesh network, a grid network, and the IEEE 118-bus network. Results indicated that splitting has the potential to be effective on problems for which standard simulation may be infeasible. In all numerical experiments, geometrically increasing levels gave an improvement over equally-spaced levels. This is due to the cascading nature of the blackouts. In some cases, the improvement in efficiency was several orders of magnitude. We also applied two splitting schemes – an equal-allocation scheme and a modified-allocation scheme [22]. The modified scheme attempts to adjust the number of runs at each stage to account for differences in the level probabilities. In the examples considered, the modified-allocation scheme gave an improvement over the equal-allocation scheme, though the improvement was sometimes minor and this result may not universally hold.

To put the results of this paper in the proper context, the underlying blackout model in this paper is relatively simple. Many complexities of real blackouts are ignored. For example, transient effects are ignored (all computations assume steady state), equipment failures are ignored (only transmission-line failures are considered), generation failures are ignored, human errors and issues of situational awareness are ignored, and so forth. The objective here is not to conduct a detailed risk analysis of a real network. Rather the objective is to demonstrate the potential of using rare-event techniques on simulating cascading blackouts and to gain insights on the application of these techniques using simple models.

Future work may involve application of the concepts in this paper to more complex models to address some of the current limitations. Because the splitting methodology is implemented independently from the blackout model, more complex blackout models can be easily integrated with the splitting methodology. Future work may involve modeling various aspects of a smart grid, such as an increased presence of renewable energy sources, demand management, and self-healing controls in the grid. The spirit would be to capture representative behavior of key elements, rather than low-level complexities, and to understand the resulting impact on blackout dynamics. Another potential application of the methodology in this paper is capacity expansion. Consider the problem of determining the optimal expansion of transmission lines in order to minimize the probability of a large-scale blackout. Because the set of potential expansion decisions might be large, and because each decision requires a separate evaluation of a rare-event probability, a fast rare-event model is needed to find the best solution. The methodology in this paper can be directly imbedded within a stochastic optimization framework to address this problem.

Acknowledgments

This material is based upon work supported by the Department of Energy under Award Number DE-SC0002223.[3]

[3] This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

[1] Anghel M, Werley K, Motter A. Stochastic model for power grid dynamics. In: Fortieth Hawaii international conference on system sciences; 2007.

[2] Bae K, Thorp J. A stochastic study of hidden failures in power system protection. Decis Support Syst 1999;24(3–4):259–68.

[3] Blom H, Bakker G, Krystul J. Rare event estimation for a large-scale stochastic hybrid system with air traffic application. In: Rubino G, Tuffin B, editors. Rare event simulation using Monte Carlo methods. Chichester (UK): Wiley; 2009. p. 193–214.

[4] Carreras B, Lynch V, Dobson I, Newman D. Critical points and transitions in an electric power transmission model for cascading failure blackouts. Chaos 2002;12(4):985–94.

[5] Carreras B, Lynch V, Newman D, Dobson I. Blackout mitigation assessment in power transmission systems. In: Hawaii international conference on system science; 2003.

[6] Carreras B, Lynch V, Sachtjen M, Dobson I, Newman D. Modeling blackout dynamics in power transmission networks with simple structure. In: 34th Hawaii international conference on system sciences; 2001.

[7] Chen J, Thorp J, Dobson I. Cascading dynamics and mitigation assessment in power system disturbances via a hidden failure model. Electr Power Energy Syst 2005;27(4):318–26.

[8] Chen Q, Mili L. Risk-based composite power system vulnerability evaluation to cascading failures using importance sampling. In: Power and energy society general meeting. IEEE.

[9] Dobson I, Carreras B, Newman D. A loading-dependent model of probabilistic cascading failure. Probab Eng Inform Sci 2005;19(1):15–32.

[10] Dzobo O, Gaunt CT, Herman R. Investigating the use of probability distribution functions in reliability-worth analysis of electric power systems. Electr Power Energy Syst 2012;37(1):110–6.

[11] Garvels M. The splitting method in rare event simulation. Ph.D. thesis, University of Twente, The Netherlands; 2000.

[12] Glasserman P, Heidelberger P, Shahabuddin P, Zajic T. Multilevel splitting for estimating rare event probabilities. Oper Res 1999;47(4):585–600.

[13] Glynn PW. Efficiency improvement technique. Ann Oper Res 1994;53(1):175–97.

[14] González A, Echavarren F, Rouco L, Gomez T, Cabetas J. Reconfiguration of large-scale distribution networks for planning studies. Electr Power Energy Syst 2012;37(1):86–94.

[15] Heidelberger P. Fast simulation of rare events in queueing and reliability models. ACM Trans Model Comput Simul 1995;5(1):43–85.

[16] L'Ecuyer P, Demers V, Tuffin B. Splitting for rare-event simulation. In: Perrone LF, Wieland FP, Liu J, Lawson BG, Nicol DM, Fujimoto RM, editors. Proceedings of 2006 winter simulation conference. Piscataway (NJ): IEEE; 2006. p. 137–48.

[17] L'Ecuyer P, Gland FL, Lezaud P, Tuffin B. Splitting techniques. In: Rubino G, Tuffin B, editors. Rare event simulation using Monte Carlo methods. Chichester (UK): Wiley; 2009. p. 39–62.

[18] L'Ecuyer P, Simard R, Chen EJ, Kelton WD. An object-oriented random-number package with many long streams and substreams. Oper Res 2002;50(6):1073–5.

[19] Lin Y-J. Prevention of transient instability employing rules based on back propagation based ANN for series compensation. Electr Power Energy Syst 2011;33(10):1776–83.

[20] Newman D, Carreras B, Lynch V, Dobson I. Exploring complex systems aspects of blackout risk and mitigation. IEEE Trans Reliab 2011;60(1):134–43.

[21] Rubino G, Tuffin B, editors. Rare event simulation using Monte Carlo methods. Chichester (UK): Wiley; 2009.

[22] Shortle J, Chen CH, Crain B, Brodsky A, Brod D. Optimal splitting for rare-event simulation. IIE Transactions 2012;44(5):352–67.

[23] Shortle J, L'Ecuyer P. Introduction to rare-event simulation. In: Cochran J, editor. Wiley encyclopedia of operations research and management science. Wiley; 2010.

[24] Theodoro E, Benedito R, London J, Alberto L. Algebraic-graph method for identification of islanding in power system grids. Electr Power Energy Syst 2012;35(1):171–9.

[25] US-Canada Power System Outage Task Force. Final report on the August 14, 2003 blackout in the United States and Canada: Causes and recommendations. Tech. rep. <http://www.nerc.com>.

[26] Wang SP, Chen A, Liu CW, Chen CH, Shortle J. Rare-event splitting simulation for analysis of power system blackouts. In: IEEE power engineering society general meeting. Paper 678, 2011.