

SENSITIVITY ANALYSIS OF RARE-EVENT SPLITTING APPLIED TO CASCADING BLACKOUT MODELS

John Shortle
Chun-Hung Chen

Systems Engineering & Operations Research
George Mason University
Fairfax, VA 22030, USA

ABSTRACT

Splitting is a technique that can be used to improve the efficiency in simulating rare events. The basic idea is to create separate copies (splits) of the simulation whenever it gets close to the rare event. To implement splitting, several decisions must be made – for example, choosing a level function, choosing the number of simulation runs for each level, etc. This paper analyzes the sensitivity of the variance of the rare-event estimator to several parameters used within the splitting framework. We specifically consider a two-level fixed-effort variation of splitting for which analytic results can be derived. Results are applied to a simple model of cascading blackouts. The results illustrate that a good choice for the locations of levels may be more important than a good choice for the importance function for these types of problems.

1 INTRODUCTION

Estimating rare-event probabilities using standard Monte-Carlo simulation can be intractable due to the huge number of replications needed to obtain reasonable confidence intervals. Splitting is a technique that can be used to reduce the variance in such problems (e.g., L'Ecuyer et al. 2009). The basic idea of splitting is to interrupt the simulation whenever it gets “close” to the rare event and then to split the simulation into multiple independent replications. In this way, more computer effort is spent on runs that are more likely to hit the rare event. To implement splitting, the analyst must generally make several decisions, such as: (a) What variant of splitting should be used (e.g., fixed-effort, fixed splitting, etc.)? (b) How many levels should be used? (c) Where should the levels be located? (d) How many simulation replications should be allocated to each level? (e) What importance function should be used? The importance function maps the state of the system, which is typically multi-dimensional, to a single real value, denoting a measure of proximity to the rare event. The choice of the importance function is generally regarded as the most important decision to be made when implementing splitting (e.g., L'Ecuyer et al. 2009).

In this paper, we analyze the sensitivity of the estimator variance to several of the parameters used within a splitting framework. We specifically consider a *two-level fixed-effort* variation of splitting. With only two levels, it is possible to determine a relatively simple analytic formula for the variance of the rare-event estimator, which we can then apply to several test problems to yield an analytic comparison of results. We compare the relative impact of the following parameters – the locations of the single intermediate level, the number of replications per level, and the choice of the importance function. That is, with respect to the five previous questions, we fix the answer to questions (a) and (b), and then investigate the sensitivity of the estimator variance to questions (c), (d), and (e).

We consider test problems in the application domain of power-grid blackouts. In these test cases, we observe that the choice of the importance function does not seem to be as important as the choice of the intermediate level. A similar observation was made in more complex simulations of power grids

(Shortle 2013). In such problems, the system tends to move in small “jumps” as several power lines may simultaneously fail in each iteration of the model. Thus, no matter where one draws the contours of the importance function in the state space, the system may “jump over” a given contour in one step, rather than landing exactly on the contour. This makes it difficult to define an importance function with the desired property that the probability of reaching the next level does not depend on the starting state from that level.

2 PRELIMINARIES AND NOTATION

This section describes the splitting methodology and notation used in this paper. For an introduction and overview of splitting, see L’Ecuyer, Demers, and Tuffin (2006), L’Ecuyer et al. (2009) and Shortle and L’Ecuyer (2011).

Let $X \equiv \{X_t, t \geq 0\}$ be a Markov process (possibly multi-dimensional) with state space χ . Let $h(\cdot) : \chi \rightarrow \mathbb{R}^+$ be a map of the state space to the “level” of the process. The rare-event set is defined as a set of states that are at or above some particular level l – that is, $\mathcal{R} \equiv \{x \in \chi : h(x) \geq l\}$. The probability to estimate is $\gamma \equiv \Pr\{X_t \text{ enters } \mathcal{R} \text{ before } X_t \text{ returns to level } 0\}$. For example, if $h(X_t)$ denotes the total number of customers in a queueing system at time t , then γ is the probability that – starting from an empty system – the system reaches some critical threshold of customers before returning to empty. To implement splitting, define a sequence of m levels $l_0 \equiv 0 < l_1 < \dots < l_m \equiv l$. Let p_j be the probability that $h(X_j)$ up-crosses l_j before returning to 0 given that $h(X_j)$ has up-crossed l_{j-1} before returning to 0. Figure 1 shows an example with a two-dimensional state space and two levels.




Figure 1: Example level function

This paper considers a variation of splitting called *fixed-effort* splitting in which the numbers of runs at each level are pre-selected, but not necessarily equal. This works as follows. Let A_0 denote a set of possible starting states (often there is only one possible starting state). Draw n_1 states at random, with replacement, from the set A_0 . Start n_1 independent runs from these starting points and simulate each of them until $h(X_t)$ up-crosses l_1 or returns to 0. Let Q_1 be the number of those runs that up-cross l_1 before returning to 0. Then $\hat{p}_1 \equiv Q_1/n_1$ is an unbiased estimator of p_1 . The end states of the Q_1 simulations that up-cross l_1 are collected into a set A_1 , and these states become the starting states for the next stage of simulation. (A_1 may include duplicate copies of the same state, so that $|A_1| = Q_1$.) Similarly, draw n_2 states at random, with replacement, from the set A_1 and conduct n_2 independent runs start from these points until $h(X_t)$ up-crosses l_2 or returns to 0. Simulation at the higher stages proceeds in an analogous manner, though this paper only considers a two level implementation ($m = 2$). An unbiased estimator for the rare-event probability is $\hat{\gamma} \equiv \hat{p}_1 \hat{p}_2 \dots \hat{p}_m$.

Let $p_2(X)$ be the probability of reaching the rare-event set from state X (before returning to 0). A key challenge in splitting is that there may be significant differences between $p_2(X)$ for different starting states X in A_1 . For example, in Figure 1 the system may be much more likely to reach the rare-event set starting from a_1 than from a_2 ($p_2(a_1) \gg p_2(a_2)$). On the other hand, the system may be much more likely to reach a_2 in the first place, compared to a_1 . Thus, the set A_1 would contain many more copies of the

state a_2 than a_1 . In this case, splitting would tend to emphasize paths that have a relatively low probability of reaching the rare event. This problem is discussed in Glasserman et al. (1998).

Figure 2 shows a discrete representation of two-level splitting. In this figure, we ignore the detailed sample paths of X_t and focus on the intermediate states at which the system may up-cross level 1. Suppose that there are k possible entrance states to level 1. (All test cases in this paper have discrete state spaces; this representation could otherwise be generalized to a continuous state space.) The probability of up-crossing level 1 via state i (before $h(X_t) = 0$) is assumed to be $p_1 q_i$. Without loss of generality, we assume that $\sum_{i=1}^k q_i = 1$. Then the probability of up-crossing level 1 (before $h(X_t) = 0$) is $p_1 = \sum_{i=1}^k p_1 q_i$, consistent with the notation defined previously. And q_i can be interpreted as the conditional probability of up-crossing level 1 through state i , given that level 1 has been up-crossed. Starting from the intermediate state i , the probability of reaching the rare-event set (before $h(X_t) = 0$) is assumed to be $p_2(1 + c_i)$. Without loss of generality, assume that $\sum_{i=1}^k q_i c_i = 0$. (If $\sum_{i=1}^k q_i c_i = a$, relabel the transition probabilities as $p'_2(1 + c'_i)$ where $p'_2 \equiv (1 + a)p_2$ and $c'_i \equiv (c_i - a)/(1 + a)$. Then $p'_2(1 + c'_i) = p_2(1 + c_i)$ and $\sum_{i=1}^k q_i c'_i = 0$.) With this convention, the probability of reaching the rare-event set, given that level 1 has been up-crossed is:

$$\sum_{i=1}^k q_i \cdot p_2(1 + c_i) = p_2 \left(\sum_{i=1}^k q_i + \sum_{i=1}^k q_i c_i \right) = p_2 + p_2 \sum_{i=1}^k q_i c_i = p_2,$$

consistent with the previous definition of p_2 . Intuitively, c_i represents the deviation of the state-based probability $p_2(a_i)$ from the overall value p_2 , which is the aggregate probability of advancing from level 1 to the rare event, starting from an *arbitrary* state chosen from the entrance distribution to level 1.




Figure 2: General case with 2 levels

In an ideal situation, the probability of reaching the rare event is the same starting from any of the k intermediate states – that is, $c_i = 0$ for all i . It is potentially problematic when, for example, state 1 is the most likely entrance state to level 1 ($q_1 \approx 1$), but it is otherwise a dead-end (e.g., $c_1 = -1$ so $p_2(1 + c_1) = 0$). The severity of this problem can be quantified by $\text{var}[p_2(S)]$, where S is a random state drawn from the entrance distribution to level 1. Using the notation from Figure 2,

$$\begin{aligned} \text{var}[p_2(S)] &= E[p_2^2(S)] - E^2[p_2(S)] = \left(\sum_{i=1}^k q_i [p_2(1 + c_i)]^2 \right) - p_2^2 = p_2^2 \left(-1 + \sum_{i=1}^k q_i (1 + 2c_i + c_i^2) \right) \\ &= p_2^2 \sum_{i=1}^k q_i c_i^2. \end{aligned}$$

3 OPTIMAL ALLOCATION FOR 2-LEVEL SPLITTING

To obtain analytic results in a simplified setting, a common assumption in the literature (e.g., Villén-Altamirano and Villén-Altamirano 1994, Garvels and Kroese 1998, Garvels 2000, p. 19, L’Ecuyer, Demers, and Tuffin 2006, Shortle et al. 2012) is the following:

Assumption 1 The success probability of a stage- j run is independent of the starting state in A_{j-1} .

Under the notation from the previous section, this assumption means that $c_i = 0$ for all i and $\text{var}[p_2(S)] = 0$. Under this assumption, in a fixed-effort splitting scheme, the number of hits to level j is a binomial random variable with parameters n_j and p_j . Further, \hat{p}_i and \hat{p}_j ($i \neq j$) are independent. It is straightforward to derive the variance of $\hat{\gamma}$:

$$\text{var}[\hat{\gamma}] = E[\hat{p}_1^2 \cdots \hat{p}_m^2] - E^2[\hat{p}_1 \cdots \hat{p}_m] = \left(p_1^2 + \frac{p_1(1-p_1)}{n_1} \right) \cdots \left(p_m^2 + \frac{p_m(1-p_m)}{n_m} \right) - p_1^2 \cdots p_m^2.$$

In the two-level case, this reduces to:

$$\text{var}[\hat{\gamma}] = \gamma^2 \left[\frac{1-p_1}{n_1 p_1} + \frac{1-p_2}{n_2 p_2} + \frac{(1-p_1)(1-p_2)}{n_1 n_2 p_1 p_2} \right]. \tag{1}$$

For problems in which the state space is multi-dimensional, Assumption 1 does not typically apply (e.g., Figure 1). In such cases, it is generally hard to derive an explicit expression for $\text{var}[\hat{\gamma}]$. One special case for which analytical results exist is the two-level case ($m = 2$). Garvels and Kroese (1998) gave the variance of \hat{p}_2 conditional on the number of hits Q_1 to level 1. From this, it is relatively straightforward to derive the overall variance of $\hat{\gamma} = \hat{p}_1 \hat{p}_2$ (proof in Appendix).

Lemma 1 (Garvels and Kroese 1998)

$$\text{var}[\hat{p}_2|Q_1] = \frac{1}{n_2} \left(p_2 - p_2^2 + \frac{n_2 - 1}{Q_1} \text{var}[p_2(S)] \right). \tag{2}$$

Theorem 1 In two-level fixed-effort splitting,

$$\text{var}[\hat{\gamma}] = \gamma^2 \left[\frac{1-p_1}{n_1 p_1} + \frac{1-p_2}{n_2 p_2} + \frac{(1-p_1)(1-p_2)}{n_1 n_2 p_1 p_2} + \frac{(1-1/n_2)\text{var}[p_2(S)]/p_2^2}{n_1 p_1} \right]. \tag{3}$$

Alternatively, using the discrete formulation in Figure 2, equation (3) can be written as:

$$\text{var}[\hat{\gamma}] = \gamma^2 \left[\frac{1-p_1}{n_1 p_1} + \frac{1-p_2}{n_2 p_2} + \frac{(1-p_1)(1-p_2)}{n_1 n_2 p_1 p_2} + \frac{(1-1/n_2)\sum_{i=1}^k q_i c_i^2}{n_1 p_1} \right]. \tag{4}$$

From these theorems, it is possible to derive the optimal allocations of runs (n_1 and n_2) that minimizes $\text{var}[\hat{\gamma}]$ subject to a fixed total number of runs $n_1 + n_2 = T$.

Theorem 2 (Garvels and Kroese 1998) Under Assumption 1, the optimal allocation of runs to each level approximately satisfies (when n_1, n_2, \dots, n_m are large):

$$n_1 \sqrt{\frac{p_1}{1-p_1}} = n_2 \sqrt{\frac{p_2}{1-p_2}} = \cdots = n_m \sqrt{\frac{p_m}{1-p_m}}. \tag{5}$$

Shortle et al. (2012) gave an extension to this result in the non-asymptotic case (i.e., without requiring that each n_i be large). The analogous result when Assumption 1 is relaxed, for a two-level problem, is:

Theorem 3 In two-level fixed-effort splitting (without Assumption 1), the optimal allocation of runs to each level approximately satisfies (when n_1 and n_2 are large):

$$n_1 \sqrt{\frac{p_1}{1-p_1 + \sum_{i=1}^k q_i c_i^2}} = n_2 \sqrt{\frac{p_2}{1-p_2}}. \tag{6}$$

The derivation of Theorem 3 is analogous to the derivation of Theorem 2. Specifically, one can compare the forms of (1) and (3) and ignore terms with $n_1 n_2$ in the denominator. As in Theorem 2, we do not attempt to define what “approximately” means or what “ n_1 and n_2 are large” means. Intuitively, the correction term $\text{var}[p_2(S)]/p_2^2 = \sum_{i=1}^k q_i c_i^2$ in (6) increases the relative allocation of level-1 runs n_1 compared to the allocation rule in (5). In other words, as the system deviates from Assumption 1, more runs must be conducted in the first stage to ensure that a good sample distribution is obtained for starting states in the second stage.

While (6) is a theoretically better allocation rule than (5), it is practically much more difficult to implement. This is because (6) requires obtaining estimates for q_i and c_i for each value of i . Equation (5) only requires obtaining estimates for two parameters, p_1 and p_2 . Thus, the required warmup period may be very long for (6). Nevertheless, we can still evaluate the theoretical benefit of (6) on analytic test problems to understand the potential benefit that could be achieved with perfect information.

4 CASCADING BLACKOUTS

Splitting has been applied to cascading blackout models in several recent papers (Kim, Bucklew, and Dobson 2013, Shortle 2013, Wang et al. 2011). This section applies the theory of splitting from the previous section to a simple test problem given in Shortle (2013). The test problem can be solved analytically. A similar type of analytical blackout model was given in Dobson, Carreras, and Newman (2005).

Consider a power network with N lines. The total load (power demand) on the network is L . This load is evenly distributed among the N lines. Initially, each line has a carried load of L/N . The capacity of each line is individually set according to a fixed but random value with cumulative distribution function F , conditional on the capacity being at least L/N . The stochastic nature of capacity is meant to model variations in weather and other factors (such as vegetation growth) that impact line capacity. The system begins with the failure of one line. When a line fails, the total load on the network is evenly distributed among the remaining working lines. (In general, this assumption is unrealistic. One network topology for which this assumption holds is a network of two buses connected by N identical lines in parallel. When a line fails, the total load between the buses is evenly divided among the remaining working lines. However, for general network topologies, the redistribution of flows depends on which line failed, which lines are currently working, the laws of physics and so forth. These more complex models were also considered in Shortle (2013) for a variety of network topologies including a grid network, a mesh network, and a small real network.) If the resulting load on any working line exceeds its capacity, the line fails. At each step, it is possible for multiple lines to fail simultaneously. The process continues until there is an iteration with no new line failures. The objective is to find the probability that a blackout, starting with the failure of one line, results in the failure of all N lines.

If F is an exponential distribution, the model can be solved analytically as a continuous-time Markov chain. If F is non-exponential, the model can be solved numerically via a set of iterative equations (Shortle 2013). The model is relatively easy to solve because the topological properties of the network are not modeled – every line has the same structural role in the model as every other. The state of the system can be specified by two parameters (m, n) , where n is the *present* number of failed lines and $m \leq n$ is the *previous* number of failed lines in the previous iteration. Under this formulation, it is possible to calculate: (a) the probability that the system ever reaches state (m, n) and (b) the probability that the system reaches the rare-event set starting from (m, n) (Shortle 2013).

Figure 3 shows two candidate definitions for the level function $h(m, n)$. The obvious choice is to let $h(m, n) = n$, the number of presently failed lines (left figure). We call this the “flat” level function. The rare-event set is $\mathcal{R} \equiv \{(m, n) : h(m, n) \geq N\}$. An alternate choice is suggested by Garvels, Kroese, and VanOmeren (2002). They state that the level function should be chosen so that the probability of reaching the next level does not depend on the starting state within the current stage. This suggests the following level function, which we call the “modified” level function: $h(m, n) = \text{Pr}\{\text{system reaches rare event from state } (m, n)\}$. The right figure shows sample contours of this level function (for an example with $N = 100$, $L = 100$, and

F is an exponential distribution with a mean of about 3.034; the rare-event probability is about $\gamma \approx 10^{-10}$. The contours of this function are curved because the system is more likely to reach the rare-event set from (m, n) than from $(m + a, n)$ where $a > 0$. This is because the first state represents $n - m$ line failures in the most recent iteration of the model, while the second state represents $n - m - a < n - m$ line failures. A larger number of line failures in one step corresponds to a greater increase in load on the remaining lines (even though the total load is the same in the two cases), resulting in a larger marginal failure rate of the working lines.




Figure 3: Two level functions for power model (from Shortle (2013)).

Figure 4 shows sample evaluations of the work-normalized variance $(n_1 + n_2)\text{var}[\hat{\gamma}]$ using (3). In this example, $N = 30$, $L = 10$, and F is an exponential distribution with a mean of 1. The resulting rare-event probability is about 5×10^{-5} . The left figure shows the impact of using the two different level functions. The right figure shows the impact of using two different run-allocation schemes – equal allocation ($n_1 = n_2$) and the allocation rule in (5) based on Assumption 1. (Assumption 1 does not hold for this problem. Thus, the allocation rule is not technically optimal. However, the rule can still be applied.) The x -axis shows the location of level 1, ranging from 2 to 29. (The system starts with one failed component, so the first level must be at least 2; there are 30 lines, so the first level can be no more than 29.) Several observations are made from this figure:

1. The optimal location of l_1 is around 5. This is the “midway” point in the sense that $p_1 \approx p_2$. That is, the system is (roughly) equally likely to go from 1 failed line to 5 failed lines as it is to go from 5 failed lines to 30 failed lines. The first few lines are relatively unlikely to fail, but as more lines fail, there are fewer lines to carry the fixed load L , so the failure probabilities of individual lines increase as the blackout progresses.
2. If the level location is chosen appropriately ($l_1 \approx 5$), there is no benefit in using the optimal allocation scheme (right figure). This is because the optimal allocation scheme of (5) reduces to the equal-allocation scheme when $p_1 = p_2$. On the other hand, if a poor choice is made for l_1 , then the optimal allocation provides a good benefit. This includes a choice of $l_1 = 15$, which might be considered “reasonable” in the sense that it is halfway between 0 and 30. Thus, in this example, the optimal allocation provides somewhat of a mitigation against a poor choice for l_1 .
3. The modified level function provides some benefit (left figure), particularly for intermediate locations of l_1 . Little benefit is achieved near the end points, including near the optimal choice for l_1 around 5.

Figure 5 shows combined results for six different splitting schemes, including the results from Figure 4 in one graph. These results include two different level functions (flat levels and modified levels) and three different run allocation schemes: (a) Equal allocation ($n_1 = n_2$), (b) optimal allocation based on (5) with Assumption 1, and (c) optimal allocation based on (6) without Assumption 1. The y -axis is shown on

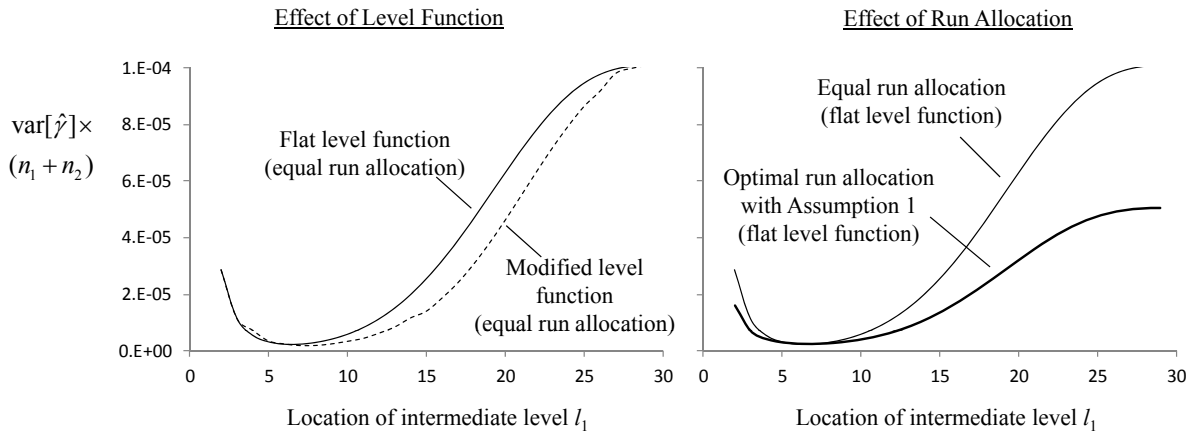


Figure 4: Impact of splitting parameters on work-normalized variance, $N = 30, L = 10$

a log-scale in this figure. The differences between the allocations in (5) and (6) are small and hard to see visually from the figure. Overall, the figure shows that, in this example, the most important decision to get right is to put l_1 near 5. A good level function is helpful as is a good allocation scheme, but the benefit is dominated by the choice of l_1 . The likely reason for this outcome is that the system can move in jumps, rather than in single steps. That is, it is possible for more than one line to fail simultaneously in each iteration of the model. Thus, it is impossible to create a level function following the rule in Garvels, Kroese, and VanOmmeren (2002) which states that the probability of reaching the rare-event set from every possible entrance state to level 1 should be the same. No matter where the contour of the level function is drawn, it is possible for the system to “jump over” the contour as well as to hit it exactly, meaning that the probabilities cannot all be the same for all possible entrance states. This may work to reduce the benefit of choosing a good importance function.

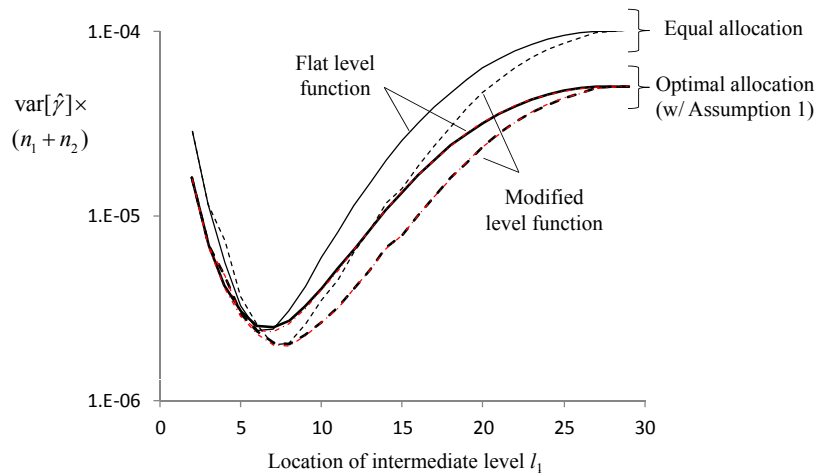


Figure 5: Impact of splitting parameters on work-normalized variance, $N = 30, L = 10$

Figure 6 shows results for a problem with $N = 60$ and $L = 15$ (as before, F is an exponential distribution with a mean of 1). The rare-event probability is about 7×10^{-11} . This is a larger network than before with a lower rare-event probability. Qualitatively, Figures 5 and 6 show similar results. Again, the modified level function is helpful in reducing the variance as is a good allocation scheme, but the benefit is dominated by the choice of l_1 . The differences between the allocations in (5) and (6) are again small, though slightly more discernable in this figure.

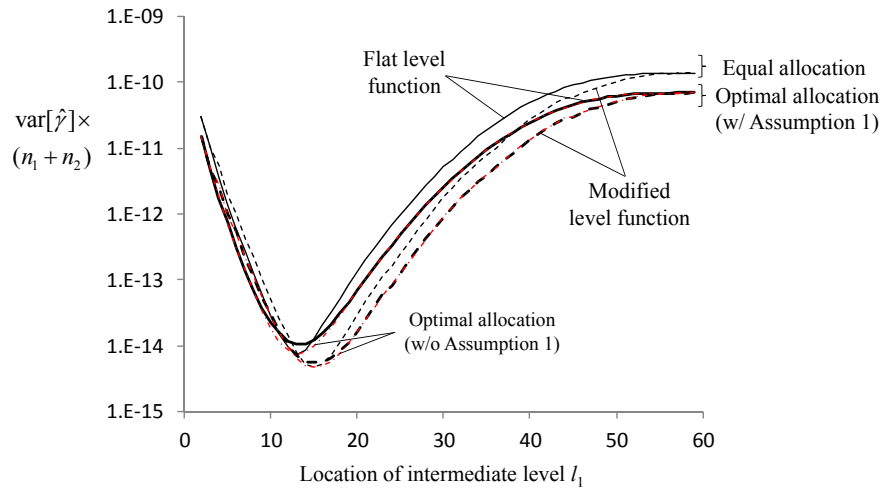


Figure 6: Impact of splitting parameters on work-normalized variance, $N = 60$, $L = 15$.

5 CONCLUSIONS

This paper analyzed the sensitivity of the variance of a rare-event estimator to several parameters used within a splitting framework on a simple model of cascading blackouts. Specifically, we compared the relative impact of the following parameters – the locations of an intermediate level, the number of replications per level, and the choice of the importance function. The results illustrate that a good choice for the locations of levels may be more important than a good choice for the importance function for these types of problems. A key reason may be that the system can move in small jumps rather than in single steps. Thus it is impossible to construct an “optimal” level function in a manner suggested in Garvels, Kroese, and VanOmmeren (2002). These types of results, of course, may be specific to the model and test cases given in this paper. Nevertheless, the results are consistent with experimental results applying splitting to more complicated blackout models that incorporate network topologies and the physics of power flows (Shortle 2013). In those experiments, the locations of the levels also seemed to be critical in reducing the variance of the estimator. In other types of test problems (e.g., the tandem queue in Glasserman et al. 1998), the choice of the level function is expected to have the largest impact.

ACKNOWLEDGMENTS

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

A APPENDIX

Proof of Theorem 1: Condition on the number of hits Q_1 to level 1:

$$\begin{aligned}
\text{var}[\hat{p}_1\hat{p}_2] &= \text{E}[\text{var}[\hat{p}_1\hat{p}_2|Q_1]] + \text{var}[\text{E}[\hat{p}_1\hat{p}_2|Q_1]] \\
&= \text{E}\left[\frac{Q_1^2}{n_1^2}\text{var}[\hat{p}_2|Q_1]\right] + \text{var}\left[\frac{Q_1}{n_1}\text{E}[\hat{p}_2|Q_1]\right] \\
&= \text{E}\left[\frac{Q_1^2}{n_1^2}\frac{1}{n_2}\left(p_2 - p_2^2 + \frac{n_2 - 1}{Q_1}\text{var}[p_2(S)]\right)\right] + \text{var}\left[\frac{Q_1}{n_1}p_2\right] \\
&= \frac{\text{E}[Q_1^2](p_2 - p_2^2) + \text{E}[Q_1](n_2 - 1)\text{var}[p_2(S)]}{n_1^2 n_2} + \frac{p_2^2 \text{var}[Q_1]}{n_1^2} \\
&= \frac{[n_1 p_1 (1 - p_1) + n_1^2 p_1^2](p_2 - p_2^2) + n_1 p_1 (n_2 - 1)\text{var}[p_2(S)]}{n_1^2 n_2} + \frac{p_2^2 n_1 p_1 (1 - p_1)}{n_1^2} \\
&= p_1^2 p_2^2 \left[\frac{(1 - p_1 + n_1)(1 - p_2)}{n_1 n_2 p_1 p_2} + \frac{(n_2 - 1)\text{var}[p_2(S)]}{n_1 n_2 p_1 p_2^2} + \frac{1 - p_1}{n_1 p_1} \right] \\
&= \gamma^2 \left[\frac{1 - p_1}{n_1 p_1} + \frac{1 - p_2}{n_2 p_2} + \frac{(1 - p_1)(1 - p_2)}{n_1 n_2 p_1 p_2} + \frac{(1 - 1/n_2)\text{var}[p_2(S)]/p_2^2}{n_1 n_2 p_1} \right].
\end{aligned}$$

The third equality uses Lemma 1 and the fact that $\text{E}[\hat{p}_2|Q_1] = p_2$. The fifth equality uses the fact that Q_1 is a binomial random variable.

REFERENCES

- Dobson, I., B. Carreras, and D. Newman. 2005. "A loading-dependent model of probabilistic cascading failure". *Probability in the Engineering and Informational Sciences* 19 (1): 15–32.
- Garvels, M. 2000. *The splitting method in rare event simulation*. Ph. D. thesis, University of Twente, The Netherlands.
- Garvels, M., and D. Kroese. 1998. "A comparison of RESTART implementations". In *Proceedings of the 1998 Winter Simulation Conference*, 601–608. Piscataway, NJ: IEEE.
- Garvels, M., D. Kroese, and J. VanOmmeren. 2002. "On the importance function in splitting simulation". *European Transactions on Telecommunications* 13 (4): 363–371.
- Glasserman, P., P. Heidelberger, P. Shahabuddin, and T. Zajic. 1998. "A large deviations perspective on the efficiency of multilevel splitting". *IEEE Transactions on Automatic Control* 43 (12): 1666–1679.
- Kim, J., J. A. Bucklew, and I. Dobson. 2013. "Splitting method for speedy simulation of cascading blackouts". *IEEE Transactions on Power Systems*. To appear.
- L'Ecuyer, P., V. Demers, and B. Tuffin. 2006. "Splitting for rare-event simulation". In *Proceedings of 2006 Winter Simulation Conference*, edited by L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 137–148. Piscataway, NJ: IEEE.
- L'Ecuyer, P., F. L. Gland, P. Lezaud, and B. Tuffin. 2009. "Splitting techniques". In *Rare Event Simulation using Monte Carlo Methods*, edited by G. Rubino and B. Tuffin, 39–62. Chichester, U.K.: Wiley.
- Shortle, J. 2013. "Efficient simulation of blackout probabilities using splitting". *Electrical Power and Energy Systems* 44 (1): 743–751.
- Shortle, J., C. H. Chen, B. Crain, A. Brodsky, and D. Brod. 2012. "Optimal splitting for rare-event simulation". *IIE Transactions* 44 (5): 352–367.
- Shortle, J., and P. L'Ecuyer. 2011. "Introduction to rare-event simulation". In *Wiley Encyclopedia of Operations Research and Management Science*, edited by J. Cochran.
- Villén-Altamirano, M., and J. Villén-Altamirano. 1994. "RESTART: A straightforward method for fast simulation of rare events". In *Proceedings of the 1994 Winter Simulation Conference*, edited by J. D. Tew, S. Manivannan, D. A. Sadowski, and A. F. Seila, 282–289. Piscataway, NJ: IEEE.

Wang, S.-P., A. Chen, C.-W. Liu, C.-H. Chen, and J. F. Shortle. 2011. "Rare-event splitting simulation for analysis of power system blackouts". In *Proceedings of the IEEE Power and Energy Society General Meeting*. Paper 678.

AUTHOR BIOGRAPHIES

JOHN F. SHORTLE is an Associate Professor of Systems Engineering and Operations Research at George Mason University. His research interests include simulation and queueing applications in air transportation, telecommunications, and energy. Previously, he worked at US WEST Advanced Technologies. He received his doctorate degree in industrial engineering and operations research from UC Berkeley. His email address is jshortle@gmu.edu.

CHUN-HUNG CHEN is a Professor of Systems Engineering and Operations Research at George Mason University. Dr. Chen has led research projects in stochastic simulation and optimization, systems design under uncertainty, and air traffic management, which are sponsored by NSF, FAA, and NASA. He served as Co-Editor of the Proceedings of the 2002 Winter Simulation Conference and Program Co-Chair for 2007 INFORMS Simulation Society Workshop. He is serving on the editorial boards of IEEE Transactions on Automatic Control, IIE Transactions, Journal of Simulation Modeling Practice and Theory, and International Journal of Simulation and Process Modeling. He received his Ph.D. degree from Harvard University in 1994. His email address is cchen9@gmu.edu.