

IMPORTANCE SAMPLING THE UNION OF RARE EVENTS  
WITH AN APPLICATION TO POWER SYSTEMS ANALYSIS

By

Art B. Owen  
Yury Maximov  
Michael Chertkov

Technical Report No. 2017-10  
October 2017

Department of Statistics  
STANFORD UNIVERSITY  
Stanford, California 94305-4065



IMPORTANCE SAMPLING THE UNION OF RARE EVENTS  
WITH AN APPLICATION TO POWER SYSTEMS ANALYSIS

By

Art B. Owen  
Stanford University

Yury Maximov  
Michael Chertkov  
Los Alamos National Laboratory

Technical Report No. 2017-10  
October 2017

**This research was supported in part by  
National Science Foundation grants  
DMS 1521145 and DMS 1407397.**

Department of Statistics  
STANFORD UNIVERSITY  
Stanford, California 94305-4065

<http://statistics.stanford.edu>

# Importance sampling the union of rare events with an application to power systems analysis

Art B. Owen                      Yury Maximov  
Stanford University          Los Alamos National Laboratory

Michael Chertkov  
Los Alamos National Laboratory

October 2017

## Abstract

This paper presents a method for estimating the probability  $\mu$  of a union of  $J$  rare events. The method uses  $n$  samples, each of which picks one of the rare events at random, samples conditionally on that rare event happening and counts the total number of rare events that happen. We call it ALORE, for ‘at least one rare event’. The ALORE estimate is unbiased and has a coefficient of variation no larger than  $\sqrt{(J + J^{-1} - 2)/(4n)}$ . The coefficient of variation is also no larger than  $\sqrt{(\bar{\mu}/\mu - 1)/n}$  where  $\bar{\mu}$  is the union bound. Our motivating problem comes from power system reliability, where the phase differences between connected nodes have a joint Gaussian distribution and the  $J$  rare events arise from unacceptably large phase differences. In the grid reliability problems even some events defined by 5772 constraints in 326 dimensions, with probability below  $10^{-22}$ , are estimated with a coefficient of variation of about 0.0024 with only  $n = 10,000$  sample values.

## 1 Introduction

In this paper we consider a mixture importance sampling strategy to estimate the probability that one or more of a set of rare events takes place. Our main method is to repeatedly choose a rare event at random, and then sample the system conditionally on that one event taking place. The result is an importance sampler with a sampling standard deviation that is no more than some modest multiple of the rare event probability, even in settings where that probability tends to zero.

Our motivating context is the reliability of the electrical grid when subject to random inputs, such as variable demand by users and variable production, as occurs at wind farms. The rare events describe unacceptably large electrical phase differences at pairs of connected nodes in the grid.

It is common to use a simplified linear direct current (DC) model of the electrical grid, because the equations describing alternating current (AC) are significantly more difficult to work with, and some authors (e.g., Van den Bergh et al. (2014)) find that there is little to be gained from the complexity of an AC model. This DC model is presented in Sauer and Christensen (1984) and Stott et al. (2009). It is also common to model the randomness in the grid as Gaussian, especially over short time horizons. We make both of these simplifications, linearity and Gaussianity, for this paper.

The probability we consider can be written

$$\mu = \Pr(\cup_{j=1}^J H_j), \quad H_j = \{\mathbf{x}^\top \omega_j \geq \tau_j\}, \quad \text{where } \mathbf{x} \sim \mathcal{N}(\eta, \Sigma). \quad (1)$$

Section 2 introduces more notation for the event (1) and develops an importance sampling estimator  $\hat{\mu}$  for the generally intractable quantity  $\mu$ . An upper bound for  $\mu$  is  $\bar{\mu} = \sum_{j=1}^J \Pr(H_j)$ , the union bound. Theorem 1 proves that the importance sampling estimator has variance at most  $\mu(\bar{\mu} - \mu)/n$  when  $n$  IID samples are used. Section 3 discusses some further sampling properties of our estimator that hold without the Gaussian assumption. When there are  $J$  events, the variance of  $\hat{\mu}$  is at most  $(J + 1/J - 2)\mu^2/(4n)$  when the system is sampled  $n$  times. Section 4 compares our importance sampler to a state of the art code (Genz et al., 2017) for estimating the probability that a multivariate Gaussian of up to 1000 variables with arbitrary covariance belongs to a given hyperrectangle. Our algorithm is simpler and extends to higher dimensions. When we studied rare event cases, our algorithm was more accurate. In our examples that are not rare events, the other algorithm was more accurate. Section 5 describes the power system application. Section 6 contains some discussions. The appendix proves Theorem 1 for any set of  $J$  events, not just those given by a Gaussian distribution. The theorem applies so long as we can sample conditionally on any one event  $H_j$  and then determine which other events  $H_\ell$  also occur.

We call our method ALORE, for “at least one rare event”, because the sampler only samples from  $\cup_j H_j$ . One common way for rare event sampling to become inaccurate is that we might fail to obtain any points where the rare event happens. That leads to a severe under-estimation of the rare event probability. In ALORE, the corresponding problem is the failure to sample any points where two or more of the rare constituent events occur. In that case ALORE will return the union bound as the estimated rare event probability instead of zero. That is also a setting where the union bound is likely to be a good approximation. So ALORE is robust against severe underestimates of the rare event probability. The second common problem for rare event sampling is an extreme value of the likelihood ratio weighting applied to the observations. In ALORE, the largest possible weight is only  $J$  times as large as the smallest one.

Our sampler is closely related to instanton methods in power systems engineering (Chertkov et al., 2011a,b; Kersulis et al., 2015). Out of all the configurations of random inputs to a system, the most probable one causing the failure is called the instanton. When there are thousands of failure types there

are correspondingly thousands of instantons, each one a conditional mode of the distribution of  $\mathbf{x}$ . Our initial thought was to do importance sampling from a mixture of distributions, with each mixture component defined by shifting the Gaussian distribution’s mean to an instanton. What we do instead is sample conditionally on the rare event happening. By sampling in a half-space we avoid wasting samples outside the failure region. By conditioning instead of shifting, we get better control over the likelihood ratio in the importance sampler. Each instanton remains the most probable point in its half-space under this sampling.

Our sampler is a form of multiple importance sampling. Multiple importance sampling originated in computer graphics (Lafortune and Willems, 1993; Veach and Guibas, 1994). Owen and Zhou (2000) found a useful way to combine it with control variates defined by the mixture components. Elvira et al. (2015a,b) investigate computational efficiency of some mixture importance sampling and weighting strategies.

We do not consider self-normalized importance sampling (SNIS) in this paper. SNIS is useful in settings where we can compute an unnormalized version of our target density but cannot sample from it efficiently, if at all. This is common in Bayesian applications (Liu, 2001, Chapter 2). For a recent adaptive version of SNIS, see Cornuet et al. (2012). In estimating rare event probabilities, the optimal sampler for SNIS can be shown to place half of its probability in the rare event and half outside the rare event, while the optimal plain IS estimator places all of its probability on the rare event. The best coefficient of variation for rare event estimation under SNIS is asymptotically  $2/\sqrt{n}$ . Ordinary importance sampling can attain much smaller variances, and so we focus on it for the rare event problem.

## 2 Gaussian case

For concreteness, we present the algorithm first for Gaussian random variables. We let  $\mathbf{x} \in \mathbb{R}^d$  have the standard Gaussian distribution,  $\mathcal{N}(0, I)$ , deferring general Gaussians to Section 2.1. We are interested in computing the probability that  $\mathbf{x}$  lies outside a polytope  $\mathcal{P}$ . In our motivating applications, the interior of the polytope defines a safe operating region and we assume that  $\mathbf{x} \notin \mathcal{P}$  is a rare event. For  $j = 1, \dots, J$ , define half-spaces

$$H_j = \{\mathbf{x} \mid \omega_j^\top \mathbf{x} \geq \tau_j\}$$

where each  $\tau_j \in \mathbb{R}$  and  $\omega_j \in \mathbb{R}^d$ , with  $\omega_j^\top \omega_j = 1$ . Then  $\mathcal{P} = \bigcap_{j=1}^J H_j^c$  and we want to find  $\mu = \Pr(\mathbf{x} \in H)$  where  $H = \bigcup_{j=1}^J H_j = \mathcal{P}^c$ . The set  $\mathcal{P}$  is convex and not necessarily bounded. Ordinarily  $\tau_j > 0$ , because we are interested in rare events.

The setting is illustrated in Figure 1 for  $J = 6$  half-spaces. In that example, two of the half-spaces have their conditional modes inside the union of the other half-spaces. One of those half-spaces is entirely included in the union of the others.




Figure 1: The solid circles contain 10%, 20% up to 90% of the  $\mathcal{N}(0, I)$  distribution. The dashed circles contain all but  $10^{-k}$  of the probability for  $3 \leq k \leq 7$ . The six solid lines denote half-spaces. The solid points are the corresponding conditional modes (instantons). The rare event of interest is  $\mathbf{x}$  in the shaded region, when  $\mathbf{x} \sim \mathcal{N}(0, I)$ .

Letting  $P_j = \Pr(\mathbf{x} \in H_j) = \Phi(-\tau_j)$ , we know that

$$\max_{1 \leq j \leq J} P_j =: \underline{\mu} \leq \mu \leq \bar{\mu} := \sum_{j=1}^J P_j.$$

The right hand side is the union bound. Sometimes that bound is very conservative, and other times it is quite accurate. For a given high dimensional problem we might not know whether the bound is close. Therefore we seek a more accurate estimate of  $\mu$ .

We will need to use some inclusion-exclusion formulas, so some notation for these follows. For any  $u \subseteq 1:J \equiv \{1, 2, \dots, J\}$ , let  $H_u = \cup_{j \in u} H_j$ , so  $H_j = H_{\{j\}}$  and by convention  $H_\emptyset = \emptyset$ . We identify the set  $H_u$  with the function  $H_u(\mathbf{x}) = \mathbf{1}\{\mathbf{x} \in H_u\}$ . Next define  $P_u = \mathbb{E}(H_u(\mathbf{x}))$  for  $\mathbf{x} \sim \mathcal{N}(0, I)$ . Let  $S(\mathbf{x}) = \sum_{j=1}^J H_j(\mathbf{x})$  count the number of rare events that happen. For  $s = 0, 1, \dots, J$ , let  $T_s = \Pr(S = s)$  give the distribution of  $S$ . We use  $|u|$  for the cardinality of  $u$ .

Our estimand is

$$\mu = P_{1:J} = \sum_{|u|>0} (-1)^{|u|-1} P_u, \quad (2)$$

by inclusion-exclusion. Our approach to estimating  $\mu$  is through mixture importance sampling. For  $j = 1, \dots, J$ , let  $q_j = \mathcal{L}(\mathbf{x} \mid \omega_j^\top \mathbf{x} \geq \tau_j)$  and take

$q_0 = p = \mathcal{N}(0, I)$ , with  $P_0 = 1$  and  $H_0(\mathbf{x}) = 1$ . The probability density functions for  $j > 0$  are  $q_j(\mathbf{x}) = p(\mathbf{x})H_j(\mathbf{x})/P_j$ , while  $q_0 = p$ .

Let  $\alpha_0, \alpha_1, \dots, \alpha_J$  be nonnegative numbers summing to 1, and  $q_\alpha = \sum_{j=0}^J \alpha_j q_j$ . The mixture importance sampling estimate of  $\mu$  based on  $n$  draws  $\mathbf{x}_i \sim q_\alpha$  is

$$\hat{\mu}_\alpha = \frac{1}{n} \sum_{i=1}^n \frac{p(\mathbf{x}_i)H_{1:J}(\mathbf{x}_i)}{\sum_{j=0}^J \alpha_j q_j(\mathbf{x}_i)} = \frac{1}{n} \sum_{i=1}^n \frac{H_{1:J}(\mathbf{x}_i)}{\sum_{j=0}^J \alpha_j H_j(\mathbf{x}_i) P_j^{-1}}. \quad (3)$$

Here the  $\alpha_0$  term represents a defensive mixture component (Hesterberg, 1995). Although the inclusion-exclusion formula (2) contains  $2^J - 1$  nonzero terms, each summand in the unbiased estimate in (3) can be computed at cost  $O(J)$ .

It is interesting here to discard the defensive component, taking  $\alpha_0 = 0$ . Now with  $\alpha_j = \alpha_j^* = P_j/\bar{\mu}$ , we get

$$\hat{\mu}_{\alpha^*} = \frac{\bar{\mu}}{n} \sum_{i=1}^n \frac{H_{1:J}(\mathbf{x}_i)}{\sum_{j=1}^J H_j(\mathbf{x}_i)} = \frac{\bar{\mu}}{n} \sum_{i=1}^n \frac{1}{S(\mathbf{x}_i)}, \quad \mathbf{x}_i \stackrel{\text{iid}}{\sim} q_{\alpha^*}, \quad (4)$$

because  $H_{1:J}(\mathbf{x}) = 1$  always holds for  $\mathbf{x} \sim q_{\alpha^*}$ . The estimate (4) is a multiplicative adjustment to the union bound  $\bar{\mu}$ . The terms  $S(\mathbf{x}_i)^{-1}$  range from 1 to  $1/J$  and so we will never get  $\hat{\mu}_{\alpha^*}$  larger than the union bound or smaller than  $\bar{\mu}/J$ . This is convenient because  $\bar{\mu} \geq \mu \geq \bar{\mu}/J$  always holds.

**Theorem 1.** *Let  $\hat{\mu}_{\alpha^*}$  be given by (4). Then*

$$\mathbb{E}(\hat{\mu}_{\alpha^*}) = \mu, \quad (5)$$

and

$$\text{Var}(\hat{\mu}_{\alpha^*}) = \frac{1}{n} \left( \bar{\mu} \sum_{s=1}^J \frac{T_s}{s} - \mu^2 \right) \leq \frac{\mu(\bar{\mu} - \mu)}{n}. \quad (6)$$

*Proof.* See the appendix, where this is proved for a general set of  $J$  events, not necessarily from Gaussian half-spaces.  $\square$

The union bound can be written

$$\bar{\mu} = \sum_{j=1}^J \Pr(H_j(\mathbf{x})) = \sum_{j=1}^J \sum_{u \subseteq 1:J} \mathbb{E}(H_u(\mathbf{x})H_{-u}^c(\mathbf{x})) 1_{j \in u} = \sum_{s=1}^J sT_s.$$

That is  $\bar{\mu} = \mathbb{E}(S(\mathbf{x})) = \mu \mathbb{E}(S(\mathbf{x}) \mid S(\mathbf{x}) > 0)$  and so we may write (6) as

$$\text{Var}(\hat{\mu}_{\alpha^*}) = \frac{\mu^2}{n} \left( \mathbb{E}(S \mid S > 0) \mathbb{E}(S^{-1} \mid S > 0) - 1 \right). \quad (7)$$

## 2.1 General Gaussians

Now suppose that we are given  $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\eta}, \Sigma)$  and the half-spaces are defined by  $\boldsymbol{\gamma}_j^\top \mathbf{y} \geq \kappa_j$ . We assume that  $\Sigma$  is nonsingular. If it is not, then we can reduce  $\mathbf{y}$  to a subset of components whose variance is nonsingular, and write the other components as linear functions of this reduced set. We also assume that we can afford to take a matrix square root  $\Sigma^{1/2}$ . Now  $\mathbf{x} = \Sigma^{-1/2}(\mathbf{y} - \boldsymbol{\eta}) \sim \mathcal{N}(0, I)$ , and  $\mathbf{y} = \boldsymbol{\eta} + \Sigma^{1/2}\mathbf{x}$ . Then the half-spaces are given by

$$\boldsymbol{\omega}_j^\top \mathbf{x} \geq \tau_j, \quad \text{where} \quad \boldsymbol{\omega}_j = \frac{\boldsymbol{\gamma}_j^\top \Sigma^{1/2}}{\sqrt{\boldsymbol{\gamma}_j^\top \Sigma \boldsymbol{\gamma}_j}}, \quad \text{and} \quad \tau_j = \frac{\kappa_j - \boldsymbol{\gamma}_j^\top \boldsymbol{\eta}}{\sqrt{\boldsymbol{\gamma}_j^\top \Sigma \boldsymbol{\gamma}_j}},$$

for  $\mathbf{x} \sim \mathcal{N}(0, I)$ . For rare events, we will have  $\kappa_j > \boldsymbol{\gamma}_j^\top \boldsymbol{\eta}$ .

## 2.2 Sampling algorithms

We want to sample  $\mathbf{x} \sim \mathcal{N}(0, I)$  conditionally on  $\mathbf{x}^\top \boldsymbol{\omega} \geq \tau$  for a unit vector  $\boldsymbol{\omega}$  and scalar  $\tau$ . We can use the following steps:

- 1) Sample  $\mathbf{z} \sim \mathcal{N}(0, I)$ .
- 2) Sample  $u \sim \mathbb{U}(0, 1)$ .
- 3) Let  $\mathbf{y} = \Phi^{-1}(\Phi(\tau) + u(1 - \Phi(\tau)))$ .
- 4) Deliver  $\mathbf{x} = \boldsymbol{\omega} \mathbf{y} + (I - \boldsymbol{\omega} \boldsymbol{\omega}^\top) \mathbf{z}$ .

Step 3 generates a truncated Gaussian by a ‘complementary inversion’ method where ordinary inversion of the conditional CDF would use  $1 - u$  in place of  $u$ . If a random number generator produces values below the machine epsilon then complementary inversion has better resolution of extreme  $\mathbf{y}$  values than inversion has.

This algorithm above can be problematic numerically when  $\Phi(\tau)$  is close to 1 as it will be for very rare events. For instance, in the R language (R Core Team, 2015),  $\Phi(10)$  yields 1 and then  $\Phi^{-1}(\Phi(10) + u(1 - \Phi(10)))$  yields  $\infty$  for any  $u$ . Some of our electrical grid examples have  $\max_j \tau_j > 10^{10}$ . That is, some of the potential failure modes are virtually impossible.

Because  $\tau > 0$  might be quite large, we get better numerical stability by sampling  $\mathbf{x} \sim \mathcal{N}(0, I)$  conditionally on  $\mathbf{x}^\top \boldsymbol{\omega} \leq -\tau$  and then delivering  $-\mathbf{x}$ . The steps are as follows:

- 1) Sample  $\mathbf{z} \sim \mathcal{N}(0, I)$ .
- 2) Sample  $u \sim \mathbb{U}(0, 1)$ .
- 3) Let  $\mathbf{y} = \Phi^{-1}(u\Phi(-\tau))$ .
- 4) Let  $\mathbf{x} = \boldsymbol{\omega} \mathbf{y} + (I - \boldsymbol{\omega} \boldsymbol{\omega}^\top) \mathbf{z}$ .
- 5) Deliver  $\mathbf{x} = -\mathbf{x}$ .

Even a very small  $u = 10^{-12}$  combined with  $\tau = 10$  yields

$$\Phi^{-1}(10^{-12} \times \Phi(-10)) \doteq \Phi^{-1}(7.62 \times 10^{-36}) \doteq -12.44$$

without any underflow in the R language (R Core Team, 2015). In cases with some  $\tau_j \geq 10^{10}$  we will ordinarily get  $P_j = 0$  and then never sample conditionally on the corresponding  $H_j$ . We compute step 4 via  $\mathbf{x} = \boldsymbol{\omega} \mathbf{y} + \mathbf{z} - \boldsymbol{\omega}(\boldsymbol{\omega}^\top \mathbf{z})$  to avoid a potentially expensive multiplication  $(I - \boldsymbol{\omega} \boldsymbol{\omega}^\top) \mathbf{z}$ .



### 3 Importance sampling properties

As shown in the Appendix, Theorem 1 holds more generally than the Gaussian case. In this more general setting, we have  $J$  events,  $H_j$ , on a common sample space  $\mathcal{X}$  where  $\mathbf{x} \in \mathcal{X}$  has probability density  $p$ . Event  $H_j$  has probability  $P_j$ . We want  $\mu = \Pr(H)$  where  $H = \cup_j H_j$ . The union bound is  $\mu \leq \bar{\mu} = \sum_j P_j$ . We assume that  $0 < \bar{\mu} < \infty$ . The upper bound only has to be checked if  $J = \infty$ . If  $\bar{\mu} = 0$ , then we know  $\mu = 0$  without any sampling.

When we sample, we ensure that at least one rare event takes place every time, by first picking an event  $H_j$  with probability proportional to  $P_j$ . Then we sample  $\mathbf{x} \in \mathcal{X}$  conditionally on  $H_j$  and find  $S(\mathbf{x}) = \sum_{\ell=1}^J H_\ell(\mathbf{x})$ , the total number of events that occur. This includes  $H_j$  and so our sample values always have  $S(\mathbf{x}_i) \geq 1$ . The importance sampling estimate  $\hat{\mu}_{\alpha^*}$  averages  $\bar{\mu}/S(\mathbf{x}_i)$  over  $n$  independent replicates. As in the prior section, we use

$$T_s = \Pr(S(\mathbf{x}) = s) = \int_{\mathbb{R}^d} 1\{S(\mathbf{x}) = s\} p(\mathbf{x}) d\mathbf{x},$$

for the probability of exactly  $s$  events happening. Then the variance of  $\hat{\mu}$  is given by (6).

The optimal importance sampling distribution for estimating  $\mu$  is uniform on  $H = \{\mathbf{x} \mid H(\mathbf{x}) = 1\}$ . Sampling from this distribution would yield an estimate with variance zero. Not surprisingly, we are seldom able to do that in applications. Our ALORE sampler takes  $\mathbf{x} \in \mathcal{X}$  with probability proportional to  $S(\mathbf{x})$ , so it has support set  $H$ .

We think that many applications will have events  $H_j$  that rarely co-occur. In that case  $S(\mathbf{x})$  is nearly constant at 1 for  $\mathbf{x} \in H$ , and then our sampler is close to the optimal one. Other applications may have a few near duplicated events that co-occur often. One extreme setting has a common cause that triggers all  $J$  events at once and those events almost never arise outside of that common situation. In that case  $S(\mathbf{x})$  is again nearly constant on  $H$ , this time usually equal to  $J$ , and our sampler is again nearly optimal.

The variance bound  $\mu(\bar{\mu} - \mu)/n$  from (6) can be conservative. It stems from  $T_s/s \leq T_s$ , when  $s \geq 1$ . If  $\Pr_{\alpha^*}(S > 1)$  is appreciably large then the variance is less than our bound. Because  $\bar{\mu} \leq J\mu$  we also have

$$\text{Var}(\hat{\mu}_{\alpha^*}) \leq (J - 1) \frac{\mu^2}{n}. \quad (8)$$

Equality in (8) only happens if  $\bar{\mu} = J\mu$ . Then the  $J$  events are almost surely the same and  $\text{Var}(\hat{\mu}_{\alpha^*}) = 0$ . Therefore strict inequality always holds in equation (8). The trivial exception from  $\mu = 0$  is not possible because we have  $\bar{\mu} > 0$ . We can sharpen the bound (8) by using the following lemma.

**Lemma 1.** *Let  $S$  be a random variable supported on  $\{1, 2, \dots, J\}$  for  $J \in \mathbb{N}$ . Then*

$$\mathbb{E}(S)\mathbb{E}(S^{-1}) \leq \frac{J + J^{-1} + 2}{4} \quad (9)$$

with equality if and only if  $S \sim \mathbb{U}\{1, J\}$ .

*Proof.* See the appendix.  $\square$

Lemma 1 tells us that for  $J \geq 2$ , our worst case setting is one where half of the time that one or more events happen, exactly one happens and half of the time, all  $J$  of them happen. We might have guessed this, because with  $\bar{\mu} = \mathbb{E}(S)$  fixed at some level, the variance of  $S^{-1}$  is maximized when  $S^{-1}$  takes only the two extreme values 1 and  $J^{-1}$  with equal probability. The situation is a bit more complicated because  $\mathbb{E}(S)$  and  $\mathbb{E}(S^{-1})$  both depend on the distribution of  $S$  and we need to bound the product  $\mathbb{E}(S)\mathbb{E}(S^{-1})$ , where both expectations are under  $q_{\alpha^*}$ .

From Theorem 2 and Lemma 1, we get

$$\text{Var}(\hat{\mu}_{\alpha^*}) = \frac{\mu^2}{n} \left( \left( \sum_{s=1}^J s \frac{T_s}{\mu} \right) \left( \sum_{s=1}^J s^{-1} \frac{T_s}{\mu} \right) - 1 \right) \leq \frac{\mu^2}{n} \frac{J + J^{-1} - 2}{4}, \quad (10)$$

because  $T_s/\mu$  is a probability distribution on  $\{1, 2, \dots, J\}$ . This bound describes relative error without reference to the union bound  $\bar{\mu}$ .

A rare event estimator has bounded relative error if  $\text{Var}(\hat{\mu})/\mu^2$  remains bounded as one takes the limit in a sequence of problems (Asmussen and Glynn, 2007, Chapter VI). The sequence is typically one where the event of interest becomes increasingly rare. Equation (10) shows a bounded relative error property for ALORE in any sequence with  $J/n$  uniformly bounded and  $\mu > 0$ .

When we want an upper bound in terms of  $\bar{\mu}$ , then the right hand side of (6) is maximized over  $\mu$  at  $\mu = \bar{\mu}/2$ , yielding

$$\text{Var}(\hat{\mu}_{\alpha^*}) \leq \frac{\bar{\mu}^2}{4n}. \quad (11)$$

Using  $\mu \leq \bar{\mu}$  in (10) yields

$$\text{Var}(\hat{\mu}_{\alpha^*}) \leq \frac{\bar{\mu}^2}{4n} (J + J^{-1} - 2).$$

This is only an improvement on (11) when  $J = 2$ . The inequality  $\mu = \bar{\mu}$  cannot actually hold under the conditions where equality holds in (10).

Now suppose that  $f(\mathbf{x})$  is supported on  $H$  and we want to estimate  $\nu(f) = \nu = \int f(\mathbf{x})p(\mathbf{x}) d\mathbf{x} = \int_H f(\mathbf{x})p(\mathbf{x}) d\mathbf{x}$ . We may use ALORE, via

$$\hat{\nu} = \frac{\bar{\mu}}{n} \sum_{i=1}^n \frac{f(\mathbf{x}_i)}{S(\mathbf{x}_i)}, \quad \mathbf{x}_i \stackrel{\text{iid}}{\sim} q_{\alpha^*}.$$

Then by the same arguments used in the Appendix,

$$\mathbb{E}(\hat{\nu}) = \nu \quad \text{and} \quad \text{Var}(\hat{\nu}) = \frac{1}{n} \left( \bar{\mu} \int_H \frac{f(\mathbf{x})^2 p(\mathbf{x})}{S(\mathbf{x})} d\mathbf{x} - \nu^2 \right).$$

If  $f(\mathbf{x}) \in \{0, 1\}$ , then  $\text{Var}(\hat{\nu}) \leq \nu(\bar{\mu} - \nu)/n$ . That is, when  $f$  describes a rare event that can only occur if one or more of the  $H_j$  also occur, we can reduce its Monte Carlo variance from  $\nu(1 - \nu)/n$  to at most  $\nu(\bar{\mu} - \nu)/n$ , in cases where  $\bar{\mu} < 1$ . In the Gaussian context, such an  $f$  is only nonzero outside the polytope  $\mathcal{P}$ .

Now suppose that  $f(\mathbf{x}) = g(S(\mathbf{x}))$  only depends on  $S$  and  $g(0) = 0$ . For instance, in some applications we might want the probability that  $S(\mathbf{x}) > k$  where  $k \geq 1$  is the largest acceptable number of failure events. Then

$$\text{Var}(\hat{\nu}) = \frac{1}{n} \left( \bar{\mu} \sum_{s=1}^J \frac{g(s)^2 T_s}{s} - \nu^2 \right).$$

## 4 Comparisons

Here we consider some numerical examples comparing ALORE to `pmvnorm` from the R package `mvtnorm` (Genz et al., 2017). This package can make use of special properties of the Gaussian distribution, and it works in high dimensions.

We begin by describing `mvtnorm` based on Genz and Bretz (2009) and a personal communication from Alan Genz. The program computes

$$\Pr(\mathbf{a} \leq \mathbf{y} \leq \mathbf{b}) \equiv \Pr(a_j \leq y_j \leq b_j, j = 1, \dots, d)$$

for  $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\eta}, \Sigma)$ , where  $-\infty \leq a_j \leq b_j \leq \infty$  for  $j = 1, \dots, d$ , and  $\Sigma$  can be rank deficient. We can use it to compute  $\mu = \Pr(\sum_{j=1}^J \mathbf{1}\{\boldsymbol{\omega}_j^\top \mathbf{x} \geq \tau_j\} > 0)$  for  $\mathbf{x} \sim \mathcal{N}(0, I)$  via

$$1 - \mu = \Pr(\boldsymbol{\Omega}^\top \mathbf{x} \leq \mathcal{T}) = \Pr(\mathbf{y} \leq \mathcal{T}), \quad \mathbf{y} \sim \mathcal{N}(0, \boldsymbol{\Omega}^\top \boldsymbol{\Omega}).$$

The code can handle dimensions up to 1000. In our context, that means at most  $J = 1000$  half-spaces. The dimension  $d$  can be higher. The related `pmvt` function handles multivariate  $t$  random variables. The code has three different algorithms in it. One from Genz (2004) handles two and three dimensional semi-infinite regions, one from Miwa et al. (2003) is for dimensions up to 20 and the rest are handled by an algorithm from Genz and Bretz (2009). This latter algorithm uses a number of methods. It uses randomized Korobov lattice rules as described by Cranley and Patterson (1976) for the first 100 dimensions, in conjunction with antithetic sampling. There are usually 8 randomizations. For more than 100 dimensions it applies a method from Niederreiter (1972). There are a series of increasing sample sizes in use, and the method provides an estimated error (3.5 standard errors) based on the randomization. The approach is via sequential conditional sampling, after strategically ordering the variables (e.g., putting unconstrained ones first). The R package calls a FORTRAN program for the computation, so it is very fast. We use the default implementation which uses up to 25,000 quadrature points.

The main finding is that importance sampling is more effective when the polytope of interest is the complement of a rare event. This is not meant to

be a criticism of `pmvnorm`. That code was not specifically designed to compute the complement of a rare event. The comparison is relevant because we are not aware of alternative code tuned for the high dimensional rare event cases that we need, and `pmvnorm` is a well regarded and widely available general solution, that seemed to us like the best off-the-shelf tool.

## 4.1 Circumscribed polygon

Let  $\mathcal{P}(J, \tau) \subset \mathbb{R}^2$  be the regular polygon of  $J \geq 3$  sides circumscribed around the circle of radius  $\tau > 0$ . This polygon is the intersection of  $H_j^c$  where  $H_j = \{\mathbf{x} \in \mathbb{R}^2 \mid \omega_j^\top \mathbf{x} \geq \tau\}$  where  $\omega_j^\top = (\sin(2\pi j/J), \cos(2\pi j/J))$ , for  $j = 1, \dots, J$ . We want  $\mu = \Pr(\mathbf{x} \in \mathcal{P}^c)$  for  $\mathbf{x} \sim \mathcal{N}(0, I)$ . Here we know that  $\mu \leq \Pr(\chi_{(2)}^2 \geq \tau^2) = \exp(-\tau^2/2)$ . Also, the gap between the circle of radius  $\tau$  and the circumscribed polygon has area  $G(J, \tau) = (J \tan(\pi/J) - \pi)\tau^2$ . The bivariate Gaussian density in this gap is at most  $\exp(-\tau^2/2)/(2\pi)$ . Therefore

$$\Pr(\mathbf{x} \in \mathcal{P}^c) \geq \exp(-\tau^2/2) - G(J, \tau) \exp(-\tau^2/2)/(2\pi)$$

that is

$$1 \geq \frac{\Pr(\mathbf{x} \in \mathcal{P}^c)}{\exp(-\tau^2/2)} \geq 1 - \frac{G(J, 1)\tau^2}{2\pi} \doteq 1 - \frac{\pi^2\tau^2}{6J^2},$$

for large  $J$ .

For  $J = 360$  and  $\tau = 6$ , we have  $\mu \leq \exp(-18) \doteq 1.52 \times 10^{-8}$ . The lower bound is about 0.9995 times the upper bound, so we treat the upper bound as exact. Figure 2 shows histograms of 100 simulations of  $\hat{\mu}/\mu$  using ALORE and using `pmvnorm`. In this case ALORE is much more accurate. The mean square relative error  $\mathbb{E}((\hat{\mu}/\mu - 1)^2)$  is about 800-fold smaller for ALORE than `pmvnorm`. We also see that `pmvnorm` has high positive skewness and the histogram of estimates has most of its mass well below the mean.

Table 1 shows summary results for this problem with different values of  $\tau$ . We see that `pmvnorm` is superior when the event is not rare but ALORE is superior for rare events. The large error for `pmvnorm` with  $\tau = 5$  stemmed from a small number of outliers among the 100 trials.

The upper bound in equation (6) is  $\text{Var}(\hat{\mu}) \leq \mu(\bar{\mu} - \mu)/n$ , from which  $\mathbb{E}((\hat{\mu}/\mu - 1)^2) \leq (\bar{\mu}/\mu - 1)/n$ . For  $\tau = 6$  this yields about 0.022, which is over 20 times the actual mean squared relative error from Table 1.

It is possible that this example is artificially easy for importance sampling, due to the symmetry. Whichever half-space  $H_j$  we sample, the distribution of overlapping half-spaces  $H_k$  for  $k \neq j$  is the same. Two half-spaces differ from  $H_j$  by a one degree angle, two differ by a two degree angle and so on. To get a more varied range of overlap patterns, we replaced angles  $2\pi j/360$  by angles  $2\pi \times p(j)/360$  where  $p(j)$  is the  $j$ 'th prime among integers up to 360. There are 72 of them, of which the largest is 359. With  $\tau = 6$  and 100 replications using  $n = 1000$  points in importance sampling, we have variance of  $\hat{p}/\exp(-18)$  equal to 0.00077. The comparable figure for `mvtnorm` is 8.5. There were a few outliers there including one that was more than 6 times the union bound. The




Figure 2: Results of 100 estimates of the  $\Pr(\mathbf{x} \notin \mathcal{P}(360, 6))$ , divided by  $\exp(-6^2/2)$ . Left panel: ALORE. Right panel: pmvnorm.

gap between the prime angle polygon and the inscribed circle is larger than the one formed by the full polygon. Pooling all the importance sampling runs leaves an estimate of about  $0.94 \times \exp(-18)$  for  $\mu$ . In this example, we see importance sampling working quite well without symmetry.

## 4.2 High dimensional half-spaces

The previous example was low dimensional and each of the half-spaces sampled had numerous similar ones, differing in angle by a small number of degrees. Thus  $\mu$  was quite a bit smaller than  $\bar{\mu}$ . Here we consider a high dimensional setting where the half-spaces have less overlap.

Two uniform random unit vectors  $\omega_1$  and  $\omega_2$  in  $\mathbb{R}^d$  are very likely to be nearly orthogonal for large  $d$ . Then  $\mathbf{x}^\top \omega_j > \tau_j$  are nearly independent events. For independent events, we would have

$$\Pr(\mathbf{x} \notin \mathcal{P}) = 1 - \prod_{j=1}^J (1 - P_j).$$

To make  $\mathbf{x} \notin \mathcal{P}$  a rare event, the  $P_j$  must be small and then the probability above will be close to the union bound. Theorem 2 predicts good performance for importance sampling here.

For this test 200 sample problems were constructed. The dimensions were chosen by  $d \sim \mathbb{U}\{20, 50, 100, 200, 500\}$ . Then there were  $J \sim \mathbb{U}\{d/2, d, 2d\}$  constraints chosen with uniform random unit vectors  $\omega_j \in \mathbb{R}^d$ . The threshold  $\tau$

$\tau$	$\mu$	$\mathbb{E}((\hat{\mu}_{\text{ALORE}}/\mu - 1)^2)$	$\mathbb{E}((\hat{\mu}_{\text{MVN}}/\mu - 1)^2)$
2	$1.35 \times 10^{-01}$	0.000399	$9.42 \times 10^{-08}$
3	$1.11 \times 10^{-02}$	0.000451	$9.24 \times 10^{-07}$
4	$3.35 \times 10^{-04}$	0.000549	$2.37 \times 10^{-02}$
5	$3.73 \times 10^{-06}$	0.000600	$1.81 \times 10^{+00}$
6	$1.52 \times 10^{-08}$	0.000543	$4.39 \times 10^{-01}$
7	$2.29 \times 10^{-11}$	0.000559	$3.62 \times 10^{-01}$
8	$1.27 \times 10^{-14}$	0.000540	$1.34 \times 10^{-01}$

Table 1: Results from 100 computations of  $\Pr(\mathbf{x} \notin \mathcal{P}(360, \tau))$  for various  $\tau$ . The true mean  $\mu$  is very nearly  $\exp(-\tau^2/2)$ . Importance sampling is more accurate for large  $\tau$  (rare events), while `pmvnorm` is more accurate for small  $\tau$ .

was chosen so that  $\log_{10}$  of the union bound was  $\mathbb{U}[4, 8]$ , followed by rounding to two significant figures. Then  $\hat{\mu}$  was computed by importance sampling with  $n = 1000$  samples, and by `pmvnorm`. Figure 3 shows the results. The importance sampling value was always very close to the union bound which in turn is essentially equal to what one would see for independent events. The values from `pmvnorm` were usually too small. By construction the intersection probabilities are quite rare. In importance sampling, 77.5% of the simulations had no intersections among 1000 trials and the others had only a few intersections. Therefore it is clear that the probabilities should be close to the union bounds here. We also see evidence of strong positive skewness for `pmvnorm`.

## 5 Power system infeasibility

### 5.1 Model

Our power system models are based on a network of  $N$  nodes, called busses. Some busses put power into the network and others consume power. The  $M$  edges in the network correspond to power lines between busses. The network is ordinarily sparse, with  $M$  a small multiple of  $N$ .

The power production at bus  $i$  is  $p_i$ , with negative values indicating consumption. For some busses,  $p_i$  is tightly controlled and deterministic in the relevant time horizon. Other busses have random  $p_i$  corresponding, for example, to variable consumption levels, that we treat as independent. Busses corresponding to wind farms have random power production levels with meaningfully large correlations. Our models contain one special bus  $S$ , called the slack bus, at which the power is  $p_S = -\sum_{i \neq S} p_i$ . The total power in the system is zero because transmission power losses are ignored in the DC approximation that we use.

The power at all busses can be represented by the vector  $p = (p_F^\top, p_R^\top, p_S)^\top$  corresponding to fixed busses, ordinary random busses (including any correlated ones) and the slack bus. There are  $N_F$  fixed busses,  $N_S = 1$  slack bus and




Figure 3: Results of 200 estimates of the  $\mu$  for varying high dimensional problems with nearly independent events.

$N_R = N - N_F - N_S$  random busses apart from the slack bus. We will use  $1_R$  to denote a column vector of  $N_R$  ones, and  $I_R$  to denote the identity matrix of size  $N_R$  and similarly for  $1_F$  and  $I_F$ .

The power  $p_i$  at bus  $i$  must satisfy the constraints

$$\underline{p}_i \leq p_i \leq \bar{p}_i. \quad (12)$$

The vector  $p$  has a Gaussian distribution, determined entirely by the random components  $p_R \sim \mathcal{N}(\eta_R, \Sigma_{RR})$ . Therefore in the present context,  $p_R$  is the Gaussian random variable  $\mathbf{x}$  from Section 2. The fixed components satisfy  $p_F = \eta_F$  and then the slack bus satisfies  $p_S \sim \mathcal{N}(\eta_S, \Sigma_{SS})$  where  $\eta_S = -1_R^T \eta_R - 1_F^T \eta_F$  and  $\Sigma_{SS} = 1_R^T \Sigma_{RR} 1_R$ . Because all of the randomness comes from  $p_R$ , we will abbreviate  $\Sigma_{RR}$  to  $\Sigma$ .

The node to node inductances in the network form a Laplacian matrix  $B$  where  $B_{ij} \neq 0$  if busses  $i$  and  $j$  are connected with  $B_{ii} = -\sum_{j \neq i} B_{ij}$  (up to rounding). The Laplacian is symmetric and has one eigenvalue of zero for a connected network. It has a pseudo-inverse  $B^+$ . We partition  $B$  and  $B^+$  as follows

$$B = \begin{pmatrix} B_{RR} & B_{RF} & B_{RS} \\ B_{FR} & B_{FF} & B_{FS} \\ B_{SR} & B_{SF} & B_{SS} \end{pmatrix}, \quad \text{and} \quad B^+ = \begin{pmatrix} B^{RR} & B^{RF} & B^{RS} \\ B^{FR} & B^{FF} & B^{FS} \\ B^{SR} & B^{SF} & B^{SS} \end{pmatrix}.$$

We also group  $B^+$  into three sets of columns via  $B^+ = (B^{\bullet R} \quad B^{\bullet F} \quad B^{\bullet S})$ .

The phase at bus  $i$  is denoted  $\theta_i$ . In our DC approximation of AC power flow, the phases approximately satisfy  $B\theta = p$ . Given the power vector  $p$ , we

take

$$\theta = B^+ p = \begin{pmatrix} B^{RR} & B^{RF} & B^{RS} \\ B^{FR} & B^{FF} & B^{FS} \\ B^{SR} & B^{SF} & B^{SS} \end{pmatrix} \begin{pmatrix} p_R \\ p_F \\ p_S \end{pmatrix}.$$

The phase constraints on the network are

$$|\theta_i - \theta_j| \leq \bar{\theta}_{ij}, \quad \text{for } i \neq j \text{ and } B_{ij} \neq 0. \quad (13)$$

In our examples, all  $\bar{\theta}_{ij} = \bar{\theta}$  for a single value  $\bar{\theta}$  such as  $\pi/6$  or  $\pi/4$ .

Let  $D \in \{-1, 1\}^{M \times N}$  be the incidence matrix. Each edge in the network is represented by one row of  $D$  with an entry of  $+1$  for one of the busses on that edge and  $-1$  for the other. The phase constraints are  $|D\theta| \leq \bar{\theta}$  componentwise. Now

$$D\theta = D \begin{pmatrix} B^{\bullet R} & B^{\bullet F} & B^{\bullet S} \end{pmatrix} \begin{pmatrix} p_R \\ p_F \\ p_S \end{pmatrix} = D(B^{\bullet R} p_R + B^{\bullet F} p_F + B^{\bullet S} p_S).$$

The constraint that  $D\theta \leq \bar{\theta}$  for every edge  $ij$  can be written

$$DB^{\bullet R} p_R \leq \bar{\theta} - DB^{\bullet F} p_F - DB^{\bullet S} p_S.$$

Now  $p_S = -1_R^T p_R - 1_F^T p_F$  and  $p_F = \eta_F$ , so the constraint on  $p_R$  is

$$D(B^{\bullet R} - B^{\bullet S} 1_R^T) p_R \leq \bar{\theta} - D(B^{\bullet F} - B^{\bullet S} 1_F^T) \eta_F. \quad (14)$$

We also have constraints  $D\theta \geq -\bar{\theta}$  which can be written

$$D(B^{\bullet S} 1_R^T - B^{\bullet R}) p_R \leq \bar{\theta} + D(B^{\bullet F} - B^{\bullet S} 1_F^T) \eta_F. \quad (15)$$

Equations (14) and (15) supply  $2M$  constraints on the random vector  $p_R$ . We have also the two slack bus constraints  $p_S \leq \bar{p}_S$  and  $-p_S \leq -\underline{p}_S$ , that is

$$-1_R^T p_R \leq \bar{p}_S + 1_F^T \eta_F \quad \text{and} \quad 1_R^T p_R \leq -\underline{p}_S + 1_F^T \eta_F. \quad (16)$$

Finally, there are individual constraints on the random busses

$$p_R \leq \bar{p}_R, \quad \text{and} \quad -p_R \leq -\underline{p}_R,$$

componentwise.

When we combine all of the constraints on  $p_R$ , we get a matrix  $\Gamma$  with rows  $\gamma_j$  and a vector of upper bounds  $\mathcal{K}$  with entries  $\kappa_j$  for which the constraints are  $\Gamma \times p_R \leq \mathcal{K}$  componentwise. Here those matrices are

$$\Gamma = \begin{pmatrix} I_R \\ -I_R \\ 1_R^T \\ -1_R^T \\ D(B^{\bullet R} - B^{\bullet S} 1_R^T) \\ -D(B^{\bullet R} - B^{\bullet S} 1_R^T) \end{pmatrix}, \quad \text{and} \quad \mathcal{K} = \begin{pmatrix} \bar{p}_R \\ -\underline{p}_R \\ -\underline{p}_S + 1_F^T \eta_F \\ \bar{p}_S + 1_F^T \eta_F \\ \bar{\theta} - D(B^{\bullet F} - B^{\bullet S} 1_F^T) \eta_F \\ \bar{\theta} + D(B^{\bullet F} - B^{\bullet S} 1_F^T) \eta_F \end{pmatrix}.$$



$\bar{\omega}$	$\hat{\mu}$	se/ $\hat{\mu}$	$\underline{\mu}$	$\bar{\mu}$
$\pi/4$	$3.7 \times 10^{-23}$	0.0024	$3.6 \times 10^{-23}$	$4.2 \times 10^{-23}$
$\pi/5$	$2.6 \times 10^{-12}$	0.0022	$2.6 \times 10^{-12}$	$2.9 \times 10^{-12}$
$\pi/6$	$3.9 \times 10^{-07}$	0.0024	$3.9 \times 10^{-07}$	$4.4 \times 10^{-07}$
$\pi/7$	$2.0 \times 10^{-03}$	0.0027	$2.0 \times 10^{-03}$	$2.4 \times 10^{-03}$

Table 2: Rare event estimates for the winter peak grid.  $\bar{\omega}$  is the phase constraint,  $\hat{\mu}$  is the ALORE estimate, se is the estimated standard error,  $\underline{\mu}$  is the largest single event probability and  $\bar{\mu}$  is the union bound.

These are the linear constraints on  $p_R$ . There are  $2M$  phase constraints and there are two constraints for all of the non-fixed busses, including the slack bus. These constraints can be turned into constraints  $\Omega$  and  $\mathcal{T}$  on a  $\mathcal{N}(0, I_R)$  vector as described in Section 2.1.

## 5.2 Examples

We considered several model electrical grids included in the MATPOWER distribution (Zimmerman et al., 2011). In each case we modeled violations of the phase constraints, and used  $n = 10,000$  samples. For some cases we found that, under our model, phase constraint violations were not rare events. In some other cases, the rare event probability was dominated by one single phase condition:  $\underline{\mu} = \max_{j=1}^J P_j \approx \sum_{j=1}^J P_j = \bar{\mu}$ . For cases like this there is no need for elaborate computation because we know  $\mu$  is within a narrow interval  $[\underline{\mu}, \bar{\mu}]$ . The interesting cases were of rare events not dominated by a single failure mode. We investigate two of them.

The first is the Polish winter peak grid of 2383 busses. There were  $d = 326$  random (uncontrolled) busses and  $J = 5772$  phase constraints. We varied  $\bar{\omega}$  as shown in Table 2. For  $\bar{\omega} = \pi/7$  constraint violations are not very rare. At  $\bar{\omega} = \pi/4$  they are quite rare. The estimated coefficient of variation is nearly constant over this range.

The second interesting case is the Pegase 2869 model of Fliscounakis et al. (2013). This has  $d = 509$  uncontrolled busses and  $J = 7936$  phase constraints. It is described as “power flow for a large part of the European system”. The results are shown in Table 3. We include an unrealistically large bound  $\bar{\omega} = \pi/2$  in that table, to test the limits of our approach. For  $\bar{\omega} = \pi/2$ , the standard error given is zero. One half-space was sampled 9408 times, another was sampled 592 times but in no instance were there two or more phase violations. The estimate reverts to the union bound. Getting 0 doubletons ( $S = 2$ ) among  $n = 10,000$  tries is compatible with the true probability of a doubleton being as high as  $3/n$ . Even if  $T_1 = .9997$  and  $T_2 = .0003$  then we would have  $\mu = (1 - (3/2) \times 10^{-4})\bar{\mu}$  instead of  $\bar{\mu}$ . We return to this issue in the discussion.

In addition to the examples above we investigated IEEE case 14, IEEE case 300, and Pegase 1354, which were all dominated by one failure. We considered

$\bar{\omega}$	$\hat{\mu}$	se/ $\hat{\mu}$	$\underline{\mu}$	$\bar{\mu}$
$\pi/2$	$3.5 \times 10^{-20}$	0*	$3.3 \times 10^{-20}$	$3.5 \times 10^{-20}$
$\pi/3$	$8.9 \times 10^{-10}$	$5.0 \times 10^{-5}$	$7.7 \times 10^{-10}$	$8.9 \times 10^{-10}$
$\pi/4$	$4.3 \times 10^{-06}$	$1.8 \times 10^{-3}$	$3.5 \times 10^{-06}$	$4.6 \times 10^{-06}$
$\pi/5$	$2.9 \times 10^{-03}$	$3.5 \times 10^{-3}$	$1.8 \times 10^{-03}$	$4.1 \times 10^{-03}$

Table 3: Rare event estimates for the Pegase 2869 model. The columns are as in Table 2. \*The se was 0 because there were never two or more failures in any sample. See text for discussion.

a system which included random and correlated wind power generators, but phase failure was not a rare event in that system. Pegase 13659 was too large for our computation. The Laplacian matrix has 37,250 rows and columns and we use the SVD to compute the generalized inverses we need. Pegase 9241 was large enough to be very slow and it did not have rare failures.

## 6 Discussion

We have introduced a version of mixture importance sampling for problems with multiple failure modes. The sample values are constrained to have at least one failure and we obtain bounded relative error.

The ALORE importance sampler is more accurate than a state of the art code for computing high dimensional Gaussian probabilities in our rare event setting, but not otherwise.

We have noticed two areas where ALORE can be improved. First, if we never see  $S \geq 2$  concurrent rare events, ALORE will return the union bound  $\hat{\mu} = \bar{\mu}$ , with an estimated variance of zero. That variance estimate could be undesirable even when  $\mu/\bar{\mu} \approx 1$ . Because  $S(\mathbf{x})$  is supported on  $\{1, 1/2, \dots, 1/J\}$  we can get an interval estimate of  $\mu$  by putting a multinomial prior on this support set and using the posterior distribution given the sample.

A second and related issue is that while  $\bar{\mu} \geq \hat{\mu} \geq \bar{\mu}/J$  always occurs, it is possible to get  $\hat{\mu} < \underline{\mu} = \max_{1 \leq j \leq J} P_j$ . We have seen this in cases where  $\bar{\mu} \approx \underline{\mu}$  because one of the  $P_j$  dominates all of the others combined. In such cases  $\mu$ ,  $\bar{\mu}$  and  $\underline{\mu}$  are all very close together and  $\hat{\mu}$  has small relative standard deviation. Improving these two issues is outside the scope of this paper. They are both things that happen in cases where we already have a very good idea of the magnitude of  $\mu$ .

In large problems the algebra can potentially be reduced by ignoring the very rarest events and simply adding their probabilities to the estimate. This will provide a mildly conservative bound. There is also the possibility of exploiting many generalized upper and lower bounds on the probability of a union. See for instance the survey by Yang et al. (2014).

## Acknowledgments

We thank Alan Genz for providing details of the `mvtnorm` package. We thank Yanbo Tang and Jeffrey Negrea for noticing that Lemma 1 could be proved by Cauchy-Schwarz, which is shorter than our original proof and also establishes necessity. ABO thanks the Center for Nonlinear Studies at Los Alamos for their hospitality while he visited. The work of YM and MC was funded by DOE/GMLC 2.0 project: “Emergency Monitoring and controls through new technologies and analytics”. ABO was supported by the NSF under grants DMS-1521145 and DMS-1407397.

## References

- Asmussen, S. and Glynn, P. W. (2007). *Stochastic simulation: algorithms and analysis*, volume 57. Springer Science & Business Media.
- Chertkov, M., Pan, F., and Stepanov, M. G. (2011a). Predicting failures in power grids: The case of static overloads. *IEEE Transactions on Smart Grid*, 2(1):162–172.
- Chertkov, M., Stepanov, M., Pan, F., and Baldick, R. (2011b). Exact and efficient algorithm to discover extreme stochastic events in wind generation over transmission power grids. In *2011 50th IEEE Conference on Decision and Control and European Control Conference*, pages 2174–2180.
- Cornuet, J., Marin, J.-M., Mira, A., and Robert, C. P. (2012). Adaptive multiple importance sampling. *Scandinavian Journal of Statistics*, 39(4):798–812.
- Cranley, R. and Patterson, T. N. L. (1976). Randomization of number theoretic methods for multiple integration. *SIAM Journal of Numerical Analysis*, 13(6):904–914.
- Elvira, V., Martino, L., Luengo, D., and Bugallo, M. F. (2015a). Efficient multiple importance sampling estimators. *IEEE Signal Processing Letters*, 22(10):1757–1761.
- Elvira, V., Martino, L., Luengo, D., and Bugallo, M. F. (2015b). Generalized multiple importance sampling. *arXiv preprint arXiv:1511.03095*.
- Fliscounakis, S., Panciatici, P., Capitanescu, F., and Wehenkel, L. (2013). Contingency ranking with respect to overloads in very large power systems taking into account uncertainty, preventive, and corrective actions. *IEEE Transactions on Power Systems*, 28(4):4909–4917.
- Genz, A. (2004). Numerical computation of rectangular bivariate and trivariate normal and t probabilities. *Statistics and Computing*, 14(3):251–260.
- Genz, A. and Bretz, F. (2009). *Computation of Multivariate Normal and t Probabilities*. Springer-Verlag, Berlin.

- Genz, A., Bretz, F., Miwa, T., Mi, X., Leisch, F., Scheipl, F., and Hothorn, T. (2017). *mvtnorm: Multivariate Normal and t Distributions*. R package version 1.0-6.
- Hesterberg, T. C. (1995). Weighted average importance sampling and defensive mixture distributions. *Technometrics*, 37(2):185–192.
- Kersulis, J., Hiskens, I., Chertkov, M., Backhaus, S., and Bienstock, D. (2015). Temperature-based instanton analysis: Identifying vulnerability in transmission networks. In *2015 IEEE Eindhoven PowerTech*, pages 1–6.
- Lafortune, E. P. and Willems, Y. D. (1993). Bidirectional path tracing. In *Proceedings of CompuGraphics*, pages 95–104.
- Liu, J. S. (2001). *Monte Carlo strategies in scientific computing*. Springer, New York.
- Miwa, T., Hayter, A., and Kuriki, S. (2003). The evaluation of general non-centred orthant probabilities. *Journal of the Royal Statistical Society, Series B*, 65(1):223–234.
- Niederreiter, H. (1972). On a number-theoretical integration method. *Aequationes Math*, 8(3):304–311.
- Owen, A. B. and Zhou, Y. (2000). Safe and effective importance sampling. *Journal of the American Statistical Association*, 95(449):135–143.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Sauer, P. W. and Christensen, J. P. (1984). Active linear DC circuit models for power system analysis. *Electric machines and power systems*, 9(2-3):103–112.
- Stott, B., Jardim, J., and Alsaç, O. (2009). DC power flow revisited. *IEEE Transactions on Power Systems*, 24(3):1290–1300.
- Van den Bergh, K., Delarue, E., and Dhaeseleer, W. (2014). Dc power flow in unit commitment models. Technical Report WP EN2014-12, KU Leuven.
- Veach, E. and Guibas, L. (1994). Bidirectional estimators for light transport. In *5th Annual Eurographics Workshop on Rendering*, pages 147–162.
- Yang, J., Alajaji, F., and Takahara, G. (2014). A short survey on bounding the union probability using partial information. Technical report, University of Toronto.
- Zimmerman, R. D., Murillo-Sánchez, C. E., and Thomas, R. J. (2011). MATPOWER: steady-state operations, planning, and analysis tools for power systems research and education. *IEEE Transactions on power systems*, 26(1):12–19.

## Appendix: Proofs

### Proof of Theorem 2

Our motivating problem involves probabilities defined by Gaussian content of half-spaces. The approach generalizes to the estimating the probability of the union of any finite set of events. We can also consider a countable number of events when the union bound is finite; see remarks below. We will assume that each event has positive probability, but that condition can also be weakened to a positive union bound, as described at the end of this section.

For definiteness, we define our sets in terms of indicator functions of a random variable  $\mathbf{x} \in \mathbb{R}^d$  with probability density  $p$ . The same formulas work for general sample spaces and the density can be with respect to an arbitrary base measure.

We cast our notation into this more general setting as follows. For  $J \geq 1$ , and  $j = 1, \dots, J$ , let the subset  $H_j \subset \mathbb{R}^d$  define both the event  $H_j = \mathbf{1}\{\mathbf{x} \in H_j\}$  and the indicator function  $H_j(\mathbf{x}) = \mathbf{1}\{\mathbf{x} \in H_j\}$ . For  $u \subseteq \{1, \dots, J\}$  we let  $H_u = \cup_{j \in u} H_j$  and  $H_u(\mathbf{x}) = \max_{j \in u} H_j(\mathbf{x})$ , with  $H_\emptyset(\mathbf{x}) = 0$ . As before  $P_u = \mathbb{E}(H_u(\mathbf{x}))$ , the number of events is  $S(\mathbf{x}) = \sum_{j=1}^J H_j(\mathbf{x})$  and  $\Pr(S = s) = T_s$ .

Next we introduce some further notation. We use  $-u$  for complements with respect to  $1:J$ , especially within subscripts, and  $H_u^c(\mathbf{x})$  for the complementary outcome  $1 - H_u(\mathbf{x})$ . Then  $H_u(\mathbf{x})H_{-u}^c(\mathbf{x})$  describes the event where  $\mathbf{x}_j \in H_j$  if and only if  $j \in u$ .

If  $P_j > 0$ , then the distribution  $q_j$  of  $\mathbf{x}$  given  $H_j$  is well defined:  $q_j(\mathbf{x}) = p(\mathbf{x})H_j(\mathbf{x})/P_j$ . If  $\min_j P_j > 0$  then we can define the mixture distribution

$$q_{\alpha^*} = \sum_{j=1}^J \alpha_j^* q_j, \quad \alpha_j^* = P_j/\bar{\mu}, \quad \bar{\mu} = \sum_{j=1}^J P_j. \quad (17)$$

For  $n \geq 1$ , our estimator of  $\mu = \Pr(S(\mathbf{x}) > 0)$  is

$$\hat{\mu}_{\alpha^*} = \frac{\bar{\mu}}{n} \sum_{i=1}^n \frac{1}{S(\mathbf{x}_i)}, \quad \mathbf{x}_i \stackrel{\text{iid}}{\sim} q_{\alpha^*}. \quad (18)$$

In this section, some equations include both randomness due to  $\mathbf{x}_i \sim q_{\alpha^*}$  and randomness due to  $\mathbf{x} \sim p$ . For section only, we use  $\Pr_*$ ,  $\mathbb{E}_*$  and  $\text{Var}_*$  when the randomness is from observations  $\mathbf{x}_i \sim q_{\alpha^*}$ , while  $\mathbb{E}$ ,  $\Pr$  and  $\text{Var}$  are with respect to  $\mathbf{x} \sim p$ .

**Theorem 2.** *If  $1 \leq J < \infty$  and  $\min_j P_j > 0$  and  $n \geq 1$ , then  $\hat{\mu}_{\alpha^*}$  defined by (17) satisfies  $\Pr_*(\bar{\mu}/J \leq \hat{\mu}_{\alpha^*} \leq \bar{\mu}) = 1$ ,*

$$\mathbb{E}_*(\hat{\mu}_{\alpha^*}) = \mu, \quad (19)$$

and

$$\text{Var}_*(\hat{\mu}_{\alpha^*}) = \frac{1}{n} \left( \bar{\mu} \sum_{s=1}^J \frac{T_s}{s} - \mu^2 \right) \leq \frac{\mu(\bar{\mu} - \mu)}{n}. \quad (20)$$

*Proof of Theorem 2.* Let  $H(\mathbf{x}) = \max_{1 \leq j \leq J} H_j(\mathbf{x})$  and  $H = \{\mathbf{x} \mid H(\mathbf{x}) = 1\}$ . If  $\mathbf{x}_i \sim q_{\alpha^*}$ , then  $\mathbf{x}_i \in H$  always holds. Then  $1 \leq S(\mathbf{x}_i) \leq J$  holds establishing the bounds on  $\hat{\mu}_{\alpha^*}$ . Next

$$\mathbb{E}_* \left( \left( \sum_{j=1}^J H_j(\mathbf{x}_1) \right)^{-1} \right) = \sum_{\ell=1}^J \frac{P_\ell}{\bar{\mu}} \int_H \frac{H_\ell(\mathbf{x}) P_\ell^{-1} p(\mathbf{x})}{\sum_{j=1}^J H_j(\mathbf{x})} d\mathbf{x} = \frac{1}{\bar{\mu}} \int_H p(\mathbf{x}) d\mathbf{x} = \frac{\mu}{\bar{\mu}},$$

establishing (19).

Because  $\hat{\mu}_{\alpha^*}$  is unbiased its variance is

$$\frac{1}{n} \left( \bar{\mu}^2 \mathbb{E}_* \left( \left( \frac{H(\mathbf{x}_1)}{\sum_{j=1}^J H_j(\mathbf{x}_1)} \right)^2 \right) - \mu^2 \right). \quad (21)$$

Next

$$\begin{aligned} \mathbb{E}_* \left( \left( \frac{H(\mathbf{x}_1)}{\sum_j H_j(\mathbf{x}_1)} \right)^2 \right) &= \sum_{j=1}^J \frac{P_j}{\bar{\mu}} \int_H \left( \frac{H(\mathbf{x})}{1 + \sum_{\ell \neq j} H_\ell(\mathbf{x})} \right)^2 \frac{p(\mathbf{x}) H_j(\mathbf{x})}{P_j} d\mathbf{x} \\ &= \bar{\mu}^{-1} \int_H \frac{p(\mathbf{x})}{\sum_{j=1}^J H_j(\mathbf{x})} d\mathbf{x} \\ &= \bar{\mu}^{-1} \sum_{|u|>0} \frac{1}{|u|} \int H_u(\mathbf{x}) H_{-u}^c(\mathbf{x}) p(\mathbf{x}) d\mathbf{x} \\ &= \bar{\mu}^{-1} \sum_{s=1}^J \frac{1}{s} T_s, \end{aligned}$$

which, with (21), establishes the equality in (20). Next

$$\bar{\mu}^{-1} \sum_{s=1}^J \frac{T_s}{s} \leq \bar{\mu}^{-1} \sum_{s=1}^J T_s = \bar{\mu}^{-1} (1 - T_0) = \bar{\mu}^{-1} \mu.$$

Finally  $\bar{\mu}^2 (\bar{\mu}^{-1} \mu) - \mu^2 = (\bar{\mu} - \mu) \mu$ , establishing the upper bound in (20).  $\square$

We can generalize the previous theorem to higher moments. Our estimate is an average of  $\bar{\mu}/S(\mathbf{x}_i)$ . The  $k$ 'th moment of this quantity is

$$\begin{aligned} \mathbb{E}_* \left( \left( \frac{\bar{\mu}}{S(\mathbf{x}_1)} \right)^k \right) &= \bar{\mu}^k \sum_{j=1}^J \alpha_j \int_H S(\mathbf{x})^{-k} q_j(\mathbf{x}) d\mathbf{x} \\ &= \bar{\mu}^k \sum_{j=1}^J \frac{P_j}{\bar{\mu}} \int_H S(\mathbf{x})^{-k} \frac{p(\mathbf{x}) H_j(\mathbf{x})}{P_j} d\mathbf{x} \\ &= \bar{\mu}^{k-1} \int_H S(\mathbf{x})^{1-k} d\mathbf{x} \\ &= \sum_{s=1}^J T_s \left( \frac{\bar{\mu}}{s} \right)^{k-1}. \end{aligned}$$

**Remark 1.** Suppose that one of the  $P_j = 0$  but  $\bar{\mu} > 0$ . In this case  $q_j$  is not well defined. However  $q_{\alpha^*}$  places probability 0 on  $q_j$ , so we may delete the  $q_j$  component without changing the algorithm and then sampling from  $q_{\alpha^*}$  is well defined.

**Remark 2.** Next suppose that there are infinitely many events, one for each  $j \in \mathbb{N}$ . If  $\bar{\mu} \in (0, \infty)$ , then  $q_{\alpha^*}$  is well defined. The same proof goes through, only now sums over  $1:J$  must be replaced by sums over  $\mathbb{N}$ .

### Proof of Lemma 1

From the Cauchy-Schwarz inequality,

$$1 - \mathbb{E}(S)\mathbb{E}(S^{-1}) = \text{Cov}(S, S^{-1}) \geq -\sqrt{\text{Var}(S)\text{Var}(S^{-1})}. \quad (22)$$

Now  $\text{Var}(S) \leq (J-1)^2/4$  and  $\text{Var}(S^{-1}) \leq (1-J^{-1})^2/4$  because the support of  $S$  is in  $[1, J]$ . Therefore

$$\mathbb{E}(S)\mathbb{E}(S^{-1}) \leq 1 + \frac{(J-1)(1-J^{-1})}{4} = \frac{J+J^{-1}+2}{4}.$$

Finally, the unique distribution for which  $\text{Var}(S)$ ,  $\text{Var}(S^{-1})$  and  $-\text{Corr}(S, S^{-1})$  all attain their maxima is  $\mathbb{U}\{1, J\}$ .  $\square$

### Generalization

The lemma generalizes. If  $\Pr(a \leq X \leq b) = 1$  for  $0 < a \leq b < \infty$  then the same argument yields  $\mathbb{E}(X)\mathbb{E}(X^{-1}) \leq (a/b + b/a + 2)/4$ .