

ESTIMATION OF SMALL FAILURE PROBABILITIES IN HIGH DIMENSIONS BY ADAPTIVE LINKED IMPORTANCE SAMPLING

L.S. Katafygiotis¹ and K.M. Zuev¹

¹ Hong Kong University of Science and Technology,
Department of Civil Engineering,
Clear Water Bay, Kowloon, Hong Kong, China,
email: {lambros, zuev}@ust.hk

Key words: Reliability, Importance sampling, Subset simulation, Bridging, MCMC.

Abstract. *A novel simulation approach, called Adaptive Linked Importance Sampling (ALIS), is proposed to compute small failure probabilities encountered in high-dimensional reliability analysis of engineering systems.*

It was shown by Au and Beck (2003) that Importance Sampling (IS) does generally not work in high dimensions. A geometric understanding of why this is true when one uses a fixed importance sampling density (ISD) was given in Katafygiotis and Zuev (2006). In this paper we introduce an algorithm referred to as Adaptive Linked Importance Sampling (ALIS). The basic idea of ALIS is instead of using a fixed ISD, as done in standard IS, to use a family of intermediate distributions that will converge to the target optimal ISD corresponding to the conditional probability given the failure event. We show that Subset Simulation (SS), which was introduced by Au and Beck (2001), does correspond to a special case of ALIS where the intermediate ISD's are chosen to correspond to the conditional distributions given adaptively chosen intermediate nested failure events. However, the general formulation of ALIS allows for a much richer choice of intermediate ISD's. As the concept of subsets is not a central feature of ALIS, the failure probability is not any longer expressed as a product of conditional failure probabilities as in the case of subset simulation. Starting with a one-dimensional example we demonstrate that ALIS can offer drastic improvements over subset simulation. In particular, we show that the choice of intermediate ISD's prescribed by subset simulation is far from optimal. The generalization to high-dimensions is discussed. The accuracy and efficiency of the method is demonstrated with numerical examples.

1 INTRODUCTION

In reliability engineering our task is to calculate the reliability or equivalently the probability of failure of a system under given loading conditions. The mathematical models of the input load x and the system output $f(x)$ are random vectors $x \in \mathbb{R}^N$ with joint probability density function (PDF) π_0 and function $f : \mathbb{R}^N \rightarrow \mathbb{R}_+$ correspondingly. For example, if our system corresponds to a tall building, the input may represent wind velocities along the building height and the output may represent the maximum roof displacement or the maximum interstory drift (absolute value) under the given wind load. Define the failure domain F as the set of inputs that lead to the exceedance of some prescribed critical threshold $b \in \mathbb{R}_+$:

$$F = \{x \in \mathbb{R}^N | f(x) > b\} \quad (1)$$

We are interested in the probability of failure :

$$p_F = P(x \in F) = \int_F \pi_0(x) dx = \int_{\mathbb{R}^N} I_F(x) \pi_0(x) dx = E_{\pi_0}[I_F] \quad (2)$$

where I_F is the indicator function ($= 1$ if $x \in F$, $= 0$ otherwise) and E_{π_0} denotes expectation with respect to the distribution π_0 , where π_0 denotes the PDF of the input random variables x .

Among all procedures developed for estimation of p_F , a prominent position is held by stochastic simulation methods. The expression of p_F as a mathematical expectation (2) renders standard Monte Carlo method [8] directly applicable, where p_F is estimated as a sample average of I_F over independent and identically distributed samples of x drawn from the PDF π_0 :

$$\hat{p}_{mc} = \frac{1}{n} \sum_{k=1}^n I_F(x^{(k)}), \quad x^{(k)} \sim \pi_0 \quad (3)$$

This estimate is unbiased and the coefficient of variation (COV), serving as a measure of the statistical error, is

$$\delta_{mc} = \sqrt{\frac{(1 - p_F)}{np_F}} \quad (4)$$

Although standard Monte Carlo is independent of the dimension N of the parameter space, it is inefficient in estimating small probabilities because it requires a large number of samples ($\sim 1/p_F$) to achieve an acceptable level of accuracy. For example, if $p_F = 10^{-4}$ and we want to achieve an accuracy of $\delta_{mc} = 10\%$ we need approximately 10^6 samples. Therefore, standard Monte Carlo becomes computationally prohibitive for our problems of interest involving small failure probabilities.

Importance Sampling is a fundamental technique in stochastic simulation that tries to reduce the COV of the Monte Carlo estimate. In statistical physics literature this procedure is also called ‘‘simple importance sampling’’ and ‘‘free energy perturbation’’. The basic idea of Importance Sampling is to generate more samples in the important region of the failure domain, i.e., in the region of the failure domain that contains most of the probability mass and, therefore, contributes mostly to the integral (2). Roughly speaking standard Monte Carlo does not work because the vast majority of terms in the sum (3) are zero and only very few are equal to one. Using Importance Sampling we

want instead of estimating p_F as the average of a vast majority of 0's and, occasionally some (if any) 1's, to calculate it as the average of less zeros and many more nonzero small numbers, each being ideally of the order of p_F .

Specifically, let π_1 be any PDF on the parameter space \mathbb{R}^N such that its support (domain where π_1 is not zero) contains the intersection of failure domain with the support of π_0 :

$$\text{supp } \pi_1 \supset F \cap \text{supp } \pi_0 \quad (5)$$

Then we can rewrite the probability integral (2) as follows :

$$p_F = \int_{\mathbb{R}^N} I_F(x) \pi_0(x) dx = \int_{\mathbb{R}^N} I_F(x) \frac{\pi_0(x)}{\pi_1(x)} \pi_1(x) dx = E_{\pi_1} \left[I_F \frac{\pi_0}{\pi_1} \right] \quad (6)$$

Suppose that we are able to generate random samples from π_1 , called importance sampling density (ISD), and compute the value of $\pi_1(x)$ easily for any given x . Then similarly to (3) we have :

$$\hat{p}_{is} = \frac{1}{n} \sum_{k=1}^n I_F(x^{(k)}) \frac{\pi_0(x^{(k)})}{\pi_1(x^{(k)})}, \quad x^{(k)} \sim \pi_1 \quad (7)$$

Note that the standard Monte Carlo method is a special case of Importance Sampling when $\pi_1 = \pi_0$. The estimate \hat{p}_{is} in (7) has the same statistical properties as \hat{p}_{mc} in (3), i.e., it is unbiased and it converges to p_F with probability 1. Choosing π_1 in some appropriate way we hope to be able to reduce the COV.

It was shown by Au and Beck [2] that Importance Sampling does generally not work in high dimensions. A geometric understanding of why this is true was given by Katafygiotis and Zuev in [5]. In that paper it is shown that Importance Sampling is inapplicable in high dimensions because it is practically impossible to select a fixed ISD such that the samples generated by it cover the important region of the failure domain given the fact that the geometry of this important region is generally very complex and unknown a priori.

In this paper we introduce an algorithm called Adaptive Linked Importance Sampling (ALIS) that is based on the methodology developed in [9]. We start from the description of all constructions that are needed.

2 PROPOSED METHODOLOGY

We proceed from the classical Importance Sampling that was described above. It is easy to see that the theoretically optimal choice of π_1 is the conditional PDF :

$$\pi_1^{opt}(x) = \pi_0(x|F) = \frac{\pi_0(x)I_F(x)}{p_F} \quad (8)$$

For this ISD the estimate (7) is $p_{is} = p_F$ for any number of samples n (even for $n = 1$) with resulting COV $\delta_{is} = 0$. We can conclude that the problem of estimating the failure probability is equivalent to estimating the normalizing constant for the following "optimal" non-normalized density function :

$$p_1^{opt}(x) = \pi_0(x)I_F(x) \quad (9)$$

Now we will reformulate the problem of estimating the failure probability in that of estimating the ratio of the normalizing constants of two distributions.

2.1 Reformulation of the problem

Consider two PDFs π_0 and π_1 on the same space \mathbb{R}^N :

$$\pi_0(x) = \frac{p_0(x)}{Z_0}, \quad \pi_1(x) = \frac{p_1(x)}{Z_1} \quad (10)$$

where Z_0 and Z_1 are the corresponding normalizing constants ensuring the total probability is equal to one. Suppose that we are not able to directly compute π_0 and π_1 since we do not know the normalizing constants Z_0 and Z_1 , but p_0 and p_1 are known pointwise. Our goal is to find a Monte Carlo estimate for the ratio

$$r = \frac{Z_1}{Z_0} \quad (11)$$

Assuming this problem can be tackled, the estimation of failure probability easily follows by appropriately choosing p_0 and p_1 :

$$\begin{aligned} p_0(x) &= \pi_0(x) & Z_0 &= 1 \\ p_1(x) &= \pi_0(x)I_F(x) & Z_1 &= p_F \end{aligned} \Rightarrow r = p_F \quad (12)$$

This observation links the problem of estimating failure probabilities to two other extensively researched problems that can also be considered as problems of estimating the ratio of normalizing constants :

- Finding the free energy of a physical system (Statistical Physics),
- Finding the marginal likelihood of a Bayesian statistical model (Bayesian Statistics)

In this paper we will address the problem of reliability estimation as a special case of calculating the ratio r of normalizing constants of two distributions.

Using (10), (11) it follows that

$$E_{\pi_0} \left[\frac{p_1}{p_0} \right] = \int_{\mathbb{R}^N} \frac{p_1(x)}{p_0(x)} \pi_0(x) dx = \int_{\mathbb{R}^N} \frac{Z_1}{Z_0} \pi_1(x) dx = \frac{Z_1}{Z_0} \quad (13)$$

This means that we can estimate r by

$$\hat{r}_{mc(sis)} = \frac{1}{n} \sum_{k=1}^n \frac{p_1(x^{(k)})}{p_0(x^{(k)})}, \quad x^{(k)} \sim \pi_0 \quad (14)$$

In the context of failure probabilities, i.e., using (12), one can easily recognize that (14) becomes exactly the expression for estimating p_F using standard Monte Carlo. In statistical physics literature (14) is also known as the Simple Importance Sampling (SIS) estimate. Now we turn to the idea of intermediate distributions.

2.2 Introduction of intermediate distributions

If π_0 and π_1 are not close enough, the estimate (14) will be improper: the variance of this estimate will be very large. In such a situation, it may be possible to obtain a good estimate by introducing artificial intermediate distributions (AID's).

Let us define a sequence $p_{\alpha_0}, \dots, p_{\alpha_m}$ of non-normalized (or unnormalized) density functions (UDF) using $0 = \alpha_0 < \alpha_1 < \dots < \alpha_{m-1} < \alpha_m = 1$, so that the first and last functions in the sequence are p_0 and p_1 with the intermediate UDF's interpolating between them. Let $Z_\alpha = \int p_\alpha(x)dx$ and $\pi_{\alpha_0}, \dots, \pi_{\alpha_m}$ denote the corresponding sequence of normalizing constants and PDFs, respectively. Then we can represent the ratio $r = Z_1/Z_0$ as follows :

$$\frac{Z_1}{Z_0} = \frac{Z_{\alpha_1}}{Z_0} \frac{Z_{\alpha_2}}{Z_{\alpha_1}} \dots \frac{Z_{\alpha_{m-1}}}{Z_{\alpha_{m-2}}} \frac{Z_1}{Z_{\alpha_{m-1}}} \quad (15)$$

Now, if π_{α_j} and $\pi_{\alpha_{j+1}}$ are sufficiently close and we can sample from π_{α_j} , we can accurately estimate each factor $Z_{\alpha_{j+1}}/Z_{\alpha_j}$ using simple importance sampling (14), and from these estimates obtain an estimate for r :

$$\hat{r}_{gis} = \prod_{j=0}^{m-1} \left(\frac{1}{n} \sum_{k=1}^n \frac{p_{\alpha_{j+1}}(x_j^{(k)})}{p_{\alpha_j}(x_j^{(k)})} \right), \quad x_j^{(k)} \sim \pi_{\alpha_j} \quad (16)$$

We will refer to (16) as Generalized Importance Sampling (GIS).

A novel procedure for evaluating failure probabilities called Subset Simulation was suggested by Au and Beck in [1].

The main idea of this method is as follows. Given a failure domain F let $\mathbb{R}^N = F_{\alpha_0} \supset F_{\alpha_1} \supset \dots \supset F_{\alpha_m} = F$ be a filtration, in other words a sequence of failure events such that $F_{\alpha_k} = \bigcap_{j=0}^k F_{\alpha_j}$. In terms of the representation (1), the failure events F_{α_j} can be chosen as follows :

$$F_{\alpha_j} = \{x \in \mathbb{R}^N | f(x) > \alpha_j b\} \quad (17)$$

Using the definition of conditional probability it can be shown that

$$p_F = \prod_{j=0}^{m-1} P(F_{\alpha_{j+1}} | F_{\alpha_j}) \quad (18)$$

The main observation is that, even if p_F is small, by choosing m and $F_{\alpha_j}, j = 1, \dots, m-1$ appropriately, the conditional probabilities can be made large enough for efficient evaluation using simulations. Each factor in (18) is approximated by

$$\hat{P}(F_{\alpha_{j+1}} | F_{\alpha_j}) = \frac{1}{n} \sum_{k=1}^n I_{F_{\alpha_{j+1}}}(x^{(k)}), \quad x^{(k)} \sim \pi_0(\cdot | F_{\alpha_j}) \quad (19)$$

It follows that Subset Simulation is exactly Generalized Importance Sampling with specific intermediate distributions :

$$\pi_{\alpha_j}(x) = \pi_0(x | F_{\alpha_j}), \quad j = 0, \dots, m \quad (20)$$

Intuition suggest that it is rather unlikely for the family of conditional intermediate distributions (20) with sharp boundaries to be optimal for estimating failure probability. Indeed, our numerical examples will show that this family of AID's is far from optimal and, therefore, there is room to further improve efficiency beyond that offered by Subset Simulation.

Here we propose two smooth families of AIDs designed for reliability problems. First, returning to the definition of the failure domain (1), define the limit state function :

$$\Phi(x) = b - f(x), \quad (21)$$

so that failure domain F is defined as the subset of \mathbb{R}^N where Φ is negative. Now we are define the following two families of non-normalized AIDs:

$$p_\alpha^I(x) = \pi_0(x) \min\{e^{-\alpha\Phi(x)}, 1\} \quad (22)$$

$$p_\alpha^{II}(x) = \frac{\pi_0(x)}{1 + e^{\alpha\Phi(x)}} \quad (23)$$

It is easy to check, that for both distributions we have:

$$\lim_{\alpha \rightarrow +\infty} p_\alpha(x) = p_0(x)I_F(x) = p_\infty(x) \quad (24)$$

Although before it was assumed that $\alpha \in [0, 1]$, in (22), (23) the α belongs in the ray $\alpha \in [0, +\infty]$. Obviously this is not of fundamental importance. We will refer to (22), (23) as to the distributions of the first and the second types correspondingly. Note, that for the first type distributions $p_0 = \pi_0$ and $Z_0 = 1$ as in (12), but for the second ones we have $p_0 = \pi_0/2$ and therefore $Z_0 = 0.5$. Again this can be easily incorporated in the expression for the estimate of p_F . Specifically, as can be easily derived from (11) and (12), one simply has to multiply the estimate for r by Z_0 , i.e., $p_F = 0.5r$.

2.3 Bridging

There is a potential problem one needs to be aware of when using intermediate distributions and that is the following. If the distributions $\{\pi_\alpha\}$ are not nested, one can not hope to estimate $Z_{\alpha_{j+1}}/Z_{\alpha_j}$ by sampling just from π_{α_j} . Indeed, this can be easily demonstrated with the following example involving uniform distributions.

Let $p_{\alpha_j} = I_{[0,4]}$ and $p_{\alpha_{j+1}} = I_{[3,5]}$, so that $Z_{\alpha_j} = 4$, $Z_{\alpha_{j+1}} = 2$ and $Z_{\alpha_{j+1}}/Z_{\alpha_j} = 1/2$. Suppose we have n points drawn uniformly from $[0, 4]$, i.e. from π_{α_j} . The fraction of these that lie in $[3, 5]$, i.e. the simple importance sampling estimate (14), will obviously not give an estimate of $Z_{\alpha_{j+1}}/Z_{\alpha_j}$ since it converges to $1/4$ when $n \rightarrow \infty$ rather than $1/2$. Instead it will be an estimate of Z_*/Z_{α_j} , where Z_* is the length of $[3, 4] = [0, 4] \cap [3, 5]$. Suppose now that we also have n samples drawn uniformly from $[3, 5]$, i.e., from $\pi_{\alpha_{j+1}}$. The fraction of these that lie in $[0, 4]$ will be an estimate of $Z_*/Z_{\alpha_{j+1}}$. Taking the ratio of these two estimates $(Z_*/Z_{\alpha_j})/(Z_*/Z_{\alpha_{j+1}})$ gives an estimate of $Z_{\alpha_{j+1}}/Z_{\alpha_j}$.

This idea can be generalized and non-nested distributions can be worked up by “bridging”. In this method we replace the simple importance sampling estimate for $Z_{\alpha_{j+1}}/Z_{\alpha_j}$ by a ratio of estimates for Z_*/Z_{α_j} and $Z_*/Z_{\alpha_{j+1}}$, where Z_* is the normalizing constant for a “bridge distribution”, $\pi_*(x) = p_*(x)/Z_*$, which is chosen so that it is overlapped by both π_{α_j} and $\pi_{\alpha_{j+1}}$. Using simple importance sampling estimates for Z_*/Z_{α_j} and $Z_*/Z_{\alpha_{j+1}}$, we can obtain the estimate for $Z_{\alpha_{j+1}}/Z_{\alpha_j}$:

$$\frac{Z_{\alpha_{j+1}}}{Z_{\alpha_j}} = \frac{E_{\pi_{\alpha_j}} \left[\frac{p_*}{p_{\alpha_j}} \right]}{E_{\pi_{\alpha_{j+1}}} \left[\frac{p_*}{p_{\alpha_{j+1}}} \right]} \approx \frac{\frac{1}{n} \sum_{k=1}^n \frac{p_*(x_j^{(k)})}{p_{\alpha_j}(x_j^{(k)})}}{\frac{1}{n} \sum_{k=1}^n \frac{p_*(x_{j+1}^{(k)})}{p_{\alpha_{j+1}}(x_{j+1}^{(k)})}}, \quad \begin{array}{l} x_j^{(k)} \sim \pi_{\alpha_j} \\ x_{j+1}^{(k)} \sim \pi_{\alpha_{j+1}} \end{array} \quad (25)$$

This technique was introduced by Bennett [3], Lu, Singh and Kofke in [6] give a recent review.




Figure 1: Estimate of ratio of normalizing constants for Gaussian distributions using SIS and geometrical bridge.

One natural choice for the bridge distribution is the “geometric” bridge:

$$p_*^{geo}(x) = \sqrt{p_{\alpha_j}(x)p_{\alpha_{j+1}}(x)}, \quad (26)$$

that is in some sense an average between π_{α_j} and $\pi_{\alpha_{j+1}}$. It was shown in [3], that an asymptotically optimal choice of bridge distribution is

$$p_*^{opt}(x) = \frac{p_{\alpha_j}(x)p_{\alpha_{j+1}}(x)}{\frac{n_j}{n_{j+1}} \frac{Z_{\alpha_{j+1}}}{Z_{\alpha_j}} p_{\alpha_j}(x) + p_{\alpha_{j+1}}(x)} \quad (27)$$

In general we cannot use this bridge distribution in practice, since we do not know $Z_{\alpha_{j+1}}/Z_{\alpha_j}$ beforehand. But we can use a preliminary guess of $Z_{\alpha_{j+1}}/Z_{\alpha_j}$ to define a bridge distribution (27) and then iterate until convergence is observed.

Next we discuss the meaning of “nested” for non-uniform distributions. Again let us start with an easy example. Let x_1, \dots, x_n be samples following the standard one-dimensional Gaussian distribution. It is easy to check that for $n = 10^4$ the probability that at least one sample lies outside the interval $[-6, 6]$ is smaller than 10^{-4} . Although, strictly speaking, the support of the standard Gaussian distribution is the entire real line (theoretical support), in practice the support may be regarded as just a segment (practical support). So if π_{α_j} and $\pi_{\alpha_{j+1}}$ are two Gaussian distributions with different mean values we cannot hope to estimate $Z_{\alpha_{j+1}}/Z_{\alpha_j}$ by sampling just from π_{α_j} . The reason is the same as in the case of uniform distributions: their practical supports are not nested. As an example, the dependence between the estimate of Z_α/Z_0 , where $\pi_\alpha = \mathcal{N}_{(\alpha,1)}$ is a Gaussian distribution centered at α and having unit variance is plotted against α in Figure 1. Two methods are used: i) SIS according to (14) and ii) using geometrical bridge according to (25). Clearly the SIS estimate deteriorates as α increases, i.e., as the practical supports of the two distributions become further apart. On the contrary, the results using bridge sampling remain relatively unaffected.

So, in order to apply AIDs for estimating the of ratio of normalizing constants we have to use the concept of “bridge”, especially if these AIDs are not nested as is generally the

case. Thus, for each couple of adjacent intermediate distributions p_{α_j} and $p_{\alpha_{j+1}}$ we have to construct a bridge p_{j^*j+1} using, for example, (26) or (27) and then for each multiplier in (15) use the bridged approximation (25). Note, that in Subset Simulation the conditional intermediate distributions (20) are nested and, therefore, bridging makes no sense.

2.4 Markov Chain Monte Carlo

As we can see from the previous sections we need some procedure to sample independently from an intermediate PDF π_α . Although we assume we can always do this directly for π_0 , since usually this is taken as the standard Gaussian, it can not be done for other distributions with $\alpha \neq 0$. Recall that we can only evaluate the target distribution π_α up to a normalizing constant Z_α , i.e. $\pi_\alpha = p_\alpha/Z_\alpha$. Here we assume that the non-negative function p_α is known pointwise.

A practical framework for sampling under the above context is that provided by discrete time Markov chains underlying the Markov Chain Monte Carlo (MCMC) methods. The main idea of all MCMC methods is the following: although it is difficult to efficiently generate independent samples according to the target PDF, it is possible, using a specially designed Markov Chain, to efficiently generate dependent samples that are at least asymptotically (as the number of Markov steps increases) distributed as the target PDF. So, for each α we define a Markov chain whose equilibrium distribution is π_α . Simulating such chain for some time, and discarding the early portion, provides dependent points sampled approximately according to π_α . It is importance to mention that usually it is difficult to tell whether or not the chain reaches equilibrium in a reasonable time. Furthermore, even if it has reached equilibrium, it is hard to find out for sure.

The Metropolis-Hastings Algorithm (MH) is a general method for constructing a Markov chain with equilibrium PDF $\pi_\alpha = p_\alpha/Z_\alpha$, dating back over thirty years [4]. If the current state of the chain is x , we move to the next state x' as follows:

- 1) Generate a candidate state x^* according to some proposal distribution with PDF $S_\alpha(\cdot|x)$.
- 2) Compute the acceptance ratio : $a(x, x^*) = \frac{p_\alpha(x^*)S_\alpha(x|x^*)}{p_\alpha(x)S_\alpha(x^*|x)}$
- 3) Set $x' = x^*$ with probability $\min\{1, a(x, x^*)\}$ and set $x' = x$ with the remaining probability.

One can show that such updates leave π_α invariant, and hence the chain will eventually converge to π_α as its equilibrium distribution. Note, that the MH algorithm does not require information about the normalizing constant.

The proposal PDF $S_\alpha(\cdot|x)$ governs the distribution of the candidate and affects the transition of the chain from one state to another. The practical support of $S_\alpha(\cdot|x)$, i.e., the region that contains almost all probability mass of $S_\alpha(\cdot|x)$, affects how fast the chain can explore the parameter space and also the dependence among the generated samples. If the practical support is too small, the acceptance rate is high, but the next sample will be near the current one, increasing dependence. If the practical support is too large, the acceptance rate can be low, making the next sample identical to the current one, increasing correlation. Therefore, the optimal choice of the proposal PDF is a trade-off between acceptance rate and spatial correlation of samples. A Gaussian

PDF or a uniform PDF centered at the current sample x are commonly used. When the proposal PDF is symmetric, as in the case of Gaussian or uniform distributions, one has $S_\alpha(x^*|x) = S_\alpha(x|x^*)$ and the MH algorithm reduces to the original Metropolis algorithm [7].

The Modified Metropolis (MM) algorithm for sampling from conditional distributions (20) in high dimensions was introduced in [1]. It has been found that when the dimension is large, the probability to obtain repeated sample in standard Metropolis algorithm is close to one, making the corresponding Markov chain useless. A geometric explanation of why this happens is given in [5]. The MM algorithm differs from the original one in the way the candidate state is generated. Instead of using an N -dimensional proposal PDF, in the modified algorithm a sequence of 1-dimensional proposals are used updating one component at a time. For more details see [1] or [5].

4 EXAMPLES

The ALIS algorithm is applied to calculate the failure probability in two examples. In these examples we use intermediate distributions of both types (22), (23) and optimal bridge (27). ALIS will be compared with Subset Simulation and standard Monte Carlo in terms of the coefficient of variation (COV) of the failure probability estimates.

4.1 One-dimensional ray

Consider the simplest possible one-dimensional example. Let the failure domain F be a ray on the real line :

$$F = \{x \in \mathbb{R} | x > b\}, \quad (28)$$

for some threshold b , where x is a standard Gaussian random variable, $x \sim \mathcal{N}_{(0,1)}$. Obviously, the failure probability in this case is

$$p_F = 1 - \Phi(b), \quad (29)$$

where Φ is the CDF of the standard Gaussian distribution. We apply the described simulation methods for estimating the following values of failure probability : $p_F = 10^{-k}$, $k = 2, \dots, 10$. The corresponding thresholds are :

$$b_k = \Phi^{-1}(1 - 10^{-k}). \quad (30)$$

The COV of the estimates obtained by ALIS, Subset Simulation and Monte Carlo are given in Fig. 2 using 100 runs. For each intermediate distribution $n = 10^4$ samples were drawn for ALIS and SS. The comparison between the different methods is made for the same total computational effort. It can be clearly seen that ALIS clearly outperforms SS. Furthermore, the computational effort required by ALIS in order to achieve a certain level of accuracy is found to increase only linearly with the order of target failure probability while this is not the case for SS.

4.2 High-dimensional sphere

The next example is high-dimensional. Let $x \in \mathbb{R}^N$ be a standard Gaussian random vector. As discussed in [5], the vast majority of the probability mass in the N -dimensional standard Gaussian space belongs in an Important Ring (= practical support),

$$\sqrt{N} - r < R < \sqrt{N} + r, \quad (31)$$

where r depends on the amount of probability mass that we want to consider. For example, if $N = 10^3$ and $r = 3.46$ the probability of the corresponding Important Ring is more than $1 - 10^{-6}$. Thus, a sample $x \in \mathbb{R}^N$ distributed according to this high-dimensional standard Gaussian distribution will lie with extremely large probability in this Important Ring.

Consider in \mathbb{R}^N the interior of a cone with axis a and angle φ :

$$\mathbb{C}_{a,\varphi} = \{x \in \mathbb{R}^N \mid \widehat{x, a} < \varphi\} \quad (32)$$

and define the failure domain to be the intersection of $\mathbb{C}_{a,\varphi}$ and the hypersphere $\mathbb{S}_{\sqrt{N}}^{N-1}$ with radius \sqrt{N} and center in the origin (i.e., the middle hypersphere in the Important Ring) :

$$F = \mathbb{C}_{a,\varphi} \cap \mathbb{S}_{\sqrt{N}}^{N-1}. \quad (33)$$

Instead of a Gaussian distribution in \mathbb{R}^N , we consider here a uniform distribution on the hypersphere $\mathbb{S}_{\sqrt{N}}^{N-1}$ with failure domain defined by (33). Unlike in the one-dimensional example, where an input load x with sufficiently large “energy” causes failure, in this example all inputs have the same energy. In this case safety or failure depends on the “direction” of the input : the closer the direction of x to the cone axis a , the more unsafe the system is. In this sense this example quite differs from the previous one.

We generate a candidate state in MH update for sampling from π_α in the following way. If the current state of the chain is $x \in \mathbb{S}_{\sqrt{N}}^{N-1}$, we first generate a sample $y \sim \mathcal{N}_{(x,1)}$ from the multidimensional Gaussian distribution centered in x . The candidate state is then the projection of y on the sphere $\mathbb{S}_{\sqrt{N}}^{N-1}$, $x^* = \text{proj}_{\mathbb{S}} y$. Obviously, the corresponding proposal PDF is symmetric, $S(x^*|x) = S(x|x^*)$, and the acceptance ratio $a(x, x^*) = p_\alpha(x^*)/p_\alpha(x)$. Since everything is taking place on a sphere, the probability to obtain a repeated sample during these updates is not close to 1, and the corresponding Markov chain is able to explore the entire parameter space.

Choosing φ we can vary the failure probability. We consider the following values of failure probability : $p_F = 10^{-k}$, $k = 2, \dots, 7$. The COV of the failure probability estimates obtained by ALIS, Subset Simulation and Monte Carlo are given in Fig. 3 and Fig. 4 for dimensions $N = 100$ and $N = 1000$ respectively, using 100 runs each time. For each intermediate distribution $n = 10^4$ samples were drawn. The comparison between the different methods is made for the same total computational effort. It can be clearly seen that ALIS clearly again outperforms SS although the computational effort does not quite increase linearly with the order of failure probability in this case.

5. CONCLUSIONS AND FUTURE RESEARCH

In this paper we have introduced a novel simulation approach, called Adaptive Linked Importance Sampling (ALIS). This methodology generalizes the Subset Simulation algorithm introduced in [1]. Instead of using conditional distributions (20) ALIS allows for a much richer choice of artificial intermediate distributions (AID's). The case of non-nested distributions can be overcome by the concept of bridging. The accuracy and efficiency of the ALIS is demonstrated with numerical examples. In particular, we show that the choice of AID's prescribed by Subset Simulation is far from optimal.

ACKNOWLEDGMENTS

This research has been supported by the Hong Kong Research Grants Council under grants 614305. This support is gratefully acknowledged.

REFERENCES

- [1] Au S.K., Beck J.L., (2001), *Estimation of small failure probabilities in high dimensions by subset simulation*, Probab. Engineering Mechanics; 16(4): 263-277. 1.
- [2] Au S.K., Beck J.L., (2003), *Importance Sampling in high dimensions*, Structural Safety, 25(2): 139-163.
- [3] Bennett C.H., (1976), *Efficient estimation of free energy differences from Monte Carlo data*, Journal of Computational Physics, vol. 22, pp. 245-268.
- [4] Hastings W.K., (1970), *Monte Carlo sampling methods using Markov chains and their applications*, Biometrika, 57:97-109.
- [5] Katafygiotis L.S., Zuev K.M., (2006), *Geometric insight into the challenges of solving high-dimensional reliability*, Proceedings Fifth Computational Stochastic Mechanics Conference, Rhodos Greece, 21-23 June 2006.
- [6] Lu N., Singh J.K., Kofke D.A., (2003), *Appropriate methods to combine forward and reverse free-energy perturbation averages*, Journal of Chemical Physics, vol. 118, pp. 2977-2984.
- [7] Metropolis N. et al., (1953), *Equation of State Calculations by Fast Computing Machines*, J. Chemical Physics, vol. 21, pp. 1087-1092.
- [8] Metropolis N., Ulam S., (1949), *The Monte Carlo method*, J. Amer. statistical assoc., vol. 44, N. 247, pp. 335-341.
- [9] Neal R.M., (2005), *Estimating ratios of normalizing constants using Linked Importance Sampling*, Technical Report No. 0511, Dept. of Statistics, University of Toronto.




Figure 2: The COV of estimates obtained by ALIS, Subset Simulation and Monte Carlo in one-dimensional problem

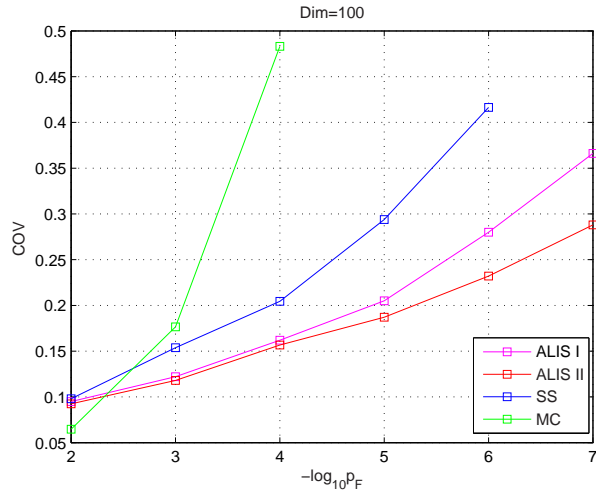


Figure 3: The COV of estimates obtained by ALIS, Subset Simulation and Monte Carlo in 100-dimensional problem




Figure 4: The COV of estimates obtained by ALIS, Subset Simulation and Monte Carlo in 1000-dimensional problem