Available online at www.sciencedirect.com
ScienceDirect
Probabilistic Engineering Mechanics 23 (2008) 208–218
PROBABILISTIC ENGINEERING MECHANICS
www.elsevier.com/locate/probengmech

Geometric insight into the challenges of solving high-dimensional reliability problems

L.S. Katafygiotis, K.M. Zuev*

Department of Civil Engineering, HKUST, Hong Kong, China

Received 12 December 2007; accepted 17 December 2007
Available online 31 December 2007

Abstract

In this paper we adopt a geometric perspective to highlight the challenges associated with solving high-dimensional reliability problems. Adopting a geometric point of view we highlight and explain a range of results concerning the performance of several well-known reliability methods.

We start by investigating geometric properties of the N-dimensional Gaussian space and the distribution of samples in such a space or in a subspace corresponding to a failure domain. Next, we discuss Importance Sampling (IS) in high dimensions. We provide a geometric understanding as to why IS generally does not work in high dimensions [Au SK, Beck JL. Importance sampling in high dimensions. Structural Safety 2003;25(2):139–63]. We furthermore challenge the significance of "design point" when dealing with strongly nonlinear problems. We conclude by showing that for the general high-dimensional nonlinear reliability problems the selection of an appropriate fixed IS density is practically impossible.

Next, we discuss the simulation of samples using Markov Chain Monte Carlo (MCMC) methods. Firstly, we provide a geometric explanation as to why the standard Metropolis–Hastings (MH) algorithm does "not work" in high-dimensions. We then explain why the modified Metropolis–Hastings (MMH) algorithm introduced by Au and Beck [Au SK, Beck JL. Estimation of small failure probabilities in high dimensions by subset simulation. Probabilistic Engineering Mechanics 2001;16(4):263–77] overcomes this problem. A study of the correlation of samples obtained using MMH as a function of different parameters follows. Such study leads to recommendations for fine-tuning the MMH algorithm. Finally, the MMH algorithm is compared with the MCMC algorithm proposed by Katafygiotis and Cheung [Katafygiotis LS, Cheung SH. Application of spherical subset simulation method and auxiliary domain method on a benchmark reliability study, Structural Safety 2006 (in print)] in terms of the correlation of samples they generate.

Keywords: Reliability; Importance sampling; Subset simulation; Markov Chains; Dynamic analysis

1. Geometric properties of high-dimensional Gaussian space

In practical applications random parameters of a system can usually be generated by appropriate transformation of independent standard Gaussian variables. So let us start with the investigation of the geometric properties of the N-dimensional standard Gaussian space and the distribution of samples in such a space.

Let $\theta = (\theta_1, \ldots, \theta_N)$ be a random vector in $\mathbb{R}^N$ where each of its components follows the standard Gaussian distribution:

$$\theta_i \sim \mathcal{N}(0,1), \quad i = 1, \ldots, N. \tag{1}$$

By definition the square of the Euclidean norm of θ is distributed according to the chi-square distribution with N degrees of freedom:

$$R^2 = \sum_{i=1}^{N} \theta_i^2 \sim \chi_N^2. \tag{2}$$

As N tends to infinity, the distribution of R^2 tends to normality (by Central Limit Theorem). In particular, it becomes more and more symmetric. However, the tendency is slow: the skewness is $\sqrt{8/N}$. The probability density function (PDF) and cumulative distribution function (CDF) of R^2 for $N = 10^3$ are plotted in Fig. 1.

* Corresponding author.
E-mail addresses: lambros@ust.hk (L.S. Katafygiotis), zuev@ust.hk (K.M. Zuev).

0266-8920/$ - see front matter © 2007 Elsevier Ltd. All rights reserved.
doi:10.1016/j.probengmech.2007.12.026

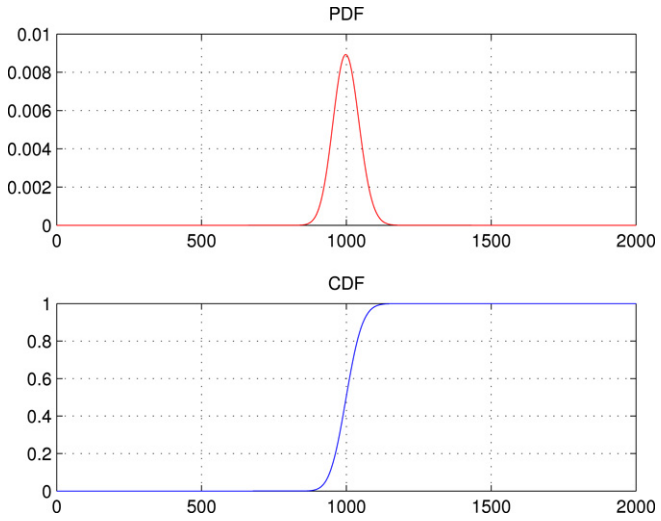

Fig. 1. Probability Density Function and Cumulative Distribution Function of chi-square distribution with $N = 10^3$ degrees of freedom.

It can be shown [3,4] that $\sqrt{2R^2}$ is approximately normally distributed with mean $\sqrt{2N-1}$ and unit variance. Hence the norm of the random vector $\|\theta\| = R$ is also approximately a Gaussian random variable:

$$R \overset{\text{app}}{\sim} \mathcal{N}\left(\sqrt{N-1/2}, 1/2\right) \approx \mathcal{N}\left(\sqrt{N}, 1/2\right), \tag{3}$$

when $N \to \infty$. This means that the huge part of probability mass in the N-dimensional standard Gaussian space belongs to a spherical ring, so-called Important Ring,

$$\sqrt{N} - r < R < \sqrt{N} + r, \tag{4}$$

where r depends on the amount of probability mass that we want to contain inside the Important Ring. For example, if $N = 10^3$ and $r = 3.46$ the probability of the corresponding Important Ring $28.16 < R < 35.08$ is more than $1 - 10^{-6}$. Thus, any sample $\theta \in \mathbb{R}^N$ distributed according to the high-dimensional standard Gaussian distribution will lie with extremely large probability in the Important Ring.

Now let us fix one particular direction (here direction means a ray passing through zero), say $e = (1, 0, \ldots, 0)$, and explore the distribution of the angle $\alpha = \widehat{\theta e}$ between this direction and a random vector θ. By definition

$$\begin{aligned} f_\alpha(\alpha_0)\mathrm{d}\alpha &= P(\alpha_0 \le \alpha \le \alpha_0 + \mathrm{d}\alpha) \\ &= P(\alpha_0 \le \widehat{\theta e} \le \alpha_0 + \mathrm{d}\alpha), \end{aligned} \tag{5}$$

where f_α is the PDF of α. Since the Gaussian space is isotropic (there are no preferable directions) and any point that lies along a particular ray forms the same angle with e, we can simplify the problem by considering the $(N-1)$-dimensional sphere of unit radius $\mathbb{S}_1^{N-1}$. Note that herein a sphere is defined as the set of all points located at distance equal to R from a given fixed point corresponding to the center of the sphere. Thus, in a three-dimensional space, according to our definition, a sphere is a two-dimensional surface, while the interior of the sphere, comprised of all points at distance smaller than R from the center, is a three-dimensional set. Clearly, on the sphere $\mathbb{S}_1^{N-1}$ all points are uniformly distributed. Therefore, $f_\alpha(\alpha_0)$ is proportional to the geometric volume of part of this sphere:

$$f_\alpha(\alpha_0) \sim Vol(\Omega_{\alpha_0}), \tag{6}$$

$$\Omega_{\alpha_0} = \{\theta \in \mathbb{S}_1^{N-1} : \widehat{\theta e} = \alpha_0\}. \tag{7}$$

If $\langle \cdot, \cdot \rangle$ denotes the standard scalar product in $\mathbb{R}^N$ then $\langle \theta, e \rangle = \theta_1$. On the other hand $\langle \theta, e \rangle = \|\theta\|\|e\| \cos \widehat{\theta e}$. Therefore

$$\Omega_{\alpha_0} = \{\theta \in \mathbb{S}_1^{N-1} : \theta_1 = \cos \alpha_0\}. \tag{8}$$

One can rewrite (8) as the intersection of the hypersphere $\mathbb{S}_1^{N-1}$ with the hyperplane

$$\pi_{\alpha_0}^{N-1} = \{\theta \in \mathbb{R}^N : \theta_1 = \cos \alpha_0\}. \tag{9}$$

Thus,

$$\Omega_{\alpha_0} = \mathbb{S}_1^{N-1} \cap \pi_{\alpha_0}^{N-1}. \tag{10}$$

This intersection can be easily evaluated.

$$\begin{cases} \theta_1^2 + \cdots + \theta_N^2 = 1 \\ \theta_1 = \cos \alpha_0. \end{cases} \quad \Leftrightarrow \tag{11}$$

$$\theta_2^2 + \cdots + \theta_N^2 = \sin^2 \alpha_0. \tag{12}$$

Thus, the region in which we are interested is a $(N-2)$-dimensional sphere of radius $\sin \alpha_0$:

$$\Omega_{\alpha_0} = \mathbb{S}_{\sin \alpha_0}^{N-2}. \tag{13}$$

It is well known that the volume of a k-dimensional sphere is proportional to the radius in the power k. So we have:

$$f_\alpha(\alpha_0) \sim Vol(\Omega_{\alpha_0}) \sim \sin^{N-2} \alpha_0. \tag{14}$$

Finally, we can obtain that the PDF and CDF of α (Fig. 2) are correspondingly equal to

$$f_\alpha(\alpha) = \frac{\sin^{N-2} \alpha}{\int_0^\pi \sin^{N-2} \alpha \mathrm{d}\alpha} \tag{15}$$

$$F_\alpha(\alpha) = \frac{\int_0^\alpha \sin^{N-2} \alpha \mathrm{d}\alpha}{\int_0^\pi \sin^{N-2} \alpha \mathrm{d}\alpha}. \tag{16}$$

From this result and by plotting these distributions for large N it follows that if we fix a particular direction e then a sample $\theta \in \mathbb{R}^N$ that is distributed according to the high-dimensional standard Gaussian distribution will be with high probability almost perpendicular to e. Although we proved this result for the specific direction $e = (1, 0, \ldots, 0)$ this also holds for arbitrary e since the Gaussian space is isotropic.

This also can be argued in a more intuitive way. If θ is such a sample then

$$\cot^2 \alpha = \frac{\theta_1^2}{\theta_2^2 + \cdots + \theta_N^2}. \tag{17}$$

Since θ_i are independent and identically distributed random variables we have that expectation:

$$E[\cot^2 \alpha] = \frac{1}{N-1} \to 0 \quad \text{as } N \to \infty. \tag{18}$$

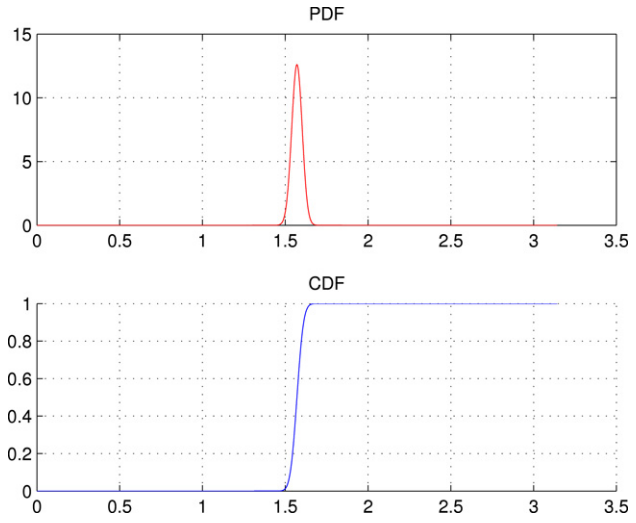

Fig. 2. Probability Density Function and Cumulative Distribution Function of angle α.

Therefore,

$$\alpha \to \pi/2 \quad \text{as } N \to \infty. \tag{19}$$

2. Importance sampling

Consider the problem of computing failure probability. The probability of failure can be simply expressed as:

$$p_F = \int_{\mathbb{R}^N} I_F(\theta) q(\theta) \mathrm{d}\theta, \tag{20}$$

where $F \subseteq \mathbb{R}^N$ is the failure domain, I_F is the indicator function (=1 if $\theta \in F$, = 0 otherwise) and θ represents the random parameters of the system with joint PDF q.[1]

Among all procedures developed for the estimation of p_F, a prominent position is held by simulation methods. In standard Monte Carlo simulations the integral (20) is interpreted as mathematical expectation:

$$p_F = E_q[I_F] \tag{21}$$

and, keeping in mind the Law of Large Numbers, the following estimator is used:

$$p_F^{mc} = \frac{1}{n} \sum_{k=1}^{n} I_F(\theta^{(k)}), \quad \theta^{(k)} \sim q(\cdot). \tag{22}$$

The statistical error of unbiased estimators is most appropriately measured by the coefficient of variation. In the case of standard Monte Carlo:

$$\delta_{mc} = \sqrt{\frac{(1 - p_F)}{n p_F}}. \tag{23}$$

The main advantage of the standard Monte Carlo is that its efficiency does depend neither on the dimensionality N of the random vector θ nor on the complexity of the failure domain. Its main disadvantage is its inefficiency in estimating small failure probabilities because of the large number of samples ($\sim 1/p_F$) needed to achieve an acceptable level of accuracy. Thus, standard Monte Carlo is not practically feasible with current computational power in the case of small failure probabilities (even in a one-dimensional case). Importance Sampling is a variance reduction method aiming to improve the standard Monte Carlo method.

The basic idea of Importance Sampling is to generate samples that lie more frequently in the important region of the failure domain, i.e., the region that accounts for most of the failure probability. Roughly speaking the inefficiency of standard Monte Carlo is due to the fact that most of the samples it generates lie in the safe domain so that the vast majority of the terms in the sum (22) are zero and only very few, if any, are equal to one. Rather than calculating the failure probability estimate as an average of many 0's and very few 1's, Importance Sampling calculates it as the average of a set of numbers involving many more nonzero terms where ideally these nonzero terms are small numbers of the same order as the target failure probability.

Specifically, for any PDF q_{is} we can write:

$$\begin{aligned} p_F &= \int_{\mathbb{R}^N} I_F(\theta) q(\theta) \mathrm{d}\theta \\ &= \int_{\mathbb{R}^N} \frac{I_F(\theta) q(\theta)}{q_{is}(\theta)} q_{is}(\theta) \mathrm{d}\theta = E_{q_{is}}\left[\frac{I_F q}{q_{is}}\right] \end{aligned} \tag{24}$$

and similarly to (22) we have:

$$p_F^{is} = \frac{1}{n} \sum_{k=1}^{n} \frac{I_F(\theta^{(k)}) q(\theta^{(k)})}{q_{is}(\theta^{(k)})}, \quad \theta^{(k)} \sim q_{is}(\cdot). \tag{25}$$

Note that standard Monte Carlo method is a special case of Importance Sampling when $q_{is} = q$.

The most important task in applying Importance Sampling is the construction of the importance sampling density (ISD) q_{is}. If it is "good" then we can get great improvement in efficiency.

What does it mean: "good" ISD ? Let us consider an example. Suppose we know that a certain vector x belongs to the failure domain $x \in F$. It is natural to assume that in the neighborhood of x there are more points from F. Thus, we can consider q_{is} to be a Gaussian PDF centered at x:

$$q_{is}(\theta) = \frac{1}{(\sqrt{2\pi})^N} \exp\left(-\frac{\|\theta - x\|^2}{2}\right). \tag{26}$$

As discussed in the previous section, the main part of probability mass is concentrated inside the Important Ring. This means that we can restrict the sample space and consider only the part of failure domain that belongs to the Important

[1] Without limiting the generality the random vector $\theta = (\theta_1, \ldots, \theta_N)$ will be considered to be standard Gaussian with independent components (any random vector can be appropriately transformed to a vector that is at least approximately distributed like this). Thus,

$$q(\theta) = \prod_{i=1}^{N} q_i(\theta_i) = \prod_{i=1}^{N} \mathcal{N}_{(0,1)}(\theta_i) = \frac{1}{(\sqrt{2\pi})^N} \exp\left(-\frac{\|\theta\|^2}{2}\right).$$

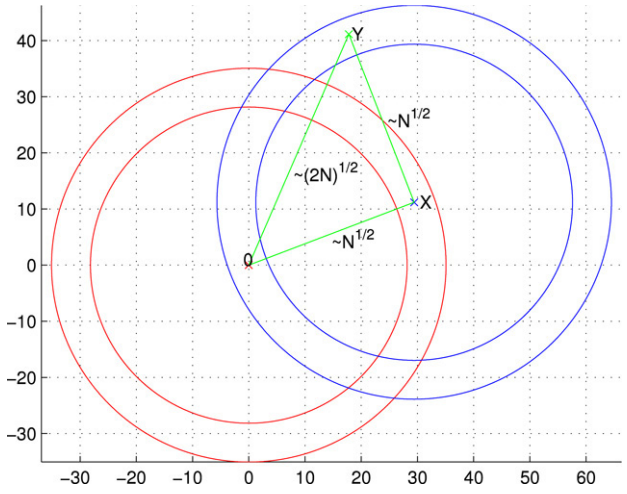

Fig. 3. Vectors $x, y \in \mathbb{R}^N$ are drawn from standard Gaussian and Gaussian centered at x distributions, respectively, $N = 1000$.

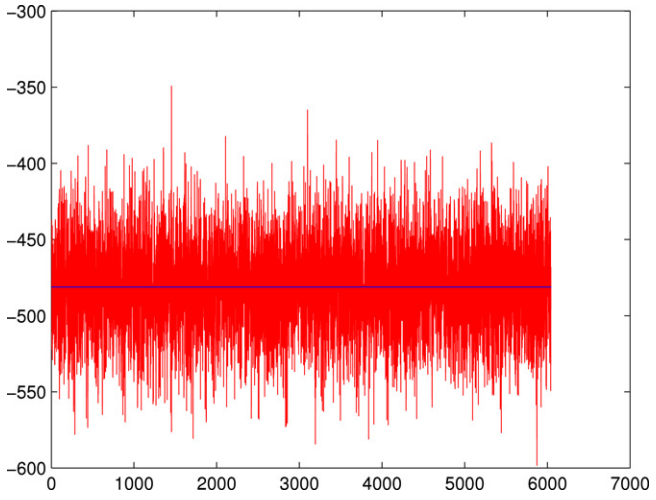

Fig. 4. The ratio $\frac{q(y^{(k)})}{q_{is}(y^{(k)})}$ plotted in log scale against the sample number, for $n = 10^4$ samples and for $N = 1000$.

Ring, since the contribution of the remaining part of the failure domain to the probability p_F is comparatively very small. So we can consider that x belongs in the Important Ring. Again, from the previous section we know that the angle between the vectors x and $(y - x)$, where y is a random vector drawn from q_{is}, is approximately equal to $\pi/2$. Moreover y will lie in the Important Ring corresponding to q_{is}. This is schematically shown in Fig. 3.

It follows that in high dimensions we will have:

$$\|x\| \approx \sqrt{N}, \qquad \|y - x\| \approx \sqrt{N}. \tag{27}$$

Therefore

$$\|y\| \approx \sqrt{2N}. \tag{28}$$

Now the estimator (25) gives:

$$p_F^{is} = \frac{1}{n}\sum_{k=1}^{n} \frac{I_F(y^{(k)})q(y^{(k)})}{q_{is}(y^{(k)})}, \tag{29}$$

since each $y^{(k)}$ is drawn from the Gaussian distribution centered at x and, therefore, satisfies (28), it follows from (27), (28) that the ratio

$$\frac{q(y^{(k)})}{q_{is}(y^{(k)})} \approx \mathrm{e}^{-N/2}, \tag{30}$$

which is extremely small in high dimensions. Thus, Importance Sampling leads to underestimation of p_F. Here we provided a geometrical explanation as to why this happens; specifically, we showed that the simulation of samples distributed according to the chosen importance sampling density is very unlikely to yield samples that lie within the Important Ring centered at zero. Note that the important region of any failure domain should be a subset of this important ring, as the probability volume outside the important ring can be considered to be negligible. Therefore, we can conclude that the chosen importance sampling density fails to generate samples in the important region of the failure domain rendering Importance Sampling inapplicable, severely underestimating the true failure probability.

The above results are next confirmed using simulations. A linear failure domain with reliability index $\beta = 3$ is considered and Importance sampling is applied using a Gaussian IS density with unit variance centered at one of the failure points, i.e., q_{is} is given by (26). Fig. 4 shows for a single run the ratio $\frac{q(y^{(k)})}{q_{is}(y^{(k)})}$ plotted in log scale against the sample number, for $n = 10^4$ samples and for $N = 1000$. Note that this figure shows only the ratio for those samples corresponding to failure points, in this case for 6042 points. The horizontal line shows the mean value of the log of these values which is calculated to be equal to -481. This is consistent with the rough estimate -500 given by (30). Fig. 5 shows the failure probability estimate p_F^{is} in log scale, as a function of the dimension N. Here, $n = 10^4$ samples were used for each run; the plotted value of p_F^{is} corresponds to the average value of 50 runs. Confirming the earlier discussion, the failure probability is found to be severely underestimated, the underestimation becoming worse as the dimension N increases. Note that p_F^{is} is found to be of order not quite as low as $\exp(-N/2)$. The reason is that the ratios $\frac{q(y^{(k)})}{q_{is}(y^{(k)})}$ are, when plotted in linear scale, varying tremendously, so that the order of p_F^{is} as calculated by (29) is governed by the largest of these terms. For example, in the case of the run corresponding to Fig. 4 the maximum term was of the order of $\exp(-350)$. This explains why the dependency on N is not quite as bad as $\exp(-N/2)$.

3. Design points and nonlinear problems

The next question we discuss is the significance of design points when dealing with strongly nonlinear problems.

Basically, the calculation of the probability of failure p_F in (20) is the evaluation of the total probability volume corresponding to the failure domain $F \subset \mathbb{R}^N$ defined by $g(\theta) < 0$, where $g(\cdot)$ denotes the limit-state function (LSF). The design point is defined as the point θ^* on the limit-state surface $\{\theta : g(\theta) = 0\}$ that is nearest to the origin when the random variables are assumed to have been transformed to

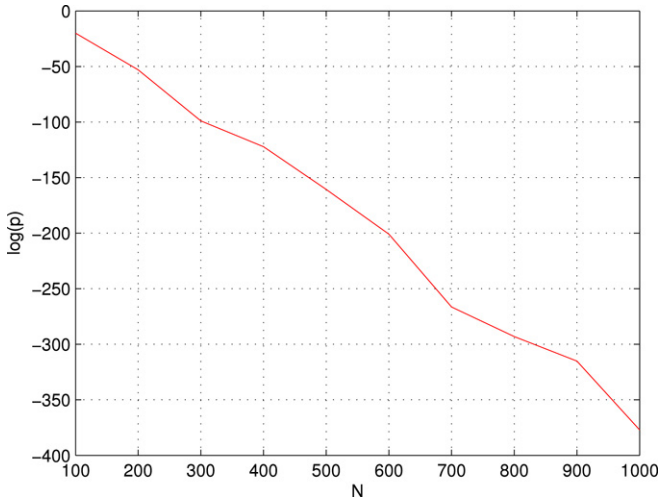

Fig. 5. Failure probability estimate p_F^{is} plotted in log scale, as a function of the dimension N.

the standard Gaussian space. Due to the rotational symmetry of the standard Gaussian distribution, the design point is the most likely single realization of the random variables that gives rise to the failure event, i.e., it is the failure point with the largest probability density. The norm of the design point, i.e., its distance from the origin, is referred to as the reliability index β. Note that in the case of high-dimensional reliability problems it is always true that $\beta \ll R = \sqrt{N}$. This leads to the important, although often misunderstood, point that the design point does not belong in the important ring and, therefore, according to our earlier discussions, it is practically impossible during simulations to obtain a sample in its neighborhood. Thus, although the design point is the point with the maximum likelihood, a sample in its vicinity will practically never be realized through simulations. The reason for this is that the geometric volume describing the vicinity of the design point is extremely small in high dimensions, making the probability mass associated with the neighborhood of the design point negligible. In other words, the relatively smaller probability density corresponding to the failure points located in the important ring is overcompensated by the vast number of these failure points so that if one tries to simulate a failure realization it is almost certain that he will obtain a sample belonging in the important ring rather than in the vicinity of the design point.

It is considered to be that design points play very important role in solving reliability problems. Without doubt this is true in the linear case. Consider a linear reliability problem with LSF expressed in terms of the standard Gaussian vector θ as follows:

$$g(\theta) = a^T\theta + b, \tag{31}$$

where $a \in \mathbb{R}^N$ and b are fixed coefficients. The design point θ^* is then the point on the plane $g(\theta) = 0$ that is located closest to the origin and can be easily calculated in terms of a and b as follows:

$$\theta^* = -\frac{b}{\|a\|^2}a \tag{32}$$

whose reliability index is given by:

$$\beta = \|\theta^*\| = \frac{b}{\|a\|}. \tag{33}$$

It is well known that the failure probability corresponding to this linear failure domain is given in terms of β by the expression:

$$p_F = P(\theta : g(\theta) < 0) = 1 - \Phi(\beta), \tag{34}$$

where Φ denotes the CDF of the standard Gaussian variable. So, the failure probability in the linear case is completely defined by the design point. In the case where the failure domain is almost linear the first-order reliability method (FORM) based on (34) for the given design point provides a good approximation of p_F.

In the case where the LSF can be approximated by a second-order polynomial function, the design point still plays a very important role and is the basis of the second-order reliability method (SORM). However in nonlinear case the significance of the design point(s) is not as clear and is in need of research.

In the next section we consider a nonlinear dynamic problem to show that the design point does not provide sufficient information to describe the complex geometry of the corresponding nonlinear failure domain.

3.1. Duffing oscillator subjected to white-noise

This nonlinear elastic system is taken from Koo et al. [6]. Consider the Duffing oscillator defined by

$$m\ddot{x}(t) + c\dot{x}(t) + k[x(t) + \gamma x(t)^3] = f(t) \tag{35}$$

with $m = 1000$ kg, $c = 200\pi$ N s/m, $k = 1000(2\pi)^2$ N/m, and $\gamma = 1$ m^{-2}, and assume the input is white-noise with intensity $S_0 = 10^6$ N^2 s/rad. Then in the discrete form $f(t)$ is a vector of pulses $f = [f_1, \ldots, f_N]^T = \sigma\theta$, where $\sigma = \sqrt{2\pi S_0/\Delta t}$ and θ is a standard Gaussian random vector. As described by Der Kiureghian and Li [2] most events of interest in random vibration analysis can be represented in terms of the instantaneous failure event

$$E_\tau = \{x(\tau) > x_0\}, \tag{36}$$

i.e., the event that the response at a specified time τ exceeds a specified threshold x_0. The corresponding failure domain is:

$$F_\tau = \{\theta : x(\tau) > x_0\}. \tag{37}$$

We consider the specific time instance $\tau = 12$ s, three different thresholds $x_0 = K\sigma_0$, where $K = 3, 4, 5$ and $\sigma_0^2 = \pi S_0/ck$ is the stationary response variance for the linear case ($\gamma = 0$). We use $\Delta\tau = 0.01$ so that the dimension $\dim\theta = N = \tau/\Delta\tau + 1 = 1201$.

If the threshold is too large, then the probability $p_F = P(F_\tau)$ is small and the intersection of F_τ with a random two-dimensional plane is zero. Let θ^* and θ_L^* denote the design points for the nonlinear and for the linear ($\gamma = 0$) problems correspondingly. Then, the two-dimensional plane $\pi(\theta^*, \theta_L^*)$ formed by θ^* and θ_L^* contains failure points for sure. Note, that

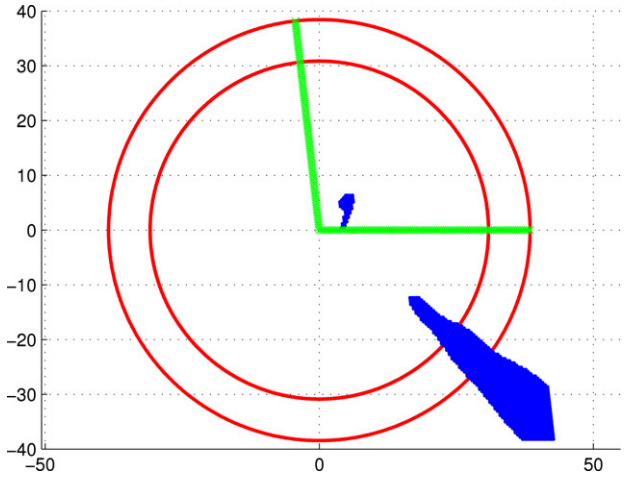

Fig. 6. Intersection of failure domain F_τ with the two-dimensional plane $\pi(\theta^*, \theta_L^*)$ for $K = 3$.

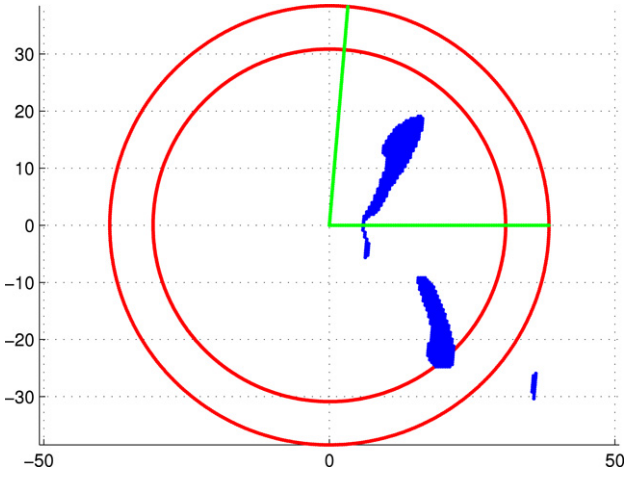

Fig. 7. Intersection of failure domain F_τ with the two-dimensional plane $\pi(\theta^*, \theta_L^*)$ for $K = 4$.

θ^* can be found using a mirror-image excitation [6]. In Figs. 6–8 the intersections $F_\tau \cap \pi(\theta^*, \theta_L^*)$ for $K = 3, 4, 5$ are shown. The concentric circles are the intersection of the Important Ring with this plane and the rays shown are the directions of θ^* (the ray in the horizontal direction) and θ_L^*. The reliability index for this nonlinear problem is $\beta = 3.76$, with corresponding failure probability $1 - \Phi(\beta) = 8.5 \times 10^{-5}$. Note, that the true failure probability is equal to 0.0092 (Monte Carlo with 10^4 samples).

Let f_s^* denote an excitation along the design point direction $f^* = \sigma\theta^*$, i.e.

$$f_s^* = sf^* = s\sigma\theta^*, \tag{38}$$

where s is a scaling factor allowing to realize scaled versions of the design excitation. The corresponding response $x(\tau)$ at a specified time τ is shown in Fig. 9 plotted against the norm of f_s^*. When $s = 1$ we have an excitation that correspond to the design point itself and the response reaches the threshold x_0. Then there is a small period of exceedance of the threshold and after that the response begins to decrease so that around the Important Ring ($s \sim \sqrt{N} = 34.64$) we have a safe region.

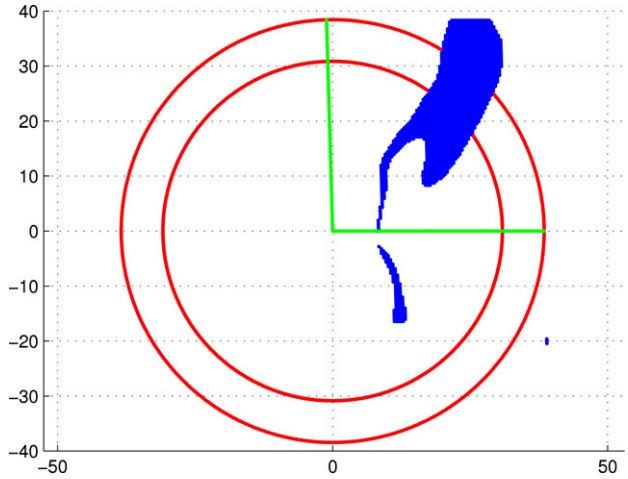

Fig. 8. Intersection of failure domain F_τ with the two-dimensional plane $\pi(\theta^*, \theta_L^*)$ for $K = 5$.

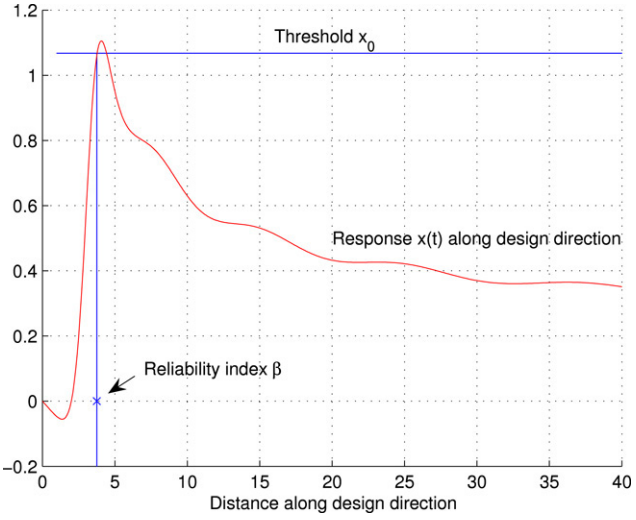

Fig. 9. Response of the Duffing oscillator along design point direction.

These figures show that the design point itself as well as the direction of the design point can be of no consequence when searching the main parts of the failure domain (intersections of failure domain with the Important Ring).

3.2. Nonlinear failure domain of parabolic shape

The design points are usually used in conjunction with Importance Sampling for calculating the failure probability. For the first passage problem the failure domain F corresponding to a time duration T can be represented as a union of elementary failure regions F_i, defined similarly to (37) as exceedance at time $\tau_i = i\Delta\tau$:

$$F_i = \{\theta : x(\tau_i) > x_0\}, \quad i = 0, \ldots, N = \frac{T}{\Delta\tau}. \tag{39}$$

When the design points θ_i^* are known, the ISD can be constructed as a weighted sum of Gaussian PDFs centered at the design points, i.e.:

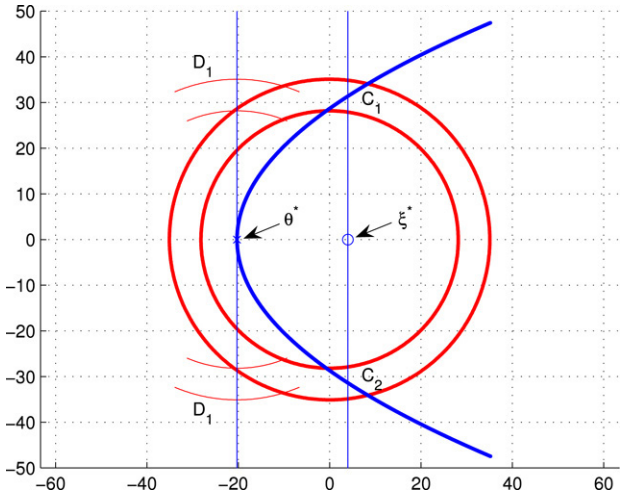

Fig. 10. Failure domain of parabolic shape.

$$q_{is}(\theta) = \sum_{i=1}^{N} w_i \mathcal{N}_{(0,1)}(\theta - \theta_i^*). \tag{40}$$

Our goal is to demonstrate that when the failure domain is strongly nonlinear Importance Sampling with ISD (40) can be feeble.

We consider a paraboloid in N-dimensional space defined as follows:

$$P : x_1 = a \sum_{i=2}^{N} x_i^2 - b, \tag{41}$$

and define failure domain as the interior of this paraboloid:

$$F = \left\{ x \in \mathbb{R}^N : x_1 > a \sum_{i=2}^{N} x_i^2 - b \right\}. \tag{42}$$

The intersection of this high-dimensional failure domain with an arbitrary plane containing the x_1 direction is shown in Fig. 10. In this example $a = 0.025$, $b = 20.27$ and $N = 1000$. The probability of this parabolic failure domain calculated using standard MC simulation (10^4 samples and 100 runs) is equal to $p_F = 0.00074$ with c.o.v. $\delta = 0.41$.

However if we use Importance Sampling method with Gaussian ISD centered at the design point $\theta^* = (-b, 0, \ldots, 0)$ we will get underestimation. For 10^4 samples and 20 runs we got $p_F^{is} = 0$. This happens because, as was explained in Section 2, all samples generated from the Gaussian density centered at θ^* will lie in the Important Ring centered at θ^* and will be almost perpendicular to the fixed x_1-direction. In the two-dimensional picture of Fig. 10 these regions are denoted as D_1 and D_2 which clearly do not correspond to failure regions. Thus, all samples generated by such an IS density yield zero indicator function and, therefore, the resulting IS estimate of the failure probability turns out to be zero.

Suppose now that we know the shape of the failure domain. Then from the previous discussion a good choice of ISD would be a Gaussian PDF $\mathcal{N}(\xi^*, \sigma^2)$ centered at ξ^*, where ξ^* is a point along the x_1-direction, such that

$$P \cap \pi_{\xi^*}^{\perp} = P \cap \mathbb{S}_{\sqrt{N}}^{N-1}, \tag{43}$$

where $\pi_{\xi^*}^{\perp}$ is a hyperplane passing through ξ^* and normal to the x_1 direction, $\mathbb{S}_{\sqrt{N}}^{N-1}$ is a hypersphere of radius $\sqrt{N}$ (middle hypersphere in the Important Ring). In Fig. 10 regions C_1 and C_2 belong to the intersection described by (43). The value of ξ^* can be easily calculated as:

$$\xi^* = \frac{\sqrt{4a^2N - 4ab + 1} - 1}{2a}. \tag{44}$$

The variance σ^2 of this Gaussian ISD has to be such that

$$E\|x - \xi^*\|^2 = \|\xi^* C_1\|^2, \quad x \sim \mathcal{N}(\xi^*, \sigma^2). \tag{45}$$

Since

$$E\|x - \xi^*\|^2 = N\sigma^2, \qquad \|\xi^* C_1\|^2 = N - (\xi^*)^2 \tag{46}$$

we finally obtain:

$$\sigma^2 = \frac{N - (\xi^*)^2}{N}. \tag{47}$$

Now if we will use the described ISD then almost all samples will lie in the important region of the failure domain, that is, in the intersection of the failure domain and the important ring and as a result the estimate of the failure probability obtained will be unbiased and accurate. The result obtained from simulations is $p_F = 0.00076$ with c.o.v. $\delta = 0.83$ (10^4 samples and 100 runs), showing that the obtained failure probability estimate is unbiased. However, the coefficient of variation is larger than that obtained from MC.

Note that finding ξ^* assumes knowledge not only of the axis of the paraboloid but also of its curvature. Furthermore, in non-toy problems the geometry of the important region of the failure domain is by far not as simple as the $(N-2)$-dimensional sphere described by (43). As has become evident by now, in order for IS to be successful its ISD must be chosen such that its corresponding important ring contains the important region of the failure domain; the latter is a subset of the important ring corresponding to the standard Gaussian distribution. Thus, the question of finding an appropriate ISD is analogous to finding a sphere (corresponding to the ISD), such that its intersection with a fixed sphere (corresponding to the standard Gaussian density) is a given set of fixed, generally complex, geometry. Clearly selecting such a sphere is impossible, unless the specified intersection is a sphere in itself (of one dimension lower) as in the paraboloid example considered earlier.

The above point is made again more formally as follows. Let F denote the failure domain and S the important spherical ring corresponding to the standard Gaussian. Then, the important region of the failure domain is: $\tilde{F} = F \cap S$. Based on the earlier discussions: $p_{\tilde{F}} = p_F$. Let now S_{is} denote the important ring corresponding to the ISD. A necessary condition for IS to work is that it must satisfy: $\tilde{F} \subset S_{is}$. Therefore, since $\tilde{F} \subset S$ also the following must hold: $\tilde{F} \subset (S_{is} \cap S)$. The intersection of two N-dimensional spherical rings has very specific geometry.

In particular, if we idealize the spherical rings by spheres, this intersection is a sphere itself of one dimension less. In the general case where $\tilde{F}$ is of arbitrary geometry, the necessary condition $\tilde{F} \subset (S_{is} \cap S)$ can therefore, not be satisfied, unless $S \equiv S_{is}$ which reduces to standard Monte Carlo.

Therefore, we conclude that IS is not applicable in general nonlinear high-dimensional reliability problems of practical interest.

4. Subset simulation

In the previous section we saw that getting information about the failure domain is quite difficult. Another approach for evaluating failure probability is Subset Simulation introduced by Au and Beck [1].

The main idea of this method is as follows. Given the original failure domain F let $F_1 \supset \cdots \supset F_n = F$ be a filtration, in other words a sequence of failure events so that $F_k = \cap_{i=0}^{k} F_i$. Using the definition of conditional probability it can be shown that,

$$p_F = P(F_1) \prod_{i=1}^{n-1} P(F_{i+1}|F_i). \tag{48}$$

The main observation is that, even if p_F is small, by choosing n and $F_i, i = 1, \ldots, n-1$ appropriately, the conditional probabilities can be made large enough for efficient evaluation by simulation.

In engineering applications the failure event usually can be expressed in terms of exceedance of some demand-capacity ratio and, therefore, the probability of failure can be written in the form $p_F = P(X > b)$, where X is the response variable and b is a critical threshold. The sequence of intermediate failure events $\{F_1, \ldots, F_{n-1}\}$ can then be chosen as $F_i = \{X > b_i\}$ for some intermediate thresholds $b_1 < \cdots < b_n = b$. During Subset Simulation, $b_1, \ldots, b_{n-1}$ are adaptively chosen such that all probabilities $P(F_1), P(F_2|F_1), \ldots, P(F_n|F_{n-1})$ are equal to, say, $p_0 = 0.1$.

Let us briefly recall how Subset Simulation works. We start by simulating n samples $\{\theta_0^k, k = 1, \ldots, n\}$ by standard Monte Carlo simulation. Perform n system analyses to obtain the corresponding response values $\{X(\theta_0^k), k = 1, \ldots, n\}$. Then, the first intermediate threshold b_1 is adaptively chosen as the $((1-p_0)n-1)$th value in the ascending list of response values, so that the sample estimate for $P(F_1) = P(X > b_1)$ is equal to p_0. There are np_0 samples among $\{\theta_0^k, k = 1, \ldots, n\}$ whose response X is greater than b_1, and hence lie in F_1. These samples are distributed as $q(\cdot|F_1)$ and provide "seeds" for simulating additional samples. Starting from each of these samples Markov chain Monte Carlo simulation (MCMC) is used to obtain an additional $(1-p_0)n$ samples, making up a total of n conditional samples $\{\theta_1^k, k = 1, \ldots, n\}$ distributed according to $q(\cdot|F_1)$. The intermediate threshold b_2 is then adaptively chosen as the $((1-p_0)n-1)$th value in the ascending list of $\{X(\theta_1^k), k = 1, \ldots, n\}$, and it defines the next intermediate failure event $F_2 = \{X > b_2\}$. The sample estimate for $P(F_2|F_1)$ is automatically equal to p_0. Repeating this process, one can generate conditional samples for higher-conditional levels until the target failure probability level has been reached.

4.1. Original metropolis algorithm

The Metropolis algorithm belongs to the class of very powerful techniques, called Markov chain Monte Carlo simulations, for simulating samples according to an arbitrary distribution. In these methods samples are simulated as the states of a Markov chain which has the target distribution as its stationary distribution.

The significance of the Metropolis algorithm in Subset simulation is that it allows to construct a Markov chain with $q(\cdot|F_i)$ as its stationary distribution. Therefore, we can use this algorithm for simulating new samples starting from "seeds" that were obtained in the previous step of Subset Simulation. Even if the current sample is not distributed as $q(\cdot|F_i)$, the limiting distribution property of Markov chain guarantees that the distribution of simulated samples will tend to $q(\cdot|F_i)$ as the number of Markov steps increases.

Given a current sample $\theta^{(1)}$ (a "seed" point) the original Metropolis algorithm works as follows.

Let $p(\xi|\theta)$, called proposal PDF, be a N-dimensional PDF for ξ centered at θ with symmetry property $p(\xi|\theta) = p(\theta|\xi)$. Generate a sequence of samples $\{\theta^{(1)}, \theta^{(2)}, \ldots\}$ starting from a given sample $\theta^{(1)}$ by computing $\theta^{(k+1)}$ from $\theta^{(k)}$ as follows:

1. *Generate candidate state* $\tilde{\theta}$. Simulate ξ according to $p(\cdot|\theta^{(k)})$. Compute the ratio $r = q(\xi)/q(\theta^{(k)})$. Set $\tilde{\theta} = \xi$ with probability $\min\{1, r\}$ or set $\tilde{\theta} = \theta^{(k)}$ with the remaining probability.

2. *Accept/Reject* $\tilde{\theta}$. If $\tilde{\theta} \in F_i$ accept it as a next step, i.e. $\theta^{(k+1)} = \tilde{\theta}$; otherwise reject it and take the current sample as a next state of the Markov chain, i.e., $\theta^{(k+1)} = \theta^{(k)}$.

Au and Beck [1] realized that the original Metropolis algorithm does not work in high dimensions.

The geometric reason of this inapplicability is that the same effect as that explained by Fig. 3 arises, where the symbols X and Y in this Figure now represent $\theta^{(k)}$ and ξ, respectively. Therefore, for the reasons explained previously, in each step of the original Metropolis algorithm the ratio $r = q(\xi)/q(\theta^{(k)})$ will be extremely small, broadly speaking of the order of $\exp(-N/2)$. Therefore, with extremely high probability one obtains repeated samples. Thus, a chain of practically meaningful length may consist of as few as a single sample. This renders Subset Simulation practically inapplicable.

4.2. Modified metropolis algorithm

The modified Metropolis algorithm [1] differs from the original Metropolis algorithm in the way the candidate state $\tilde{\theta}$ is generated. Rather than using a N-dimensional proposal to directly obtain the candidate state $\tilde{\theta}$, in the modified algorithm a sequence of one-dimensional proposals $p(\xi_j|\theta_j^{(k)}), j = 1, \ldots, N$ is used. Specifically, each component of $\tilde{\theta}_j$ of

the candidate state $\tilde{\theta}$ is generated separately using a one-dimensional proposal centered at $\theta_j^{(k)}$. Thus, being at state $\theta^{(k)}$ the candidate $\tilde{\theta}$ for the next state θ^{k+1} of the Markov chain is generated as follows:

For each component $j = 1, \ldots, N$ simulate ξ_j from $p(\cdot|\theta_j^{(k)})$. Compute the ratio $r_j = q_j(\xi_j)/q_j(\theta_j^{(k)})$. Set $\tilde{\theta}_j = \xi_j$ with probability $\min\{1, r_j\}$ and set $\tilde{\theta}_j = \theta_j^{(k)}$ with the remaining probability.

Once the candidate state $\tilde{\theta}$ has been generated according to the above procedure, the second step of the standard Metropolis algorithms is performed involving the acceptance/rejection of $\tilde{\theta}$ as the next state of the Markov chain. Thus, if $\tilde{\theta}$ is found to be a failure point it is accepted as the next state θ^{k+1}; otherwise, $\theta^{(k+1)} = \theta^{(k)}$, i.e., one obtains a repeated sample.

It can be easily seen that the modified Metropolis algorithm overcomes the deficiency of standard Metropolis algorithm by producing distinct candidate states, rather than repeated samples. Specifically, for large N it is highly unlikely that the candidate state $\tilde{\theta}$ is equal to the current state $\theta^{(k)}$, because this would mean that all N components ξ_j were rejected as candidate state components, which is highly unlikely.

However, it can be shown that the above modified algorithm is not perfect, in the sense that it does not guarantee for a chain starting within the important ring that it will remain within this ring at all times. Specifically, we will show that there is a speculative chance that such a chain may exit the important ring for a while. This is demonstrated by the following example. Ignore for a moment the Acceptance/Rejection step, in other words assume that the entire space $\mathbb{R}^N$ is a failure domain. Consider a Markov chain starting from a random vector $\hat{\theta}$ drawn from the N-dimensional standard Gaussian distribution and governed by the modified Metropolis algorithm with proposal PDF $p(\cdot|\theta_j) = \mathcal{N}(\theta_j, 1)$. In Fig. 11 the evolution of the Euclidean norm (length) of the Markov chain state is shown. The horizontal lines show the inner and outer radii of the Important Ring. This figure demonstrates that all states of the generated chain belong in the important ring.

Now consider the same Markov chain but starting not from a random point within the important ring but from a point along one of the coordinate axes. For example, consider as starting point the point $\hat{\theta} = (\sqrt{N}, 0, \ldots, 0)$ which clearly also belongs to the important ring. The evolution of the Euclidean norm of the states of such a Markov chain is shown in Fig. 12. As can be seen from this figure the Markov chain jumps out of the Important Ring and only after approximately 50 steps returns to it. Thus, using a chain of small length, as is done in practical applications, will result into misleading conclusions. This is because such a chain has anything but reached its underlying stationary distribution.

This happens because the modified Metropolis algorithm assumes some structure while the standard Gaussian space is isotropic and has no preferred directions. However, in the modified Metropolis algorithm each coordinate axis direction is a special direction. The closer the initial state of a Markov chain is to one of the x_j-directions the worse the modified Metropolis

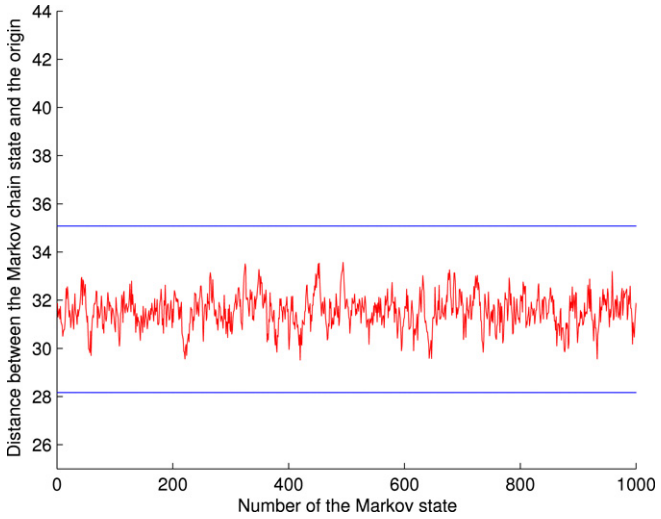

Fig. 11. Markov chain using MMA starting from the random Gaussian vector ($N = 10^3$).

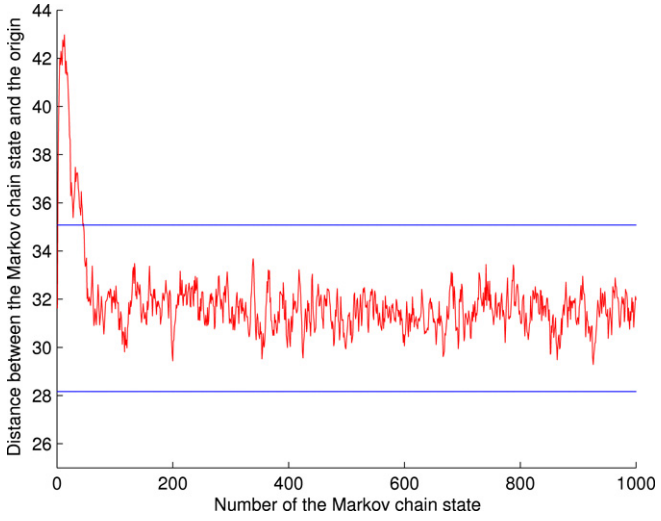

Fig. 12. Markov chain using MMA starting from $\hat{\theta} = (\sqrt{N}, 0, \ldots, 0)$, $N = 10^3$.

algorithm works. Of course the probability to obtain an initial state closely aligned to one of the coordinates is very small.

To demonstrate the effect of the above behavior when calculating reliability estimates consider a linear problem with LSF (31) with corresponding failure probability (34). Let us apply Subset Simulation with modified Metropolis algorithm for evaluating the failure probabilities of two "equivalent" linear failure problems. The design point of the first problem is chosen as $\theta_1^* = \beta \frac{x}{\|x\|}$, where x is drawn from N-dimensional standard Gaussian distribution, while that of the second problem is along one of the "preferred" directions, namely, $\theta_2^* = (\beta, 0, \ldots, 0)$. These failure domains have equal probabilities, since they have the same reliability index β. In Fig. 13 the coefficients of variation (c.o.v.) for corresponding estimators for different values of n (number of samples in each intermediate Subset Simulation level) are shown. Here $\beta = 3$, $N = 10^3$ and 50 runs of the algorithm are used. Clearly in the second case the modified Metropolis algorithm works worse.

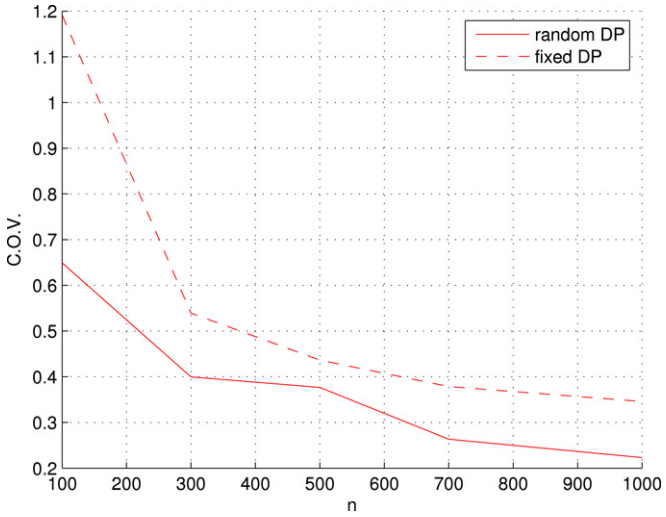

Fig. 13. C.o.v. of p_F estimator for linear failure domain with random and fixed design points.

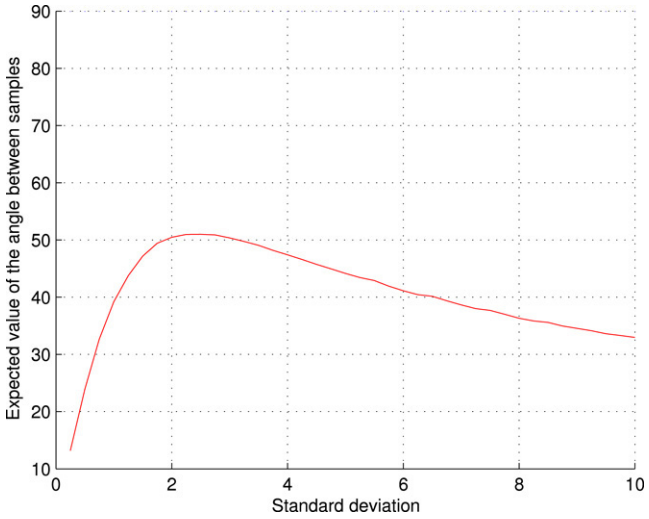

Fig. 14. Expected value of angle between samples plotted against standard deviation of proposal ($N = 10^3$).

This is due to the existence of preferred directions and the earlier discussed deficiency of Markov Chains generated from points closely aligned with these directions.

In the implementation of Subset Simulation the choice of proposal distribution is very important since it governs the efficiency in generating the samples by MCMC. Let us again consider a Markov chain simulated by the modified Metropolis algorithm when the entire space is a failure domain, i.e., when there is no Acceptance/Rejection step. For each component we use a Gaussian proposal as before but instead of using unit variance we consider it as a parameter:

$$p(\cdot|\theta_j) = \mathcal{N}(\theta_j, \sigma^2), \quad j = 1, \ldots, N. \tag{49}$$

The question we consider is how σ affects the correlation between samples. The expected value of the angle α between two consecutive samples is a good parameter for measuring the correlation of the Markov chain states as it provides a measure of the separation of samples that is independent of the dimension N while alternative measures, such as the Euclidean distance between samples, depend on N. The larger the angle, the more independent are the samples and, therefore, the more ergodic the chain. Clearly, standard Monte Carlo provides the most uncorrelated samples satisfying:

$$E[\alpha] = \pi/2. \tag{50}$$

For the described one-parameter family of Markov chains the expected value of α is shown in Fig. 14. As we can see the "optimal" standard deviation is somewhere between 2 and 3.

Of course in real engineering applications the failure domain is only a small subset of the parameter space and, therefore, adopting too large standard deviation in the proposal PDF may lead to too many rejections in the second step of the modified Metropolis algorithm yielding a highly-correlated chain.

4.3. Spherical subset simulation method (S^3)

A new MCMC algorithm was proposed by Katafygiotis and Cheung [5]. Given a sample $\theta = Ru$, where $u = \theta/\|\theta\| \in \mathbb{S}^{N-1}$ is a vector of unit length and $R = \|\theta\|$, this algorithm consists of two steps. The first step involves sampling an intermediate failure point $\theta' = Ru'$, having the same length R as the current sample θ. That is, in the first step we simulate a new failure point lying on the sphere of radius R. Clearly, all failure points on this sphere are uniformly distributed as they have equal corresponding PDF values. Therefore, the selection of one such failure point is performed by first selecting a random plane $\pi(u, \tilde{u})$ containing the direction u and a uniformly distributed random direction $\tilde{u}$ realized as $\tilde{u} = Z/\|Z\|$ where Z is distributed according to the standard Gaussian. Next, we consider the circle C of radius R defined as the intersection of the sphere of radius R and the random plane $\pi(u, \tilde{u})$, i.e.,

$$C = \left\{\theta : \theta = R\frac{(u\cos\phi + \tilde{u}\sin\phi)}{\|u\cos\phi + \tilde{u}\sin\phi\|}\right\}, \tag{51}$$

where $\phi \in [0, 2\pi]$. Finally, we simulate uniformly distributed points on this circle, by drawing ϕ uniformly from the interval $[0, 2\pi]$ until a failure point Ru' is reached.

The second step involves radial sampling along the direction u' to obtain R'. The distinct sample $\theta'' = R'u'$ is then the next state of the chain, i.e. it is accepted with probability one. For more details, refer to [5]. From a geometric point of view the advantage of this method is that it is consistent with the nature of the Gaussian space, namely there are no preferred directions as in the modified Metropolis algorithm. If we assume that the entire domain corresponds to failure it is clear that the expectation of the angle between samples generated by the S^3 method is equal to $\pi/2$ as in standard Monte Carlo.

Next, we compare the angular distances between failure samples generated by standard Monte Carlo, Modified Metropolis algorithm (MMA) with $\sigma = 1, 2$ and the above described S^3 Markov Chain algorithm in the case of a linear failure domain, where the Markov Chain is initiated at one of the failure points obtained by Monte Carlo simulations. The results are presented in Fig. 15. We can see that in terms of the angular distances between samples, the S^3 method is much

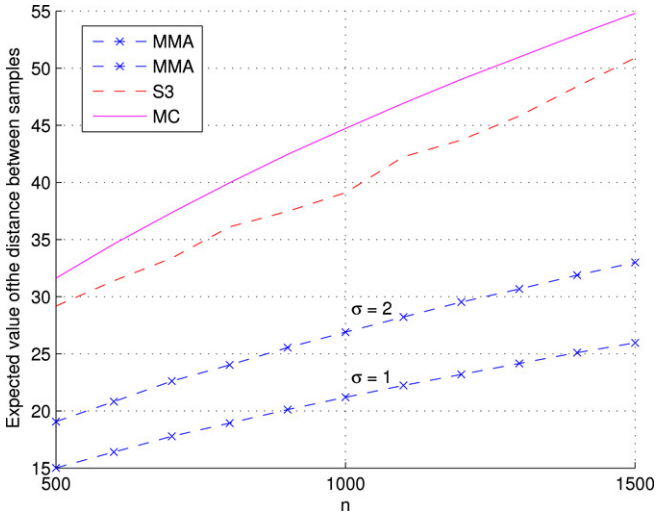

Fig. 15. Expected value of angular distance between samples plotted against dimension N for MMA, S^3 and MC.

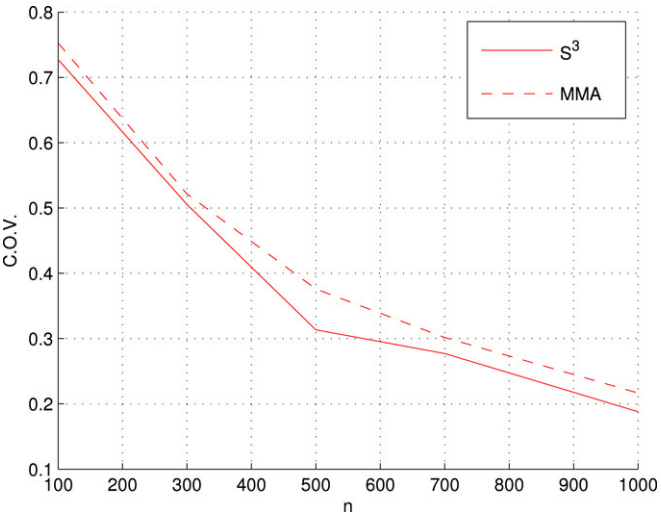

Fig. 16. C.o.v. of p_F estimator for linear failure domain using S^3 and MMA.

closer to Monte Carlo simulation and, therefore, one can argue that it provides a better quality Markov chain.

Finally Fig. 16 shows the c.o.v of the failure probability estimates obtained using Subset simulation employing two alternative Markov Chain algorithms: a) the S^3 Markov Chain algorithm described in this Section and b) MMA ($\sigma = 1$). Here reliability index $\beta = 3$ and 50 runs of each algorithms were used. As can be seen, the S^3 Markov Chain algorithm provides better estimates due to the higher-quality (more uncorrelated) chain it produces.

5. Conclusions

In this paper a geometric perspective has been adopted to provide insight into the computational difficulties arising in high-dimensional reliability problems. Several important results are explained following this perspective. An explanation is provided as to why Importance Sampling using a fixed importance sampling density is inapplicable in moderate to strongly nonlinear high-dimensional reliability problems. The reason why the Metropolis Algorithm is inapplicable in high dimensions and is bound to produce repeated samples is also given. Potential deficiencies of the Modified Metropolis Algorithm are revealed and an alternative algorithm that overcomes these deficiencies is presented. In summary, the paper provides a deeper understanding of the geometric features of high-dimensional reliability problem. This understanding is invaluable in developing novel efficient algorithms for treating such problems.

Acknowledgments

This research has been supported by the Hong Kong Research Grants Council under grants 614305 and CA04/05.EG01. This support is gratefully acknowledged.

References

[1] Au SK, Beck JL. Estimation of small failure probabilities in high dimensions by subset simulation. Probabilistic Engineering Mechanics 2001;16(4):263–77.
[2] Der Kiureghian A, Li C-C. Nonlinear random vibration analysis through optimization. In: Proceedings of the seventh IFIP WG 7.5 conference on optimization of structural systems. 1996. p. 197–206.
[3] Luisa Canal. A normal approximation for the chi-square distribution. Computational Statistics & Data Analysis 2005;48:803–8.
[4] Fisher RA. On the interpretation of χ^2 from contingency tables and calculations of P. Journal of the Royal Statistical Society A 1922;85: 87–94.
[5] Katafygiotis LS, Cheung SH. Application of spherical subset simulation method and auxiliary domain method on a benchmark reliability study, Structural Safety 2006 (in print).
[6] Koo H, Der Kiureghian A, Fujimura K. Design-point excitation for nonlinear random vibrations. Probabilistic Engineering Mechanics 2005;20(2): 136–47.