Targeted Random Sampling for Reliability Assessment: A Demonstration of Concept

Michael D. Shields
Johns Hopkins University, michael.shields@jhu.edu

Abstract. Monte Carlo simulation provides the benchmark for computational methods in reliability assessment in terms of both accuracy and robustness. However, Monte Carlo methods are well-known to be computationally expensive and, in many cases, intractable. Several improvements have been proposed to increase their computational efficiency for reliability analysis including Latin hypercube sampling; see (Olsson et al., 2003), importance sampling; see (Scheuller and Stix, 1987), subset sampling; see (Au and Beck, 2001), and line sampling; see (Koutsourelakis et al., 2004), among others. The primary drawback of Monte Carlo methods is that, even using the most advanced methods, a large proportion of the samples are of little, if any, use to the reliability assessment.

This work demonstrates a new Monte Carlo concept, referred to as Targeted Random Sampling (TRS), that is rooted in the recently developed Refined Stratified Sampling method; see the forthcoming work (Shields et al., 2014). TRS enables the selection of random samples from specific strata of the space (with known probability of occurrence) that are identified, based on statistical information available from existing samples, as being particularly important to estimation of the statistical quantity of interest (here probability of failure). In this work, the concept is developed and it is demonstrated how the method can be used to concentrate random samples in the vicinity of the failure surface to facilitate very rapid convergence of the probability of failure estimate.

Keywords: Monte Carlo Simulation; Variance Reduction; Stratified Sampling; Refined Stratified Sampling; Structural Reliability

1. Introduction

Generically speaking, reliability analysis refers to the evaluation or estimation of the probability of failure of a (deterministic or stochastic) system under specified (deterministic or stochastic) conditions. Stated mathematically, this probability of failure, p_f is evaluated as:

$$p_f = \int_{\Omega_f} f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} \tag{1}$$

where Ω_f denotes the failure domain defined as $g(\mathbf{x}) \leq 0$ for the limit state function $g(\mathbf{x})$ with random vector $\mathbf{X}$ describing the stochasticity of the system and possessing joint probability density $f_{\mathbf{X}}(\mathbf{x})$. The evaluation of this integral becomes exceedingly difficult for many practical applications. The most common difficulty arises when the limit state function $g(\mathbf{x})$ is not known explicitly and takes a complex form possessing, for example, strong nonlinearities, non-monotonicity, and/or discontinuities (resulting in disjoint failure regions).

An expansive body of research has been developed over the past several decades across numerous disciplines aimed specifically at developing methods to compute or estimate this integral expression for various systems. The methods that have been employed in the field of structural reliability can be broadly classified into two groups: approximate methods and exact methods. In general, the approximate methods aim to approximate the function $g(\mathbf{x})$ with some simpler representation. The simplest of the approximate methods is the First Order Reliability Method (FORM) (Hasofer and Lind, 1974), which expands the limit state function as a first-order Taylor Series around the so-called design point - thus linearizing the limit-state function. Despite its low order, FORM approximations have proven effective for many problems particularly when the limit state function is smooth and not strongly curved. Numerous higher-order approximations are available beginning with the Second Order Reliabilty Method (SORM) (Breitung, 1984) that uses the Taylor Series expansion as in FORM while keeping the second order terms and continuing up to those utilizing a response surface as a surrogate for the system itself (Faravelli, 1989) or for the limit state function (Hurtado, 2004). These approximate methods will not be discussed further here.

The exact methods for reliability analysis aim to determine a solution to Eq. (1) directly. Using these methods, the probability of failure integral is reformulated as:

$$p_f = \int_{\mathbb{R}^n} I(\mathbf{x} \in \Omega_f) f_{\mathbf{X}}(\mathbf{x}) d\mathbf{x} \tag{2}$$

where $I(\cdot)$ denotes the indicator function. There are two general approaches to solving this integral exactly: numerical integration and Monte Carlo simulation. Numerical integration is intractable for most practical problems because failure probabilities are typically very small and, consequently, a very large number of integration points (computationally expensive function evaluations) are necessary for accurate estimation. Monte Carlo simulation is perhaps the most widely used method and, when it can be achieved, it is the most accurate and robust means of evaluating reliability. Monte Carlo simulation evaluates Eq. (2) as:

$$p_f = \sum_{i=1}^{N} I(\mathbf{x}_i \in \Omega_f) p(\mathbf{x}_i) \tag{3}$$

where $I(\cdot)$ again denotes the indicator function at a *randomly selected* point $\mathbf{x}_i$, $p(\mathbf{x}_i)$ denotes the probability weight assigned to sample $\mathbf{x}_i$, and N denotes the sample size. In classical MCS, the $\mathbf{x}_i$ are independent and identically distributed (*iid*) samples of the random vector $\mathbf{X}$ and their associated probability weights are given by $p(\mathbf{x}_i) = \frac{1}{N}$ with the summation converging to the true probability of failure as $N \to \infty$.

Monte Carlo simulation typically requires a very large number of samples $\mathbf{x}_i$ and, therefore is considered intractable for many problems. However, several methods have been developed with the intention of reducing the sample size while maintaining the required accuracy of the estimate. Several of these methods are discussed briefly in the following section. In this work, a new sampling method referred to as Targeted Random Sampling (TRS) is developed that relies on the stratified sampling design – and specifically the newly developed concept of Refined Stratified Sampling (RSS) – to target samples in regions of the space that are of high consequence to the probability of failure estimate (i.e. in the vicinity of the failure surface). The method operates by iterative refining a stratified design to add samples in the regions separating points identified as failure and non-failure. A simple algorithm for this refinement is provided along with some discussion of its effectiveness. The method is shown to perform remarkably well for a demonstration problem with a highly complex and strongly nonlinear limit state function.

2. Monte Carlo Methods for Reliability Analysis

Computational efficiency is the paramount concern with Monte Carlo simulation for reliability analysis. For this reason, several methods have been developed - broadly referred to as variance reduction techniques - to produce samples of the random variables for Monte Carlo simulation that minimize the number of calculations required to accurately estimate a given quantity of interest (here probability of failure). Some of the most common methods employed are Latin hypercube sampling (Olsson et al., 2003), importance sampling (Scheuller and Stix, 1987), subset sampling (Au and Beck, 2001), and line sampling (Koutsourelakis et al., 2004). A detailed discussion of these methods is beyond the scope of this work.

The importance sampling, subset sampling, and line sampling methods share a common objective but rely on fundamentally different approaches to achieve this objective. Simply stated, the objective of these methods is to produce an optimal sample set that is concentrated in the vicinity of the failure surface to reduce the cost of Monte Carlo probability of failure estimates. Note this is the precise objective of the Targeted Random Sampling method presented herein as well. By achieving this objective, the methods can significantly reduce computational cost. However, they are not without their drawbacks. Importance sampling can be highly effective for problems in relatively low dimension where a single design point can be identified around which to center the samples. But, it will struggle with problems possessing multiple or competing design points. The subset simulation method is very effective in high dimension because the convergence rate does not depend on the number of random parameters in the problem. The drawback of the subset sampling method is that the Markov Chain used to generate samples leads to dependence among the samples that produces biased estimates of the probability of failure. Lastly, the line sampling method reduces the stochastic dimension by one degree and performs a line search by identifying an "important direction." As such, this method has been demonstrated to be effective for high dimensional problems but still may necessitate a relatively large sample size if the important direction is difficult to identify. In this case, a large number of samples are required simply to identify the important direction which is followed by a significant number of samples to produce the probability of failure estimate.

The method proposed herein shares this common objective – taking an approach rooted in stratified sampling and the newly developed Refined Stratified Sampling (RSS) method introduced by the author. The following sections outline the concepts of stratified sampling and RSS needed for use in developing the Targeted Random Sampling method outlined in Section 5.

3. Stratified Sampling

Stratified sampling operates by dividing the sample space of the random variables $\mathbf{S_X}$ into a collection of M disjoint subsets (strata) $\mathbf{\Omega}_k; k = 1, 2, \ldots, M$ with $\cup_{k=1}^{M} \Omega_k = \mathbf{S_X}$ and $\Omega_p \cap \Omega_q = \emptyset; p \neq q$. Samples $\mathbf{x}_k; k = 1, 2, \ldots, M$ are then drawn at random from within each of the strata according to:

$$x_{ik} = F_{X_i}^{-1}(U_{ik}); i = 1, 2, \ldots, n; k = 1, 2, \ldots, M \tag{4}$$

where $F_{X_i}(\cdot)$ is the marginal cumulative distribution function of variable X_i, U_{ik} are iid uniformly distributed samples on $[\xi_{ik}^l, \xi_{ik}^u]$ with $\xi_{ik}^l = F_{X_i}(\zeta_{ik}^l)$ and $\xi_{ik}^u = F_{X_i}(\zeta_{ik}^u)$ and ζ_{ik}^l and ζ_{ik}^u denote the lower and upper bounds respectively of the i^{th} vector component of stratum $\mathbf{\Omega}_k$.

Samples drawn using a stratified design according to Eq. (4) and the subsequently defined strata bounds are drawn from a rectilinear grid of the space. Consider, for example, a two-dimensional random vector $\mathbf{X} = [X_1, X_2]$ where X_1 has a lognormal distribution with marginal probability density function (pdf) given by:

$$f(x_1) = \frac{1}{x_1\sqrt{2\pi}\sigma} \exp\left[-\frac{(\ln x_1 - \mu)^2}{2\sigma^2}\right] \tag{5}$$

with parameters $\mu = 0$ and $\sigma = 0.5$ and X_2 is uniformly distributed over the range $[0, 5]$ (i.e. $f(x_2) = \frac{1}{5}|_{x_2 \in [0,5]}$). Sixteen samples drawn from strata of equal probability are shown in the probability space and the mapped sample space in Figure 2. Note that other stratified designs are possible and, in

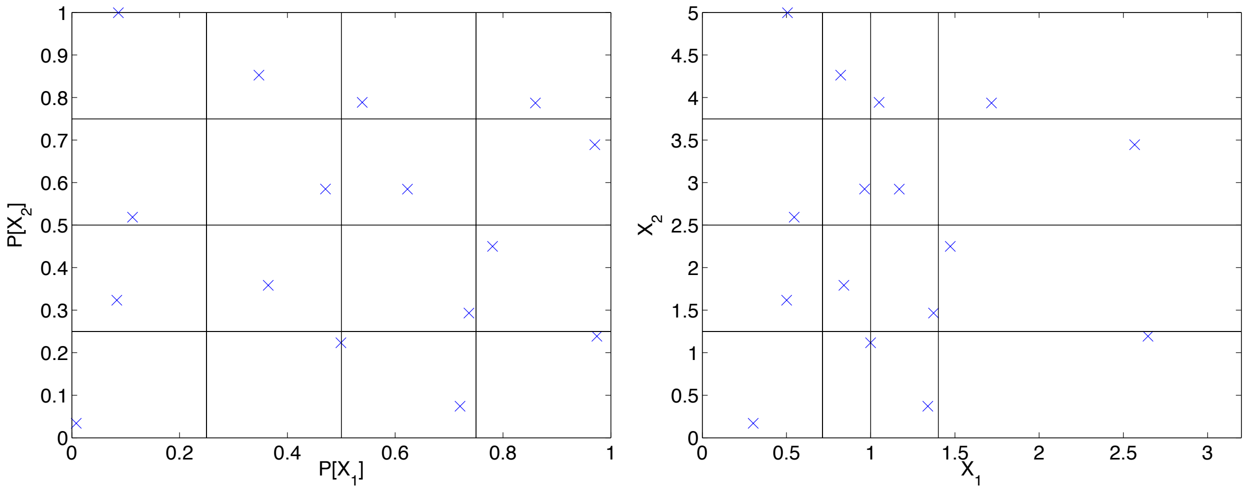

Figure 1. Sample realizations of a 2D random vector $\mathbf{X} = [X_1, X_2]$ with X_1 ∼LogN(0, 0.5) and X_2 ∼Unif(0, 5) generated using stratified sampling in the probability space (left) and sample space (right).

general, a stratified design is not required to divide the space into such a grid or even strata of equal probability.

The, in general, unequal stratum sizes necessitates the use of probabilistic weights in statistical calculations performed on samples drawn from a stratified design. These weights are assigned according to the probability of occurrence of the stratum, $\mathbf{\Omega}_k$, being sampled as:

$$w_k = \frac{P[\mathbf{\Omega}_k]}{M_k} \tag{6}$$

where M_k is the total number of samples in stratum $\mathbf{\Omega}_k$. For probability of failure estimation using Eq. (3) and assuming a single sample per stratum (i.e. $M_k = 1; \forall k$), these stratum weights correspond to the probability weights $p(\mathbf{x}_i)$. Note that for problems with non-rectilinear strata (particularly in high dimension), the calculation of the strata weights can be a difficult task.

Stratified designs possess several nice features that make them attractive, in general, for Monte Carlo simulation. First, stratified sampling has been demonstrated to *always* reduce the variance of statistical estimators when compared to *iid* random sampling. Consider the class of statistical estimators defined by:

$$T(\mathbf{Y}) = \frac{1}{N}\sum_{i=1}^{N} g(Y_i) \tag{7}$$

where $g(\cdot)$ is an arbitrary function. This form of estimators can be used to estimate statistical moments, empirical distribution functions, and other characteristics of the quantity Y. It has been shown in(McKay et al., 1979) that the difference between the variance of such an estimator using a stratified design (T_S) and the same estimator computed using *iid* random sampling (T_R) is given as follows:

$$Var(T_S) - Var(T_R) = -\frac{1}{N}\sum_{i=1}^{I} p_i(\mu_i - \tau)^2 \tag{8}$$

where p_i denotes the probability weight of stratum i computed using Eq. (6), μ_i denotes the mean value of Y computed over stratum i, and τ is the overall mean value. Notice that this quantity is always negative; hence stratified sampling produces a reduction in variance.

Next, balanced stratified designs (those with equal strata probabilities) possess greatly enhanced space-filling properties when compared to other variance reduction techniques such as Latin hypercube sampling (McKay et al., 1979). This property is not discussed in detail here but is a point of emphasis in a future work of the author's (Shields et al., 2014). Space-filling is a desirable property for many problems, but it is not necessarily desirable for others - such as probability of failure analysis where it is beneficial to sample heavily in the vicinity of the failure surface. In this case, unbalanced stratified designs (those with unequal strata probabilities) afford the flexibility to define strata that concentrate samples where they are most useful and avoids oversampling in other regions. In general, the optimal strata definitions are dependent on the nature of the operator - which may be for example a function, a set of differential equations, or some numerical computer model. That is, consider the basic transformation of the random vector $\mathbf{X}$ given by $Y = F(\mathbf{X})$. The optimal strata definition for $\mathbf{X}$ is dependent on the form of $F(\cdot)$ and its characteristics such as strong nonlinearities, periodicity, discontinuities, etc. The flexibility of stratified sampling to accommodate such varying operators provides the foundation of the Targeted Random Sampling presented herein.

4. Refined Stratified Sampling

The Targeted Random Sampling method relies on the ability to extend the sample size from a stratified design. That is, given a sample of size N produced from a stratified design it is necessary to produce a larger sample of size $N+k$ that retains the original N samples. The classical means of adding samples to a stratified design, originally proposed by (Tocher, 1963), simply adds samples to preexisting strata. By doing so, the value of M_k is increased and the associated sample weights are adjusted.

Recently, (Shields et al., 2014) have developed a new method, referred to as Refined Stratified Sampling (RSS), wherein the strata themselves are divided and samples are added to the newly defined (empty) strata. Upon refinement, the redefined strata weights (Eq. (6)) are computed for use in statistical calculations. In the forthcoming work describing RSS (Shields et al., 2014), it is proven that refinement of the strata produces a reduced variance in the statistical estimator T when compared to a design where samples are added to existing strata. Thus, any optimal variance reducing stratified design where samples are added should not contain more than one sample per stratum.

The ability to refine the strata and re-evaluate sample weights is coupled with the flexibility afforded by stratified sample designs to produce unbalanced designs that are catered to the specific application. This is the driving concept behind the Targeted Random Sampling method outlined in the following section.

5. Targeted Random Sampling

5.1. Basic TRS algorithm

Consider a system with random vector $\mathbf{X}$ (possessing uncorrelated components) describing its stochasticity and unknown limit state function $g(\mathbf{X})$. The Targeted Random Sampling method proceeds as follows and is demonstrated schematically for a two-dimensional system with $\mathbf{X} = [X_1, X_2]$ in Figure 2:

1. Draw an initial sample set using a stratified design as shown in Figure 2(a). Notes on this initial design are provided in Section 5.2;

2. Evaluate system response $Y = F(\mathbf{X})$ for each new sample, identifying each response as a failure or a non-failure.

3. Identify pairs of sample points existing in adjoining strata that correspond to different outcomes (i.e. one failure and one non-failure). Denote these pairs by $\mathbf{X}_{fsi} = [\mathbf{X}_{fi}, \mathbf{X}_{si}], i = 1, \dots, N_p$ where N_p is the number of such pairs, $\mathbf{X}_{fi}$ denotes a sample of $\mathbf{X}$ in pair i corresponding to a failure and $\mathbf{X}_{si}$ denotes a sample of $\mathbf{X}$ in pair i corresponding to a "safe" outcome. Note that any given sample $\mathbf{X}_j$ may be present in multiple "fail-safe" pairs. In Figure 2 the strata corresponding to these sample pairs are hatched. This hatched region demarcates a "region of interest" from which future samples should be drawn;

4. Compute the separation distance for these sample pairs in the probability space as:

$$D_i = \sqrt{\sum_{k=1}^{M} \left(F_{X_k}^{-1}(X_{fik}) - F_{X_k}^{-1}(X_{sik}) \right)^2} \tag{9}$$

   where X_{fik} and X_{sik} corresponds to the k component of $\mathbf{X}_{fi}$ and $\mathbf{X}_{si}$ respectively.

5. Identify the pair with the largest separation:

$$\mathbf{X}_{fs}^m = \max_i \left[\mathbf{X}_{fsi} \right] \tag{10}$$

   In Figure 2(b), this pair are shown by points A and B and the corresponding strata have been identified with a box.

6. Divide one of the pair strata at the bisecting point in the direction/component of the largest separation. In Figure 2(b), the direction of largest separation between points A and B is in the direction of X_2. The bisecting line may fall in either of the strata. The stratum division is made to the stratum in which the bisecting line falls. In 2(b), the stratum corresponding to sample A is divided as denoted by the dashed line.

7. Recompute the strata probabilities (weights) and draw a sample from the newly defined stratum as shown in Figure 2(c).

8. Return to step 2 and repeat.

The probability of failure estimate resulting from this process corresponds to the summation of the strata probabilities for all samples corresponding to a failure occurrence.

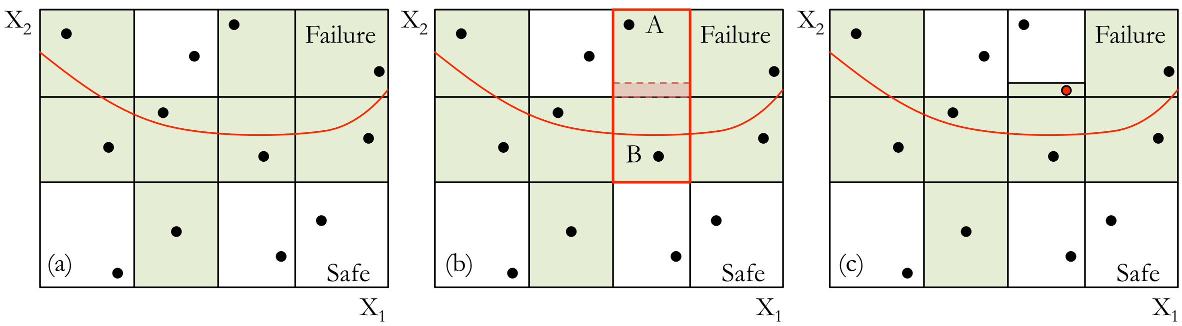

Figure 2. Conceptual demonstration of the Targeted Random Sampling method.

5.2. Identification of failure domains & initial stratified design

Clearly the above procedure relies on the identification of both failure and non-failure events. In fact, the proposed algorithm must identify at least one failure per disjoint region of the failure domain.

As such, careful consideration should be given to the initial stratified design. A space-filling design, such as will be achieved through a balanced stratification, is inappropriate for most reliability applications because probability of failure is typically very low. Consequently, a very large number of samples will be necessary just to identify a single failure occurrence - which represents a starting point for the proposed procedure. Instead, an unbalanced design is likely necessary (and perhaps even a highly unbalanced design) wherein knowledge of the system should be used to inform the initial strata definitions. For example, failure of many systems will be associated with values present in the tails (or extreme values) of the probability distributions. In this case, the strata should be defined such that a certain number of samples are drawn from the tails (and perhaps far into the tails) to ensure a failure event is realized. Similarly, failure may correspond to a combination of parameter values that is rare but not necessarily in the tails of the distribution. If the region of the space corresponding to this combination is broadly known (perhaps qualitatively), the initial stratification should concentrate samples in this region.

5.3. Additional considerations

The basic algorithm presented above does not possess the robustness that may be necessary to identify the more nuanced features of some failure domains. Rather, it is intended as a starting-point from which more sophisticated and robust algorithms will be developed. These algorithms must necessarily be able to identify subtle features in a failure domains such as the following.

- Certain seemingly disjoint regions of the failure domain may actually be connected.
- Certain seemingly connected regions of the failure domain may actually be disjoint.
- Samples near the boundary of the probability space that correspond with non-failure events may not be close enough to the boundary. That is, samples drawn from the tails of the distribution may not be drawn from sufficiently far into the tails.

This list is not intended to be comprehensive. In identifying these subtle features, it will be necessary to consider the relatively probabilities of such conditions. For example, two failure domains that appear to be disjoint but are, in fact, connected by a regime of probability that is negligibly small can be adequately treated as two distinct failure domains for practical purposes. Instead of focusing the samples on identifying the link between these domains, the samples will be better utilized by resolving the domains more completely. Such considerations should be informed by the relatively strata probabilities.

6. Example

Consider samples of a random vector $\mathbf{X} = [X_1, X_2]$ drawn from a standard bivariate normal distribution with zero mean and unit standard deviation. Failure is assumed to occur under the condition $x_2 > -\frac{x_1}{4} + \sin(5x_1) + 4$. For reference, the probability of failure for this problem was determined to be $p_f \approx 4.15e-4$ using Monte Carlo simulation with 100,000,000 samples.

The proposed targeted random sampling method was applied for 1,000 samples as demonstrated in Figure 3 producing an estimate of $p_f = 4.12e-4$. In Figure 3, samples corresponding to a failure are demarcated with an 'o' while those corresponding to a non-failure are shown with an 'x'. The figure clearly shows that samples are concentrated in the vicinity of the limit state function and, moreover the samples are concentrated near the limit state function in the regions of higher probability. That is, the samples are more heavily concentrated near the "tips" of the finger-like troughs of the limit state function that correspond to higher probabilities of **X**.

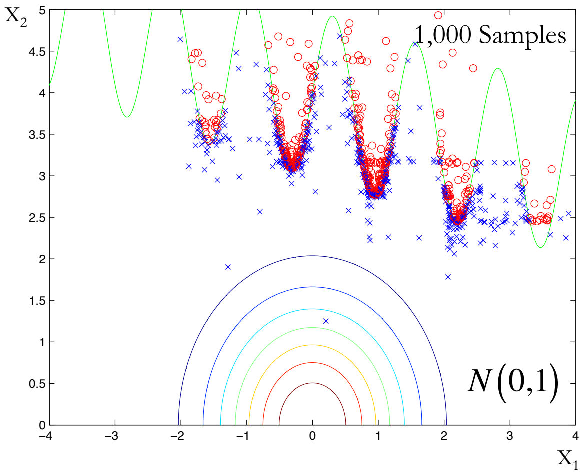

Figure 3. Example of 1,000 samples drawn using the TRS method for a 2D problem with complex failure surface.

The initial stratified design prescribed nine strata and their associated stratum probabilities as follows:

$$\begin{aligned}
\Omega_1 &= ([0.0, 1.0e-5]\,;[0.0, 1.0e-5])\,; \;\; p_1 = 1e-10 \\
\Omega_2 &= ([1.0e-5, 0.99999]\,;[0.0, 1.0e-5])\,; \;\; p_2 = 9.9998e-6 \\
\Omega_3 &= ([0.99999, 1.0]\,;[0.0, 1.0e-5])\,; \;\; p_3 = 1e-10 \\
\Omega_4 &= ([0.0, 1.0e-5]\,;[1.0e-5, 0.99999])\,; \;\; p_4 = 9.9998e-6 \\
\Omega_5 &= ([1.0e-5, 0.99999]\,;[1.0e-5, 0.99999])\,; \;\; p_5 = 0.99996 \\
\Omega_6 &= ([0.99999, 1.0]\,;[1.0e-5, 0.99999])\,; \;\; p_6 = 9.9998e-6 \\
\Omega_7 &= ([0.0, 1.0e-5]\,;[0.99999, 1.0])\,; \;\; p_7 = 1e-10 \\
\Omega_8 &= ([1.0e-5, 0.99999]\,;[0.99999, 1.0])\,; \;\; p_8 = 9.9998e-6 \\
\Omega_9 &= ([0.99999, 1.0]\,;[0.99999, 1.0])\,; \;\; p_9 = 1e-10
\end{aligned}$$

Notice that this is a *highly* unbalanced initial stratification with stratum Ω_5 encompassing 99.996% of the domain. This stratification clearly assumes that failures will occur in the tails of the distribution (note this is the only assumption made in the selection of the initial stratification). Therefore, this degree of imbalance is intentional since, under the assumption that failure occurs in the tails of the distribution, there is very little to be gained by sampling heavily outside of those tails.

The use of only 1,000 samples in this example demonstrates the efficiency of the method. The limit state function is strongly nonlinear with several regimes (albeit not disjoint) of the failure domain that must be identified. Approximate methods such as FORM and SORM will not be able to adequately resolve this failure surface and methods such as importance sampling will struggle because multiple design points need to be identified. Furthermore, the algorithm used to generate these samples is in its most basic form. It does not include logic that accounts for the considerations discussed in Section 5.3 and other logic that can improve sample placement. Improved algorithms have the potential to further reduce the number of simulations considerably (potentially as low as 200 samples or less for this example).

7. Conclusions

A new sampling method, Targeted Random Sampling, for Monte Carlo reliability estimation has been presented. The method utilizes a stratified sampling design in which samples are added (a process enabled by the Refined Stratified Sampling method) in targeted regions of the sample space that are identified as simultaneously close to the limit state and possessing a known probability of occurrence as defined by the stratum design. Progressive refinement of the strata design allows the samples to converge toward the limit state and produce an accurate estimate of the probability of failure in a small number of samples. A discussion of algorithmic features is provides and example demonstrating the capabilities for two-dimensional problem possessing a highly nonlinear failure domain is provided.

References

Olsson, A. and Sandberg, G. and Dahlblom, O. On Latin hypercube sampling for structural reliability analysis. *Structural Safety*, 25: 47–68, 2003.

Schueller, G. I. and Stix, R. A critical appraisal of methods to determine failure probabilities. *Structural Safety*, 4:293–309, 1987.

Au, S-K. and Beck, J. L. Estimation of small failure probabilities in high dimensions by subset simulation. *Probabilistic Engineering Mechanics*, 16:263–277, 2001.

Koutsourelakis, P. S. and Pradlwarter, H. J. and Schueller, G. I. Reliability of structures in high dimensions, part I: algorithms and applications. *Probabilistic Engineering Mechanics*, 19:409–417, 2004.

Shields, M. D. and Teferra, K. and Hapij, A. and Daddazio, R. P. Refined stratified sampling for efficient Monte Carlo based uncertainty quantification. (Under development).

Hasofer, A.M. and Lind, N.C. An exact and invariant first order reliability format. *ASCE Journal of the Engineering Mechanics Division*, 100: 111–121, 1974.

Faravelli, L. Response surface approach for reliability analysis. *Journal of Engineering Mechanics*, 115: 2763–2781, 1989.

Hurtado, J.E. An examination of methods for approximating implicit limit state functions from the viewpoint of statistical learning theory. *Structural Safety*, 26: 271–293, 2004.

Breitung, K. Asymptotic approximations for multinormal integrals. *Journal of Engineering Mechanics*, 110: 357–366, 1984.

McKay, M.D. and Beckman, R.J. and Conover, W.J. A comparison of three methods of selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21: 239–245, 1979.

Tocher, K.D. *The art of simulation.* The English Universities Press, 1963.