48th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference
23 - 26 April 2007, Honolulu, Hawaii

AIAA 2007-1946

Multimodal Reliability Assessment for Complex Engineering Applications using Efficient Global Optimization

B. J. Bichon[††*], M. S. Eldred[**†], L. P. Swiler[**‡], S. Mahadevan[††§], and J. M. McFarland[††¶]

[††] *Vanderbilt University, Nashville, TN 37235*

[**] *Sandia National Laboratories[‖], Albuquerque, NM 87185*

As engineering applications become increasingly complex, they are often characterized by implicit response functions that are both expensive to evaluate and nonlinear in their behavior. Reliability assessment given this type of response is difficult with available methods. Current reliability methods focus on the discovery of a single most probable point of failure, and then build a low-order approximation to the limit state at this point. This creates inaccuracies when applied to engineering applications for which the limit state has a higher degree of nonlinearity or is multimodal. Sampling methods, on the other hand, do not rely on an approximation to the shape of the limit state and are therefore generally more accurate when applied to problems with nonlinear limit states. However, sampling methods typically require a large number of response function evaluations, which can make their application infeasible for computationally expensive problems.

This paper describes the application of efficient global optimization to reliability assessment to provide a method that efficiently characterizes the limit state throughout the uncertain space. The method begins with a Gaussian process model built from a very small number of samples, and then intelligently chooses where to generate subsequent samples to ensure the model is accurate in the vicinity of the limit state. The resulting Gaussian process model is then sampled using multimodal adaptive importance sampling to calculate the probability of exceeding (or failing to exceed) the response level of interest.

By locating multiple points on or near the limit state, more complex limit states can be modeled, leading to more accurate probability integration. By concentrating the samples in the area where accuracy is important (i.e. in the vicinity of the limit state), only a small number of true function evaluations are required to build a quality surrogate model. The resulting method is both accurate for any arbitrarily shaped limit state and computationally efficient even for expensive response functions.

This new method is applied to a collection of example problems that currently available methods have difficulty solving either accurately or efficiently. The focus is on forward reliability analysis (calculating the probability of exceeding a specified response level), but some discussion of application to inverse reliability analysis (calculating the response level that corresponds to a specified probability) is included.

*Doctoral Candidate and NSF IGERT Fellow, Department of Civil and Environmental Engineering, Student Member AIAA.

†Principal Member of Technical Staff, Optimization and Uncertainty Estimation Department, Associate Fellow AIAA.

‡Principal Member of Technical Staff, Optimization and Uncertainty Estimation Department, Member AIAA.

§Professor, Department of Civil and Environmental Engineering and Mechanical Engineering, Member AIAA.

¶Graduate Student and NSF IGERT Fellow, Department of Mechanical Engineering.

‖Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under Contract DE-AC04-94AL85000.

I. Introduction

Accurate reliability assessment is a problem of great importance to the engineering community. Poor solutions lead to designs that are either unreliable or overly conservative. However, the ability to accurately quantify the uncertainty in a design becomes increasingly difficult as the analysis of the design becomes more expensive and its behavior more nonlinear.

Current methods of reliability assessment solve a local optimization problem to locate the most probable point of failure (MPP), and then quantify the reliability based on its location and an approximation to the shape of the limit state at this point. Typically, gradient-based solvers are used to solve this optimization problem, which may fail to converge for nonsmooth response functions with unreliable gradients or may converge to only one of several solutions for response functions that possess multiple local optima. In addition to these MPP convergence issues, the evaluated probabilites can be adversely affected by limit state approximations that may be inaccurate. Engineers are then forced to revert to sampling methods, which do not rely on MPP convergence or simplifying approximations to the true shape of the limit state. However, employing such methods is impractical when evaluation of the response function is expensive.

A reliability assessment method that is both efficient when applied to expensive response functions and accurate for a response function of any arbitrary shape is needed. This paper investigates the application of efficient global optimization[15] (EGO) to the search for multiple points on or near the limit state throughout the uncertain variable space. By locating multiple points on the limit state, more complex limit states can be accurately modeled, resulting in a more accurate assessment of the reliability. It should be emphasized here that these multiple points exist on a single limit state. This is significantly different (and more computationally demanding) than problems concerned with locating single MPPs on multiple limit states.

EGO was developed to facilitate the unconstrained minimization of expensive implicit response functions. The method builds an initial Gaussian process model as a global surrogate for the response function, then intelligently selects additional samples to be added for inclusion in a new Gaussian process model in subsequent iterations. The new samples are selected based on how much they are expected to improve the current best solution to the optimization problem. When this expected improvement is acceptably small, the globally optimal solution has been found. The application of this methodology to equality-constrained reliability assessment is the primary contribution of this work.

Section II describes the reliability assessment problem and traditional methods of solving it. Section III gives an overview of EGO, and outlines how it is adapted for application to reliability assessment. Section IV describes a collection of example problems and compares the performance of this EGO-based method to that of other available methods. Finally, Section V provides concluding remarks on the promise and shortcomings of this new method.

II. Reliability Assessment

The goal of reliability assessment is to determine the probability that an engineered device, component, system, etc. will fail in service given that its behavior is dependent on random inputs. This behavior is defined by a response function $g(\mathbf{x})$, where $\mathbf{x}$ represents the vector of random variables defined by known probability distributions. Failure is then defined by that response function exceeding (or failing to exceed) some threshold value $\bar{z}$. The probability of failure, p_f, is then defined by

$$p_f = \int_{g > \bar{z}} \cdots \int f_{\mathbf{x}}(\mathbf{x})\, d\mathbf{x} \tag{1}$$

where $f_{\mathbf{x}}$ is the joint probability density function of the random variables $\mathbf{x}$, and the integration is performed over the failure region where $g > \bar{z}$. In general, $f_{\mathbf{x}}$ is impossible to obtain, and even when it is available, evaluating the multiple integral is impractical.[10] Because of these complications, methods of approximating this integral are used in practice.

A. MPP Search Methods

These methods involve solving a nonlinear optimization problem to locate the point on the limit state (the contour on the response function where $g = \bar{z}$) that has the greatest probability of occurring. This point is

known as the most probable point or MPP. An approximation to the limit state is then formed at this point to facilitate the integration required to compute the probability of failure.

The MPP search is performed in uncorrelated standard normal space because it simplifies the probability integration; in this space, the distance from the origin to the MPP is equivalent to the number of input standard deviations from the mean response at which the limit state lies. This distance is known as the reliability index and is denoted by β. The transformation from correlated non-normal distributions (x-space) to uncorrelated standard normal distributions (u-space) is nonlinear in general, and possible approaches include the Rosenblatt,[19] Nataf,[5] and Box-Cox[1] transformations. The nonlinear transformations may also be linearized, and common approaches for this include the Rackwitz-Fiessler[17] two-parameter equivalent normal and the Chen-Lind[3] and Wu-Wirsching[23] three-parameter equivalent normals. This work employs the Nataf nonlinear transformation, which occurs in the following two steps. To transform between the original correlated x-space variables and correlated standard normals (z-space), the CDF matching condition is used:

$$\Phi(z_i) = F(x_i) \tag{2}$$

where $F()$ is the cumulative distribution function of the original probability distribution and $\Phi()$ is the standard normal cumulative distribution function. Then, to transform between correlated z-space variables and uncorrelated u-space variables, the Cholesky factor $\mathbf{L}$ of a modified correlation matrix is used:

$$\mathbf{z} = \mathbf{L}\mathbf{u} \tag{3}$$

where the original correlation matrix for non-normals in x-space has been modified for z-space.[5]

The forward reliability analysis algorithm for computing the probability/reliability level that corresponds to a specified response level is called the reliability index approach (RIA), and the inverse reliability analysis algorithm for computing the response level that corresponds to a specified probability/reliability level is called the performance measure approach (PMA).[21] The differences between the RIA and PMA formulations appear in the objective function and equality constraint formulations in the MPP searches. For RIA, the MPP search for achieving the specified response level $\bar{z}$ is formulated as

$$\begin{aligned} \text{minimize} \quad & \mathbf{u}^T\mathbf{u} \\ \text{subject to} \quad & G(\mathbf{u}) = \bar{z} \end{aligned} \tag{4}$$

and for PMA, the MPP search for achieving the specified probability/reliability level $\overline{p_f}, \bar{\beta}$ is formulated as

$$\begin{aligned} \text{minimize} \quad & \pm G(\mathbf{u}) \\ \text{subject to} \quad & \mathbf{u}^T\mathbf{u} = \bar{\beta}^2 \end{aligned} \tag{5}$$

where $\mathbf{u}$ is a vector centered at the origin in u-space and $G(\mathbf{u}) \equiv g(\mathbf{x})$ by definition. In the RIA case, the optimal MPP solution $\mathbf{u}^*$ defines the reliability index from $\beta = \pm\|\mathbf{u}^*\|_2$, which in turn defines the probability of failure through the probability integration. In the PMA case, the value of the response function at the optimal MPP solution $G(\mathbf{u}^*)$ defines the desired response level result.

Recent research has focused on the use of local and multipoint surrogate models to reduce the expense of the MPP search.[7,8] All of these MPP search methods employ local optimization techniques and converge to a single MPP. But the limit state of a complex engineering application may be multimodal and possess multiple significantly probable points of failure. The method for reliability assessment proposed here uses a global surrogate model and global optimization methods to reduce expense and locate multiple points along the limit state.

B. Probability Integration

For an RIA formulation, after the MPP is found and the reliability index β is known, the next step is to integrate over the failure region to calculate the probability of failure. This can be greatly simplified from Eqn. 1 by approximating the shape of the limit state with one over which it is easier to integrate. The simplest approximation is the first-order reliability method (FORM), which approximates the limit state as a linear function. Because β represents the distance from the mean response to the MPP in standard normal space, the probability integration simplifies to:

$$p_f = \Phi(-\beta) \tag{6}$$

Another alternative is the second-order reliability method (SORM), which incorporates some curvature in the limit state approximation.[2,12,13] Breitung applies a correction based on asymptotic analysis:[2]

$$p_f = \Phi(-\beta) \prod_{i=1}^{n-1} \frac{1}{\sqrt{1+\beta\kappa_i}} \tag{7}$$

where κ_i are the principal curvatures of the limit state (the eigenvalues of an orthonormal transformation of $\nabla^2_{\mathbf{u}} G$, taken positive for a convex limit state). This method essentially uses a parabolic approximation to the limit state and is more accurate for large values of β because it collapses to first-order integration at $\beta = 0$. An alternative correction in Ref. 12 is consistent with Breitung's correction in the asymptotic regime ($\beta \to \infty$) but does not approach first-order integration as $\beta \to 0$:

$$p_f = \Phi(-\beta) \prod_{i=1}^{n-1} \frac{1}{\sqrt{1+\psi(-\beta)\kappa_i}} \tag{8}$$

where $\psi() = \frac{\phi()}{\Phi()}$ and $\phi()$ is the standard normal density function. Ref. 13 applies further corrections to Eqn. 8 based on point concentration methods. Each of these methods makes a local approximation to the shape of the limit state, centered at the point of highest probability, making them inaccurate if the approximation is poor. In today's complex engineering models, the true shape of the limit state is not likely to be linear or parabolic, so more advanced methods are needed.

This nonlinearity causes a problem for the PMA formulation described in Eq. 5. For a linear approximation to the limit state, the relationship between the probability and the reliability index is easily invertible by $\bar{\beta} = -\Phi^{-1}(\overline{p_f})$, which means that constraining the MPP to be a specific distance $\bar{\beta}$ from the mean is equivalent to specifying the probability at the MPP. Ref. 7 has shown that a similar relationship can be derived for second-order approximations to the limit state. However, for more complex limit states that cannot be accurately modeled with these low-order functions, there is no $\beta \to p_f$ relationship, much less one that can be inverted to provide a constraint for the optimization. For multimodal functions, there are multiple MPP solutions with different βs and the inverse formulation of constraining to a single β loses meaning. For this reason, PMA is not an effective formulation for inverse reliability assessment of multimodal functions. Due to this lack of a proper formulation for the inverse problem, this paper focuses on direct reliability analysis.

C. Sampling Methods

An alternative to MPP search methods is to directly perform the probability integration numerically by sampling the response function. Sampling methods do not rely on a simplifying approximation to the shape of the limit state, so they can be more accurate than FORM and SORM, but they can also be prohibitively expensive because they generally require a large number of response function evaluations. Importance sampling methods reduce this expense by focusing the samples in the important regions of the uncertain space. They do this by centering the sampling density function at the MPP rather than at the mean. This ensures the samples will lie in an interesting region of the space, thus increasing the efficiency of the sampling method. Adaptive importance sampling (AIS) further improves the efficiency by adaptively updating the sampling density function. Multimodal adaptive importance sampling[6,24] is a variation of AIS that allows for the use of multiple sampling densities making it better suited for cases where multiple sections of the limit state are highly probable.

Note that importance sampling methods require the location of at least one MPP be known because it is used to center the initial sampling density. However, current gradient-based, local search methods used in MPP search may fail to converge or may converge to poor solutions, possibly making these methods inapplicable. As the next section describes, EGO is a global optimization method that does not depend on the availability of accurate gradient information, making convergence more reliable for nonsmooth response functions. Moreover, EGO has the ability to locate multiple failure points, which would provide multiple starting points and a true multimodal sampling density for the initial steps of multimodal AIS. An additional advantage of the EGO-based method proposed is that a byproduct of the search for failure points is a Gaussian process model that is accurate in the vicinity of the limit state, thereby providing an inexpensive surrogate that can be used to provide response function samples. As will be seen, using EGO to locate multiple points along the limit state, and then using the resulting Gaussian process model to provide function evaluations in multimodal AIS for the probability integration, results in an accurate and efficient reliability assessment tool.

III. Efficient Global Optimization

Efficient global optimization was originally proposed by Jones et al.[15] and has been adapted into similar methods such as sequential kriging optimization (SKO).[14] The main difference between SKO and EGO lies within the specific formulation of what is known as the expected improvement function (EIF), which is the feature that sets all EGO/SKO-type methods apart from other global optimization methods. The EIF is used to select the location at which a new training point should be added to the Gaussian process model by maximizing the amount of improvement in the objective function that can be expected by adding that point. A point could be expected to produce an improvement in the objective function if its predicted value is better than the current best solution, or if the uncertainty in its prediction is such that the probability of it producing a better solution is high. Because the uncertainty is higher in regions of the design space with few observations, this provides a balance between exploiting areas of the design space that predict good solutions, and exploring areas where more information is needed. The general procedure of these EGO-type methods is:

1. Build an initial Gaussian process model of the objective function
2. Find the point that maximizes the EIF. If the EIF value at this point is sufficiently small, stop.
3. Evaluate the objective function at the point where the EIF is maximized. Update the Gaussian process model using this new point. Go to Step 2.

To construct a parallel algorithm, the n best points could be selected and evaluated in steps 2 and 3. The following sections discuss the construction of the Gaussian process model used, the form of the EIF, and then some ideas on how that EIF could be altered for application to MPP search.

A. Gaussian Process Model

Gaussian process (GP) models are set apart from other surrogate models because they provide not just a predicted value at an unsampled point, but a full Gaussian distribution with an expected value and a predicted variance. This variance gives an indication of the uncertainty in the model, which results from the construction of the covariance function. This function is based on the idea that when input points are near one another, the correlation between their corresponding outputs will be high. As a result, the uncertainty associated with the model's predictions will be small for input points which are near the points used to train the model, and will increase as one moves further from the training points.

It is assumed that the true response function being modeled $G(\mathbf{u})$ can be described by:[4]

$$G(\mathbf{u}) = \mathbf{h}(\mathbf{u})^T \boldsymbol{\beta} + Z(\mathbf{u}) \tag{9}$$

where $\mathbf{h}()$ is the trend of the model, $\boldsymbol{\beta}$ is the vector of trend coefficients, and $Z()$ is a stationary Gaussian process with zero mean (and covariance defined below) that describes the departure of the model from its underlying trend. The trend of the model can be assumed to be any function, but taking it to be a constant value is generally sufficient.[20] The covariance between outputs of the Gaussian process $Z()$ at points $\mathbf{a}$ and $\mathbf{b}$ is defined as:

$$Cov\left[Z(\mathbf{a}), Z(\mathbf{b})\right] = \sigma_Z^2 R(\mathbf{a}, \mathbf{b}) \tag{10}$$

where σ_Z^2 is the process variance and $R()$ is the correlation function. There are several options for the correlation function, but the squared-exponential function is common, and is used here for $R()$:

$$R(\mathbf{a}, \mathbf{b}) = \exp\left[-\sum_{i=1}^{d} \theta_i (a_i - b_i)^2\right] \tag{11}$$

where d represents the dimensionality of the problem, and θ_i is a scale parameter that indicates the correlation between the points within dimension i. A large θ_i indicates a low correlation.

The expected value $\mu_G()$ and variance $\sigma_G^2()$ of the GP model prediction at point $\mathbf{u}$ are:

$$\mu_G(\mathbf{u}) = \mathbf{h}(\mathbf{u})^T \boldsymbol{\beta} + \mathbf{r}(\mathbf{u})^T \mathbf{R}^{-1} (\mathbf{g} - \mathbf{F}\boldsymbol{\beta}) \tag{12}$$

$$\sigma_G^2(\mathbf{u}) = \sigma_Z^2 - \begin{bmatrix} \mathbf{h}(\mathbf{u})^T & \mathbf{r}(\mathbf{u})^T \end{bmatrix} \begin{bmatrix} \mathbf{0} & \mathbf{F}^T \\ \mathbf{F} & \mathbf{R} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{h}(\mathbf{u}) \\ \mathbf{r}(\mathbf{u}) \end{bmatrix} \tag{13}$$

where $\mathbf{r}(\mathbf{u})$ is a vector containing the covariance between $\mathbf{u}$ and each of the n training points (defined by Eq. 10), $\mathbf{R}$ is an $n \times n$ matrix containing the correlation between each pair of training points, $\mathbf{g}$ is the vector of response outputs at each of the training points, and $\mathbf{F}$ is an $n \times q$ matrix with rows $\mathbf{h}(\mathbf{u}_i)^T$ (the trend function for training point i containing q terms; for a constant trend $q = 1$). This form of the variance accounts for the uncertainty in the trend coefficients $\boldsymbol{\beta}$, but assumes that the parameters governing the covariance function (σ_Z^2 and $\boldsymbol{\theta}$) have known values.

The parameters β, σ_Z^2 and $\boldsymbol{\theta}$ are determined through maximum likelihood estimation. This involves taking the log of the probability of observing the response values $\mathbf{g}$ given the covariance matrix $\mathbf{R}$, which can be written as:[20]

$$\log\left[p(\mathbf{g}|\mathbf{R})\right] = -\frac{1}{n}\log|\mathbf{R}| - \log(\hat{\sigma}_Z^2) \tag{14}$$

where $|\mathbf{R}|$ indicates the determinant of $\mathbf{R}$, and $\hat{\sigma}_Z^2$ is the optimal value of the variance given $\boldsymbol{\theta}$ and is defined by:

$$\hat{\sigma}_Z^2 = \frac{1}{n}(\mathbf{g} - \mathbf{F}\hat{\boldsymbol{\beta}})^T\mathbf{R}^{-1}(\mathbf{g} - \mathbf{F}\hat{\boldsymbol{\beta}}) \tag{15}$$

where $\hat{\boldsymbol{\beta}}$ is the generalized least squares estimate of $\boldsymbol{\beta}$ from:

$$\hat{\boldsymbol{\beta}} = \left[\mathbf{F}^T\mathbf{R}^{-1}\mathbf{F}\right]^{-1}\mathbf{F}^T\mathbf{R}^{-1}\mathbf{g} \tag{16}$$

Maximizing Eq. 14 gives the maximum likelihood estimate of $\boldsymbol{\theta}$.

B. Expected Improvement Function

The expected improvement function is used to select the location at which a new training point should be added. It does this by calculating the expectation that any point in the search space will provide a better solution than the current best solution based on the expected values and variances predicted by the GP model. An important feature of the EIF is that it provides a balance between exploiting areas of the design space where good solutions have been found, and exploring areas of the design space where the uncertainty is high. First, recognize that at any point in the design space, the GP prediction $\hat{G}()$ is a Gaussian distribution:

$$\hat{G}(\mathbf{u}) \sim N\left[\mu_G(\mathbf{u}), \sigma_G(\mathbf{u})\right] \tag{17}$$

where the mean $\mu_G()$ and the variance $\sigma_G^2()$ were defined in Eqs. 12 and 13, respectively. The EIF is defined as:[14,15]

$$EI\big(\hat{G}(\mathbf{u})\big) \equiv E\left[\max\Big(G(\mathbf{u}^*) - \hat{G}(\mathbf{u}), 0\Big)\right] \tag{18}$$

where $G(\mathbf{u}^*)$ is the current best solution chosen from among the true function values at the training points (henceforth referred to as simply G^*). This expectation can then be computed by integrating over the distribution $\hat{G}(\mathbf{u})$ with G^* held constant:

$$EI\big(\hat{G}(\mathbf{u})\big) = \int_{-\infty}^{G^*} (G^* - G)\ \hat{G}(\mathbf{u})\ dG \tag{19}$$

where G is a realization of $\hat{G}$. This integral can be expressed analytically as:[14,15]

$$EI\big(\hat{G}(\mathbf{u})\big) = (G^* - \mu_G)\ \Phi\left(\frac{G^* - \mu_G}{\sigma_G}\right) + \sigma_G\,\phi\left(\frac{G^* - \mu_G}{\sigma_G}\right) \tag{20}$$

where it is understood that μ_G and σ_G are functions of $\mathbf{u}$.

The point at which the EIF is maximized is selected as an additional training point. With the new training point added, a new GP model is built and then used to construct another EIF, which is then used to choose another new training point, and so on, until the value of the EIF at its maximized point is below some specified tolerance. In Ref. 14 this maximization is performed using a Nelder-Mead simplex approach, which is a local optimization method. Because the EIF is often highly multimodal[15] (particularly in the early stages of the process) it is expected that Nelder-Mead may fail to converge to the true global optimum. In Ref. 15 a branch-and-bound technique for maximizing the EIF is used, but was found to often be too expensive to run to convergence.

It is important to understand how the use of this EIF leads to optimal solutions. Eq. 20 indicates how much the objective function value at $\mathbf{x}$ is expected to be less than the predicted value at the current best solution. It contains a balance between exploiting regions of the design space where good solutions have been discovered, and exploring regions that have not been well sampled and thus have greater uncertainty. Because the GP model provides a Gaussian distribution at each predicted point, expectations can be calculated. Points with good expected values and even a small variance will have a significant expectation of producing a better solution (exploitation), but so will points that have relatively poor expected values and greater variance (exploration).

The application of EGO to reliability assessment, however, is made more complicated due to its inclusion of equality constraints (see Eqs. 4-5). For inverse reliability analysis, this extra complication is small. The response being modeled by the GP is the objective function of the optimization problem (see Eq. 5) and the deterministic constraint might be handled through the use of a merit function, thereby allowing EGO to solve this equality-constrained optimization problem. Here the problem lies in the interpretation of the constraint for multimodal problems as mentioned previously. In the forward reliability case, the response function appears in the constraint rather than the objective. Here, the maximization of the EIF is inappropriate because feasibility is the main concern. This application is therefore a significant departure from the intentions of EGO and requires a new formulation. For this problem, the expected feasibility function is introduced.

C. Expected Feasibility Function

The expected improvement function provides an indication of how much the true value of the response at a point can be expected to be less than the current best solution. It therefore makes little sense to apply this to the forward reliability problem where the goal is not to minimize the response, but rather to find where it is equal to a specified threshold value. The expected feasibility function (EFF) is introduced here to provide an indication of how well the true value of the response is expected to satisfy the equality constraint $G(\mathbf{u}) = \bar{z}$. Inspired by the contour estimation work in Ref. 18, this expectation can be calculated in a similar fashion as Eq. 19 by integrating over a region in the immediate vicinity of the threshold value $\bar{z} \pm \epsilon$:

$$EF\big(\hat{G}(\mathbf{u})\big) = \int_{z-\epsilon}^{z+\epsilon} \big[\epsilon - |\bar{z} - G|\big]\, \hat{G}(\mathbf{u})\, dG \tag{21}$$

where G denotes a realization of the distribution $\hat{G}$, as before. Allowing z^+ and z^- to denote $\bar{z} \pm \epsilon$, respectively, this integral can be expressed analytically as:

$$\begin{aligned} EF\big(\hat{G}(\mathbf{u})\big) = (\mu_G - \bar{z}) &\left[2\,\Phi\left(\frac{\bar{z} - \mu_G}{\sigma_G}\right) - \Phi\left(\frac{z^- - \mu_G}{\sigma_G}\right) - \Phi\left(\frac{z^+ - \mu_G}{\sigma_G}\right)\right] \\ - \sigma_G &\left[2\,\phi\left(\frac{\bar{z} - \mu_G}{\sigma_G}\right) - \phi\left(\frac{z^- - \mu_G}{\sigma_G}\right) - \phi\left(\frac{z^+ - \mu_G}{\sigma_G}\right)\right] \\ + \quad \epsilon &\left[\Phi\left(\frac{z^+ - \mu_G}{\sigma_G}\right) - \Phi\left(\frac{z^- - \mu_G}{\sigma_G}\right)\right] \end{aligned} \tag{22}$$

where ϵ is a constant proportional to the standard deviation of the GP predictor ($\epsilon \propto \sigma_G$). In this case, z^-, z^+, μ_G, σ_G, and ϵ are all functions of the location $\mathbf{u}$, while $\bar{z}$ is a constant. Note that the EFF provides the same balance between exploration and exploitation as is captured in the EIF. Points where the expected value is close to the threshold ($\mu_G \approx \bar{z}$) and points with a large uncertainty in the prediction will have large expected feasibility values.

IV. Computational Experiments

This new reliability assessment method has been applied to a collection of example problems to compare it to other available methods. The following process was used:

1. Generate a small number of samples from the true response function.
   (a) Only $\frac{(n+1)(n+2)}{2}$ samples are used (n is the number of uncertain variables). This initial selection is arbitrary, but the number of samples required to define a quadratic polynomial is used as a convenient rule of thumb.

(b) The samples uniformly span the uncertain space over the bounds $\pm 5\sigma$.

(c) Latin hypercube sampling (LHS) is used to generate the samples.

2. Construct an initial Gaussian process model from these samples.

3. Find the point with maximum expected feasibility.

(a) The expected feasibility function is built with $\epsilon = 2\sigma_G$.

(b) To ensure the global optimum of this multimodal function is found, either a genetic algorithm or DIRECT is used.

(c) If the maximum expected feasibility is acceptably small, go to step 6.

4. Evaluate the true response function at this point.

5. Add this new sample to the previous set and build a new GP model. Go to step 3.

6. When the maximum expected feasibility is small, the Gaussian process model is accurate in the vicinity of the limit state. This surrogate is then used to calculate the probability using multimodal adaptive importance sampling.

Computational results for all methods are from either published results or new analyses performed using the DAKOTA/UQ[22] software with solvers provided by the Coliny[11] and OPT++[16] packages. DAKOTA/UQ is the uncertainty quantification component of DAKOTA,[9] an open-source software framework for design and performance analysis of computational models on high performance computers developed at Sandia National Laboratories.

A. Multimodal Example

The first problem has a highly nonlinear response defined by:

$$g(\mathbf{x}) = \frac{\left(x_1^2 + 4\right)(x_2 - 1)}{20} - \sin\frac{5x_1}{2} - 2 \tag{23}$$

The distribution of x_1 is Normal(1.5, 1) and x_2 is Normal(2.5, 1); the variables are uncorrelated. The response level of interest for this study is $\bar{z} = 0$ with failure defined by $g > \bar{z}$. Figure 1 shows a plot in x-space of the limit state throughout the ± 5 standard deviation search space. This problem has several local optima to the forward-reliability MPP search problem (see Eq. 4) and the two most significant MPPs are shown on the plot.

This problem was solved using reliability methods available in DAKOTA/UQ that reduce the cost of the MPP search through the use of local surrogate models. Two response function approximation methods were investigated:[7] second-order iterated Advanced Mean Value (AMV^2+) and Two-point Adaptive Nonlinear Approximation (TANA).

A case using no response function approximation was also investigated. To produce results consistent with an implicit response function, numerical gradients and quasi-Newton Hessians from Symmetric Rank 1 updates were used. For each method, at the converged MPP, both first-order and second-order integration (using Eqs. 6 and 8) were used to calculate the probability.

For each method, the only algorithmic variation explored here is that the surrogate models can be built in either x-space or u-space. Determining which space is appropriate depends upon the form of the response and the space transformation. This choice can have significant effects on both accuracy and efficiency for methods that use low-order approximations to the response function. For instance, if a linear approximation is used for a response that is linear in x-space but nonlinear in u-space, then building the approximation in x-space will yield better results. Gaussian process models are not greatly affected by this choice because they do not rely on curve-fitting or any assumptions on the shape of the response. However, they can smooth out finer details in the true function, meaning they will generally provide better approximations with fewer samples if the response is well behaved. If it is known that the response function is smoother in one space than the other, the user should take advantage of that information and build the GP in that space.

Table 1 gives a summary of the results from all methods. To establish an accurate estimate of the true solution, 20 independent studies were performed using one million Latin hypercube samples per study. The

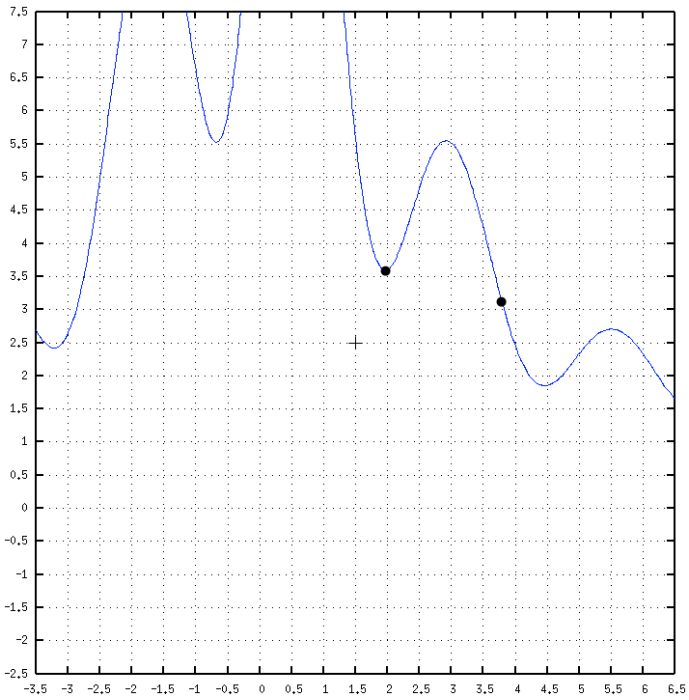

Figure 1. Limit state of the multimodal test problem.

average probability from these studies is reported as the "true" solution. Because the EGO-based method is stochastic, it was also run 20 times and the average probability are reported. To measure the accuracy of the methods, two errors are reported for the EGO results: the error in the average probability, and the average of the absolute errors from the 20 studies. For comparison, the same errors are given for the 20 LHS studies.

Table 1. Results for the multimodal test problem.

Reliability Method	Function Evaluations	First-Order p_f (% Error)	Second-Order p_f (% Error)	Sampling p_f (% Error, Avg. Error)
No Approximation	66	0.11798 (276.3%)	0.02516 (-19.7%)	—
x-space AMV^2+	26	0.11798 (276.3%)	0.02516 (-19.7%)	—
u-space AMV^2+	26	0.11798 (276.3%)	0.02516 (-19.7%)	—
x-space TANA	506	0.08642 (175.7%)	0.08716 (178.0%)	—
u-space TANA	131	0.11798 (276.3%)	0.02516 (-19.7%)	—
x-space EGO	50.4	—	—	0.03127 (0.233%, 0.929%)
u-space EGO	49.4	—	—	0.03136 (0.033%, 0.787%)
True LHS solution	1M	—	—	0.03135 (0.000%, 0.328%)

Most of the MPP search methods converge to the same MPP and thus report the same probability. These probabilities are more accurate when second-order integration is used, but still have significant errors. However, x-space TANA converges to the second MPP shown in Fig. 1, which lies in a relatively flat region of the limit state. This local lack of curvature means that first-order and second-order integration produce approximately the same probability. In isolation, this second-order result could be viewed as a verification of the first-order probability and thus provide a misguided confidence in the reliability analysis. For this multimodal problem, the new EGO-based method is more expensive than AMV^2+, but cheaper than all the other methods, and provides much more accurate results.

B. Nonlinear Example

The second example is a two-dimensional nonlinear function taken from the literature.

$$g(\mathbf{x}) = x_1^3 + x_2^3 - 18 \tag{24}$$

The distribution of x_1 is Normal(10, 5) and x_2 is Normal(9.9, 5); the variables are uncorrelated. The response level of interest for this study is $\bar{z} = 0$ with failure defined by $g < \bar{z}$. This problem was introduced by Zou et al.[24] to test a method that used a trust-region managed surrogate model to locate the MPP and then multimodal adaptive importance sampling (MAIS) to perform the probability integration.

Table 2 gives a summary of the results from the same methods investigated in the previous example plus the published results from Ref. 24. To establish an accurate estimate of the true solution, 20 independent studies were performed using one million Latin hypercube samples per study. The average probability from these studies is reported as the "true" solution. Again, two errors are reported for the EGO-based methods and LHS: the error in the average probability, and the average of the absolute errors from the 20 studies.

Table 2. Results for the nonlinear test problem.

Reliability Method	Function Evaluations	First-Order p_f (% Error)	Second-Order p_f (% Error)	Sampling p_f (% Error, Avg. Error)
No Approximation	125	0.01301 (128.2%)	0.004164 (-26.9%)	—
x-space AMV^2+	66	0.01301 (128.2%)	0.004165 (-26.9%)	—
u-space AMV^2+	66	0.01301 (128.2%)	0.004165 (-26.9%)	—
x-space TANA	21	0.01301 (128.2%)	0.004155 (-27.1%)	—
u-space TANA	36	0.01301 (128.2%)	0.004165 (-26.9%)	—
Zou et al. MAIS	560	—	—	0.005750 (0.883%, no data)
x-space EGO	40.6	—	—	0.005750 (0.837%, 2.740%)
u-space EGO	43.1	—	—	0.005652 (0.876%, 3.523%)
True LHS solution	1M	—	—	0.005700 (0.000%, 0.930%)

This problem only has one significant MPP, so all of the local search methods converge to approximately the same point (x-space TANA preconverges to a slightly sub-optimal MPP). Once again, second-order integration provides better results, but is still not an accurate approximation to the true shape of the limit state, so there are still large errors. Because this test problem is not multimodal, performing MAIS with only the MPP as a starting point is sufficient to capture the higher level of nonlinearity in the limit state and generate an excellent result. If this method had been applied to the previous test problem, it would have either been much less accurate or would require a substantial increase in cost in order to adequately locate and sample the other significantly probable regions of the space. It should also be pointed out that despite MAIS being a stochastic method, only the error for a single result is reported by Zou et al. and not an average absolute error as is included for the other sampling methods. For this nonlinear problem, the new EGO-based method is more expensive than TANA, but cheaper than all the other methods, and provides much more accurate results.

V. Conclusions

As engineering applications become increasingly complex, they are often characterized by implicit response functions that are both expensive to evaluate and nonlinear in their behavior. Current reliability methods are not adequate for these problems because they require engineers to sacrifice either accuracy or efficiency. This paper has presented a new method based on efficient global optimization that is aimed at creating a method that is both accurate for any arbitrarily shaped limit state and efficient even for computationally expensive response functions.

The example problems used to test this new method have shown that it can produce results that are far more accurate than MPP search methods, while requiring far fewer true function evaluations than sampling methods. However, the authors realize that these problems are too simple to make any definitive judgment on the performance of the new method because both problems involve only two uncorrelated normal variables.

Investigation is ongoing to determine if this EGO-based method will scale to problems with larger dimensions and non-normal distributions.

An existing limitation to this new reliability assessment method is that it has no knowledge that points far from the mean are less probable than ones near the mean. Because of this limitation, the method may waste function evaluations in an attempt to improve the accuracty of the GP model in a region of the space where additional accuracy is not important due to its low probability. Current research is investigating a way to bias the maximization of the expected feasibility function towards regions of the space that are more probable. Other additional work includes developing the parallel method alluded to at the beginning of Section III and investigating how this new reliability assessment method might be utilized in reliability-based design optimization.

VI. Acknowledgments

The authors would like to express their thanks to the Sandia Computer Science Research Institute (CSRI) for support of this collaborative work between Sandia National Laboratories and Vanderbilt University.

References

[1]Box, G.E.P. and Cox, D.R., An Analysis of Transformations, *J. Royal Stat. Soc.*, Series B, Vol. 26, 1964, pp. 211-252.

[2]Breitung, K., Asymptotic Approximation for Multinormal Integrals, *J. Eng. Mech., ASCE*, Vol. 110, No. 3, 1984, pp. 357-366.

[3]Chen, X. and Lind, N.C., Fast Probability Integration by Three-Parameter Normal Tail Approximation, *Struct. Saf.*, Vol. 1, 1983, pp. 269-276.

[4]Cressie, N.A.C., *Statistics for Spatial Data*, revised edition, 1993 (Wiley: New York).

[5]Der Kiureghian, A. and Liu, P.L., Structural Reliability Under Incomplete Probability Information, *J. Eng. Mech., ASCE*, Vol. 112, No. 1, 1986, pp. 85-104.

[6]Dey, A. and Mahadevan, S., Ductile Structural System Reliability Analysis using Adaptive Importance Sampling, *Struct. Saf.*, Vol. 20, 1998, pp. 137-154.

[7]Eldred, M.S. and Bichon, B.J., New Second-Order Formulations for Reliability Analysis and Design, *AIAA J.*, in preparation.

[8]Eldred, M.S., Agarwal, H., Perez, V.M., Wojtkiewicz, S.F., Jr., and Renaud, J.E., Investigation of Reliability Method Formulations in DAKOTA/UQ, (to appear) *Structure & Infrastructure Engineering: Maintenance, Management, Life-Cycle Design & Performance*, Taylor & Francis Group.

[9]Eldred, M.S., Giunta, A.A., Brown, S.L., Adams, B.M., Dunlavy, D.M., Eddy, J.P., Gay, D.M., Griffin, J.D., Hart, W.E., Hough, P.D., Kolda, T.G., Martinez-Canales, M.L., Swiler, L.P., Watson, J.-P., and Williams, P.J., DAKOTA, A Multilevel Parallel Object-Oriented Framework for Design Optimization, Parameter Estimation, Uncertainty Quantification, and Sensitivity Analysis. Version 4.0 User's Manual. Sandia Technical Report SAND2006-6337, Revised October 2006, Sandia National Laboratories, Albuquerque, NM.

[10]Haldar, A. and Mahadevan, S., *Probability, Reliability, and Statistical Methods in Engineering Design*, 2000 (Wiley: New York).

[11]Hart, W.E., Coliny Users Manual: Version 2.0, Technical Report SAND2006-xxxx, Sandia National Laboratories, Albuquerque, NM, 2006.

[12]Hohenbichler, M. and Rackwitz, R., Improvement of Second-Order Reliability Estimates by Importance Sampling, *J. Eng. Mech., ASCE*, Vol. 114, No. 12, 1988, pp. 2195-2199.

[13]Hong, H.P., Simple Approximations for Improving Second-Order Reliability Estimates, *J. Eng. Mech., ASCE*, Vol. 125, No. 5, 1999, pp. 592-595.

[14]Huang, D., Allen, T.T., Notz, W.I., and Zeng, N., Global Optimization of Stochastic Black-Box Systems via Sequential Kriging Meta-Models, *J. Global Opt.*, Vol. 34, 2006, pp. 441-466.

[15]Jones, D., Shonlau, M., and Welch, W., Efficient Global Optimization of Expensive Black-Box Functions, *INFORMS J. Comp.*, Vol. 12, 1998, pp. 272.283

[16]Meza, J.C., OPT++: An Object-Oriented Class Library for Nonlinear Optimization, Technical Report SAND94-8225, Sandia National Laboratories, Livermore, CA, 1994.

[17]Rackwitz, R. and Fiessler, B., Structural Reliability under Combined Random Load Sequences, *Comput. Struct.*, Vol. 9, 1978, pp. 489-494.

[18]Ranjan, P., Bingham, D., and Michailidis, G., Sequential Experiment Design for Contour Estimation from Complex Computer Codes, *Technometrics*, to appear, 2007.

[19]Rosenblatt, M., Remarks on a Multivariate Transformation, *Ann. Math. Stat.*, Vol. 23, No. 3, 1952, pp. 470-472.

[20]Sacks, J., Schiller, S.B., and Welch, W., Design for Computer Experiments, *Technometrics*, Vol. 31, 1989, pp. 41-47.

[21]Tu, J., Choi, K.K., and Park, Y.H., A New Study on Reliability-Based Design Optimization, *J. Mech. Design*, Vol. 121, 1999, pp. 557-564.

[22]Wojtkiewicz, S.F., Jr., Eldred, M.S., Field, R.V., Jr., Urbina, A., and Red-Horse, J.R., A Toolkit For Uncertainty Quantification In Large Computational Engineering Models, paper AIAA-2001-1455 in *Proceedings of the 42nd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, Seattle, WA, April 16-19, 2001.

[23]Wu, Y.-T. and Wirsching, P.H., A New Algorithm for Structural Reliability Estimation, *J. Eng. Mech., ASCE*, Vol. 113, 1987, pp. 1319-1336.

[24]Zou, T., Mourelatos, Z., Mahadevan, S., and Meernik, P., Reliability Analysis of Automotive Body-Door Subsystem, *Rel. Eng. and Sys. Saf.*, Vol. 78, 2002, pp. 315-324.