SANDIA REPORT
SAND2006-5319
Unlimited Release
Printed November 2006

Development of a New Adaptive Ordinal Approach to Continuous-Variable Probabilistic Optimization

Vicente J. Romero and Chun-Hung Chen

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

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

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd.
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online order: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2006-5319
Unlimited Release
Printed November 2006

Development of a New Adaptive Ordinal Approach to Continuous-Variable Probabilistic Optimization

Vicente J. Romero
Model Validation and Uncertainty Quantification Department
Sandia National Laboratories
Albuquerque, NM 87185-0828

Chun-Hung Chen
Department of Systems Engineering and Operations Research
George Mason University
Fairfax, VA

Abstract

A very general and robust approach to solving continuous-variable optimization problems involving uncertainty in the objective function is through the use of ordinal optimization. At each step in the optimization problem, improvement is based only on a *relative ranking* of the uncertainty effects on local design alternatives, rather than on precise quantification of the effects. One simply asks "Is that alternative better or worse than this one?" –not "HOW MUCH better or worse is that alternative to this one?" The answer to the latter question requires precise characterization of the uncertainty—with the corresponding sampling/integration expense for precise resolution. However, in this report we demonstrate correct decision-making in a continuous-variable probabilistic optimization problem despite extreme vagueness in the statistical characterization of the design options. We present a new adaptive ordinal method for probabilistic optimization in which the trade-off between computational expense and vagueness in the uncertainty characterization can be conveniently managed in various phases of the optimization problem to make cost-effective stepping decisions in the design space. Spatial correlation of uncertainty in the continuous-variable design space is exploited to dramatically increase method efficiency. Under many circumstances the method appears to have favorable robustness and cost-scaling properties relative to other probabilistic optimization methods, and uniquely has mechanisms for quantifying and controlling error likelihood in design-space stepping decisions. The method is asymptotically convergent to the true probabilistic optimum, so could be useful as a reference standard against which the efficiency and robustness of other methods can be compared—analogous to the role that Monte Carlo simulation plays in uncertainty propagation.

Intentionally Left Blank

Contents

Figures

1 Introduction

A very general and robust approach to solving optimization problems is via *ordinal optimization*, where improvement at each step in the optimization problem is based only on a *relative ranking* of response, rather than on precise quantification of response. It is much less difficult and expensive to determine whether one response is greater/less than another, than to accurately quantify their response values ("...think of holding two objects to determine which one weighs more, versus trying to estimate their actual weights [by holding them]..." [10])—much more complex and accurate machinery is required to get accurate weights. Correspondingly, the ordinal approach appears to be naturally suitable for optimization problems involving nonprobabilistic and semiquantitative descriptions of uncertainty. In fact, it can be seen as a formalism of the ordinal selection process employed by decision makers in the real world when they are deliberating over competing options whose outcomes are highly uncertain.

For continuously variable design options (continuous-variable design problems), non-gradient-based local optimizers such as simplex and pattern-search methods [9,12], and global methods such as genetic algorithms and DIRECT ([13]), can be used to search the design space, using ordinal comparison to select improvement steps. This paper considers continuous-variable (C-V) problems of optimization under uncertainty (OUU) where ordinal comparisons are made of alternatives in the design space that have probabilistic information regarding their uncertainty.[1] Fundamental concepts of probabilistic ordinal optimization are presented. In particular, efficiencies gained from exploiting spatial correlation of uncertainty and sampling in the design space are demonstrated on a low-dimensional probabilistic optimization problem. More-sophisticated implementational possibilities are discussed, along with merits compared to nonordinal approaches to probabilistic optimization.

[1] C-V OUU problems are principally different from *stochastic optimization problems* where random stochastic noise in the objective function is a primary feature of the deterministic or OUU optimization problem. Specialized optimization techniques exist for this type of random stochastic uncertainty in the objective function (see, e.g., [22]). In C-V OUU problems the uncertainty in the objective function comes principally from fundamental uncertainty about the response or behavior of the system being optimized.

2 Engineering Optimization Problem Motivating Probabilistic Ordinal Concepts

2.1 Deterministic Optimization Problem

The probabilistic optimization problem considered here is an outgrowth of a deterministic optimization problem [21] in which heating conditions were sought that put a candidate weapon subsystem design most at risk. As shown in Figure 1, the design problem is parameterized in terms of two key heating variables: 1) the radius r of a circular region of impinging fire on the top of the safing subsystem; and 2) a coordinate x that moves the center of the "blow torch" on the device's surface along the cut-plane shown. A three-dimensional (3-D) finite-element conduction/radiation thermal model (described in Ref. [21]) is used to calculate the transient temperature response of the device. Figure 1 shows the calculated temperature of the device (cut along its plane of symmetry) at some point in time for a particular radius r and location x of heating.

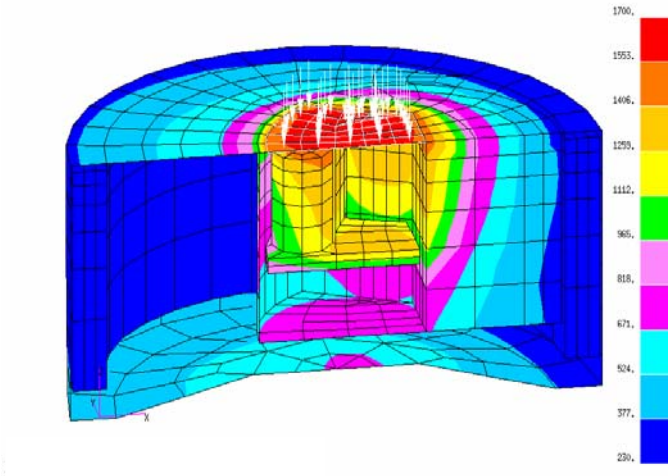

FIGURE 1. Safing device exposed to circular region of heating.

Figure 2 shows temperature histories for two safety-critical components within the safing device: a "strong link" and a "weak link." The components are termed as such because the weak link must be thermally weaker and fail at a lower temperature than the thermally stronger strong link in order to prevent inadvertent operation of the weapon.

For now focusing just on the depicted *nominal* failure temperatures of the components (nominal T_fail_{WL} = 250ºC and nominal T_fail_{SL} = 600ºC), an associated "safety margin" of time can be defined:

$$S = t\_fail_{SL} - t\_fail_{WL}, \tag{1}$$

where the subscripts SL and WL respectively denote strong link and weak link. Here t_fail_{SL} and t_fail_{WL} are the elapsed times (from time zero at the beginning of the thermal simulation) required for the links to reach their nominal failure temperatures. In Figure 2, the safety margin associated with the nominal failure temperatures of the components is positive and has a value of about S = +9 minutes (right-pointing arrow in the figure). A negative or zero safety margin indicates that the strong link is not failing after the weak link as desired, so the safing device is experiencing a

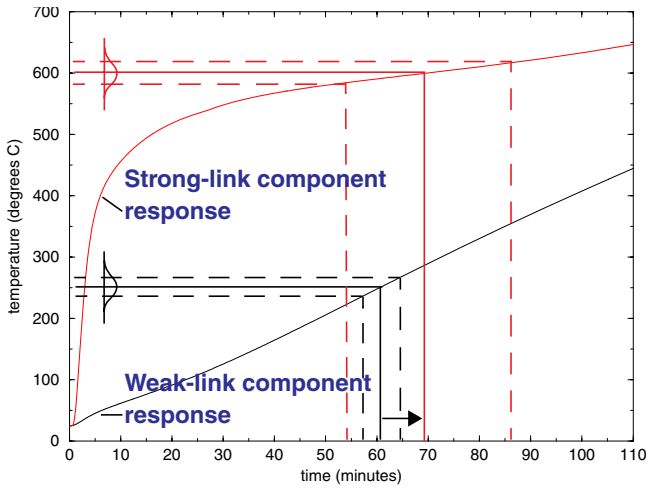

FIGURE 2. A mean realization of each component's failure temperature (solid line) reflects through component temperature response curve into a failure time. The difference in component failure times, $t\_fail_{SL} - t\_fail_{WL}$, is the corresponding *safety margin*, shown by right-pointing arrow for the positive margin here.

vulnerability for that set of heating conditions and component failure temperatures. The safety margin is a deterministic indicator of the safety of the device. A probabilistic indicator of the safety of the device will later be defined that takes into account the depicted uncertainty in the component failure temperatures.

In the optimization problem, the intent is to find the heating conditions that minimize the indicated safety of the device. This minimum corresponds to the heating conditions that most threaten the intended function of the safing device. If the vulnerability is deemed unacceptable, the design can be modified to sufficiently harden the device to these worst-case heating conditions.

In the *deterministic* optimization problem, the optimal values of the variables r and x are sought that minimize the safety margin of the device. As the $\{r, x\}$ variables are changed, the strong-link and weak-link temperature responses change, and the value of the safety margin changes accordingly (assuming fixed nominal failure temperatures for the components). Thus, the deterministic safety-margin objective function S is navigated (minimized) in the deterministic optimization problem.

The deterministic optimization problem is complicated by numerical noise resulting from discrete time and space representation in the model (see Refs. [5] and [21]). On a more global scale, the deterministic optimization problem contains mathematical difficulties associated with navigating to a global minimum at $\{r, x\} = \{1.62, 0.782\}$ on a nonconvex design surface having a fold and several local minima as described in Refs. [5] and [21].

Toward solving the deterministic local optimization problem, Newton-based nonlinear programming (NLP) methods (sequential quadratic programming, BFGS quasi-Newton, and quasi-Newton) from several different research and commercial optimization packages were tried with little success as explained in Ref. [5]. Nonconvexity and noise resulted respectively in ill-conditioning and inaccuracy of the Hessian matrix of finite-differenced second-order derivatives.

To avoid these problems with the second-order information, first-order NLP conjugate-gradient methods were employed in Ref. [5] with more success, but still exhibited some noise-induced nonrobustness. The choice of finite-difference step size for computation of sufficiently accurate gradients proved to be important. The initial starting point also proved to be important because of the noise.

The relatively high degree of model-discretization resolution required for sufficiently smooth computed response for reliable gradient-based navigation of the design space added large cost compared to the more noise-tolerant derivative-free optimization approaches applied in Ref. [6]. There, a simple Coordinate Pattern Search (CPS) method was shown to be successful with lower model resolution/cost requirements than the congugate-gradiant method. The CPS method was both more robust and less costly than the conjugate-gradient method.

More recently, a global-to-local optimizer devised with moving low-order polynomial local response surfaces to be efficient in the presence of small-scale noise was applied to this problem in Ref. [15]. The method performed considerably better than all those previously tried in terms of number of function evaluations required, cost of function evaluations (model discretization level required for accuracy), and robustness to noise level. The method was later recognized to be a type of surrogate-based *trust-region* optimization approach. Successful approaches in this vein had already been introduced elsewhere (e.g., Refs. [1] and [14]). The general advantages of surrogate-based trust-region approaches for handling noisy optimization problems are now well recognized.

2.2 Probabilistic Optimization Problem

The associated *probabilistic* optimization problem is now considered. A glance at Figure 2 indicates that strong-link and weak-link failure temperatures trace out to relatively flat portions of the temperature response curves. This is more true for the strong link than for the weak link, but regardless, it can be seen that the component failure times are very sensitive to the failure temperature values. Thus, moderate uncertainties in the component failure temperatures T_fail_{WL} and T_fail_{SL} can have a substantial impact on the uncertainty of the safety margin. In fact, when reasonable uncertainty bands 5% above and below the nominal failure temperatures are considered in Figure 2, the corresponding bands in failure times indicate that the safety margin could vary from about 28 minutes at best to about –11 minutes at worst.Thus, it is apparent that the effects of uncertainty in the component failure thresholds are very important in this problem.

2.2.1 Safety-Margin Distribution Due to Uncertain Component Failure Thresholds

For the purposes here, the strong-link and weak-link failure temperatures T_fail_{WL} and T_fail_{SL} are each assumed to be described by truncated normal distributions with means μ equal to the respective nominal failure temperatures of 600ºC and 250ºC, and standard deviations σ equal to 3% of the means, or 18ºC and 7.5ºC respectively. (The distributions are truncated at 3σ above and below their mean values, and then renormalized to integrate to unity.)

Multiple sets of weak-link and strong-link failure temperatures can be generated from standard Monte Carlo sampling, and then a safety margin can be computed for each set by using the time histories of the weak and strong links obtained from the thermal model run at the $\{r, x\}$ heating conditions. The resulting population of safety-margin realizations will be distributed with some probability density, as exemplified in Figure 3. Associated statistics, such as the mean, the standard

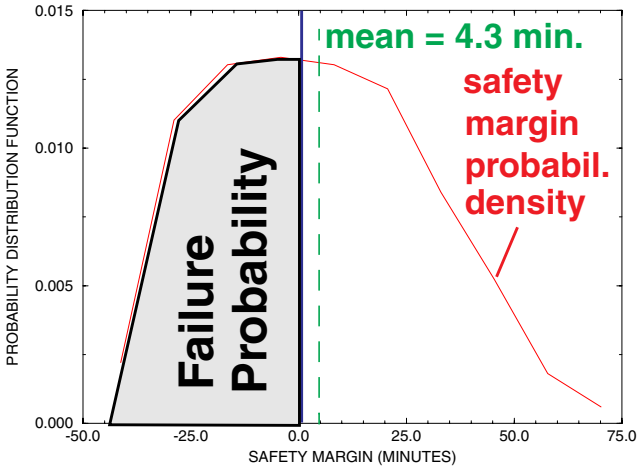

FIGURE 3. Safety-Margin probability density function at a particular set (r, x) of heating conditions for uncertain strong-link and weak-link failure temperatures.

deviation, and the probability of attaining a zero or negative safety margin, can be calculated. For example, the mean safety margin denoted in the figure is +4.3 minutes for temperature response curves obtained by running the model at the deterministically worst case set of heating conditions $\{r, x\}_{\text{opt-det}} = \{1.62, 0.782\}$.

The deterministic safety margin for this set of heating conditions (and nominal failure temperatures of T_fail_{WL} = 250ºC and T_fail_{SL} = 600ºC) is about 2.5 minutes, or 40% less than the 4.3-minute mean. This problem is clearly nonlinear in the uncertain parameters T_fail_{WL} and T_fail_{SL}; otherwise, the mean safety margin would equal the deterministic safety margin at the mean values of the uncertain parameters. This nonlinearity may result in different optimal (worst-case) heating conditions $\{r, x\}_{\text{opt}}$ if component failure uncertainty is taken into account rather than being ignored as in the deterministic optimization problem.

Figure 4 shows a 3×3 grid of points over a small subset of the design space, centered about the deterministic optima $\{r, x\}_{\text{opt-det}} = \{1.62, 0.782\}$. Figure 5 shows magnitude bars for failure probabilities calculated at the nine points of the grid by Monte Carlo sampling over the strong-link and weak-link failure temperature uncertainties. Thus, safety-margin distributions and corresponding failure-probability magnitudes (like in Figure 3) are obtained at each design point. The optimization problem now is to *maximize* the *probabilistic* objective function for failure probability in order to determine the worst-case heating values of r and x.

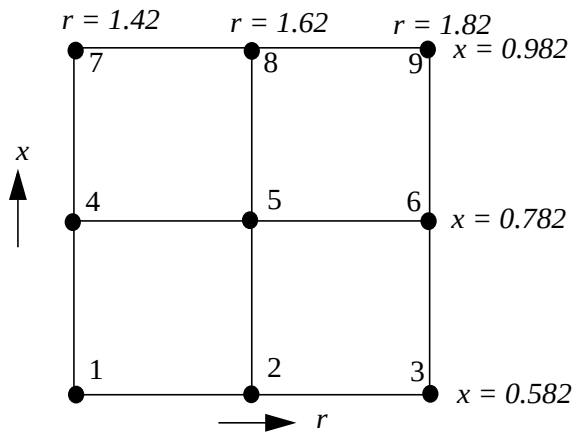

FIGURE 4. 9-point grid over important subset of 2-D optimization space, centered at point where deterministic optimum for worst-case heating occurs.

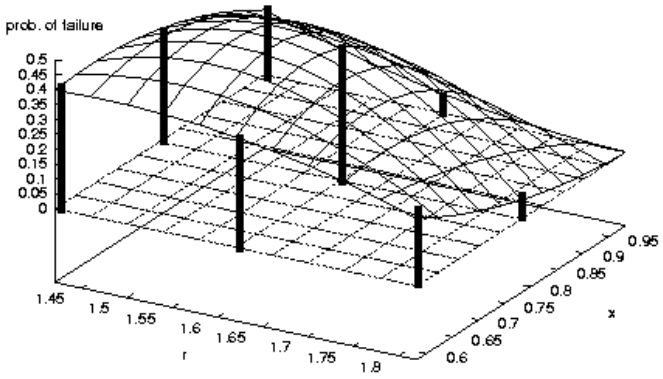

FIGURE 5. Probabilistic objective function (biquadratic response surface built from Monte Carlo point estimates of failure probabilities at the nine grid points).

2.2.2 Motivation for Probabilistic Ordinal Optimization in This Problem

The biquadratic response surface shown in Figure 5 can actually be seen to not go exactly through the end points of the probability bars because the response surface was created from probabilities obtained from 500 Monte Carlo samples, rather than from the 1000 samples that the bars are based on. The mismatch between bar height and response-surface height can be seen to vary somewhat over the design space. The mismatch reflects the fact that MC estimates from a given number of samples have confidence intervals (CI) as depicted in Figure 6 that reflect uncertainty (potential error) in the point estimates.

The following questions then arise: What effect does the number of Monte Carlo samples have on the accuracy with which a probabilistic optimum can be identified in the design space? If one wants to identify the heating variables $\{r, x\}$ that correspond to the highest probability of device failure, how can the impact of Monte Carlo sampling errors be controlled? How many samples

need to be taken at each point in the design space to shrink the confidence intervals small enough so that it is unambiguous which design point corresponds to the highest failure probability?

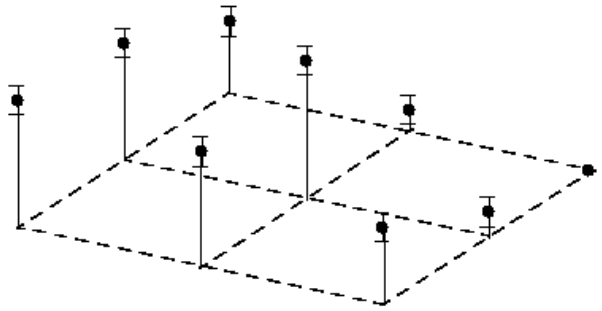

FIGURE 6. Monte Carlo point estimates of failure probabilities with associated confidence intervals shown.

3 Fundamental Concepts of Probabilistic Ordinal Optimization

In addressing the above questions, it is noted that probabilistic ordinal optimization has been studied for some time in the operations research field (e.g. Ref. [11]), to compare discrete design alternatives (as opposed to continuous-variable problems like the one in this paper). In particular, the work has addressed apportioning Monte Carlo sampling amongst multiple uncertain or stochastic discrete systems to most efficiently resolve their statistical behavior for the purpose of selecting the best option or several best options. The concepts and procedures apply whether it is desired to A) minimize the total number of samples required to reach a desired probability of correct selection of the best option or several of the best options, or B) maximize the probability of correct selection for a given budget of total samples N_T to be optimally apportioned among the various alternatives. In either case, the odds that the current identification of the best alternative(s) is correct can be estimated at every stage of the sampling.

The concepts of probabilistic ordinal optimization can also be applied to continuous-variable probabilistic optimization problems. For such problems, improvement steps in the design space can be taken based on ordinal ranking of candidate alternatives according to relative merit, rather than by attempting to resolve the actual merit value of each alternative as nonordinal approaches do. (See, e.g., [2] and [7] for brief overviews of some nonordinal approaches to OUU, such as Reliability-based and surrogate-based methods). The relaxed ranking conditions for progress in ordinal optimization allow it to be more efficient than nonordinal methods for certain problems, as demonstrated in [17]. These latter approaches can introduce sources of noise and approximation error into the probabilistic objective function that can be avoided with probabilistic ordinal approaches (see [18]). Furthermore, for C-V problems the efficiency of sampling operations for ordinal ranking can be greatly increased by consideration of local spatial correlation in the design space as explored in section 3.2. Additionally, advanced continuous-variable ordinal optimizers can be used, such as genetic algorithms, evolutionary simplex and pattern search methods ([9], [12]), and the simultaneous global/local search method DIRECT[13].

3.1 Independent (Uncorrelated) Designs

Consider two different designs or processes that each have some variability in their behavior or outputs. To compare the relative merit of the two designs, statistics of their relative merit can be compared. Thus, one can ask whether design A or design B has the higher mean output, or the lower probability of meeting or not meeting some acceptable threshold of behavior, or the smallest variance in the product produced.

The behavior or output of these independent designs can be sampled to generate statistics of their tendencies. For mean and probability, it is known from classical confidence interval theory that the values $\hat{s}_A$ and $\hat{s}_B$ of a statistic calculated from finite sampling will be realizations from normal distributions about the true statistical values S_A and S_B, with variances $\sigma^2_{\hat{s}_A}$ and $\sigma^2_{\hat{s}_B}$ that decrease as the number of samples increases. Figure 7 depicts normal distributions of the calculated statistics from two designs being compared.

Under limited sampling, the two normal distributions each have nonzero variance, which means that the distributions overlap to some degree as illustrated in Figure 7. Hence, it is possible for a calculated statistic $\hat{s}_A$ of design A to have a value greater than the calculated statistic $\hat{s}_B$ of design

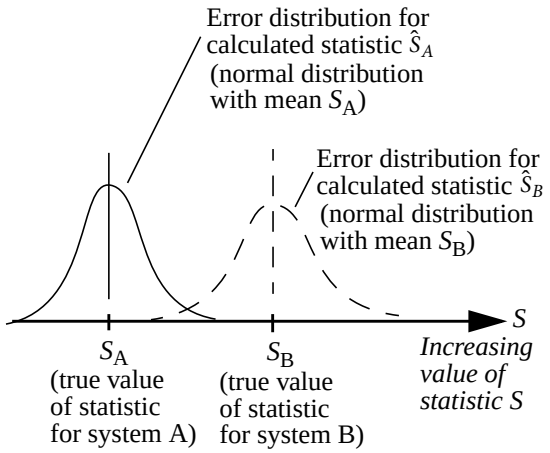

FIGURE 7. Normally distributed errors in Monte Carlo estimates of statistical behaviors of two different systems.

B, even though this is not the case for the true values S_A and S_B. If the optimization goal here is to pick the design that maximizes the critical statistic (say, maximizes the mean output rate or probability of acceptance), then there is a nonzero probability that a misleading indication would be given that design A is better than design B. The more the error distributions overlap, the less certain it is that the chosen design is indeed the better design.

As the number of Monte Carlo samples of each design increases, the critical statistics become better resolved. That is, their confidence intervals decrease such that it begins to become more and more apparent which alternative has the better statistical behavior. Hence, the probability of making a correct selection between the two alternatives (i.e., the '*probability of correct selection*' P{CS}), increases with more Monte Carlo sampling.

The P{CS} value has the property that it converges to 100% at an exponentially fast rate as the number of samples increases ([4]). This is a much faster rate than for resolving the individual alternatives' behavioral statistics. That is, the standard deviations of the individual distributions in Figure 7 only converge at a rate of 1/SQRT {number of samples of that alternative}. However, by looking at the figure it is intuitive that as the "width" (standard deviation) of either or both PDFs decrease at a fixed rate, the question of which alternative's value is greater is resolved at a much faster rate. In fact, such differentiation occurs at an exponential rate with the total number of samples expended, N_T, whether only one alternative is sampled or both are sampled: $P\{CS\} \geq 1 - \alpha e^{-\beta N_T}$ (see Ref. [4]).

Ref. [18] describes a methodology for calculating the probability of correct selection when only two design options are considered at a time. A requirement can then be posed to determine to a given level of statistical assurance P{CS} that the option picked has the better actual figure of merit. Meeting this requirement necessitates sampling the alternatives' behaviors a sufficient number of times. The question then arises: What is an adequate number of samples?

A unique answer to this question does not exist. In looking at Figure 7, a desired overlap (or nonoverlap as the present requirement is stated) may be attainable by increased Monte Carlo sampling to manipulate the spread (variance) of alternative A's distribution, of alternative B's

distribution, or some combination of both (only this last case is guaranteed to always work). If a combination of both is used, there is not a unique combination that will attain the desired result.

The way to approach the sufficiency issue, then, is to ask a second question: What is the most *efficient* way to achieve sampling sufficiency for a given P{CS} requirement? This question *does* have a unique answer, and this answer minimizes the total number of samples N_T apportioned between the two alternatives that will achieve the desired P{CS} level. In comparing two alternatives, the optimal apportionment of samples amongst the two alternatives is relatively easy to obtain, as described in Ref. [18]). In comparing more than two alternatives, the optimal solution from Ref. [3], referred to as '*optimal computing budget allocation*' (OCBA), is much more involved and difficult to obtain. Furthermore, calculation of P{CS} itself is much more difficult than the well-known approach for two alternatives.

Though the possibilities of increasing the efficiency are myriad, a simple implementation of probabilistic ordinal optimization (comparing only two alternatives at a time) was applied in Ref. [18] to the probabilistic optimization problem defined earlier.[2] A requirement of 95% P{CS} was imposed in the ordinal comparisons, and a simple CPS algorithm ([9]) was used to systematically progress toward the optimum in the continuous 2-D *r-x* design space. At the end of the process, a high confidence existed that the probabilistic optimum obtained was the true optimum among the designs considered.

To the authors' knowledge, no other methods for probabilistic optimization currently offer this type of quantitative assessment of correctness. If fact, the combination of CPS optimizer with probabilistic ordinal optimization as discussed here (with either pairwise or simultaneous comparisons), is asymptotically convergent to the true local probabilistic optimum (see the Section 5 of this report). Thus, probabilistic ordinal optimization may be usable to provide a reference standard against which the accuracy and efficiency of other OUU methods can be compared—analogous to the role that Monte Carlo simulation plays in uncertainty propagation.

3.2 Correlated Designs: Efficiency from Spatial Correlation of Uncertainty in Continuous Design Spaces

The full-variance distributions of Figure 7 apply for completely uncorrelated designs, as is often the case in comparing discrete alternatives in the realm of Operations Research. For continuous design spaces, however, closely neighboring points in the design space can have closely correlated uncertainties. The efficiency-enhancing prospects of spatial correlation of uncertainty in the design space are investigated here. Figure 8 illustrates the correlation issue. The mapping of uncertainty distributions at two neighboring design points is shown. For convenience of illustration, the uncertainties in this particular figure derive from stochastic noise in the tolerances that can be held in the design variable. This uncertainty (as a function of location in the design space) maps through

---

[2] A more sophisticated application considering multiple alternatives simultaneously on our test problem is demonstrated in this report in section 4. In comparing more than two alternatives, a '*simultaneous comparison*' approach is often more efficient than the '*sequential pairwise*' method of comparing two at a time and then dropping the "loser" and comparing the "winner" to the next alternative in the list, until all alternatives have been considered. If one of the alternatives happens to be much better than the current best point and all the other alternatives in the comparison, then simultaneous comparison will generally be more efficient than sequential pairwise comparison. The sequential pairwise method has its own appeal, however, because of its relative simplicity of implementation.

the deterministic input/output function of system behavior, as shown. Resulting output response uncertainties are depicted on the vertical axis.

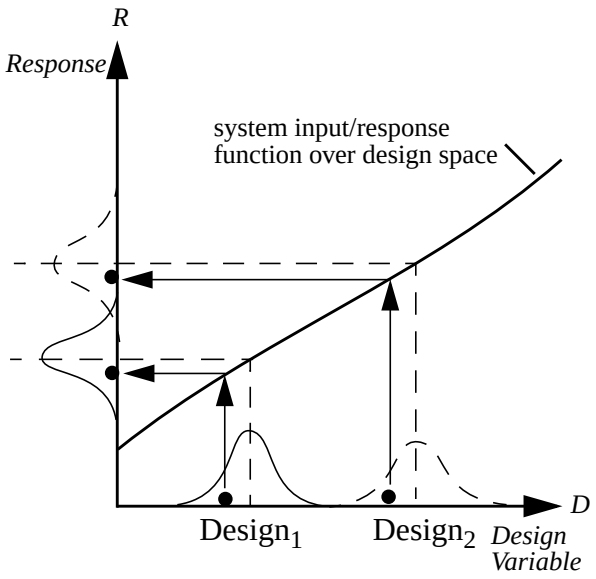

FIGURE 8. Correlated sampling of input/response distributions at two different points in design space.

Though shown on the design-variable axis, uncertainties in other (non-design-variable) inputs to the system may exist. These uncertainties would similarly map into response uncertainties. Thus, in general, the response distributions may have contributions of uncertainty from design variables and/or other system variables.

In any case, as the compared points in the design space get closer and closer, the inputs to the system or design look more and more alike. Furthermore, the response function that maps inputs in the design and uncertainty spaces to points on the response axis becomes more and more alike. This means that if Monte Carlo realizations of the involved uncertainties are taken with the exact same random-number-generator starting seed at each design point termed "*spatially correlated sampling*" in [17]), the uncertainty realizations and the response realizations will be very strongly correlated, as illustrated in Figure 8 for the *i*th sample taken at each design point. That is, on the design-variable axis, uncertainty realizations at the neighboring design points will come at similar percentile locations in the input uncertainty distributions. Likewise, the mapped response values for these input uncertainty realizations will occur at similar percentiles of the response distributions. Thus, the input/response realizations are spatially correlated (at least locally) in the design space.

3.2.1 The Spatial-Correlation Deterministic Efficiency Limit

In the spatial-correlation deterministic efficiency limit, the input uncertainties and the system response function that maps the input uncertainties to response uncertainty do not change over the design space. Therefore, the response uncertainty distribution does not change over the design space except for relative position on the response axis. Hence, it is trivial to identify which design point has the highest mean response, or probability of response exceeding some threshold value.

Only one correlated Monte Carlo sample of the uncertainty at each design point is needed to locate the relative position of the response distributions. If the single sample at each design point is evaluated at the midrange values of the input uncertainties, then this is a balanced deterministic transform of the probabilistic optimization problem. Note that with a single sample it cannot be said what the mean or probability value is at each design point, but only which design point spawns the highest or lowest such value. This is sufficient to select the best move for optimizing the probabilistic objective function.

Thus, the conditions just described enable probabilistic OUU problems to be treated as deterministic optimization problems—with the attendant simplification and cost savings that this brings. This limiting case, where perfect spatial correlation in continuous-variable probabilistic optimization problems causes them to be degenerate such that they can be recognized and treated as deterministic optimization problems, is termed the '*spatial correlation efficiency limit*' ([17]). An ordinal perspective identifies (and exploits) this determinism, whereas other OUU approaches do not.

3.2.2 Efficiency from Domination of "Uncertainty Variation" by Mean Response Trend in the Design Space

The spatial-correlation deterministic efficiency limit just described features a mechanism that allows probabilistic OUU problems to be treated as deterministic optimization problems, with the attendant cost savings. Since there is no response-uncertainty variation in the design space, the optimization problem is fully determined by the aggregate response behavior. This is a limiting case of a general mechanism in which response-uncertainty variation in the design space is substantially dominated by the mean response variation, such that effective progress in the probabilistic optimization problem can be made by solving a deterministic simplification of the optimization problem. Such dominance frequently exists in probabilistic optimization problems, particularly at stages of global and early-local search. This allows successful initial treatment as a deterministic optimization problem that enables relatively inexpensive location of an initial point in the design space (deterministic optimum) that is relatively close to the probabilistic optimum. From this starting point, OUU methods are applied to find the probabilistic optimum.

Figure 9 illustrates the issues. The middle histogram in the figure corresponds to the safety-margin uncertainty distribution (Figure 3) at the nominal deterministic optima $\{r,\ x\}_{\text{opt-det}} = \{1.62, 0.782\}$. The histograms on either side correspond to the neighboring points 4 and 6 of the 9-point grid of Figure 4. In Figure 9, the mean deterministic response is approximated by a quadratic curve through the 50th percentile locations on the safety-margin distributions. The location of the probabilistic optimum relative to the deterministic optimum is controlled by the local trends in the uncertainty and mean-deterministic behaviors. In the present problem, the mean trend of response (as indicated by the 50th percentile curve in Figure 9) rises to the left of the deterministic minimum. As the mean trend of response rises, the variance of the safety-margin distribution has to increase commensurately for the amount of the distribution beneath the $S = 0$ axis to keep increasing. At some point, even though the response variance keeps increasing in this direction, the effect of the rising mean trend of behavior overwhelms the effect of the increasing variance and the failure probability begins to decline. This is reflected in Figure 9 by the failure probability curves' maxima closely to the left of the deterministic optimum. (These curves are quadratic fits to failure probabilities at grid points 4, 5, and 6 calculated by (a) 500 Latin-

Hypercube Monte Carlo samples and (b) a 5-sample central-difference first-order second-moment "FOSM" method—see e.g. Ref. [16]).

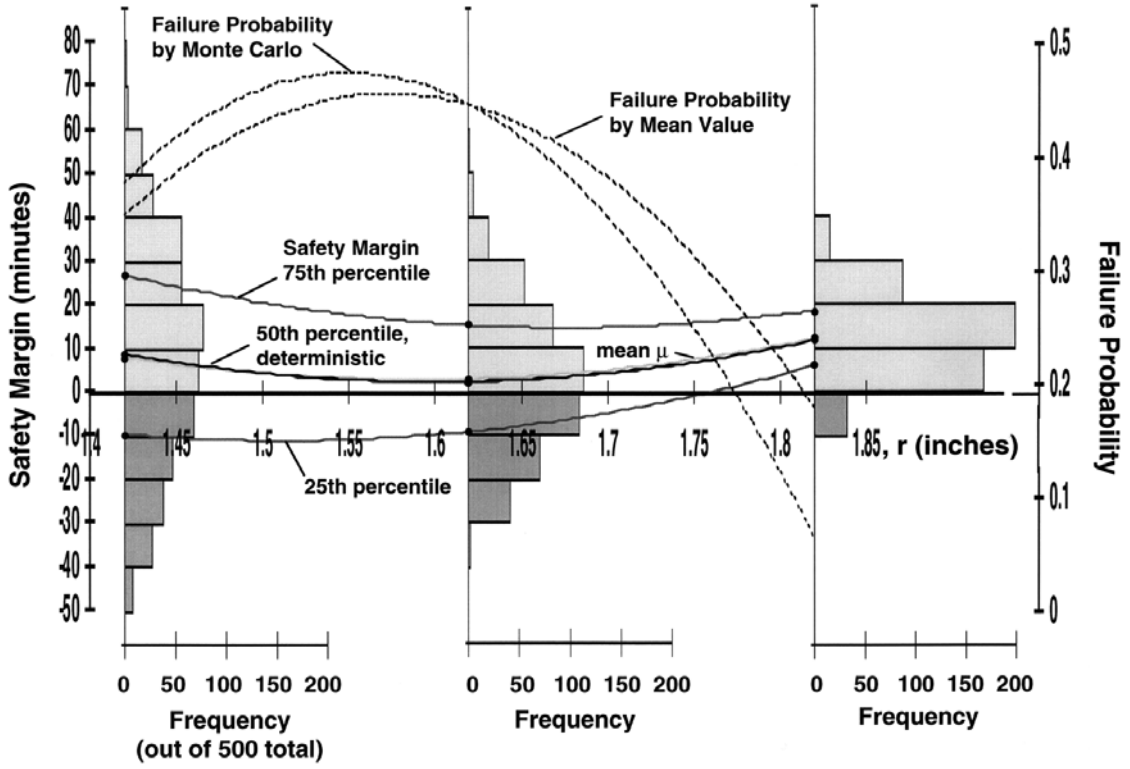

FIGURE 9. Histograms of response (safety margin) distributions at three points (#'s 4, 5, and 6 of the 9-point grid in Figure 4) along an $x = 0.782$ cut of the optimum region of the design space.

In this problem, over the early-stage probing scale defined by points 4, 5, and 6 in the design space, the failure probability is substantially determined by the trend of mean behavior, which has an optimum at the middle point #5 — very close to the probabilistic optimum (see Figure 9). Because the deterministic mean objective function is flat (has no gradient) at the deterministic optimum point #5, further progress in the probabilistic optimization problem can only be made according to the non-zero "*uncertainty gradient*" there, as explained next. Consequently, the deterministic approach to the probabilistic optimization problem (the success of which relies on dominance of the uncertainty gradient by the deterministic trend of behavior), must be abandoned in favor of a true (non-degenerate) probabilistic optimization approach that is driven by the uncertainty gradient.

At the deterministic optimum point #5 in Figure 9 an obvious *uncertainty gradient* exists. That is, the safety-margin variance increases meaningfully as the *r* coordinate is traversed leftward from the minimum point #5 of the deterministic objective function; the histograms of response and their associated 25th, 50th, and 75th percentile curves increasingly spread out, signifying increasing response variance in this direction. Since the deterministic trend curve is flat at the deterministic optimum, any nonzero uncertainty gradient there will push the probabilistic optimum off of the deterministic optimum. The direction in which the variance of the safety margin increases is the direction in which more of the distribution will fall below the $S = 0$ threshold (horizontal axis in

Figure 9) at least initially as the *r*-coordinate is varied from the deterministic optimum. This is therefore the direction in which the failure probability will increase (at least initially), and the direction in which the probabilistic optimum (maximum failure probability) will lie. Hence, if sufficiently accurate techniques are available to determine the direction of the uncertainty gradient, or if suitable ordinal optimization is used as discussed in the next subsection, then one can proceed from the local deterministic optimum to the local probabilistic optimum.

More generally beyond the optimization problem here, in a *multimodal* probabilistic optimization problem, the global optimum may not be reachable from the global deterministic optimum using *local* OUU methods (which are the subject of this paper). Nonetheless, if the uncertainty gradient is zero at a deterministic optimum, then the location of the determinstic optimum is at least a local optimum of the global probabilistic optimization problem—as there is no uncertainty gradient to push the probabilistic optimum off of the deterministic optimum. When the uncertainty gradient is zero over the entire design space, the location of the determinstic global optimum is also the location of the global probabilistic optimum. In this case, the probabilistic optimization problem can be treated as a deterministic optimization problem as was explained in the previous subsection associated with the spatial-correlation efficiency limit. Another case can be imagined where the uncertain inputs to the system, and/or the system behavior function that maps the inputs to response uncertainty, change over the design space[3] but the resulting response distribution does not. Here, too, the global deterministic and probabilistic optima occur at the same location in the design space.

If only one probabilistic optimum exists in the design space, then this global optimum can be reached from the deterministic optimum (as is the case for the optimization problem in this paper). Furthermore, for such unimodal probabilistic optimization problems, if the uncertainty gradient is zero at the deterministic optimum, then this is the location of the probabilistic optimum as well, since there is no uncertainty gradient to push the probabilistic optimum off of the deterministic optimum. In the extreme case where the uncertainty gradient is zero everywhere over the design space, it is of course also true in the unimodal problem that the deterministic and probabilistic optima occupy the same location in the design space.

3.2.3 "Point of First Separation" Ordinal Selection Efficiency Mechanism

As was established in the previous subsections, deterministic-scale efficiency in the probabilistic optimization problem can be achieved when the influence of the uncertainty gradient is overwhelmed by mean behavioral trends. In this case, a single ordinal comparison will suffice for making effective optimization decisions in the probabilistic OUU problem. Conversely, when the effect of the uncertainty gradient dominates the effect of the underlying mean trend of behavior,

[3] Here, even with correlated Monte Carlo sampling, over the design space the correspondence changes between sampled percentiles of the input uncertainties and their associated response percentiles. Thus, the response realizations at design points will become more weakly correlated in their percentiles of response as the distance between the design points increases. Hence, in contrast to the perfect-spatial-correlation special case, it is essential here that the single realization at each design point (in the deterministic analogue of the probabilistic optimization problem) occur at the midrange values of the uncertain inputs. Then the spatial discorrelation of the resulting output realizations will be minimized, and the probabilistic optimization problem can be approximated as a deterministic optimization problem. Since this sampling strategy also works for the special case of perfectly correlated uncertainty in the design space, it is a good practice to always use the midrange uncertainty values in the deterministic transform of the probabilistic optimization problem.

like ordinarily occurs at the location of the deterministic optimum, it is necessary to turn from a deterministic treatment of the OUU problem to one in which more than one correlated sample per design point is necessary for successful progress in the problem.

The next level of sampling treatment that is suitable is coined the '*Point of First Separation'* (PFS) ordinal selection method. The PFS method can be thought of as an initial-stage implementation of the regular ordinal-optimization method, that truncates the sampling at the earliest reasonable breakpoint on the expectation that more sampling will only serve to confirm more strongly (and not overturn) the earliest tangible indication of ordinal dominance of the best design option(s). This expectation assumes that reasonably strong spatial correlation of the uncertainty will exist at the relatively small design-space scales involved in the final movement from the deterministic to the probabilistic optimum. (Recall that perfect or absolute spatial correlation of uncertainty occurs in continuous-variable problems as the distance between design points approaches zero.) A later demonstration on the OUU problem indicates that the PFS method is surprisingly efficient and effective.

The PFS method is explained with the help of Figure 10. Three design alternatives (A, B, and C) are being compared in the figure to determine which has the lowest probability of failure. This illustration concerns failure-probability minimization rather than maximization, but the same principles apply whether seeking to maximize or minimize probability.

To start, a single correlated Monte Carlo sample of the response uncertainty of each design is evaluated. The sampled response of design C has a response value greater than the upper operating threshold level for acceptable system response. The first sample for the other two designs does not produce a threshold exceedence. On the assumption of closely spatially correlated sample realizations (i.e., the *n*th correlated Monte Carlo sample of each response distribution lies approximately at the same percentile on each distribution), this first sample already implies that design C will have a larger integrated failure probability than designs A and B. Because the

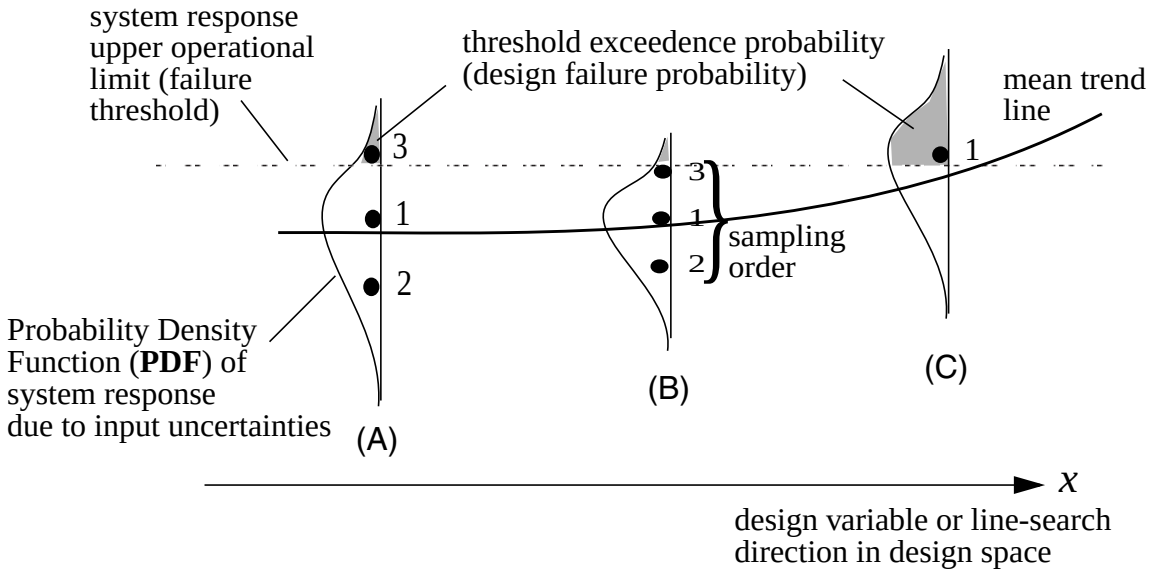

FIGURE 10. Correlated sampling of response behavior of three design alternatives at points A, B, and C in the design space.

objective is to identify the design with the lowest failure probability, design C can be immediately eliminated from further consideration.

A second correlated sample is then taken from the two designs still in contention. Neither of the remaining designs A and B indicate a failure (threshold exceedence) with this second sample. Of the two, design A has lower response values to this point, and thus would appear to be the better candidate at this point. However, here is a case where the influence of the uncertainty gradient is stronger than that of the underlying mean trend, and dominates the determination of failure-probability magnitude. Hence, selection of the best alternative from indications so far would lead to an incorrect selection.

The least-expensive strategy that can often be successful under dominant influence of the uncertainty gradient is to continue the correlated sampling of the remaining designs until one or several become distinguished from the rest by a difference in response status relative to the critical threshold in the problem. This occurs on the third correlated sample in Figure 10. The response value of design A lies above the threshold, while that of design B lies below the threshold. On the assumption of spatially correlated percentile realizations, this third sample implies that similar percentiles of response for designs A and B respectively lie above and below the threshold. Therefore, it is implied that design B will have a lower integrated failure probability than design A.

In this contrived problem, it takes a total of seven Monte Carlo samples apportioned among the three initial candidate designs to identify the correct design. With a different sample placement and ordering that would accompany a different random-number-generator initial seed, a different total number of samples might be required. If, for instance, the sampling order in the figure was reversed, then with the first sample of each alternative (sample #3 in the figure), both designs A and C would be discarded immediately. Design B would be correctly selected at a total cost of only three samples. So, the total number of samples employed in this truncated version of the full probabilistic ordinal selection process will vary depending on the specific random-number–generator starting seed. This is true for full probabilistic ordinal optimization as well.

Figure 11 shows some results of an empirical investigation of the PFS method. Calculated failure probabilities at points 4 and 5 of the 9-point grid (Figure 4) are shown versus number of correlated Monte Carlo samples of each design. With the fourth sample of each design, a separation in the calculated failure probabilities occurs. The separation wavers but generally grows as the sampling increases. Once the separation occurs, the calculated probability values never cross each other, so a correct ordinal selection was made just after the point of first separation. Any added sampling simply serves to decrease the confidence intervals about the calculated probability estimates, thereby increasing the probability, P{CS}, that the initial indicated separation is correct. The correlated sampling shown in Figure 11 takes 202 samples total (101 each) to achieve a P{CS} of 95% under equal sampling of the two alternatives.[4] Under optimally efficient (nonequal) OCBA sampling of the two alternatives, only 134 samples total are required to attain 95% P{CS}, see Ref.

---

[4]Because of the difficulty involved otherwise, all P(CS) values we state are calculated ignoring any spatial correlation of uncertainty in the design space. When such positive correlation exists, corresponding P(CS) values will be higher than when calculated on the basis of independent, uncorrelated designs (see Deng, M., Ho, Y.C., Hu, J.Q., "Effect of Correlated Estimation Errors in Ordinal Optimization," Proc. 1992 Winter Simulation Conference, eds. J.J. Swain, D. Goldsman, R.C. Crain, J.R. Wilson). Hence, the numerical values of P(CS) appearing in this paper may substantially under-state the actual probabilities of correct selection.

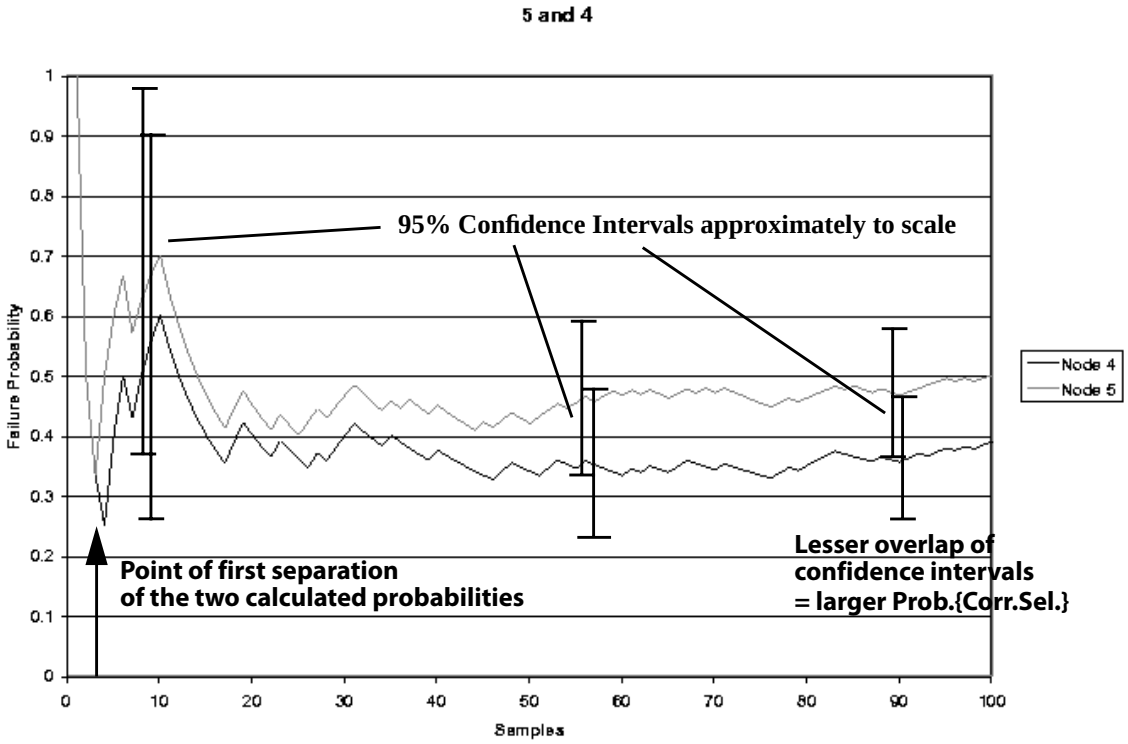

FIGURE 11. Failure probabilities by correlated sampling of response behavior at design points 4 and 5 of the 9-point–grid subspace (Figure 4).

[18]. In contrast, the point of first tangible indication in Figure 11 occurred after only eight samples (four of each alternative).

The non-crossing of probability curves (once initially separated) that is exhibited in Figure 11 also occurs for all other two-alternative comparisons possible among the points of the 9-point grid. For any problem, it appears that the more local the comparisons—as when the probabilistic optimum is being converged to in the final stages of local optimization—the more similar the sampled response percentiles should be under correlated Monte Carlo sampling, and therefore the better the chances of correct selection by the Point of First Separation mechanism.

PFS ordinal selection is used in Ref. [17] to successfully step to the probabilistic optimum of the present problem (section 2.2) with a simple CPS optimizer. The cost of solving the probabilistic optimization problem was 31 samples from the starting point. The same probabilistic optimum was found in Ref. [18] with 95% P{CS} at each CPS selection step. That result, having a very high probability of correctness, took 2721 samples to establish. Thus, PFS ordinal selection gave the same result but was tremendously cheaper. Note that the application problem here has not been contrived in any way to be favorable to the PFS ordinal selection method.

By comparison, the least-expensive non-ordinal approach possible for guiding the CPS search is to use a first-order second-moment mean-value method (FOSM, see e.g. [16]) to estimate failure probability at the various candidate points in the CPS search. This would have required 30 samples total for one-sided differencing, and 50 samples for more robust central differencing. Thus, only in

the best possible case (use of the lowest-cost probability quantification method available (linear FOSM), yet also no significant error in the calculated failure probabilities) could an explicit probability-resolution approach be as inexpensive as the 31-sample PFS ordinal method at making correct stepping decisions in the test problem. This is a reasonable first indication of the relative efficiency of the PFS ordinal OUU approach. Sections 4 and 5 present other reasons why this approach is expected to possess potential advantages over other OUU approaches in terms of robustness and cost-scaling as the number of design and/or uncertain variables increase in the OUU problem.

4 Demonstration of a Combined Adaptive Probabilistic-Ordinal OUU Method

An adaptive probabilistic-ordinal OUU method is demonstrated in this section that combines the elements described in the previous subsections. The optimization goal in this example presentation is to move in the design space to identify the optimal values of the fire heating parameters *r* and *x* that maximize the probability of failure of the safing device.

4.1 Deterministic Initial Phase of the Method

The first element of the strategy is to employ a "preconditioner" to the probabilistic optimization problem by solving a representative deterministic version. This is done by temporarily fixing the uncertainties in the problem to their mean values and then solving the corresponding deterministic optimization problem. For the present problem, this has been done with a variety of deterministic optimization approaches as reported in section 2.1. The resulting deterministic optima, $\{r,\ x\}_{\text{opt-det}} = \{1.62, 0.782\}$, define the center (point #5) of the 9-point grid in Figures 4 through 6.

4.2 Point-of-First-Separation (economical ordinal selection under uncertainty) Phase of the Method

Starting from the deterministic solution, the probabilistic optimization is begun at a suitable initial probing scale in the design space, which may or may not be the same as the final probing scale in the deterministic precursor problem. The probing scale denoted by the 9-point grid is the initial probing scale to be used for this demonstration. A simple probing/stepping algorithm is used in the following to facilitate illustration of the concepts here. Much more sophisticated ordinal optimizers could be employed, as mentioned elsewhere in this paper.

In beginning the OUU phase, the deterministic trends about the starting point #5 are known from the deterministic phase of the problem. The deterministic trend information is used for an initial indication of which direction appears most likely, based on the deterministic results, to yield the fiercest competitor to the current best point in the design space.

In the problem at hand, the closest deterministic competitor to the final deterministic optimum lies at the left-center point #4 of the 9-point grid. Hence, the PFS correlated sampling method is applied to design points (alternatives) 5 and 4 in the design space. Recall that the PFS method samples both designs equally. This is done with correlated Monte Carlo sampling until the first signal of advantage of either design. The point of first separation occurs after just 4 samples of each alternative, as established in Figure 11 and the previous subsection. Figure 12 depicts this comparison in its upper-left graphic (a). The figure notes that 8 samples total (4 for each alternative) are required for the comparison—that is, to get to the point of first indicated separation of probability values. For reference, the true failure probabilities of the alternatives being compared are printed in the figure.

With the center point having provisionally distinguished itself as having a higher failure probability than the grid point to its left, the next course of action is to determine whether the apparent trend of increasing failure probability continues to its right, in going from grid point #5 to grid point #6. As shown in Figure 12(b), only two samples (one sample for each alternative) are necessary to provisionally signify that the center point has the higher probability of failure. It is to

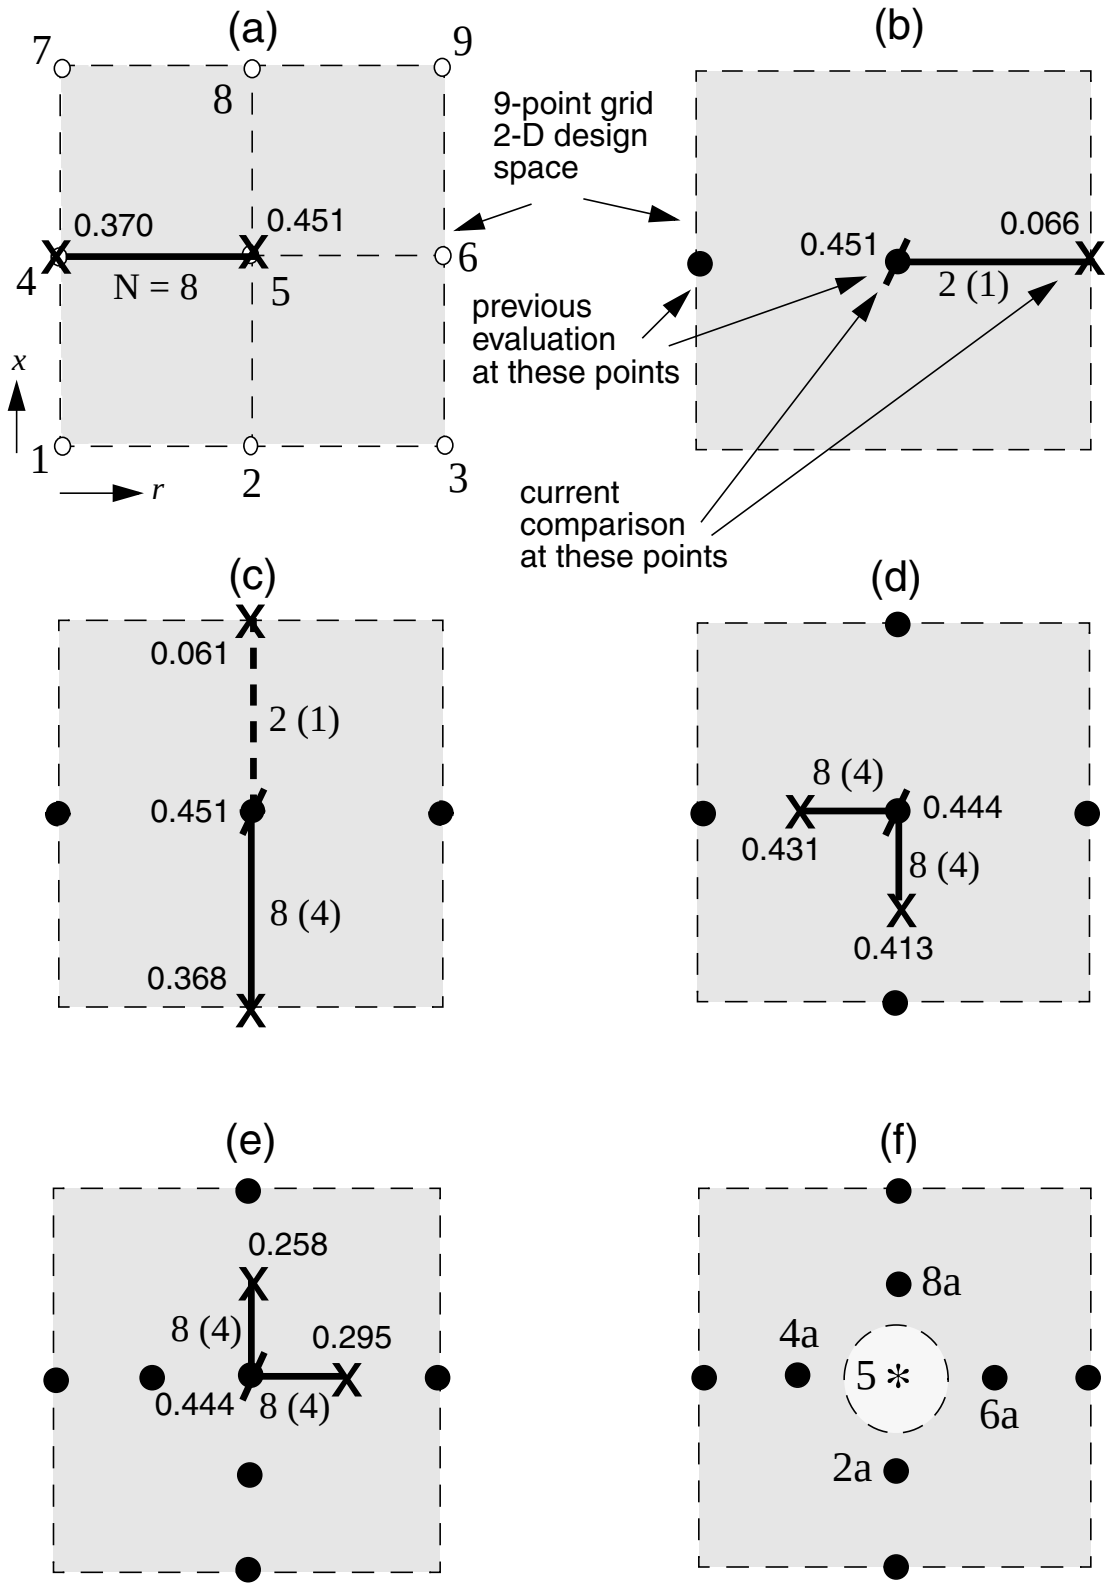

FIGURE 12. 2-D design optimization in 9-point–grid design space by simple space-searching algorithm and ordinal decision making for better designs by pairwise 'Point of First Separation' (PFS) selection criterion based on closely-related (correlated) uncertainty of nearby designs in the continuous-variable design space.

be expected that a very low number of samples would be necessary for this determination, given the large difference between the actual failure probabilities at the two points, as noted in the figure. The first safety-margin realization at point #6 is positive, or above the safety threshold, indicating an initial failure-probability estimate of zero. Conversely, the correlated first sample at the center point #5 is negative, or below the safety threshold, indicating an initial failure-probability estimate of 1. Thus, the designs immediately separate in terms of their indicated safety after just one seed-correlated sample of each. Listed in Figure 12(b) by the number 2 is the number 1 in parentheses. The number in parentheses is the number of *new* samples that must be drawn for the comparison. Only one new sample, of design #6, was needed here; the sample for the center design #5 was taken from its previous comparison against alternative #4.

Thus, for the current level of granularity (grid probing scale) of the design space, the indicated maximum failure probability is at the center point #5 in the *r* design-coordinate. We now consider the *x* design-coordinate. The deterministic trend at point #5 has the nominal safety-margin objective function decreasing in the direction from point #5 to the bottom-center point #2 of the 9-point grid. Therefore, a PFS comparison is performed between design points #5 and #2. As noted in Figure 12(c), the failure probabilities at the two design points are relatively close. A total of 8 samples (only 4 new, all at design point #2) are required to tangibly distinguish the center point #5 as having the higher failure probability. A consequent comparison of points #5 and #8, also depicted in Figure 12(c), again indicates that point #5 has the highest failure probability. This comes at a cost of only 1 new sample, of design #8.

Having exhausted all coordinate directions about the current best point at the current scale of resolution of the design space, a contraction is implemented to determine whether something off the center point at a smaller probing scale will possess an increased failure probability. Here a new variant of the search procedure is used, just to demonstrate the possibilities.

The comparisons in Figure 12(d) are driven by a desire to leverage off of the indicated directional trends of increasing failure probability in the negative *r* and *x* directions as signified by the comparisons at the coarser probing scale. From considerations of the deterministic trends, the contracted alternative #4a to the left of center is compared first. Only eight samples (just four new ones, at point #4a) are required with PFS ordinal selection to dismiss design 4a, even though its failure probability is 0.431 while the center point #5 probability is a very close 0.444. This is quite striking. Next, instead of comparing the rightward contracted point, as per the scheme of logic at the coarser probing scale, we take the option of first considering the most promising directional step in the other design variable. Only eight samples (just four new ones, at the downward contracted location #2a) are required with PFS to dismiss design 2a, even though its failure probability of 0.413 is relatively close to design 5's probability. Note that with this comparison scheme, designs 5, 2a, and 4a could be *simultaneously* sampled and compared under the PFS methodology. If one of the alternatives happens to be much better than the current best point and (all) the other alternative(s) in the comparison, then simultaneous comparison will generally be more efficient than the sequential pairwise-comparison approach used here.

Having not found a better candidate among the initially indicated most-promising designs 2a and 4a, candidate designs at opposite perturbations from the current-best point #5 are evaluated next to see if the upward trends of failure probability continue in the rightward and upward directions (as established by the initial comparisons of designs 4a and 2a respectively to point #5).

Hence, the next two candidates to be tested are designs 6a and 8a, at the rightward and upward contracted locations shown in Figure 12(e). Only eight new samples (four of each new candidate #6a and #8a) are required to provisionally dismiss these candidate designs as having lower failure probability than the reigning best design at point #5.

At this juncture, contraction can be implemented again and the optimization process be continued, or the process can be suspended for fiscal or physical reasons (i.e., when the probing step size becomes less than the level of resolution-uncertainty within which it is cared to resolve the optimal values of the design variables r and x). For the purposes of this demonstration we elect to end this PFS phase of the optimization procedure and get on to the next phase, below. Figure 12(f) conveys that in stopping here, the center point is provisionally deemed to have the highest failure probability in the space, within the resolution uncertainty indicated by the lightly shaded region about the identified optimal point (asterisk) in the design space.

An *approximate* P{CS}, or APCS (see below), can be calculated to assess the confidence that design #5 is truly the best among the five finalists 2a, 4a, 6a, 8a, and 5. The *expected value* of APCS is P{CS}, which converges to a probability of 1 as the number of samples increases (see Ref. [3]). The current realization of APCS, based on just four samples of each design's probabilistic behavior from the PFS procedure is 0.18. This value is based on just 20 samples total among the five candidate designs, so is very tentative. Even if the relatively low APCS value of 0.18 is considered acceptable for the present purposes, we will see below that APCS estimates (particular realizations) can be exceedingly unreliable at low levels of samples. Therefore, we do not put much faith in this estimate. However, this does not mean that we strongly doubt that PFS identified the best design—just that we doubt the associated value of APCS calculated with so few samples. For more concrete assurance that the PFS selection is indeed the correct one, the OCBA process can be pursued as demonstrated next.

4.3 Optimal-Computing-Budget-Allocation (efficient ordinal assurance) Final Phase of the Method

For increased confidence regarding the final optimum, the simultaneous OCBA method can be applied to the final design point and any alternatives that the user deems to be approximately competitive with it. All of the alternatives that are quasi-competitive with the PFS-identified optimum should be included in the OCBA round of sampling because if another alternative is actually better than the PFS optimum, the OCBA process will identify that alternative. There is relatively little cost-penalty for any included alternatives that are not competitive with the actual best alternative; the OCBA algorithm determines this quite early and de-emphasizes those designs fairly quickly in its optimized sampling allocation among the various designs. For these reasons, in the final OCBA step of our adaptive OUU method, it is good practice to be somewhat generous with the pool of candidates deemed to be legitimately competitive designs.

For the present demonstration purposes we include designs 2a, 4a, 6a, 8a, and 5 for the OCBA round of sampling. We start from an initial condition of 4 samples per design, from the PFS phase of sampling. Figure 13 shows the running number of samples allocated to each design, as a function of the running total number of samples allocated during the OCBA process. The more competitive designs 2a, 4a, and 5, with failure probabilities 0.413, 0.431, and 0.451 respectively, get the lionshare of the samples. The best design, # 5, gets the most samples. The OCBA algorithm devotes

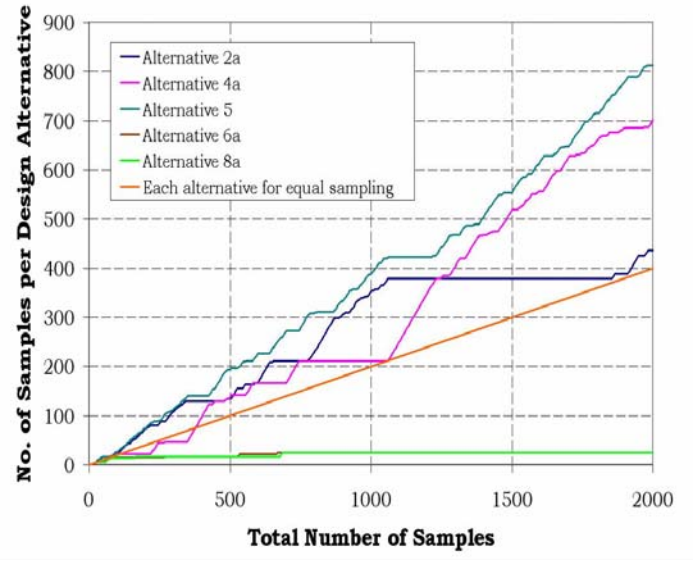

FIGURE 13. OCBA apportionment of samples to candidate design finalists 2a, 4a, 5, 6a, and 8a. In the end, the most samples go to the best designs.

most samples to the several best designs as it endeavors to isolate the highest-probability design and maximize the P{CS} probability that it is indeed the best design. Each of these three designs gets more samples than it would if all five designs were sampled equally, as shown by the 'equal sampling' curve in the figure. In stark contrast, the two non-competitive designs 6a and 8a, with failure probabilities 0.295 and 0.258 respectively, get far less than equal sampling.

Figure 14 shows the associated APCS as a function of the running total number of samples allocated. The calculated APCS value for OCBA sampling is seen to vary wildly in the early sampling, and to continue to possess significant stochastic variability or noise at higher numbers of samples. In particular, the APCS value quickly shoots up by ~0.3 from a value of 0.3 to 0.58 in

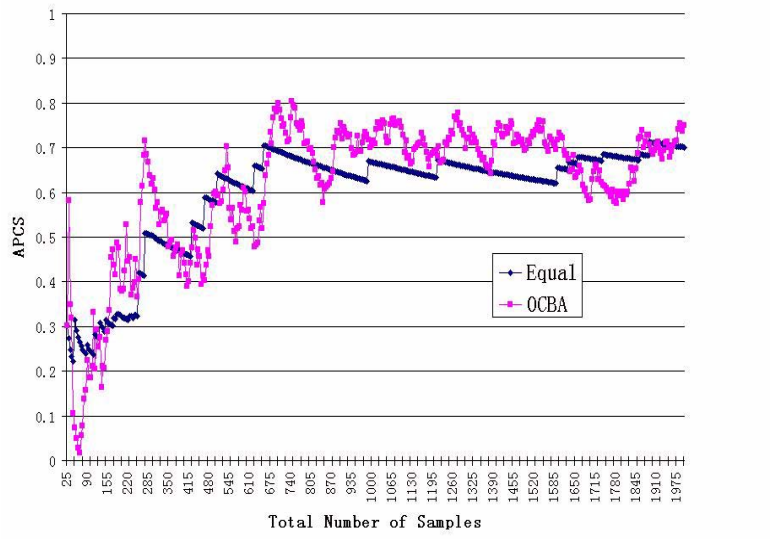

FIGURE 14. Approximate probability of correct selection, APCS, associated with the identified best design #5.

going from 25 to 30 samples. It then promptly declines to as low as 0.02 at 65 samples before beginning to go back up again in a relatively strong climb to APCS=0.72 at 275 samples. During this climb there are a couple of short periods of substantial fall-off of APCS, on the order of 0.2 drop in value. From the local peak of APCS=0.72, a ~0.3 drop-off in magnitude to 0.39 at 455 samples occurs. From this local minimum, another non-monotonic climb occurs to a local maximum of APCS=0.70 at 535 samples. A subsequent drop-off of about 0.2 occurs, to a local minimum of 0.48 at 625 samples. Then another strong but non-monotonic climb of about 0.3 in magnitude occurs, to a local peak of APCS=0.8 at 700 samples. From there the climbing trend pauses, out through 2000 samples where the sampling was stopped. In this last range, APCS actively varies within an envelope of effective lower and upper bands of 0.6 and 0.8.

We note that the calculated APCS value here applies only to the design alternatives actually included in the APCS evaluation. That is, in the APCS evaluation we could have also included designs 2, 4, 6, and 8 that were sampled in the PFS phase of the process, just to concretely affirm that these designs too are dominated by design #5 to a level of APCS that would not be much different from the present one. We also observe that Figure 9 suggests that the maximum failure probability may be slightly to the left of the center point #5. The maximum is evidently within the region of resolution-uncertainty about point #5 that is affiliated with the final probing scale of the design space. Hence, we caution that APCS is not the estimated probability that the optimum in the design space has been found, but only the probability associated with the identified optimum among the finite number of design points actually evaluated in the design space.

We also note that if we had used a different initial seed for correlated Monte Carlo sampling of all the designs, somewhat different curves than those in Figures 13 and 14 would have resulted. Yet, the *expectation* of the random processes would be the same: *i*) in the expectation, APCS rises, but at a decreasing rate, as the total number of samples increases; *ii*) the expected noise level in calculated APCS decreases as the total number of samples increases; and *iii*) the expected noise level in APCS is smaller for equal sampling, but its mean APCS value rises slower than for OCBA sampling.

The noise variation in calculated APCS values seen in this example is not out of the ordinary, except perhaps for the very large dip of 0.56 at 65 samples for OCBA APCS. Although this is uncommon in our other experiences, generally when the top designs are very close in probabilistic behavior, the variability of calculated APCS is relatively large for small numbers of samples. We note that the expectation is for OCBA APCS to converge to 1.0 faster than equal-sampling APCS, but the OCBA APCS estimate also tends to have greater variance.

Finally, we observe that a “knee” in the OCBA APCS curve seems to occur once a value of APCS=0.7 is first reached. This is also generally true in other applications that we have experience with. It seems that once OCBA APCS first reaches 0.70 there is no need to continue sampling any further, even if a very high assurance is desired that the correct design has been selected. The other argument here is that the cost of raising APCS beyond 0.6 or 0.7 begins to much more quickly accelerate beyond this point. Hence, we generally recommend OCBA APCS = 0.6 to 0.7 as the cost-effective point at which sampling can be stopped. We note here that OCBA APCS first reaches 0.7 at 275 total samples, whereas equal-sampling APCS takes 655 samples (over 2.38 times as much) to first reach 0.7. For reaching an APCS value of 0.6, OCBA costs 265 samples, whereas equal sampling costs 505 samples (not quite twice as much). Thus, even with the faster-converging OCBA method, it took over an order of magnitude more samples beyond PFS to substantiate with high certainty the correctness of the early indication.

5 Convergence, Robustness, and Cost Considerations

Our generic OUU algorithm combines simple Coordinate Pattern Search [9] for probing and stepping in the design space, combined with probabilistic ordinal optimization for decision-making about whether a candidate step with the CPS method is an improvement step in the probabilistic optimization problem. We will now consider the performance of each of these aspects separately and when combined into our OUU procedure.

When the objective function is deterministic, CPS can be proven ([23]) to converge to a local optimum under fairly standard conditions for provable convergence of gradient-based optimization methods (e.g., the objective function is bounded and continuously differentiable). The application considered in this report becomes a deterministic optimization problem that in principle meets the conditions for CPS convergence in the asymptotic case of infinite sampling of the local alternatives (those being considered at a CPS decision step). In this asymptotic case the failure probabilities become known with exact precision over the optimization design space. Thus, P{CS} among the local decision alternatives is 100% and the decision step is based on deterministic certainty. The converse is also true: to achieve local P{CS}=100%, the local alternatives' failure probabilities must be known exactly. The question then shifts to whether the probabilistic ordinal optimization method can be shown to converge (with infinite sampling) to P{CS}=100% among the local alternatives. Such convergence, at an exponentially fast rate, has been proven in [4]. This proof also holds for the optimally efficient sampling allocation variant OCBA ([3]). Therefore, asymptotically, our OUU method is *in principle* convergent to the probabilistic optimum.

Under more realistic circumstances of limited sampling, we can no longer pose the probabilistic optimization problem as a deterministic optimization problem to establish convergence. Under non-100% P{CS} at each decision point in the CPS algorithm, we cannot be absolutely certain that each step is correct so that the final optimum arrived at is the true optimum. However, once the final optimum is identified in our OUU procedure, we can perform a P{CS} calculation that includes not only the local adjacent alternatives, but all those visited during the optimization search. Then the resulting "global P{CS}" value applies to all alternatives (points in the design space) considered in the search. Of course, this leaves out all points in the design space (an infinite number of them) NOT considered in the search, so we can only claim our calculated global P{CS} is a "*conditional*" one predicated on the design points actually visited in the optimization procedure.

Although we can only lay claim to a conditional global probability that the true optimum is found by our procedure, what we can claim about our procedure appears to be somewhat more than what can be claimed for other C-V OUU procedures under realistic conditions of industrial applications. For instance, all such methods will have to successfully contend with small-scale objective-function noise coming from computational-physics models, which are characteristically noisy ([16]). As explained in Section 2.1, model-induced noise in the deterministic version of our optimization problem was effectively filtered out by the noise-tolerant CPS optimizer employed also in our OUU method. In contrast, gradient-based optimizers were not nearly as tolerant to the model noise.

Other OUU approaches than ours (such as reliability-based and surrogate-based approaches, see e.g., [2] and [7]) have to contend with more than just noise from the model. These approaches have additional sources of noise and approximation bias-error that our OUU procedure does not have

under our choice of using the same initial seed for MC sampling of all alternatives ("spatially correlated sampling", section 3.2). These potential optimizer-fooling effects can come from various sampling, iteration/optimization, and approximation-building operations that can be found in other OUU methods. Certainly, other methods do not presently (perhaps never will) have as explicit and rigorous of a basis for error assessment and control regarding decision making of improvement direction in the OUU problem.

In terms of cost, we have already mentioned several specific instances where our OUU procedure is more efficient than reliability-based and surrogate-based approaches. We also noted that a local P{CS} criterion of 60% - 70%, on the lower part of the cost-growth curve, is considered reasonable (especially in view of competing OUU approaches which do not have a reliable basis for estimation and control of error potential in local decision-making under uncertainty). Furthermore, we have empirically found here and in [17] when starting from a different point in the design space, that the exceedingly economical PFS truncated sampling method yields correct stepping decisions in the continuous-variable design space even though the associated P{CS} values are very low, often less than 0.1. Whether this type of economical performance can be expected in general will take many years and a diversity of applications to empirically answer. However, we are cautiously optimistic based on results and theoretical considerations so far.

We also note that the relative efficiency of the probabilistic ordinal approach to identify the optimum among a set of candidates generally increases as the number of alternatives increases. That is, under optimal allocation of samples amongst the alternatives, ordinal selection cost generally scales only very weakly as the number of alternatives to be considered increases. (See [3] for a variety of empirical tests.) CPS and all other non-gradient based optimizers like Nelder-Mead, DIRECT[13], and Genetic Algorithms exhibit relatively high cost-scaling (in terms of the number of alternatives investigated) with the number of design variables in the problem. Hence, the “curse of dimensionality” label attached to these optimizers. Fortunately, their cost-scaling in an OUU context is substantially mitigated by the weakness of scaling of our ordinal selection cost with the number of alternatives to be considered.

Another important property is how slowly the ordinal selection cost scales with the number of uncertain variables in the OUU problem. Ordinal selection cost generally scales only weakly with the number of uncertain variables ([3]). This is largely because the statistics are resolved by MC sampling, and the number of samples required for a given resolution (confidence-interval size) does not directly depend on the number of uncertain variables. For instance, failure-probability magnitude itself determines the size of the CI for a given number of samples –it doesn’t matter whether 1 or 100 uncertain inputs contribute to this failure probability. The same cannot be said of other methods for calculating failure statistics, such as response-surface methods and Reliability methods. These cost-scale strongly with the number of uncertain variables ([16])—although for a given accuracy their cost may be less than MC if the probability magnitudes involved are small enough. In that case, non-MC methods *might* be preferable for ordinal selection. This would depend on the added cost to assess and control (to sufficiently small levels) any noise and point-to-point varying approximation-bias error. Such assessment and control would require more function evaluations than for just a nominal point-estimate of response, and might also imply greater expense per function evaluation (i.e., models with more refined space/time discretization are usually required than with MC sampling approaches, see [16]).

In general, if using a non-gradient optimizer for OUU, an ordinal approach for selection of improvement steps is the most efficient approach. This avoids unnecessarily precise numerical quantification of response statistics. Whether using MC, response-surface methods, or Reliability approaches, calculation cost grows precipitously as the desired numerical precision of the computed results increases. It is much less difficult and expensive to determine whether one response is greater/less than another, than to accurately quantify their response values. Thus, non-gradient optimizers only require uncertainty propagation techniques to resolve the computed responses to an adequate level of precision to distinguish the best option from among the considered alternatives. At present, it appears that only MC sampling approaches feature practical mechanisms for the necessary quantification and control of precision of the computed response statistics for efficient probabilistic ordinal selection.

On the other hand, if using a gradient-based optimizer for C-V OUU, the number of evaluated points in the design space may scale with the number of design variables at a much slower rate than for a non-gradient optimizer. However, precise numerical quantification of uncertainty statistics is required. The quantification has to be relatively precise in order to get sufficiently accurate uncertainty gradients. Such preciseness can be exceedingly difficult and expensive to achieve (with any propagation method) as discussed above. Hence, the formation of such gradients is a subject of some contention regarding accuracy, robustness, and cost. To minimize the effects of imprecise gradients and objective-function noise from the numerical operations in the propagation methods, so-called “trust-region surrogate-based” OUU methods (e.g., [7]) are being developed. It remains to be seen how the complexity, cost, and accuracy of these approaches compare to probabilistic ordinal approaches on industrial-type OUU problems.

6 Concluding Remarks

Several fundamental concepts of continuous-variable probabilistic ordinal optimization have been introduced and demonstrated in this report. These are: 1) *ordinal* selection or *ranking* of design alternatives based on statistical figures of merit of the probabilistic behaviors of the candidate designs; 2) the corresponding probability, P{CS}, of correctly selecting the best (or several best) design(s) among the alternatives considered; 3) Optimal Computing Budget Allocation, OCBA, for optimally efficient sample allocation for maximizing P{CS} given a fixed number of total samples to be distributed among the alternatives, or for minimizing the total number of samples required to attain a stipulated P{CS} level; and 4) tremendous efficiencies to be gained from exploiting spatial correlation of uncertainty in continuous-variable design problems.

Though only very simple and elementary implementations of ordinal optimization concepts have been demonstrated here, promising efficiency versus non-ordinal OUU approaches has been demonstrated on the two design-variable, two uncertain-variable probabilistic optimization problem solved here. Much more sophisticated and efficient implementational possibilities are foreseen. Certainly, there are many areas to be researched in the future involving (*i*) optimization constraints; (*ii*) advanced ordinal optimizers for searching and progressing in the design space; and (*iii*) increased sampling efficiencies from advanced sampling methods which exhibit significant variance reduction in calculated statistics (therefore faster resolution of P{CS}). For instance, Ref. [20] reviews some low-variance sampling methods in the course of introducing some promising new ones. Specialized uncertainty sampling schemes for small-probability values such as Importance Sampling (e.g., [8]) may also be applicable, as well as sampling schemes for epistemic uncertainties (e.g., [24]). The many possibilities have only just begun to be identified here.

Ultimately, efficient implementations of probabilistic ordinal optimization may prove to be a reference standard to which the efficiency and accuracy of other OUU approaches can be compared and evaluated. This is analogous to the position that Monte Carlo sampling holds among uncertainty propagation approaches.

Lastly, we note that this SAND report is a slight revision of Ref. [19], which does not have Section 5 of this report, but instead points to an analogous section in reference [18]. During the journal refereeing process for Ref. [19], we encountered and answered a reviewer's comment which has important implications that are likely of interest to other readers of this report. We paraphrase the reviewer as saying that an ordinal approach might not be able to solve a typical reliability-based design optimization problem such as: minimize design cost subject to reliability constraints (requirements) on design performance. While we are not certain that this cannot be done procedurally (above we list 'optimization constraints' as an area for future research), we do not believe that treating reliability as a constraint in model-based optimization is a realistic way to proceed anyway. In view of the magnitudes of model verification and validation type errors in typical computational physics models, it is concluded in Ref. [16] that reliability at small probability levels, say less than 10^{-3}, cannot be reasonably shown (validated) to be accurately estimable with models (please see footnote [5] on the following page for essential context regarding this statement). Nevertheless, it may be that *ordinal* or *ranking* accuracy is often sufficient that reliability predictions can be used in a relative sense (as is done in our application problem) to discriminate which of several neighboring design options has the lowest failure probability. Thus,

we believe OUU problems can be realistically posed in terms of maximizing reliability subject to e.g. cost constraints, but not in terms of minimizing cost subject to reliability constraints. We realize that this is likely a very controversial stance, but one that we have seen much evidence for and no empirical evidence against.

[5]However, a potential *satisficing* approximation to validation can be employed as discussed in "Some Issues in Quantification of Margins and Uncertainty (QMU) for Phenomenologically Complex Coupled Systems," paper AIAA-2006-1989 by V. Romero, 8th AIAA Non-Deterministic Approaches Conference, Newport, RI, May 1-4, 2006. Nonetheless, it is difficult to see how this subjective inverse "Info Gap" procedure might be packaged into an optimization formalism that adequately validates calculated reliability levels in a "reliability based" design optimization (RBDO) problem. Certainly, no such validation procedures seem to exist in the many RBDO frameworks and demonstrations that we have seen in the literature to date.

References

[1] Alexandrov, N., Dennis, J.E. Jr., Lewis, R.M., and Torczon, V., A trust region framework for managing the use of approximation models in optimization. *Structural Optimization*, 1998, **15**(1), 16–23.

[2] Ba-abbad, M.A., Kapania, R.K., and Nikolaidis, E., A new approach for system reliability-based design optimization. paper 2005-01-0348 in *Reliability and Robust Design in Automotive Engineering 2005*, proceedings SP-1956 of the SAE 2005 World Congress, Detroit, MI, April 11–15, 2005.

[3] Chen, C.H., Lin, J., Yucesan, E., and Chick, S.E., Simulation budget allocation for further enhancing the efficiency of ordinal optimization. *Journal of DEDS*, 2000, **10**(3), 251–270.

[4] Dai, L., Convergence properties of ordinal comparison in the simulation of discrete event dynamic systems. *Journal of Optimization Theory and Applications*, 1996, **91**(2), 363–388.

[5] Eldred, M.S., Outka, D.E., Bohnhoff, W.J., Witkowski, W.P., Romero, V.J., Ponslet, E.J., and Chen, K.S., Optimization of complex mechanics simulations with object-oriented software design. *Computer Modeling and Simulation in Engineering*, 1996, **1**(3), 323–352.

[6] Eldred, M.S., Hart, W.E., Bohnhoff, W.J., Romero, V.J., Hutchinson, S.A., and Salinger, A.G., Utilizing object-oriented design to build advanced optimization strategies with generic implementation. In *Proceedings of the Sixth Multi-Disciplinary Optimization Symposium*, Belleview, WA, Sept. 4–6, 1996.

[7] Eldred, M.S., Giunta, A.A., Wojtkiewicz, S.F., and Trucano, T.G., Formulations for surrogate-based optimization under uncertainty. Paper AIAA-2002-5585. In *Proceedings of the 9th AIAA/ISSMO Symposium of Multidisciplinary Analysis and Optimization*, Atlanta, GA, Sept. 4–6, 2002.

[8] Harbitz, A., Efficient and accurate probability of failure calculation by use of the importance sampling technique. In *Proceedings of the 4th International. Conference on Application of Statistics and Probability in Soils and Structural Engineering*, ICASP-4, Pitagora Editrice, Bologna, Italy, 1983.

[9] Hart, W. E., Evolutionary pattern search algorithms. Sandia National Laboratories report SAND95-2293, September 1995.

[10] Ho, Y.C., An explanation of ordinal optimization: Soft computing for hard problems. *Information Sciences*, 1999, 113, 169–192.

[11] Ho, Y.C., Screenivas, R., Vakili, P., Ordinal optimization of discrete event dynamical systems. *J. of DEDS*, 1992, **2**(2), 61–88.

[12] Hough, P.D., Kolda, T.G., and Torczon, V.J., Asynchronous parallel pattern search for nonlinear optimization. *SIAM J. Scientific Computing*, 2002, **23**(1), 134–156.

[13] Jones, D.R., Perttunen, C.D., and Stuckman, B.E., Lipschitzian optimization without the Lipschitz constant. *Journal of Optimization Theory and Applications*, 1993, **79**(1), 157–181.

[14] Rodriguez, J.F., Renaud, J.D., and Watson, L.T., Trust region augmented Lagrangian methods for sequential response surface approximation and optimization. ASME paper

#97-DETC/DAC-3773. In *Proceedings of the ASME Design Engineering Technical Conference*, ISBN 0-7918-1243-X, Sacramento, CA, Sept. 14–17, 1997.

[15] Romero, V. J., Efficient global optimization under conditions of noise and uncertainty – A multi-model multi-grid windowing approach. In *Proceedings of the 3rd WCSMO (World Congress of Structural and Multidisciplinary Optimization) Conference*, Amhearst, NY, May 17–21, 1999.

[16] Romero, V.J., Characterization, costing, and selection of uncertainty propagation methods for use with large computational physics models. Paper AIAA-2001-1653, *42nd Structures, Structural Dynamics, and Materials Conference*, Seattle, WA, April 16–19, 2001. Updated revision available from the author.

[17] Romero, V.J., Efficiencies from spatially correlated uncertainty and sampling in continuous-variable ordinal optimization. Submitted to *Engineering Optimization.*

[18] Romero, V.J., Ayon, D., Chen, C.H., Demonstration of probabilistic ordinal optimization concepts for continuous-variable optimization under uncertainty. *Optimization and Engineering* Vol. 7, No. 3, Sept. 2006, pp. 343-365.

[19] Romero, V.J., and Chen, C.H., A new adaptive ordinal approach to continuous-variable probabilistic optimization. Paper AIAA-2006-1826, *8th AIAA Non-Deterministic Analysis Conference*, Newport, RI, May 1–4, 2006. Accepted (2006) for *AIAA Journal.*

[20] Romero, V.J., Burkardt, J.S., Gunzburger, M.D., Peterson, J.S., Comparison of pure and "Latinized" centroidal Voronoi tessellation against various other statistical sampling methods. *Reliability Engineering and System Safety* 91 (2006) pp. 1266-1280.

[21] Romero, V.J., Eldred, M.S., Bohnhoff, W.J., and Outka, D.E., Application of optimization to the inverse problem of finding the worst-case heating configuration in a fire. In *Numerical Methods in Thermal Problems*, Vol. IX, Part 2, edited by R.W. Lewis and P. Durbetaki, 1022–1033 (Pineridge Press) (*Proceedings of the 9th Int'l. Conf. on Numerical Methods in Thermal Problems*, Atlanta, GA., July 17–21, 1995.)

[22] Spall, J.C., *Introduction to Stochastic Search and Optimization: Estimation, Simulation, and Control*. Wiley & Sons, 2003.

[23] Torczon, V., On the convergence of pattern search methods. *SIAM Journal of Optimization*, 1997, 7, 1–25.

[24] Yager, R.R., Fedrizzi, M., and Kacprzyk, J. (eds.), *Advances in the Dempster-Shafer Theory of Evidenc*e, Wiley & Sons, 1994.

DISTRIBUTION

1	MS0384	1500	A.C. Ratzel
1	MS0828	1530	A. L. Thornton
1	MS0828	1533	M. Pilch
1	MS0828	1533	A. A. Giunta
10	MS0828	1533	V. J. Romero
1	MS0828	1533	W. R. Witkowski
1	MS0847	1411	S. A. Mitchell
1	MS0847	1411	M. S. Eldred
1	MS0370	1411	L. P. Swiler
1	MS9159	8962	M. L. Martinez-Canales
2	MS9018	8944	Central Technical Files
2	MS0899	4536	Technical Library