Approaches for the Efficient Probabilistic Calculation of Surge Hazard

Gabriel R. Toro
Risk Engineering, Inc.
3 Farmers Row Acton, MA 01720 USA
toro@riskeng.com

Alan W. Niedoroda and Chris Reed
URS Group
1625 Summit Lake Drive, Tallahassee, FL 32317 USA

1 INTRODUCTION

The Joint-Probability Method or JPM (Myers, 1975; Ho and Meyers, 1975) has become the standard approach for the evaluation of surge inundation probabilities from hurricanes. The JPM provides a rigorous--yet practical--mathematical framework for combining the probabilistic characterization of hurricane occurrences and hurricane parameters in the region of interest with the results from numerical models that calculate the surge inundation generated by a hurricane of given characteristics, to calculate surge inundation probabilities. In essence, JPM considers all possible combinations of storm characteristics at landfall, calculates the surge effects for each combination, and then combines these results considering the combinations' associated probabilities. The result is the annual probability of exceeding any desired storm stage. Mathematically, this calculation is represented as a multi-dimensional integral (the JPM integral).

Because the JPM must consider the effects of many possible combinations of hurricane parameters, traditional implementations of the JPM have required numerical wind, wave, and surge calculations for many synthetic storms (typically several thousands). At the same time, these numerical calculations have become computationally intensive with the introduction of more accurate algorithms, the ability to incorporate detailed data describing bathymetry and topography at smaller spatial scales, and the need to incorporate wave set-up. As a result, it has been necessary to develop schemes that require a smaller number of artificial storms; these schemes have come to be called JPM-Optimal Sampling or JPM-OS.

Two JPM-OS schemes have been developed and applied in recent probabilistic hurricane-studies performed by teams led by URS and by USACE for the central Gulf of Mexico coast. We will refer to these schemes as the *quadrature* and *response-surface* JPM_OS schemes, respectively, The quadrature JPM-OS keeps the number of synthetic storms small by employing an algorithm, that selects the parameter combinations in an optimal manner and assigning the appropriate weight to each synthetic storm, transforming the JPM multi-dimensional integral into a weighted summation in the process. The response-surface JMP-OS, on the other hand, interpolates between the surge results obtained for a carefully selected set of synthetic storms. Both approaches take advantage of the smooth variation of the calculated surge η as the hurricane parameters are varied.

This paper begins with a description of the hurricane parameters used in most probabilistic surge studies, followed by a summary of the JPM formulation, as it is currently applied. This is followed by a description of the quadrature and response-surface JPM-OS approaches, as well as a discussion of other approaches that require further investigation and may prove useful in probabilistic surge calculations.

A forthcoming paper will provide detailed comparisons the results from the quadrature and response-surface approaches, when they are applied to multiple sites, using the same probabilistic storm characterization and the same numerical surge model.

2 HURRICANE PARAMETERS

For the purposes of the JPM formulation, we describe the storm as it approaches the coast in terms of the following parameters (see Figure 1): the pressure deficit ΔP (representing hurricane intensity), the radius of the exponential pressure profile[1] R_p (representing hurricane size), the forward velocity V_f, the storm heading[2] θ, and the landfall location (or, equivalently, the minimum distance from the track to a reference point along the coast). These parameters are illustrated in Figure 1. These parameters represent the main hurricane characteristics affecting storm surge; they are treated as random variables in the JPM formulation. Other storm characteristics, including parameter B (Holland, 1980), are treated as constant at landfall or are not considered explicitly. Although hurricanes are much more complex than this parameterization allows for, and substantially more information is available for well-studied recent hurricanes, it is necessary at present to utilize this simple storm parameterization for the probabilistic characterization of future storms. The differences between real hurricanes and this simple parameterization are not ignored in the JPM formulation employed in these studies. They are included in a statistical sense by means of the ε terms used in the JPM calculations, and then incorporated analytically into the results, as will be explained in Section 3.

The development of the probability distributions for these hurricane parameters and the calculation of the associated annual occurrence rate for the URS/FEMA Mississippi Coastal Flooding Hazard Project--including the data selection and the statistical analysis--is documented in URS (2007) and Risk Engineering (2007). It is also summarized in another paper in this volume (Niedoroda et al. 2007)..

3 THE JOINT PROBABILITY METHOD

The JPM formulation combines the following inputs:

(1) The annual rate of storms of interest λ. Typically, it is also assumed that the occurrences of

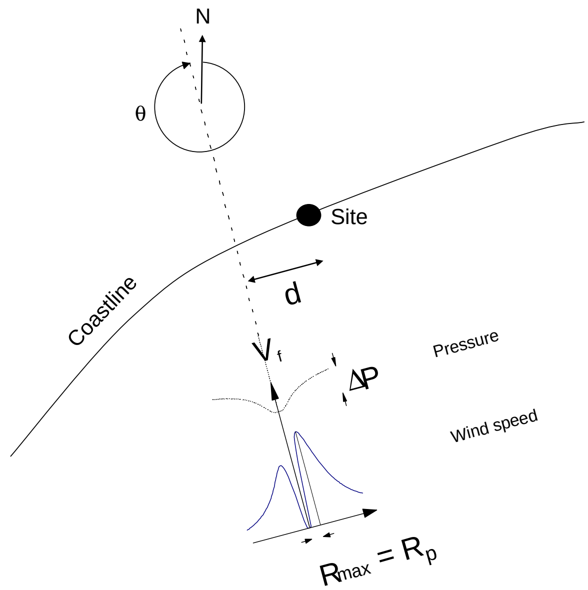

Figure 1. Characterization of storm as it approaches the coast.

these storms in time represents a Poisson random process (Parzen, 1962)3

(2) The joint probability distribution $f_{\underline{X}}(\underline{x})$ of the storm characteristics for storms of interest. These characteristics are defined very broadly at first and narrowed down later to make the approach practical.

(3) The storm-generated surge[4] $\eta(\underline{X})$ at the site of interest, given the storm characteristics.

The combined effect of these three inputs is expressed by the multiple integral

$$P[\eta_{\max(1\,yr)} > \eta] = \lambda \int \ldots \int_{\underline{x}} f_{\underline{X}}(\underline{x}) P[\eta(\underline{x}) > \eta] d\underline{x} \quad (1)$$

[1] In the URS/FEMA study, it is assumed that the radius of the exponential pressure profile R_p and the radius to maximum winds $R_{\max}$ (sometimes written as RMW) are identical or at least linked by a one-to-one relationship.

[2] Direction to, measured clockwise from North.

[3] In practice, the Poisson assumption is not necessary. Weaker assumptions are sufficient when calculating the probabilities of rare events, as will be discussed below.

[4] In this definition, the term surge represents the peak total inundation, including the surge itself, wave setup, astronomical tide, etc.

where $P[\eta(\underline{x}) > \eta]$ is the conditional probability that a storm of certain characteristics $\underline{x}$ will generate a flood elevation in excess of an arbitrary value η. This probability would be a Heaviside step function $H[\eta - \eta(\underline{X})]$ if vector $\underline{X}$ contained a complete characterization of the storm and if we had a perfect tool for the calculation of surge given $\underline{x}$, but these conditions cannot be satisfied in practice. The integral above considers all possible storm characteristics from the population of storms of interest and calculates the fraction of these storms that produce surges in excess of the value of interest η.

The right hand side in Equation 1 actually represents the mean annual rate of storms that exceed η at the site, but it also provides a good approximation to the annual exceedence probability[5].

Equation 1 defines a smooth function of η that can be used to determine the flood levels associated with any annual probability of exceedence. Typical values of interest include the 10%, 2%, 1% and 0.2 % annual probabilities. These are often referred to as the 10-, 50-, 100- and 500-yr annual exceedence levels, respectively. Unfortunately, the concept of return periods is often misunderstood.

As noted by Resio (2007), some approximations are necessary in practice for the evaluation of Equation 1. Firstly, it would be impossible to calculate the surge exactly, even if the storm's wind and pressure fields as a function of time were known exactly. To this effect, we write the actual elevation $\eta(\underline{X})$ in terms of the model-calculated elevation $\eta_m(\underline{X})$ as $\eta(\underline{X}) = \eta_m(\underline{X}) + \varepsilon_m$, where ε_m is a modeling-error term, which will be treated as a random quantity uncorrelated with $\underline{X}$. If the numerical surge model is unbiased, ε_m has a mean value of zero. Using the above representation, one can write the actual conditional probability as $P[\eta(\underline{x}) > \eta]$ as

$$P[\eta(\underline{X}) > \eta] = P[\eta_m(\underline{X}) + \varepsilon_m > \eta] \qquad (2)$$

In addition, it would be impossible to provide a complete characterization of the storm (i.e., the wind and pressure fields as a function of time). Thus, it is convenient to partition the vector of storm characteristics $\underline{X}$ into two parts, as follows: (1) a vector of principal quantities $\underline{X}_1 = (\Delta P, Rp, V_{f,} \text{landfall location}, \theta)$, whose probability distributions are represented explicitly and whose effects are also represented explicitly in the model calculations, and (2) a vector of secondary quantities $\underline{X}_2 = (B, \text{tide}...)$, whose distributions (relative to their base-case values) and effects are jointly represented in an approximate manner by random terms $(\varepsilon_B, \varepsilon_{\text{tide}}, ...)$ (which have units of elevation). These secondary quantities are ignored or set to their base-case values in the numerical surge calculations, but their effects are carried by their corresponding ε terms. Although these epsilons are conceptually different from the modeling error ε_m introduced earlier, they are combined operationally into one random quantity $\varepsilon = \varepsilon_m + \varepsilon_B + \varepsilon_{tide} + ...$.

Incorporating these simplification, Equation 1 transforms into

$$P[\eta_{\max(1\,yr)} > \eta] = \lambda \int ... \int_{\underline{x}_1} f_{\underline{X}_1}(\underline{x}_1) P[\eta_m(\underline{x}_1) + \varepsilon > \eta] \, d\underline{x}_1 \qquad (3)$$

where $\underline{X}_1 = (\Delta P, Rp, V_{f,} \text{landfall location}, \theta)$. The subscript 1 [as in X_1] will be dropped in the remainder of this paper for the sake of simplicity, resulting in

$$P[\eta_{\max(1\,yr)} > \eta] = \lambda \int ... \int_{\underline{x}} f_{\underline{X}}(\underline{x}) P[\eta_m(\underline{x}) + \varepsilon > \eta] \, d\underline{x} \qquad (4)$$

The quantification of the standard deviations for the various components of ε is not described here (see URS, 2007, or Niedoroda et al. 2007, this volume).

[5] The derivation to show that the annual probability for a rare event is approximately equal to annual rate is usually made by assuming that event occurrences represent a Poisson process and then linearizing the resulting exponential expression. The same result may be obtained under weaker assumptions. It is sufficient to assume that the probability of two or more of these rare events in one year is much lower than the probability of one event. This condition is satisfied for hurricane-generated surges and for the exceedence probabilities of interest in these studies (e.g, 5% per year or lower).

This quantification is done using a variety of approaches, such as modeling followed by uncertainty propagation, comparisons of observed surges to surges obtained using "best winds", and comparisons of surges obtained using the limited parameterization employed in the production JPM calculations to the surges obtained using the best winds. By integrating over the distributions of these ε's, we are incorporating (in an approximate manner) the effects of limitations in the numerical surge models and in our parameterization of hurricanes.

4 THE QUADRATURE JPM-OS APPROACH

As indicated earlier, evaluation of the JPM integral (Equation 4) using conventional numerical-integration approaches is impractical for the following two reasons: (1) each evaluation of the integrand involves evaluation of $\eta_m(\underline{x})$ for one value of $\underline{x}$ (i.e., one synthetic storm), which requires computationally intensive numerical calculations of wind, waves, surge, wave setup, etc.; and, (2) numerical evaluation of the 5-dimensional integral in Equation 4 using conventional approaches requires that the integrand be evaluated a large number of times (this is the so-called curse of dimensionality).

The quadrature JPM-OS approach approximates the integral in Equation 4 as a weighted summation, i.e.:

$$
\begin{aligned}
&P[\eta_{\max(1\,yr)} > \eta] = \\
&\lambda \int \ldots \int_{\underline{x}} f_{\underline{X}}(\underline{x}) P[\eta_m(\underline{x}) + \varepsilon > \eta]\, d\underline{x} \approx \\
&\sum_{i=1}^{n} \lambda_i P[\eta_m(\underline{x}_i) + \varepsilon > \eta]
\end{aligned}
\qquad (5)
$$

where each $\underline{x}_i = (\Delta P_i, Rp_i, V_{f,i,} \text{landfall location}_i, \theta_i)$ may be interpreted as a synthetic storm[6], $\lambda_i = \lambda p_i$ may be interpreted as the annual occurrence rate for that storm, and $\eta_m(\underline{x}_i)$ may be interpreted as the numerical-model's estimates of the storm elevation generated by that storm. For this approach to be practical, one must be able to specify the storm characteristics $\underline{x}_i$ and their rates λ_i so that the integral can be approximated with sufficient accuracy (for all η values of interest), using a reasonably small value of n (i.e., a reasonably small number of synthetic storms and corresponding numerical-model runs).

The approach used to define the synthetic storms and their rates uses a combination of well-known and more sophisticated techniques and may be summarized by the following three steps:

1. Discretize the distribution of ΔP into broad slices. In the URS/FEMA work, we use three slices, roughly corresponding to hurricane Categories 3, 4, and 5.

2. Within each ΔP slice, discretize the joint probability distribution of $\Delta P(\mathit{within\ slice})$, Rp, Vf, and θ using the multi-dimensional optimal-sampling procedure known as Bayesian Quadrature (Diaconis, 1988; O'Hagan, 1999; Minka, 2000; see Section 4.1). This procedure represents the response portion of the integrand (i.e., the term $P[\eta_m(\underline{x}) + \varepsilon > \eta]$) as a random function of $\underline{x}$ with certain correlation properties, and determines the optimal values of $\Delta P_i, Rp_i, V_{f,i,} \theta_i$, and the associated probability, so that the variance of the integration error is minimized. The correlation properties of the random function (which take the form of correlation distances) depend on how sensitive the response is to each variable (shorter correlation distances for the more important variables). These correlation distances were set based on judgment and on the results of the sensitivity tests described in URS (2007)

3. Discretize the distribution of landfall location by offsetting each of the synthetic storms defined in the previous two steps. The offset is equal to R_p (measured perpendicular to the storm track). Sensitivity studies indicated that a spacing of R_p is small enough to capture the peak surge at all grid locations.

Finally, one computes the probability p_i assigned to each synthetic storm as the product of the probabilities resulting from the three steps. Equivalently, one can compute the rate λ_i assigned to

[6] More precisely, $(\Delta P_i, Rp_i, V_{f,i,} \text{landfall location}_i, \theta_i)$ represent the characteristics of the synthetic storm at landfall.

each synthetic storm as the product of the probabilities from the first two steps times the rate per unit length from Section 3 times the storm spacing.

It is useful to discuss some possible variations to this scheme, which may apply to other situations.

- If one were performing calculations for a single site, one could treat location (or distance to the site) as one of the quantities in the Bayesian Quadrature (step 2). To improve the efficiency of this scheme, one could use *importance sampling* (e.g., Melchers, 1999) to sample more heavily at distances near R_p on the strong side of the storm. When performing calculations for a large number of sites, it is considered more convenient to sample distance using constant spacing (step 3 above).

- One could include ε as one of the random quantities in the Bayesian Quadrature, instead of treating it as part of the effects term. This change would have two detrimental effects on the efficiency of the Bayesian Quadrature scheme, as follows: (1) the number of dimensions is increased, and (2) the integrand becomes less smooth. On the other hand, one can think of situations where this change may be required (e.g., if the hydrographs $\eta(t)$ for the synthetic storms are to be used as inputs to a computationally intensive calculation that is nonlinear in $\eta(t)$).

4.1 BAYESIAN QUADRATURE APPROACH

4.1.1 BACKGROUND

The word *Quadrature* is often used to denote numerical techniques to approximate an integral of the form

$$I = \int_A f(x)p(x)dx \qquad (6)$$

over some domain A, as a weighted sum of the form

$$I \approx \sum_{i=1}^{n} w_i \, p(x_i) \qquad (7)$$

where $f(x)$ is typically a probability density function (i.e., it is positive and it integrates to unity) and $p(x)$ represents a function belonging to a certain family of functions.

In our case, *A* represents a four-dimensional domain, $f(x)$ represents the joint probability distribution of storm characteristics, and $p(x) = P[\eta_m(x) + \varepsilon > \eta]$ represents the "surge effects" portion of the JPM integral (for many possible locations and for many possible values of η).

The design of a quadrature involves specification of the number of nodes *n*, and selection of the node values x_i and associated weights w_i. These quantities depend on the functional form of $f(x)$ and of the $p(x)$ family, and on the desired characteristics (e.g., accuracy) of the approximation. In our case, each node becomes one synthetic storm.

Classical (Gaussian) quadrature chooses the number of nodes, node values, and weights so that the summation will integrate the function exactly if $p(x)$ is a polynomial of a certain degree and $f(x)$ is a particular function (e.g., a standard normal probability density). This technique is used frequently in one dimension (see Miller and Rice, 1983 for details and for results for a variety of probability distributions)[7]. Extension of these zero-error rules to more than one dimension is problematic. The number of required nodes increases rapidly with the number of dimensions (see Minka, 2000). Furthermore, some of the weights often become negative (see Genz and Keister, 1996), which leads to less stable results and makes it impossible to interpret the weights as probabilities.

Bayesian quadrature, in contrast, considers a much broader probabilistically defined family of functions (i.e., random functions with a certain correlation structure), and minimizes the integration error in a mean-squared sense. Conceptually, it is straightforward to extend Bayesian quadrature to multiple dimensions, although it becomes somewhat computationally demanding for more than six or seven dimension. According to Diaconis (1988), the

[7] Classical Quadrature is used later in this paper to construct an accurate, but inefficient, reference-case JPM formulation (which we call "JPM-Heavy"). In this formulation, Classical Quadrature will be used to represent the probability distribution of each individual storm characteristic.

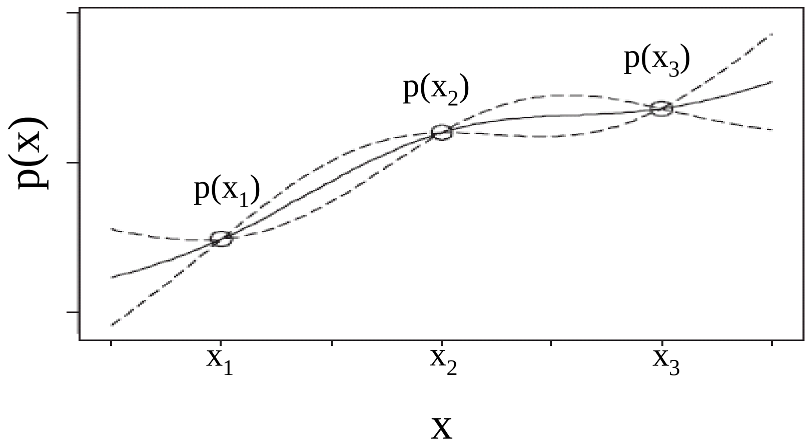

Figure 2. Illustration of the conditional distribution of random function $p(x)$ at intermediate points between sampling nodes. The function $p(x)$ has been sampled at 3 nodes x_1, x_2, x_3. The solid line displays the conditional mean value. The dashed lines display the conditional mean $\pm$ standard deviation range; the width of this range depends on the distance to the nodes and on the autocovariance function $k(x, y)$.

approach dates back to the work of Poincare in 1896. It is also closely related to the technique known as *Kriging* (e.g., Journel and Huijbregts, 1978), and even to least-squares regression.

4.1.2 DERIVATION

The development of Bayesian Quadrature begins by idealizing the function $p(x)$ in Equation 6 as a Gaussian random process[8] in *m*-dimensional space with mean zero and with autocovariance function $k(x, y) = E[p(x)p(y)]$ [9] [10], where $E[\cdot]$ denotes mathematical expectation (i.e., $E[z]$ the average value of quantity z over all possible realizations of $p(x)$), and *x* and *y* are any two arbitrary values of *x*. The autocovariance function contains information about the degree of continuity or smoothness of realizations of $p(x)$, at both small and large scales (e.g., Vanmarcke, 1983).

Let $D = [x_1, x_2, ..., x_n]$ denote *n* nodes[11] for which $p(x)$ has been evaluated, so that we know the values of $p(D) = [p(x_1), p(x_2), ..., p(x_n)]$. Because we know the value of $p(x)$ at these *n* points, we also know the conditional mean and the conditional variance of $p(x)$ at all other values of *x*. This is illustrated in Figure 2. The mean $E[p(x) | p(D)]$ is a linear combination of the known nodal values $p(D) = [p(x_1), p(x_2), ..., p(x_n)]$, with coefficients that depend on the values of the covariance function $k(x, y)$ between *x* and each nodal point and between each pair of nodal points. The conditional variance $Var[p(x) | p(D)]$ depends on the values of the covariance function $k(x, y)$ only.

[8] The assumption of the process being Gaussian is not strictly necessary. One may obtain the same results using weaker assumptions.

[9] *m* is the number of dimensions for the integral in Equation 6. In this derivation, *x* and *y* are m-dimensional vectors, but we will not use underline or boldface for them for the sake of simplicity

[10] Details on the functional form of $k(x, y)$ will be considered later. The only requirement at this stage is that the required integrals involving $k(x, y)$ and $f(x)$ do not diverge.

[11] In our *m*-dimensional integration space, each node location x_i is a vector containing *m* nodal coordinates.

O'Hagan (1991) proposes an approximation to Equation 6 in which one uses the conditional mean $E[p(x) \mid p(D)]$ in place of $p(x)$ (whose value we know at only a few points), i.e.,

$$I = \int_A f(x)p(x)dx \approx \int_A f(x)E[p(x) \mid p(D)]dx \qquad (8)$$

Because the mean $E[p(x) \mid p(D)]$ is a linear combination of the nodal values $p(D) = [p(x_1), p(x_2),..., p(x_n)]$, the above approximation is also a linear combination of the nodal values, which means that it has the same functional form of Equation 7 (i.e., a weighted sum of the values of $p(x)$ at the nodes), with weights that depend on $k(x, y)$ and on integrals involving $k(x, y)$ and $f(x)$ (namely, $W^T = [w_1, w_2,..., w_n] = U(x,D)K(D,D)^{-1}$), where

$$U(D) = \begin{bmatrix} \int_A k(x, x_1) f(x)dx \\ \int_A k(x, x_2) f(x)dx \\ \vdots \\ \int_A k(x, x_n) f(x)dx \end{bmatrix} \qquad (9)$$

and

$$K(D,D) = \begin{bmatrix} k(x_1,x_1) & k(x_1,x_2) & \dots & k(x_1,x_n) \\ k(x_2,x_1) & k(x_2,x_2) & \cdots & k(x_2,x_n) \\ \vdots & \vdots & \ddots & \vdots \\ k(x_n,x_1) & k(x_n,x_2) & \cdots & k(x_n,x_n) \end{bmatrix} \qquad (10)$$

Minka (2000) shows that one can arrive at the same weights with a least-squares formulation; i.e., by determining the weights that minimize the variance—over all possible realizations of $p(x) \mid p(D)$ -- of the difference between the exact integral and the approximation given by Equation 7.

The resulting weights from this approach may or may not add to unity, depending on the choice of $k(x, y)$ and of *D*. To ensure that the weights always add to unity—which is required because we want to interpret $p(D) = [p(x_1), p(x_2),..., p(x_n)]$ and the weights as an *m*-dimensional discrete probability distribution—we introduce the constraint $\sum_{i=1}^{n} w_i = 1$ into Minka's (2000) least-squares representation of the problem. We do this by means of a Lagrange multiplier, obtaining the following:

$$\left[\begin{array}{cccc|c} k(x_1,x_1) & k(x_1,x_2) & \dots & k(x_1,x_n) & 1 \\ k(x_2,x_1) & k(x_2,x_2) & \cdots & k(x_2,x_n) & 1 \\ \vdots & \vdots & \ddots & \vdots & 1 \\ k(x_n,x_1) & k(x_n,x_2) & \cdots & k(x_n,x_n) & 1 \\ \hline 1 & 1 & \cdots & 1 & 0 \end{array}\right] \left[\begin{array}{c} w_1 \\ w_2 \\ \vdots \\ w_n \\ \hline -\mu \end{array}\right] = \left[\begin{array}{c} U(D) \\ \hline 1 \end{array}\right] \qquad (11)$$

where μ is the Lagrange multiplier (see, e.g., Journel and Huijbregts, 1978). We solve this system of linear equations to obtain weights that add to unity.

The associated estimation variance is given by the expression

$$Var\left[\int_A f(x)p(x)dx - \sum_{i=1}^{n} w_i\, p(x_i)\right] = u - W^T U(D) + \mu \qquad (12)$$

where $W^T = [w_1, w_2,..., w_n]$ is the vector of weights obtained by solving Eq. 11 and

$$u = \int_A \int_A k(x, y) f(x) f(y)dxdy \qquad (13)$$

It may also be possible to force the weights to sum to unity by assuming that $p(x)$ has an unknown (but generally non-zero) mean.

So far in this discussion, we have treated the node locations $D = [x_1, x_2,..., x_n]$ as known. The development of an efficient quadrature rule requires finding the optimal node locations *D* that minimize the variance in Equation 12 [12], for a pre-specified

[12] Note that now we have two nested optimizations, both of which seek to minimize the variance of the integration error. At the inner level of nesting, there is the optimization to determine the best weights (for given nodal locations *D*). This is done analytically, by solving Equation 11. At the outer level, there is the search for the best set of nodal locations *D*. This

value of n. Note that each x_i in D represents the coordinates of a node in m-dimensional space. Thus, determination of the optimal D is an $(m \times n)$-dimensional optimization problem. We provide details on the algorithm used in Section 4.2.4

We have not shown that all the weights are non-negative, and in fact negative weights do arise when the nodal vector $D = [x_1, x_2, ..., x_n]$ is specified arbitrarily (i.e., without optimization). On the other hand, it is reasonable to expect (and one may be able to prove) that optimization of $D = [x_1, x_2, ..., x_n]$ so as to minimize the variance in Equation 12, forces all the weights to be positive. This argument is related to the concept of quadrature stability factor (i.e., the sum of the absolute values of the weights) employed by Genz and Keister (1996). In our practical applications, the approach introduced here has always led to positive weights when D is optimized. Negative weights, if they happen to arise, may be eliminated by employing a numerical optimization scheme with inequality constraints on the weights.

4.2 IMPLEMENTATION OF BAYESIAN QUADRATURE FOR JPM-OS

This portion of the paper provides a number of details on how Bayesian Quadrature procedure is implemented as part of the Quadrature JPM-OS formulation used in the URS/FEMA study. These details were left out of the derivation above, for the sake of generality and simplicity.

4.2.1 PROBABILITY DISTRIBUTION: CHOICE FOR $f(x)$

The Quadrature JPM-OS formulation must be flexible enough to accommodate the probability distributions typically used for $\Delta P, R_p$, etc. On the other hand, implementation of the Bayesian Quadrature formulation described above using a general form for $f(x)$ would require the repeated evaluation of the integrals in Equations 9 and 13, which are likely more complex than the JPM integrals we are trying to solve in the first place.

Instead, we formulate the problem in m-dimensional standard normal-distribution space, determine the coordinates of the integration nodes in that space, and then convert these nodal coordinates to the "physical" space of $\Delta P, \overline{R}p$, etc., in a manner that takes into account their joint probability distribution. This approach for the transformation of multivariate probability distributions is commonly used in structural reliability theory (e.g., Madsen et al., 1986; Melchers, 1999) and is built into structural reliability software (e.g., Gollwitzer et al., 2006).

One can also achieve the distribution transformations by altering the weights, or by using a combination of both approaches.

4.2.2 CORRELATION STRUCTURE OF $p(x)$: THE CHOICE FOR $k(x, y)$ AND THE SPECIFICATION OF CORRELATION DISTANCES

We use the correlation structure of $p(x)$, as represented by the covariance function $k(x, y)$, to specify the importance of the corresponding physical quantity in the surge calculations. If the physical quantity corresponding to the j-th component of x is important, correlation decays faster in that direction than in the direction corresponding to a less important quantity.

One of the consequences of formulating and solving the Bayesian Quadrature problem in normal space is that we must also define $k(x, y)$ in normal space. Also, we must choose a functional form that facilitates analytical evaluation of the integrals in Equations 9 and 13.

We choose the double-exponential functional form for the covariance function, i.e.,

$$k(x, y) = E[p(x)p(y)] = \sigma^2 \prod_{j=1}^{m} \exp\left[-\left(\frac{x_j - y_j}{c_j} \right)^2 \right] \qquad (14)$$

where c_j controls how quickly the correlation decays in the direction of a particular component [13] [14] [15]. c_j

is done numerically (more details will be provided in the next section).

[13] In this section, the subscripts denote the coordinates of one point in m-dimensional space (j=1,…m); x and y denote two points in that m-dimensional space.

is related to the corresponding correlation distance or scale of fluctuation d_j (Vanmarcke, 1983) by the relation $d_j = \sqrt{\pi} c_j$. Because the algorithm operates in standard normal space, c_j and d_j have no physical units.

One of the critical steps in the Quadrature JPM-OS analysis is the specification of the correlation distances d_j associated with the various hurricane characteristics. This is made more difficult because these correlation distances are specified in normal space, not in physical space. The following discussion provides some guidance to facilitate this step.

In a relative sense, the Quadrature JPM-OS algorithm tends to spread the sampling nodes more faithfully along those directions for which $p(x)$ has lower correlation distances, providing a closer match to the marginal probability distributions in those directions. Thus, it is important to specify correlation distances that relate to the importance of the various physical quantities, in order to obtain an optimal allocation of effort among the various directions.

In an absolute sense, numerical experiments in one dimension show that low values of the correlation distance cause the algorithm to be more cautious and tend towards equal weights, while high values provide a wide range of weights, approaching those obtained by Gaussian quadrature. The ideal choice is between these two extremes.

The following values provide preliminary guidance for the choice of correlation distances:

- Sensitive (important): d_j =1 to 3
- Insensitive (unimportant): d_j = 4 to 6

In principle, one could calculate appropriate values for these correlation distances analytically, using as inputs the results from sensitivity runs such as those documented in URS (2007) or the predictions of parametric models (e.g., Irish et al., 2007). This calculation should take into account the distribution transformation described in Section 4.2.1. This calculation of correlation distances has not been done to date, relying instead on choices made on the basis of judgment and on verification against reference JPM results, using an inexpensive numerical surge model.

4.2.3 POSSIBLE REFINEMENTS IN DISTRIBUTION SHAPE AND CORRELATION STRUCTURE

It is also possible to construct $f(x)$ as a mixed product of probability-distribution shapes (e.g., normal in some directions, uniform in others, exponential an possibly Weibull in others), chosen in such a way that these distribution shapes are closer to the distribution of the corresponding physical quantities. This will require a somewhat different functional form for the autocovariance function in the non-normal directions, in order to permit analytical evaluation of the required integrals. The effect of this refinement is anticipated to be better performance of the Bayesian quadrature for distribution shapes such as the Weibull, possibly eliminating the need for pre-slicing of the distribution of ΔP.

4.2.4 OPTIMIZATION ALGORITHM

As was indicated earlier, determination of n optimal sampling points in m dimensions constitutes an ($m \times n$)-dimensional optimization problem. We perform this optimization using an algorithm developed by Powell (2004), which does not require derivatives. We choose the starting points at random. Convergence is fast for the number of dimensions and nodes considered in this paper.

4.2.5 NEED FOR VALIDATION

Because we make a number of assumptions regarding the functional form and parameters of the autocovariance function $k(x, y)$ of $p(x) = P[\eta_m(x) + \varepsilon > \eta]$, it is important to validate the Quadrature JPM-OS scheme (i.e., the number of nodes and the correlation distances).

In the URS/FEMA study, we validate the JPM-OS scheme by comparing the cumulative distributions of surge obtained using the JPM-OS scheme and a Base Case or reference JPM scheme (denoted JPM-Heavy in this paper and JPM Gold Standard in URS, 2007), using the SLOSH software to compute the surge in

---

[14] This correlation model implies that the random field is homogeneous and twice differentiable (in a second-order sense).

[15] The variance σ^2 cancels out in the results of section 4.1.2 and will be omitted in the material that follows.

both schemes. These comparisons are shown later, in Section 4.2.6.2. One can also perform this validation using a parametric surge model (e.g., Irish et al., 2007).

4.2.6 APPLICATION TO HURRICANES IN MISSISSIPPI

This section documents the application of the Quadrature JPM-OS approach to hurricanes affecting the Mississippi coast. The example presented here considers hurricanes with $\Delta P > 45$ mb (what we call the greater storms in URS, 2007 and Risk Engineering, 2007)[16].

4.2.6.1 SELECTED JPM-OS SCHEME

This scheme (which we call the JPM-OS 6 scheme), is the scheme adopted for the final calculations, after experimenting with a number of other schemes having different numbers of nodes and somewhat different correlation distances. Table 1 shows the various slices of the ΔP distribution, their probabilities, and the number of nodes used in the Bayesian-Quadrature discretization for each slice. Table 2 shows the correlation distances used in the Bayesian-Quadrature procedure. These values were chosen based on the extensive sensitivity results presented in URS (2007), and then refined so that they preserve the marginal moments of the most important quantities.

Table 1. Discretization of ΔP into Slices in JPM-OS 6 Scheme for Greater Storms

Slice	Cat .3	Cat. 4	Cat. 5
ΔP range (mb)	45-70	70-95	95-135
Probability	0.657	0.261	0.082
#of points in Bayesian Quadrature	5	7	7

Table 2. Correlation Distances in JPM-OS 6 Scheme for Greater Storms

Correlation Distance (std normal units)			
DeltaP (within slice)	Rp	Vf	Heading
4	2.5	6	5

[16] The results shown here were obtained using the preliminary rates and probability-distribution parameters. These parameters slightly different from the final values documented in Section 3 of Risk Engineering (2007).

Figure 3 provides an illustration of the parameters of the resulting synthetic storms (for one landfall location location; i.e., prior to the distance-offsetting step in Section 4). Each chart on the main diagonal shows the probability distribution of the corresponding quantity (in the form of a histogram), as represented in the JPM-OS 6 discretization. Each off-diagonal scatter diagram shows how each pair of quantities (e.g., ΔP and R_p) are jointly distributed in the JPM-OS 6 scheme, with the areas of the circles being proportional to the associated annual rate.

4.2.6.2 VALIDATION OF JPM-OS SCHEME

As indicated earlier, we validate the JPM-OS 6 scheme by comparing the cumulative distributions of surge obtained using the more efficient JPM-OS 6 scheme and a reference-case JPM scheme (denoted JPM-Heavy), using the SLOSH software to compute the surge in both schemes. This validation is feasible because SLOSH surge computer runs are relatively fast.

The parameter values and annual rates for the JPM-Heavy synthetic storms were defined as follows. The values and probabilities for each parameter were determined using the one-dimensional quadrature approach described by Miller and Rice (1983), using different numbers of quadrature points according to the importance of the parameter. For R_p (which depends on ΔP), values were drawn from the conditional distribution of radius given ΔP). In addition, each of the 360 combinations of $\Delta P - Rp - \theta - V_f$ was assigned to multiple parallel tracks with a perpendicular spacing of R_p, as described in step 3 of Section 4. For each of the resulting 2,967 synthetic storms (i.e., for each combination of parameters and track), the storm event rate was obtained by multiplying the probability associated with each of the storm parameter values (i.e., $p(\Delta P_i)\times p(R_{p,i} \mid \Delta P_i)\times p(V_{f,i})\times p(\theta_i)$) and then multiplying that result by the storm spacing (which is equal to R_p) times the annual rate of $\Delta P > 45$ mb storms occurring in the Gulf coast area. This rate was previously determined to be 3.02E-4

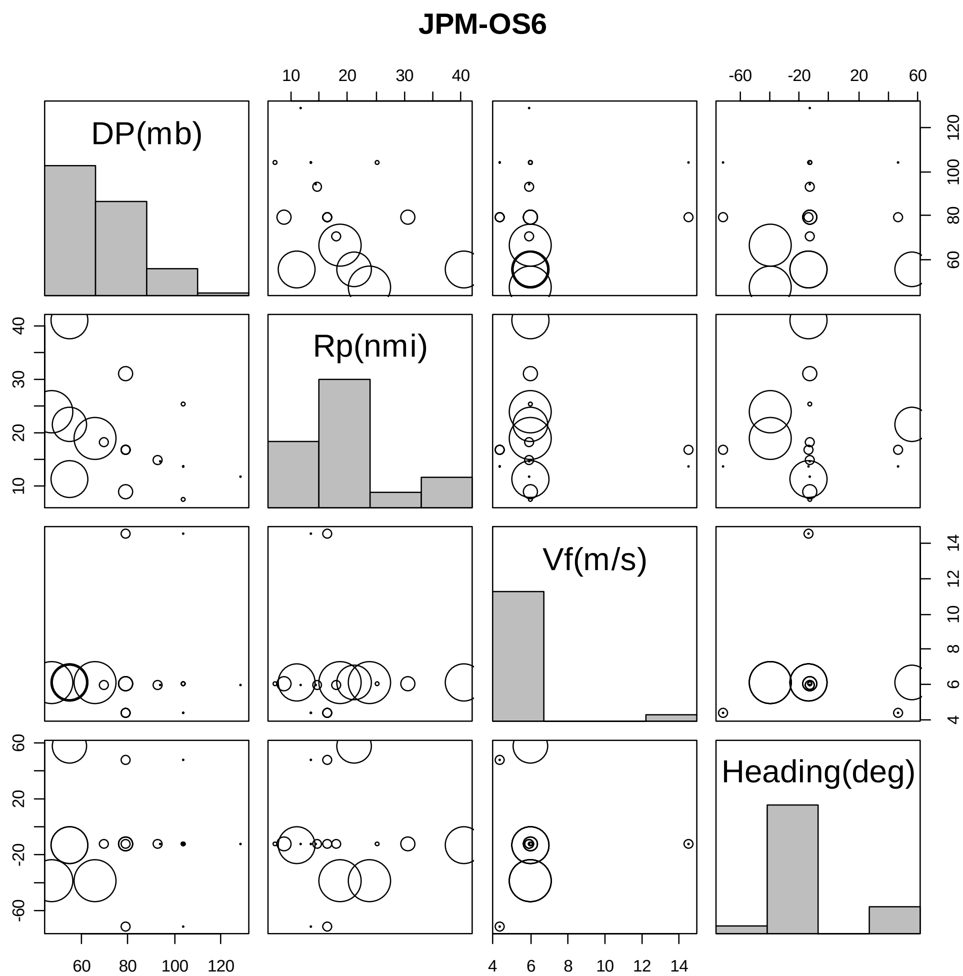

Figure 3. Graphical representation of the JPM-OS 6 scheme for one landfall location. The areas of the circles are proportional to the associated annual rate.

storms/km/yr. Table 3 shows the quadrature nodes and associated probabilities for the various hurricane parameters.

Figure 4 provides an illustration of the parameters of the resulting synthetic JPM-Heavy storms (for one landfall location), in a manner similar to Figure 3. Notice that some of the nodes receive very low weights (i.e., points that are barely visible) because they correspond to combinations of unlikely parameter values and it would be wasteful to spend significant computer resources evaluating the surge for each of these low-weight combinations. This illustrates the inefficiencies that may occur when evaluating the JPM integral using conventional approaches.

For the validation of JPM-OS 6, we calculate surges at 147 test points scattered throughout coastal Mississippi and adjacent portions of Louisiana using both the JPM-OS6 and JPM-Heavy synthetic storms, calculate the associated cumulative distributions of surge elevation for each storm set, and compare the associated 100-year surge. Figure 6 compares the complementary cumulative distribution functions obtained with the JPM-OS6 and JPM-Heavy for site 42, which is located along the shore at Biloxi, MS. The 100-year surge corresponds to an annual exceedence probability of 0.01, for which the difference between the two curves is less then 1 foot.

Table 3. Parameter discretizations for JPM-Heavy Scheme

ΔP (mb)	45.6	48.6	56.1	69.1	87.6	111.8
Probability	0.0496	0.1661	0.2844	0.2844	0.1661	0.0496

Rp (nmi) for ΔP =45.6 mb	6.94	13.43	24.38	44.28	85.71
Rp (nmi) for ΔP =48.6 mb	6.63	12.84	23.32	42.34	81.96
Rp (nmi) for ΔP =56.1 mb	5.99	11.59	21.05	38.23	74.00
Rp (nmi) for ΔP =69.1 mb	5.16	10.00	18.15	32.96	63.80
Rp (nmi) for ΔP =87.6 mb	4.36	8.44	15.33	27.84	53.88
Rp (nmi) for ΔP =111.8 mb	3.67	7.10	12.89	23.41	45.31
Probability	0.01	0.22	0.53	0.22	0.01

Heading (θ)	-73.0	-32.7	7.3	49.4
Probability	0.133	0.367	0.367	0.133

Fwd. Velocity (m/s)	2.99	6.04	12.23
Probability	0.1667	0.6666	0.1667

Note also the granularity of the cumulative distributions. Fine granularity is generally an indication that the JPM-OS scheme is adequate in the corresponding portion of the curve.

Figure 5 compares the 100-year results at the 147 test points. Note that the error in the JPM-OS6 scheme is less than 1 foot for most points. The few instances of large over-estimation (one of which is off-scale in the bottom panel of Figure 5) occur at inland or riverine points located more than 20 km away from the coast, near the edge of the of the 100-year inundation area. The instances of mild under-estimation seem to occur in areas where the coastline is complicated, such as Biloxi Bay near -88.8 longitude.

This validation was performed prior to integration over ε. The match between JPM-OS6 and JPM-Heavy would be much closer if the comparison had been performed after integration over ε because this step makes $p(x)$ smoother.

4.3 OTHER MULTI-DIMENSIONAL QUADRATURE FORMULATIONS

The simplest way to construct a multi-dimensional quadrature scheme is to apply a Classical quadrature to discretize each hurricane parameter and then generate all possible combinations of parameters. The probability assigned to each combination is the product of the probabilities for the corresponding discrete parameter values. If the parameters are not independent, this is done in a sequential manner. This approach, which is called a "product-rule" quadrature, is precisely what we used above to construct the JPM-Heavy scheme.

Unfortunately, this approach generates a very large number of parameter combinations as the number of dimensions increases. Also, some of the combinations receive very low relative probabilities, suggesting that many these combinations do not warrant separate modeling runs. Smolyka (1966) has developed a more efficient procedure to extend classical quadratures to multiple dimensions, but we have not investigated this approach.

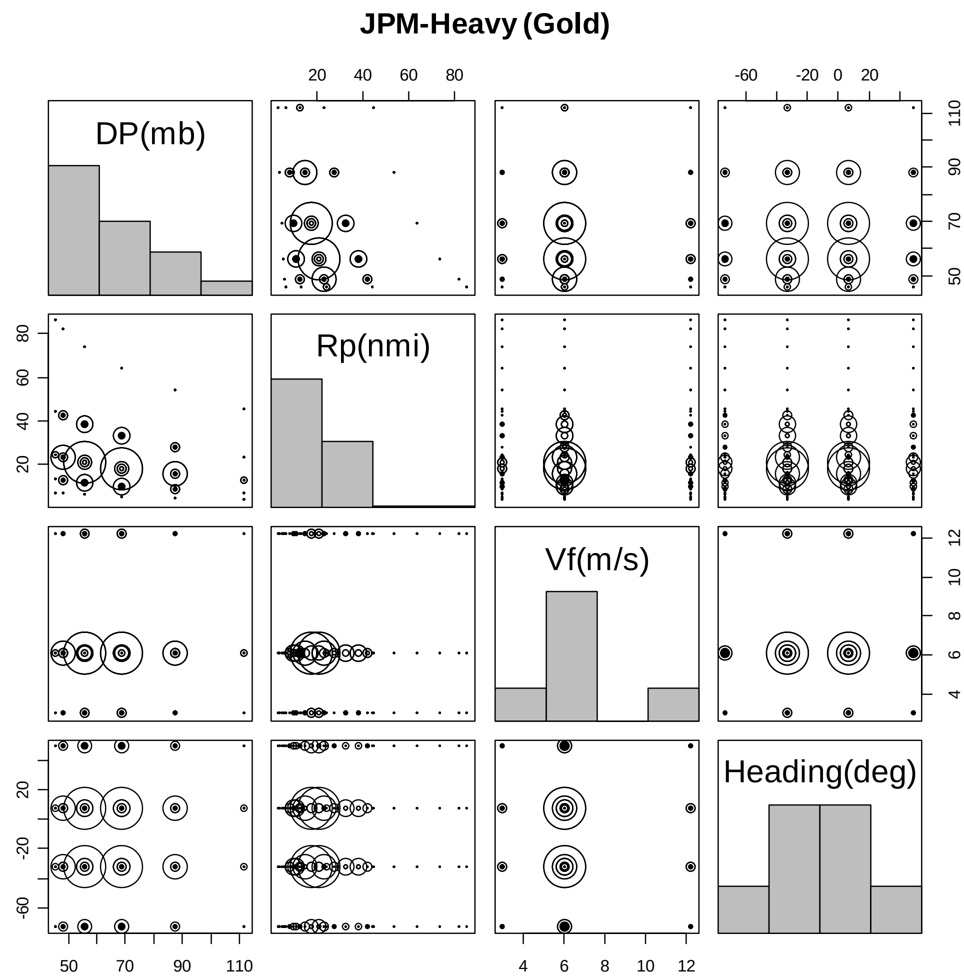

Figure 4. Graphical representation of the JPM-Heavy scheme for one landfall location. The areas of the circles are proportional to the associated annual rate.

Hong (1998) developed an approach that requires 2m+1 nodes to approximate an integral in m dimensions, while conserving moments up to order 3. Unfortunately, the weight for the central node tends to become negative as the number of dimensions increases. Negative weights preclude interpretation of the nodes and their associated weights as a discrete probability distribution leading to a representative set of synthetic storms, are usually associated with higher error variances, and may lead to other difficulties such as negative values of cumulative distributions.

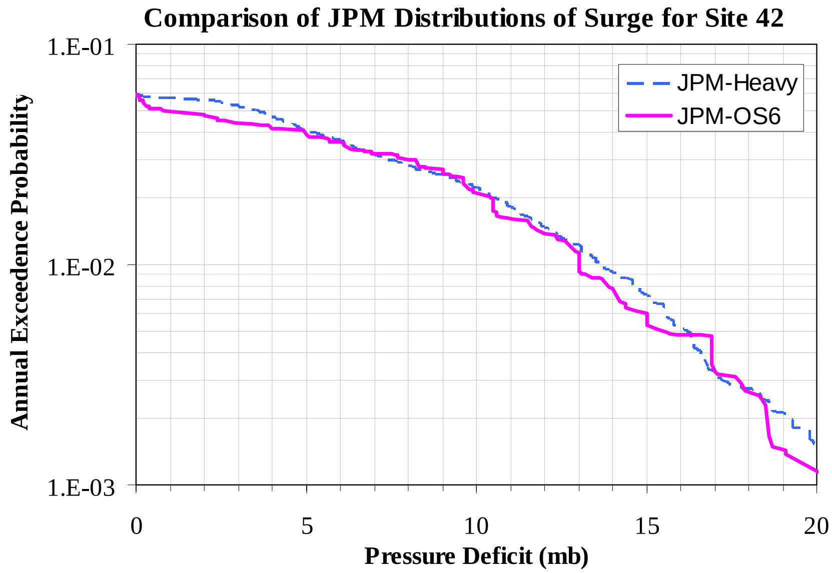

Figure 6. Comparison of cumulative distribution functions for a site along the shore at Biloxi, MS.

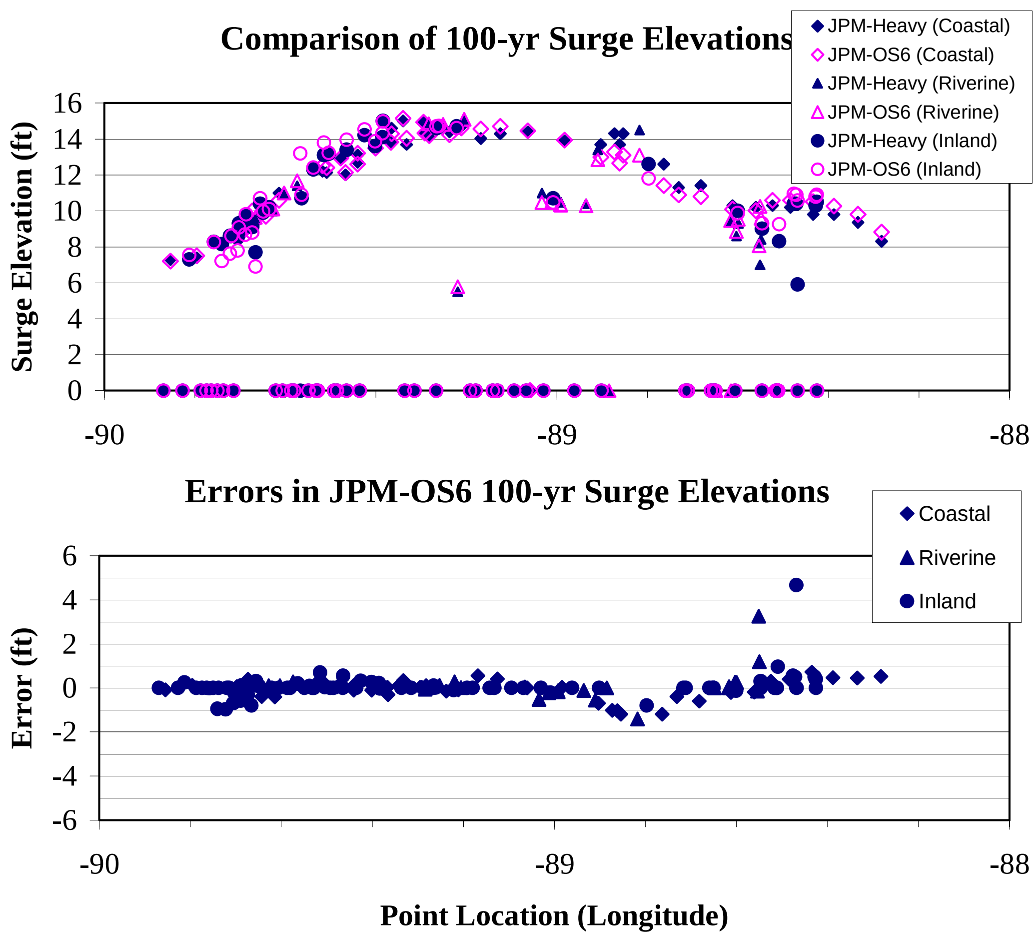

Figure 5. Comparison of 100-year surge values at 147 test points for JPM-OS6 and JPM-Heavy schemes. Positive errors indicate over-estimation by the JPM-OS6 scheme.

5 THE RESPONSE-SURFACE JPM-OS APPROACH

This approach takes advantage of the observation that the calculated surge η_m is most sensitive to ΔP, R_p, and distance (or track location). Also, sensitivity to heading angle θ and forward velocity V_f is weaker and may be approximated as linear. Furthermore, the variation of η_m as a function of these parameters is fairly smooth. These observations are confirmed by the sensitivity analyses documented in URS (2007) and by ADCIRC runs cited by Resio (2007).

As a result of these observations, it is possible to perform surge calculations for a moderate number of synthetic storms—with carefully selected combinations of parameters--and then interpolate between the calculated surge elevation (in five dimensions) to obtain the surge elevation for any desired combination of parameters. The computational cost for this interpolation is minimal. As a result, one can discretize the domain of the JPM integral very finely, even in five dimensions.

The main difficulty in the response-surface JPM-OS scheme resides in the *experimental design*; i.e., selection of the parameters combinations for the synthetic storms, in a manner that provides enough points in the five-dimensional $\Delta P - \overline{R}p - \theta - V_f - \text{distance}$ parameter space, without requiring a very large number of synthetic storms, and then implementing an accurate interpolation scheme that works reliably for all sites of interest. This selection process takes advantage of the weak sensitivity to heading angle θ and forward velocity V_f, concentrating on ΔP, R_p, and track location.

The work performed by Resio (2007) for the New Orleans area provides an illustration of the selection of parameter combinations for the synthetic storms. This study considered 3 values of ΔP, 3 values of $R_p \mid \Delta P$, 3 values of θ, 8 to 9 alternative tracks (9 for the base-case θ=0, 8 each for $\theta = \pm 45$ degrees; the spacing between tracks at landfall is ~30 km), and 3 values of forward velocity V_f. All parameter combinations would comprise a total of 675 synthetic storms. This number was reduced to 152 by reducing the number of ΔP - R_p combinations that are run for values of θ and V_f other than the base case. In particular, 9 ΔP - R_p combinations are run per primary track (every other track is a primary track, spaced ~ 60 km at landfall) when both θ and V_f are at their base case, 4 combinations are run when θ alone is changed to ± 45 degrees, 2 combinations are run when V_f alone is reduced to its low value, and 1 combination is run when both θ and V_f are at their alternative values. Figure 7 provides an illustration of the parameters of the resulting synthetic storms (for the primary tracks only), using the same format as Figure 3. A much smaller number of combinations are for the secondary tracks; these tracks provide more dense spatial sampling.

Calculations proceed by interpolating first on ΔP - R_p space for each track and each combination of θ and V_f, then interpolating linearly over θ, and then finally over V_f. Further details are provided in Resio (2007) and in Irish and Resio (2007).

6 SUMMARY AND DISCUSSION

This paper has described the quadrature and response-surface JPM-OS approaches, with emphasis on the derivation, application, and verification of the former. Experience from the URS/FEMA study for Mississippi and the USACE studies for Mississippi and Louisiana suggest that the two approaches are comparable in their accuracy and efficiency. Detailed comparisons between the two approaches will be documented in a forthcoming paper.

It may be possible, and it is certainly worthwhile, to improve the efficiency the numerical wave and surge calculations for probabilistic surge hazard studies by using somewhat more coarse computational grids and more efficient algorithms—while maintaining the necessary accuracy for studies of this kind. Nonetheless, the need for JPM-OS will remain because the hurricane parameterization for these studies will likely become more realistic and more complex, thereby increasing the number of dimensions in the JPM integral.

7 ACKNOWLEDGMENTS

The authors have benefited from extensive discussions and technical exchanges with other members of the URS/FEMA and USACE teams, particularly Don Resio, David Divoky, Vince Cardone, Andrew Cox, David Levinson, Jen Irish, Peter Vickery, and Lynda Charles.

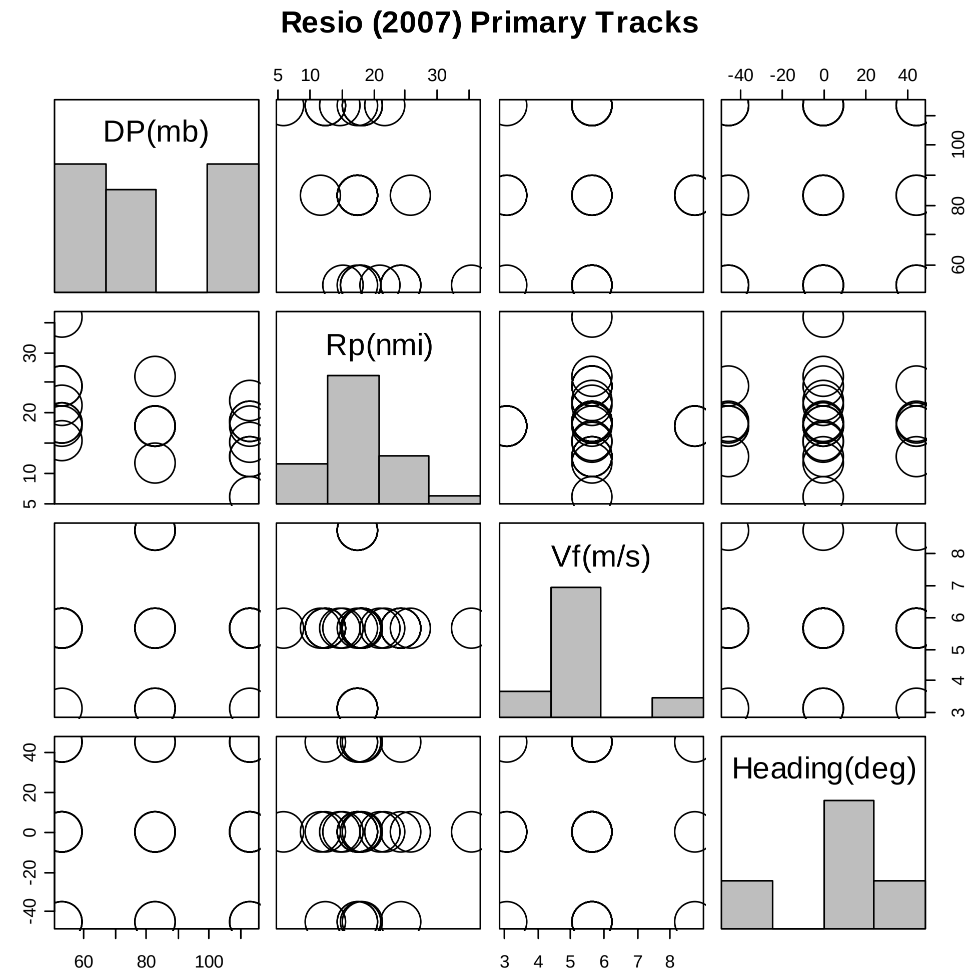

Figure 7. Graphical representation of the Resio (2007) experimental-design scheme for one landfall location. Values shown for the primary tracks only. The circle sizes and the histograms are meaningless in this case because these parameter combinations have no assigned probabilities.

This work was sponsored by FEMA under contract HSFEHQ-06-J--0018 to URS and by USACE under Purchase Order numbers W912HZ-06-P-0266 and W912BU-07-P-0211. The opinions expressed here are not necessarily those of the US Government or the sponsoring agencies.

8 REFERENCES

Benjamin, J.R. and Cornell, C.A. 1970. Probability, Statistics, and Decision for Civil Engineers. New York, New York: McGraw-Hill.

Diaconis, P. (1988). Bayesian numerical analysis. Statistical Decision Theory and Related Topics IV (pp. 163--175).

Genz, A., and B. Keister (1996). Fully Symmetric Interpolatory Rules for Multiple Integrals over Infinite Regions with Gaussian Weight. J. Comp. Appl. Math. 71 (1996), pp. 299-309

Gollwitzer, S., B. Kirchgaßner, R. Fischer, and R. Rackwitz (2006). PERMAS-RA/STRUREL system of programs for probabilistic reliability analysis. Structural Safety 28 (2006) 108–129.

Ho, F.P and V.A. Myers, 1975: Joint Probability Method of Tide Frequency Analysis applied to Apalachicola Bay and St. George Sound, Florida, NOAA Tech. Rep. WS 18, 43p.

Ho, F.P., Su, J.C., Hanevich, K.L., Smith, R.J. and F.P. Richards (1987). Hurricane Climatology for the Atlantic and Gulf Coasts of the United States, NOAA Technical Report NWS 38, Washington, D.C., April.

Holland, G.J. (1980). "An Analytic Model of the Wind and Pressure Profiles in Hurricanes." Monthly Weather Review: Vol. 108, No. 8, pp. 1212–1218.

Hong, H. P. (1998). An efficient point estimate method for probabilistic analysis. Reliability Engrg. and Sys. Safety, 59(3), 261–167.

Irish, J., D. Resio, and J. Ratcliff (2007). The influence of storm size on hurricane surge. J. of Physical Oceanography, in press.

Journel, A.G., and Ch. J. Huijbregts (1978). Minimg Geostatistics. Academic Press, London.

Madsen, H.O., S. Krenk and N.C. Lind (1986). Methods of Structural Safety, Prentice-Hall, Inc., Englewood Cliffs, NJ.

Melchers, R. E. (1999). Structural reliability: analysis and prediction. Chichester, New York, Wiley. 1999.

Miller, A.C., and T.R. Rice (1983). "Discrete approximations of probability distributions," Management Science, 29, 3, 352-361.

Minka, T. P. (2000) Deriving quadrature rules from Gaussian processes, Technical Report, Statistics Department, Carnegie Mellon University. Available on-line at http://research.microsoft.com/~minka/papers/minka-quadrature.ps.gz

Myers, V.A., 1975: Storm Tide Frequencies on the South Carolina Coast, NOAA Tech. Rep. NWS-16, 79 p.

National Oceanic and Atmospheric Administration (1979). Meteorological Criteria for Standard Project Hurricane and Probable Maximum Hurricane Wind Fields, Gulf and East Coasts of the United States, NOAA Technical Report NWS 23, Washington, D.C., September.

Niedoroda, A., D. Resio, G. Toro, D. Divoky, R. Lowe, L. Charles, H. Das, C. Reed (2007). Evaluation of the Storm Surge Hazard in Coastal Mississippi. Proc., 10th International Workshop on Wave Hindcasting and Forecasting, this volume.

O'Hagan, A. (1991). Bayes-Hermite quadrature. J Statistical Planning and Inference, 29, 245--260.

Parzen, E. 1962. Stochastic Processes. San Francisco, California: Holden Day.

Powell, M.J.D.(2004). The NEWUOA software for unconstrained optimization without derivatives. Proc., 40th Workshop on Large Scale Nonlinear Optimization, Erice, Italy.

Resio, D.T. (2007). White Paper on Estimating Hurricane Inundation Probabilities. Version 11, May 7.

Risk Engineering, Inc. (2007). Joint Probability Analysis Of Hurricane Flood Hazards For Mississippi. Report Prepared for URS Group in support of the FEMA-HMTAP flood study of the State of Mississippi, October.

Schwerdt, R.W., Ho, F.P., and R.R. Watkins (1979). Meteorological criteria for Standard Project Hurricane and Probable Maximum Hurricane Windfields, Gulf and East Coasts of the United States, Tech. Rep. NOAA-TR-NWS-23, National Oceanic and Atmospheric Administration.

Smolyak, S.A. (1963). Quadrature and interpolation formulas for tensor products of certain classes of functions. Soviet Math.Dokl., vol. 4, 240-243.

URS (2007). Coastal Documentation Report (Technical Overview). Final Report of FEMA Task Order 018, Mississippi Coastal Flood Study, unpublished consulting report.

Vanmarcke, E.H. (1983). Random Fields: Analysis and Synthesis. MIT Press.