Guidance for Flood Risk Analysis and Mapping

Coastal Flood Frequency and Extreme Value Analysis

November 2016

FEMA

Requirements for the Federal Emergency Management Agency (FEMA) Risk Mapping, Assessment, and Planning (Risk MAP) Program are specified separately by statute, regulation, or FEMA policy (primarily the Standards for Flood Risk Analysis and Mapping). This document provides guidance to support the requirements and recommends approaches for effective and efficient implementation. Alternate approaches that comply with all requirements are acceptable.

For more information, please visit the FEMA Guidelines and Standards for Flood Risk Analysis and Mapping webpage (www.fema.gov/guidelines-and-standards-flood-risk-analysis-and-mapping). Copies of the Standards for Flood Risk Analysis and Mapping policy, related guidance, technical references, and other information about the guidelines and standards development process are all available here. You can also search directly by document title at www.fema.gov/library.

Table of Revisions

Affected Section or Subsection	Date	Description
First Publication	November 2016	Initial version of new transformed guidance. The content was derived from the Guidelines and Specifications for Flood Hazard Mapping Partners, Procedure Memoranda, and/or Operating Guidance documents. It has been reorganized and is being published separately from the standards.

Table of Contents

List of Figures

1.0 General Statistical Methods

This document summarizes the most common statistical methods used to establish the 1-percent-annual-chance flood elevation for FEMA coastal Flood Map Projects. The following approach is used when at least 30 years of observational record is available. In this *extreme value analysis* approach, the data are used to establish a *probability distribution* that is assumed to describe the flooding process and that can be evaluated by using the data to determine the flood elevation at any frequency. The observational record length influences the fit of the probability distribution such that longer records improve the fit.

Coastal: Statistical Simulation Methods (FEMA 2016a) contains guidance on an alternate approach that is used when an adequate observational record of flood levels does not exist. In this case, it may be possible to simulate the flood process using hydrodynamic models driven by meteorological or other processes for which adequate data exist. That is, the hydrodynamic model (perhaps describing storm surge or waves) provides the link between the known statistics of the generating forces and the desired statistics of flood levels.

The present discussion begins with the basic concepts of probability theory and introduces the concept of a continuous probability distribution. The distributions which are most commonly used, as well as methods to fit these to the observed record, are described along with recommendations for FEMA Flood Map Project applications. A list of suggested additional information resources is included at the end of the subsection.

1.1 Elementary Probability Theory

Probability theory deals with the characterization of random events and, in particular, with the likelihood of the occurrence of particular outcomes. The word “probability” has many meanings, and there are conceptual difficulties with all of them in practical applications such as flood studies. The common frequency notion is assumed here: the probability of an event is equal to the fraction of times it would occur during the repetition of a large number of identical *trials*. For example, if one considers an annual storm season to represent a trial, and if the event under consideration is occurrence of a flood exceeding a given elevation, then the annual probability of that event is the fraction of years in which it occurs in the limit of an infinite period of observation. Clearly, this notion is entirely conceptual and cannot truly be the source of a probability estimate.

An alternate measure of the likelihood of an event is its expected *rate of occurrence*, which differs from its probability in an important way. Whereas probability is a pure number and must lie between zero and one, a rate of occurrence is a measure with physical dimensions (reciprocal of time) that can take on any value, including values greater than one. In many cases, when one speaks of the probability of a particular flood level, one actually means its rate of occurrence; thinking in terms of physical rate can help to clarify an analysis.

To begin, a number of elementary probability rules are reviewed. If an event occurs with probability P in some trial, then it fails to occur with probability $Q = 1 – P$. This is a consequence of the fact that the sum of the probabilities of all possible results must equal unity, by the definition of total probability:

$$\sum_i P(A_i) \equiv 1 \qquad \text{(Eq. 1-1)}$$

in which the summation is over all possible outcomes of the trial.

If *A* and *B* are two events, the probability that either *A* or *B* occurs is given by:

$$P(A \text{ or } B) = P(A) + P(B) - P(A \text{ and } B) \qquad \text{(Eq. 1-2)}$$

If *A* and *B* are *mutually exclusive*, then the third term on the right-hand side is zero, and the probability of obtaining either outcome is the sum of the two individual probabilities.

If the probability of *A* is contingent on the prior occurrence of *B*, then the *conditional probability* of *A* given the occurrence of *B* is defined to be:

$$P(A \mid B) \equiv \frac{P(AB)}{P(B)} \qquad \text{(Eq. 1-3)}$$

in which *P(AB)* denotes the probability of both *A* and *B* occurring.

If *A* and *B* are *stochastically independent*, *P(A|B)* must equal *P(A)*. Then the definition of conditional probability just stated gives the probability of occurrence of both *A* and *B* as:

$$P(AB) = P(A)P(B) \qquad \text{(Eq. 1-4)}$$

This expression generalizes for the joint probability of any number of independent events, as:

$$P(ABC...) = P(A)P(B)P(C)... \qquad \text{(Eq. 1-5)}$$

As a simple application of this rule, consider the chance of experiencing at least one base flood (P = 0.01) in 100 years. This is 1 minus the chance of experiencing no such flood in 100 years. The chance of experiencing no such flood in 1 year is 0.99, and if it is granted that floods in different years are independent, then the chance of not experiencing such a flood in 100 years is 0.99^{100} according to Equation 1-5, or 0.366. Consequently, the chance of experiencing at least one base flood in 100 years is 1 – 0.366 = 0.634, or about 63 percent.

1.2 Distributions of Continuous Random Variables

A continuous random variable can take on any value from a continuous range, not just a discrete set of values. The instantaneous ocean surface elevation at a point is an example of a continuous random variable. So, too, is the annual maximum water level at a point. If such a variable is observed a number of times, a set of differing values *distributed* in some manner over a range is found; this fact suggests the idea of a *probability distribution*. The observed values are a data *sample*.

If the *probability density function*, PDF, of *x* is defined as *f(x)*, such that the probability of observing the continuous random variable x to fall between *x* and *x + dx* is *f(x) dx* then, in accordance with the definition of total probability stated above:

$$\int_{-\infty}^{\infty} f(x)\ dx = 1 \qquad \text{(Eq. 1-6)}$$

If the upper limit of integration is taken to be the level L, rather than infinity, the *cumulative distribution function*, CDF, denoted by *F(x)*, is derived which specifies the probability of obtaining a value of L or less:

$$F(x \le L) \equiv \int_{-\infty}^{L} f(x)\ dx \qquad \text{(Eq. 1-7)}$$

It is assumed that the observed set of values, the sample, is derived by random sampling from a *parent distribution*. That is, there exists some unknown function, *f(x)*, from which the observed sample is obtained by random selection. No two samples taken from the same distribution will be exactly the same. Furthermore, random variables of interest in engineering cannot assume values over an unbounded range, as suggested by the integration limits in the expressions shown above. In particular, the lower bound for flood elevation at a point can be no less than ground level, wind speed cannot be less than zero, among other physical constraints. Upper bounds also exist, but they cannot be precisely specified; whatever occurs can be exceeded, if only slightly. Consequently, the usual approximation is that the upper bound of a distribution is taken to be infinity, whereas a lower bound may need to be specified based on physical constraints.

If the nature of the parent distribution can be inferred from the properties of a sample, then the distribution provides the complete statistics of the variable. If, for example, 30 years of annual peak flood data are available, and if these data can be used to specify the underlying distribution, then one can easily obtain the 10-, 2-, 1-, and 0.2-percent-annual-chance flood levels by computing *x* such that *F(x)* is 0.90, 0.98, 0.99, and 0.998, respectively.

The entirety of the information contained in the PDF can be represented by its moments. For the normal distribution, the *mean*, μ, specifies the *location* of the distribution and is the first moment about the origin:

$$\mu = \int_{-\infty}^{\infty} x\, f(x)\ dx \qquad \text{(Eq. 1-8)}$$

Two other common measures of the location of the distribution are the mode, which is the value of *x* for which *f* is maximum, and the median, which is the value of *x* for which *F* is 0.5.

The spread of the distribution is measured by its variance, σ^2, which is the second moment about the mean:

$$\sigma^2 = \int_{-\infty}^{\infty} (x-\mu)^2 f(x)\ dx \quad \text{(Eq. 1-9)}$$

The *standard deviation*, σ, is the square root of the variance.

The third and fourth moments are called the skew and the kurtosis, respectively; still higher moments fill in more details of the distribution shape, but they are seldom encountered in practice. If the variable is measured about the mean and is normalized by the standard deviation, then the *coefficient of skewness*, measuring the asymmetry of the distribution about the mean, is:

$$\eta_3 = \int_{-\infty}^{\infty} \left(\frac{x-\mu}{\sigma}\right)^3 f(x)\ dx \quad \text{(Eq. 1-10)}$$

and the *coefficient of kurtosis*, measuring the peakedness of the distribution, is:

$$\eta_4 = \int_{-\infty}^{\infty} \left(\frac{x-\mu}{\sigma}\right)^4 f(x)\ dx \quad \text{(Eq. 1-11)}$$

These four parameters are properties of the continuous distribution, not of discrete samples. However, the sample has its own set of corresponding parameters. For example, the *sample mean* is:

$$\bar{x} = \frac{1}{n}\sum_i x_i \quad \text{(Eq. 1-12)}$$

which is the average of the sample values. The *sample variance* is:

$$s^2 = \frac{1}{n-1}\sum_i (x_i - \bar{x})^2 \quad \text{(Eq. 1-13)}$$

while the sample skew and kurtosis are:

$$C_S = \frac{n}{(n-1)(n-2)s^3}\sum_i (x_i - \bar{x})^3 \quad \text{(Eq. 1-14)}$$

$$C_K = \frac{n(n+1)}{(n-1)(n-2)(n-3)s^4}\sum_i (x_i - \bar{x})^4 \quad \text{(Eq. 1-15)}$$

Note that in some literature, the kurtosis is reduced by 3, so the kurtosis of the normal distribution becomes zero; it is then called the *excess* kurtosis.

1.3 Stationarity

Roughly speaking, a random process is said to be stationary if it is not changing over time, or if its statistical measures remain constant. Many statistical tests can be performed to help determine whether a record displays a significant trend that might indicate nonstationarity. A simple test is the Spearman Rank Order Test. This is a nonparametric test operating on the ranks of the individual values sorted in both magnitude and time. The Spearman *R* statistic is defined as:

$$R = 1 - \frac{6\sum_i \left(d_i\right)^2}{n(n^2 - 1)} \qquad \text{(Eq. 1-16)}$$

in which *d* is the difference between the magnitude rank and the sequence rank of a given value. The statistical significance of *R* computed from Equation 1-16 can be found in published tables of Spearman's *R* for *n – 2* degrees of freedom.

1.4 Correlation between Series

Two random variables may or may not be statistically dependent. Dependence means that knowing the value of one of the variables permits a degree of inference regarding the value of the other. Whether paired data (*x,y*), such as simultaneous measurements of wave height and period, are interdependent or correlated is usually measured by their *linear correlation coefficient*:

$$r = \frac{\sum_i (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_i (x_i - \bar{x})^2}\sqrt{\sum_i (y_i - \bar{y})^2}} \qquad \text{(Eq. 1-17)}$$

This correlation coefficient indicates the strength of the correlation. An *r* value of +1 or -1 indicates perfect correlation, so a cross-plot of *y* versus *x* would lie on a straight line with positive or negative slope, respectively. If the correlation coefficient is near zero, then such a plot would show random scatter with no apparent trend.

1.5 Convolution of Two Distributions

If a random variable, *z*, is the simple direct sum of the two random variables *x* and *y*, then the distribution of *z* is given by the convolution integral:

$$f_z(z) = \int_{-\infty}^{\infty} f_x(T) f_y(z-T)\ dT \tag{Eq. 1-18}$$

in which subscripts specify the appropriate distribution function. This equation can be used, for example, to determine the distribution of the sum of storm surge and tide under the assumptions that surge and tide are independent and can be linearly superimposed.

1.6 Important Distributions

Many statistical distributions are used in engineering practice. Perhaps the most familiar is the *normal* or *Gaussian* distribution. The following distributions are likely to be used in a FEMA Flood Map Project. Although the normal distribution is the most familiar, the most fundamental is the uniform distribution.

1.6.1 Uniform Distributions

The uniform distribution is defined as constant over a range and zero outside that range. If the range is from *a* to *b*, then the PDF is:

$$f(x) = \frac{1}{b-a}, \quad a \le x < b, \quad 0 \text{ otherwise} \tag{Eq. 1-19}$$

which, within its range, is a constant independent of *x*; this is also called a *top-hat* distribution.

The uniform distribution is especially important because it is used in drawing random samples from all other distributions. A random sample drawn from a given distribution can be obtained by first drawing a random sample from the uniform distribution defined over the range from 0 to 1. Then set *F(x)* equal to this value, where *F* is the cumulative distribution to be sampled. The desired value of *x* is obtained by inverting the expression for *F*.

Sampling from the uniform distribution is generally done with a random number generator, returning values on the interval from 0 to 1. Most programming languages and spreadsheets have such a function built in, as do many calculators. However, not all such standard routines are satisfactory. While adequate for drawing a small number of samples, many widely used standard routines fail statistical tests of uniformity. If a critical application requires a large number of samples, these simple standard routines may be inadequate. Further guidance on this topic is available in *Coastal: Statistical Simulation Methods* (FEMA 2016a).

1.6.2 Normal or Gaussian Distribution

The normal or Gaussian distribution, sometimes called the bell-curve, has a special place among probability distributions. Consider a large number of large samples drawn from some unknown distribution. For each large sample, compute the sample mean. The distribution of those means tends to follow the normal distribution, a consequence of the *central limit theorem*. Despite this, the normal distribution does not play a central direct role in hydrologic frequency analysis. The standard form of the normal distribution is:

$$f(x) = \frac{1}{\sigma(2\pi)^{1/2}} e^{-\frac{(x-\mu)2}{2\sigma^2}}$$

$$F(x) = \frac{1}{2} + \frac{1}{2}\operatorname{erf}\left(\frac{x-\mu}{\sqrt{2}\sigma}\right)$$

(Eq. 1-20)

1.6.3 Rayleigh Distribution

The Rayleigh distribution is important in the theory of random wind waves. Unlike many distributions, it has some basis in theory; Longuet-Higgins (1952) showed that with reasonable assumptions for a narrow banded wave spectrum, the distribution of wave height will be Rayleigh. The standard form of the distribution is:

$$f(x) = \frac{x}{b^2} e^{-\frac{x^2}{2b^2}}$$

$$F(x) = 1 - e^{-\frac{x^2}{2b^2}}$$

(Eq. 1-21)

The range of *x* is positive, and the scale parameter b > 0. In water wave applications, $2b^2$ equals the mean square wave height. The mean and variance of the distribution are respectively given by:

$$\mu = b\sqrt{\frac{\pi}{2}}$$

$$\sigma^2 = b^2\left(2 - \frac{\pi}{2}\right)$$

(Eq. 1-22)

The skew and kurtosis of the Rayleigh distribution are constants (approximately 0.63 and 3.25, respectively) but are of little interest in extreme value analysis for flooding.

An application of the Rayleigh distribution in coastal flood studies is the estimation of the 2-percent-annual-chance runup level, given the mean level computed by the RUNUP 2.0 program. There is empirical evidence that the runup is Rayleigh distributed so that the ratio between the 2-percent and 50-percent runup levels can be computed from Equation 1-21.

1.6.4 Extreme Value Distribution

Extreme value distributions are recognized as having some theoretical justification and are therefore most appropriate for extreme value analyses. Among the well-known extreme value distributions are the *Gumbel* distribution and the *Weibull* distribution. Both of these are candidates for FEMA Flood Map Project applications and have been widely used with success in similar applications. Significantly, these distributions (and others, including the Rayleigh) are subsumed under a more general distribution, the *generalized extreme value* (GEV) distribution, given by:

$$f(x) = \frac{1}{b}\left\{1 + c\left(\frac{x-a}{b}\right)\right\}^{-\frac{1}{c}-1} e^{-(1+c(x-a)/b))^{-1/c}}$$

$$\text{for } -\infty < x \le a - \frac{b}{c} \text{ with } c < 0 \text{ and } a - \frac{b}{c} \le x < \infty \text{ with } c > 0$$

$$f(x) = \frac{1}{b} e^{-e^{-(x-a)/b}} e^{-(x-a)/b} \qquad \text{for } -\infty \le x < \infty \qquad \text{with } c = 0$$

(Eq. 1-23)

The cumulative distribution is given by the expressions:

$$F(x) = e^{-(1+c(x-a)/b))^{-1/c}}$$

$$\text{for } -\infty < x \le a - \frac{b}{c} \text{ with } c < 0 \text{ and } a - \frac{b}{c} \le x < \infty \text{ with } c > 0$$

$$F(x) = e^{-e^{-(x-a)/b}} \qquad -\infty \le x < \infty \qquad \text{with } c = 0$$

(Eq. 1-24)

In these expressions, *a*, *b*, and *c* are the *location*, *scale*, and *shape* factors, respectively. This distribution includes the *Frechet* (Type 2) distribution for $c > 0$ and the Weibull (Type 3) distribution for $c < 0$. If the limit of the exponent of the exponential in the first forms of these distributions is taken as *c* goes to 0, then the simpler second forms are obtained, corresponding to the Gumbel (Type 1) distribution. Note that the Rayleigh distribution is a special case of the Weibull distribution and so is also encompassed by the GEV distribution.

The special significance of the members of the extreme value family is that they describe the distributions of the extremes drawn from other distributions. That is, given a large number of samples drawn from an unknown distribution, the extremes of those samples tend to follow one of the three types of extreme value distributions, all incorporated in the GEV distribution. This is analogous to the important property of the normal distribution that the means of samples drawn from other distributions tend to follow the normal distribution. If a year of water levels is considered to be a sample, then the annual maximum, or the largest value in the sample, is an extreme and may tend to be distributed according to the statistics of extremes.

1.6.5 Pareto Distribution

If for some unknown distribution the sample extremes are distributed according to the GEV distribution, then the set of sample values exceeding some high threshold tends to follow the Pareto distribution. Consequently, the GEV and Pareto distributions are closely related in a dual manner. The Pareto distribution is given by:

$$F(y) = 1 - \left(1 + \frac{cy}{\tilde{b}}\right)^{-1/c} \quad \text{for} \quad y = x - u$$

$$\text{with} \quad \tilde{b} = b + (u - a) \qquad \text{(Eq. 1-25)}$$

where *u* is the selected threshold. In the limit as *c* goes to zero, this reduces to the simple expression:

$$F(y) = 1 - e^{-y/\tilde{b}} \qquad \text{for} \quad y > 0 \qquad \text{(Eq. 1-26)}$$

The Pareto distribution was developed to capture tails of observed data. For example, the surge elevations associated with a single Category 5 hurricane within a record where most observed surges are from tropical storms represents a tail that other distributions fail to model.

1.6.6 Poisson Distribution

The Poisson distribution – a discrete distribution – is especially important in some applications, because it describes a process in which events occur at a known average rate but with no memory of the last occurrence. Examples include such processes as radioactive decay and, of interest here, might include the number of hurricanes occurring in a year at some site. If hurricanes occur at some long-term average rate (storms per year at the site), and if the occurrence of one storm is independent of any other occurrence, then the process may be well described by a Poisson distribution, which is given by:

$$f(k:\lambda) = \frac{e^{-\lambda}\lambda^{k}}{k!} \qquad \text{(Eq. 1-27)}$$

where *f* is the probability of experiencing exactly *k* occurrences in an interval if *λ* is the number expected to occur in that interval, equal to the interval length multiplied by the average rate. The Poisson distribution is important in coastal applications of the Empirical Simulation Technique (EST) method for estimation of surge frequency. Guidance on the EST method is available in *Coastal: Statistical Simulation Methods* (FEMA 2016a).

2.0 Data Sample and Estimation of Parameters

Several methods of estimating the best values of distribution parameters have been widely used. The most frequent are the methods of plotting positions, moments, and maximum likelihood. The methods discussed here are limited to point-site estimates. If statistically similar data are available from other sites, then it may be possible to improve the parameter estimate through the method of *regional frequency analysis*; see Hosking and Wallis (1997) for information on this method. Note that the word *regional* is used to describe the geophysical boundaries which do not necessarily following jurisdictional boundaries (e.g. FEMA Regions).

2.1 Plotting Positions

Widely used in older hydrologic applications, the method of plotting positions is based on first creating a visualization of the sample distribution and then performing a curve-fit between the chosen distribution and the sample. However, the sample consists only of the process variable; there are no associated quantiles, and so it is not clear how a plot of the sample distribution is to be constructed. The simplest approach is to rank-order the sample values from smallest to largest and to assume that the value of *F* appropriate to a value is equal to its fractional position in this ranked list, *R/N*, where *R* is the value's rank from 1 to *N*. Then, the smallest observation is assigned *plotting position 1/N*, and the largest is assigned *N/N=1*. This is clearly unsatisfactory at the upper end because instances larger than the largest observed in the sample can occur. A more satisfactory and widely used plotting position expression is *R/(N+1)*, which leaves some room above the largest observation for the occurrence of still larger elevations. A number of such plotting position formulas is encountered in practice, most involving the addition of constants to the numerator and denominator, *(R+a)/(N+b)*, in an effort to produce improved estimates at the tails of specific distributions.

Given a plot produced in this way, one might simply draw a smooth curve through the points and visually extend it to the recurrence intervals of interest. This constitutes an entirely empirical approach and is sometimes made easier by using a transformed scale for the cumulative frequency to construct the plot. The simplest such transformation is to plot the logarithm of the cumulative frequency, which flattens the curve and makes extrapolation easier.

A second approach would be to choose a distribution type and adjust its free parameters so that a plot of the distribution matches the plot of the sample. This is commonly done by least-squares fitting. Fitting by eye is also possible if an appropriate *probability paper* is adopted, on which the transformed axis is not logarithmic but is transformed in such a way that the corresponding distribution plots as a straight line; however, this cannot be done for all distributions.

These simple methods based on plotting positions, although widely used, are problematic. Two fundamental difficulties with the methods are seldom addressed. First, it is inherent in the methods that each of N quantile bins of the distribution is occupied by one and only one sample point, an extremely unlikely outcome. Second, when a least-squares fit is made for an analytical distribution form, the error being minimized is taken as the difference between the sample value and the distribution value, whereas the true error is not in the value but in its frequency position.

2.2 Methods of Moments: Conventional Moments

An alternate method that does not rely upon visualization of the empirical distribution is the method of moments, of which there are several forms. The methodology is to equate the sample moments and the distribution moments and to solve the resulting set of equations for the distribution parameters. That is, the sample moments are simple functions of the sample points, as in Section 1.2. Similarly, it may be possible to express the corresponding moments of an analytical distribution as functions of the several parameters of the distribution. If this can be done, then empirical estimates of those parameters can be obtained by equating the expressions to the sample values.

2.3 Method of Moments: Probability-weighted Moments and Linear Moments

Various versions of the method of moments overcome certain difficulties inherent in conventional methods of moments. For example, simple moments may not exist for a given distribution or may not exist for all values of the parameters. Higher sample moments may not be able to adopt the full range of possible values; for example, the sample kurtosis is constrained algebraically by the sample size.

Alternate moment-based approaches have been developed, including probability-weighted moments and the newer method of linear moments, or L-moments. L-moments consist of simple linear combinations of the sample values that convey much the same information as true moments: location, scale, shape, and those of higher order. However, being linear combinations they have certain desirable properties that make them preferable to normal moments. The theory of L-moments and their application to frequency analysis has been developed by Hosking; see, for example, Hosking and Wallis (1997).

2.4 Maximum Likelihood Method

A method based on an entirely different idea is the method of maximum likelihood. Maximum likelihood estimates are preferred because they commonly show less bias than other methods and are also conceptually appealing. However, they usually require iterative calculations to locate the optimum parameters, and a maximum likelihood estimate may not exist for all distributions or for all values of the parameters for a particular distribution. If the Mapping Partner considers alternate distributions or fitting methods, the likelihood of each fit can still be computed using the equations given above even if the fit was not determined using the maximum likelihood method. The distribution with the greatest likelihood of having described the sample could then be chosen.

Consider an observation, *x*, obtained from the density distribution *f(x)*. The probability of obtaining a value close to *x*, say within the small range *dx* around *x*, is *f(x) dx*, which is proportional to *f(x)*. Then, the *posterior* probability of having obtained the entire sample of *N* points is assumed to be proportional to the product of the individual probabilities estimated in this way, in consequence of Equation 1-5. This product is called the *likelihood* of the sample, given the assumed distribution:

$$L = \prod_{1}^{N} f(x_i) \quad \text{(Eq. 2-1)}$$

It is more common to work with the logarithm of this equation, which is the *log-likelihood*, *LL*, given by:

$$LL = \sum_{1}^{N} \log f(x_i) \quad \text{(Eq. 2-2)}$$

The primary use of the maximum likelihood method is to determine the distribution parameters that maximize the likelihood of the given sample. Because the logarithm is a monotonic function, maximizing the likelihood is equivalent to maximizing the log-likelihood. Note that because *f(x)* is always less than one, all terms of the sum for *LL* are negative; consequently, larger log-likelihoods are associated with smaller numerical values.

3.0 Extreme Value Analysis in a Flood Insurance Study (FIS) FEMA Flood Map Project

For FEMA Flood Map Project extreme value analysis, the Mapping Partner may adopt the annual maxima of the data series (runup, stillwater elevation [SWEL], and other factors pertinent to coastal flooding) as the appropriate data sample. Also acceptable is the peak-over-threshold (POT) approach, fitting all observations that exceed an appropriately high threshold to the generalized Pareto distribution. Further discussion of the POT approach can be found in references such as Coles (2001).

3.1 Atlantic, Gulf and Pacific Regions

The normal approach for the Atlantic, Gulf, and Pacific regions has been to fit the GEV distribution to the data sample using the method of maximum likelihood. The Mapping Partner can also consider distributions other than the GEV for use with the annual series. However, the final distribution selected to estimate the base flood level should be based on the total estimated likelihood of the sample. In the event that methods involve different numbers of points (e.g., POT vs. annual maxima), the comparison might be made on the basis of average likelihood per sample point.

3.2 Great Lakes Region

The Great Lakes region has applied a method of deriving probability distributions for storm responses such as stillwater level or total water level (which includes wave runup) by sampling across the record of lake level observations. Melby et al. (2012) and Nadal-Caraballo et al. (2012) describe application of the POT/Generalized Pareto Distribution (GPD) method to Great Lakes storm responses in detail. Extreme value theory suggests that the Partial Duration Series determined from the POT method should conform to the GPD not the GEV distribution. The goodness of fit is evaluated using both the CDF and a quantile-quantile (Q-Q) analysis of storm responses.

A critical issue on the Great Lakes in developing a storm sampling approach is whether the events that end up being sampled as part of the composite storm set are representative of the entire record length, in terms of the distribution of mean lake levels. Thus selection of a storm set that accurately describes extremal water-level statistics everywhere within the lake must be done judiciously in order to minimize the storm sample size but accurately reflect the extremal water-level distributions everywhere along the lake shoreline. Selection of an appropriate but smaller composite storm sample can be achieved through analysis and comparison of water-level probability distributions derived from various composite storm sets and a much larger full storm set. The Mapping Partner should conduct a lake-specific analysis to determine the appropriate ratio of water level to wave events within the composite storm set to adequately represent the storm conditions that cause extreme flood hazard responses. The following steps for storm selection are recommended:

1. Determine the number of storms necessary for the composite storm set to adequately represent storm conditions and responses throughout the lake (See Nadal-Caraballo, 2012).

2. Identify storms having the highest peak storm surge. This would typically be done using NOAA water-level measurements from all available sites and over the full record length. Rank storms based on magnitude of peak surge at each site.

3. Identify storms having highest peak wave height. This would typically be done using measured data, hindcasts, or surrogate wave calculations at spatially distributed sites. Rank storms based on magnitude of peak significant wave height.

4. Select a sufficient number of the highest ranked surge and wave events at each location such that the total number of storms is greater than the previously determined composite storm set size. This sample will be further reduced through screening to achieve the target composite storm set size.

5. Screen the storm set to eliminate duplicate storms. If storms are duplicates, reject duplicate event (or events) with the lowest site-specific rank and include the next largest event at that site.

6. Balance the number of storms selected for each site to maximize consistency in geographical and temporal coverage of selected storms; define an initial set with a sufficient number of storms as previously determined with an appropriate ratio of wave height dominated events and storm surge events. For Lake Michigan, the appropriate ratio was determined to be roughly 50 percent based on maximum wave height, and 50 percent based on peak storm surge (See Nadal-Caraballo, 2012).

7. Ice screening may then be required if ice processes prevented flooding for the selected events. Data and basic process knowledge of ice cover influences on flooding are limited. Therefore, the ice screening must be done with caution. Using regional ice maps nearest the time of each storm, determine if shore fast ice has the potential to block waves. If a storm in the initial set is a low-ranked surge event and waves are blocked due to ice coverage, consider it as a candidate to remove.

4.0 1-Percent-Annual-Chance Stillwater Levels

Stillwater is the coastal water level in response to tidal and meteorological forcing in the absence of wave energy. Generally speaking, there are two distinct stillwater components: astronomic tide and storm surge (wind and pressure setup). Runoff from rainfall represents an additional meteorological contribution to stillwater in some sheltered water bodies. More specifically, increased water levels near the mouth of tidal rivers may occur where fluvial flood waters interact with coastal tides and surge. Consequently, the rainfall runoff excess elevation may need to be considered as a third contributor to coastal stillwater elevation.

4.1 Tide Statistics

The astronomic tide is a deterministic process. Consequently, tide statistics can be generated directly from the local tidal constituents. One simple way to do this is to sample the predicted tide at random times throughout the tidal epoch. Alternatively, predictions can be used to obtain highs and lows from which corresponding statistics can be derived. It is noted that the maximum possible tide is given simply by the sum of the amplitudes of the tidal constituents.

4.2 Surge Statistics

The development of surge statistics can be approached in two general ways. First, if sufficient data are available from tide gage records, the non-tidal residual can be computed by subtracting the predicted astronomical tide from the observed water levels. Annual peak residuals or peaks over threshold values for the period of record can be fit to a GEV or other appropriate distribution using the method of maximum likelihood (or an alternate acceptable method). The Mapping Partner should keep in mind that the 1-percent level determined in this way will include the contributions of all mean water components affecting the gage, which may include both static wave setup (to the degree it exists at the gage site) and riverine rainfall runoff.

The second way in which 1-percent surge levels are determined is through numerical modeling of surge elevation using 1-D or 2-D models combined with a statistical model relating the surge simulations to storm frequency and storm parameter distributions. *Coastal: Statistical Simulation Methods* (FEMA 2016a) contains detailed guidance on the statistical models for use in determining surge statistics from model simulations.

4.3 Combined Effects: Surge Plus Tide

There are several ways to consider the effects of storm surge and tides to arrive at the 1-percent-annual-chance flood. First, if the surge and tide can be assumed to combine linearly (that is, neither is physically altered to an important degree by the presence of the other), then the simplest method is to simply add them together. If a surge episode – such as a northeaster – is of relatively long duration compared with a tidal cycle, then high tide will be certain to occur at some time for which the surge is near its peak, and a simple sum of amplitudes may be sufficiently accurate.

However, if the surge duration is short – such as may be typical for hurricanes in northern latitudes – this approximation is inadequate. The next simplest assumption, still assuming linear superposition, is based on the fact that the PDF for a sinusoid is largest at its extrema – tide is generally near a local high water, or near a local low water, and spends more time near those values than in between. It may be reasonable, then, to assume that the peak surge occurs with equal probability near a high tide or near a low tide, taking mean high and mean low as representative values. Each of the corresponding elevation sums would be assigned 50 percent of the rate associated with the particular storm (as if each storm were to occur twice, once at high tide and once at low tide), and the frequency analysis would proceed with these divided rates.

A slightly more complex approach but still assuming physical independence, is based on the convolution method mentioned in Section 1.6. In this method, the PDFs for both tide and surge without tide are used. If the probability density of the tide level Z is denoted by $p_T(Z)$ and the probability density of the surge level is $p_S(Z)$, then the probability density of the sum of the two is given by:

$$p(Z) = \int_{-\infty}^{\infty} p_T(T) p_S(Z-T)\, dT = \int_{-\infty}^{\infty} p_T(Z-S) p_S(S)\, dS \qquad \text{(Eq. 4-1)}$$

where the indicated integrations are over all tide and surge levels.

In some cases, however, the essential assumption that the tide and surge can be linearly added is not satisfied. In shallow water areas extending a large distance inland, the enhanced depth associated with tide (or surge) affects the propagation and transformation of the surge (or tide). That is, there is a nonlinear hydrodynamic interaction between the two. In such a case, more complex methods are required because the nonlinear interaction can only be accounted for by hydrodynamic considerations, not by any amount of purely statistical effort. For more information related to applying tides in storm surge modeling, please see *Coastal: Water Levels* (FEMA 2016b) and the statistical approaches are in *Coastal: Statistical Simulation Methods* (FEMA 2016a).

4.4 Combined Effects: Surge Plus Riverine Runoff

The simplest assumption is to assume that the extreme levels from coastal and riverine processes are physically independent and are not concurrent. This assumption is acceptable if storms that produce extreme rainfall and runoff are not the same as the storms that produce the greatest storm surge. Furthermore, if a single storm produces both large surge and large runoff, the runoff may be significantly delayed by the time required for overland flow, causing the runoff elevation to peak after the storm surge. When there are particular storms and locations for which these assumptions are not true, care should be taken to evaluate if they strongly influence the final statistics.

The simplified procedure, assuming independence and non-concurrence, is straightforward, beginning with development of curves or tables for rate of occurrence vs. flood level for each flood source (riverine and coastal). Rate of occurrence can be assumed equal to the reciprocal of the recurrence interval, so the 100-year flood has a rate of occurrence of 0.01 times per year. This is numerically equal to what is more loosely called the *flood elevation probability*. Then one proceeds as follows at each point of interest, P, within the mixed surge/runoff tidal reach.

1. Select a flood level Z within the elevation range of interest at point P.
2. Determine the rates of occurrence $R_{P,R}$ (Z) and $R_{P,S}$ (Z) of rainfall runoff and storm surge elevations exceeding Z at site P (number of events per year).
3. Find the total rate $R_{P,T}(Z) = R_{P,R}(Z) + R_{P,S}(Z)$ at which Z is exceeded at point P, irrespective of flood source.
4. Repeat steps (1) through (3) for the necessary range of flood elevations.
5. Plot the combined rates $R_{P,T}$ (Z) vs Z and find $Z_{P,100}$ by interpolation at $R_{P,T} \approx 0.01$.
6. Repeat steps (1) through (5) for a range of sites covering the length of the mixed tidal reach.
7. Construct the 1-percent-annual-chance composite profile passing through the several combined 1-percent-annual-chance elevation points and blending smoothly into the pure-riverine and pure-surge 1-percent-annual-chance profiles at the ends of the mixed reach.

The procedure is shown schematically in Figure 4-1 in which the combined curve has been constructed by addition of the rates at elevations of 6, 8, 10, and 12 feet. The entire procedure can be implemented in a simple hand calculator program, with the input at point P being the 10-, 2-, 1-, and 0.2-percent-annual-chance levels for both runoff and surge, as obtained from standard FIS report tables.

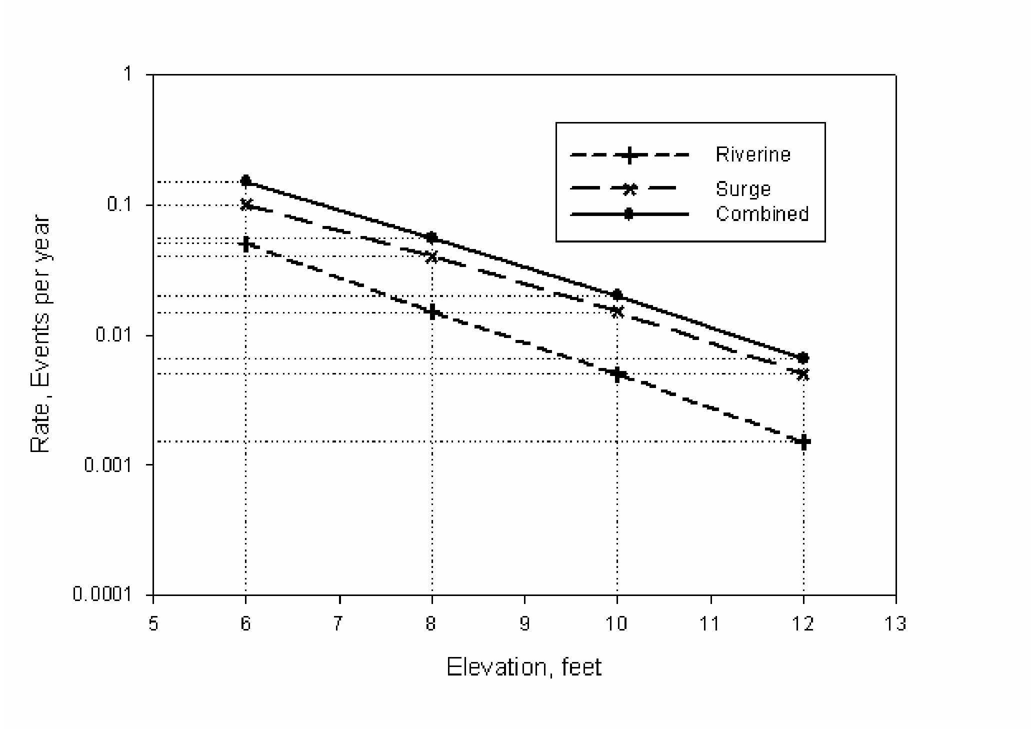

Figure 4.1 Schematic Illustration of Riverine and Surge Rate Combination

5.0 Additional Resources

The Mapping Partner may consult a large amount of literature on probability, statistics, and statistical hydrology. The following is a partial list of resources that might also be consulted for additional guidance:

Probability Theory:

An Introduction to Probability Theory and Its Applications, Third Edition, William Feller, 1968 (two volumes). This is a classic reference for probability theory, with a large number of examples drawn from science and engineering.

The Art of Probability for Scientists and Engineers, Richard Hamming, 1991. Less comprehensive than Feller but provides clear insight into the conceptual basis of probability theory.

Statistical Distributions:

Statistics of Extremes, E.J. Gumbel, 1958. A cornerstone reference for the theory of extreme value distributions.

Extreme Value Distributions, Theory and Applications, Samuel Kotz and Saralees Nadarajah, 2000. A more modern and exhaustive exposition.

Statistical Distributions, Second Edition, Merran Evans, Nicholas Hastings, and Brian Peacock, 1993. A useful compendium of distributions but lacking discussion of applications; a formulary.

An Introduction to Statistical Modeling of Extreme Values, Stuart Coles, 2001. A practical exposition of the art of modeling extremes, including numerous examples. Provides a good discussion of POT methods that can be consulted to supplement the annual maxima method.

Statistical Hydrology:

Applied Hydrology, Ven Te Chow, David Maidment, and Larry Hays, 1988. One of several standard texts with excellent chapters on hydrologic statistics and frequency analysis.

Probability and Statistics in Hydrology, Vujica Yevjevich, 1972. A specialized text with a lot of pertinent information for hydrologic applications.

General:

Numerical Recipes, Second Edition, William Press, Saul Teukolsky, William Vetterling, and Brian Flannery, 1992. A valuable and wide ranging survey of numerical methods and the ideas behind them. Excellent discussions of random numbers, the statistical description of data, and modeling of data, among much else. Includes well-crafted program subroutines; the book is available in separate editions, presenting routines in FORTRAN and C/C++.

Software:

Several open-source and commercial software packages provide tools to assist in the sorts of analyses discussed in this section. In particular, the S, S-PLUS, and R programming languages (commercial and open-source versions of a high-level statistical programming language) include comprehensive statistical tools. The R language package is available for free from the web site http://www.r-project.org/; several books discussing the use of R and S are available. Other well-known software packages include *Mathematica*, *Matlab*, *SPSS*, and *SYSTAT*.

6.0 Bibliography

Coles, Stuart. 2001. *An Introduction to Statistical Modeling of Extreme Values*. Springer, London, U.K.

Federal Emergency Management Agency. 2016a. *Coastal: Statistical Simulation Methods*. Washington, D.C.

Federal Emergency Management Agency. 2016b. *Coastal: Water Levels*. Washington, D.C.

Hosking, J.R.M. and J.R. Wallis. 1997. *Regional frequency analysis: an approach based on L-moments*. Cambridge University Press, Cambridge, U.K.

Longuet-Higgins, M.S. 1952. On the statistical distribution of the height of sea waves. *Journal of Marine Research*, Vol. 11, pp 245-266.

Melby, Jeffrey A.; Nadal-Caraballo, Noberto C.; Pagán-Albelo, Yamiresty; and Bruce Ebersole. 2012. *Wave Height and Water Level Variability on Lakes Michigan and St. Clair.* U.S. Army Corps of Engineers, Engineer Research and Development Center, Vicksburg, MS, ERDC/CHL TR-12-23.

Nadal-Caraballo, Noberto C.; Melby, Jeffrey A.; and Bruce A. Ebersole. 2012. *Statistical Analysis and Storm Sampling Approach for Lakes Michigan and St. Clair.* U.S. Army Corps of Engineers, Engineer Research and Development Center, Vicksburg, MS, ERDC/CHL TR-12-19.