Extreme Wave Height Data Analysis

Review and Recommendations

B. A. Harper

Conservation technical report No. ?

ISSN 1037-4701 RE ??? August 1996

Contents

Extreme Wave Height Data Analysis
Review and Recommendations

Technical Report: B. A. Harper
Coastal Resource Assessment Section

Introduction

This report summarises the results of a review of analysis techniques for the estimation of extreme wave heights (either significant or maximum) from measured data sets. The Department (previously as The Beach Protection Authority and The Department of Harbours and Marine) began developing wave data monitoring systems in the late 1960's and now operates more than a dozen nearshore sites along the Queensland coast. A principal reason for wave monitoring has been to establish the baseline wave climatology in various regions - an essential adjunct to the long term understanding of coastal processes. Routine wave data analyses to date have been targeted towards establishing "normal" criteria related to basic wave height and period statistics, seasonal and annual exceedance and sea state persistence.

The data base now available at some specific sites is sufficiently long (up to 20 years in some cases) to begin to consider the estimation of "extreme" conditions. Extreme wave heights are those not expected to occur, for example, more than once in several years of measurements. At the lower end this might represent an annual probability of exceedance less than 0.2 (once in every five years on average), extending out to 0.01 (once in a hundred years on average) or even 0.001 (once in a thousand years on average) or less. The lower the annual probability (or higher the average return period) the higher the expected wave height - provided other physical limits such as wave breaking do not form a limit to the process.

The utility of having such information on extreme conditions is related to a need to ensure that very severe episodic natural events can be anticipated. Such information is essential for engineering design, coastal planning and the estimation of short term coastal erosion experienced during storm events. Without access to long term data, such analyses must be based on model hindcasting of past (unwitnessed) weather events or conceptual models of possible future events based on simulation techniques. Based on recorded or hindcast data sets alone, extrapolation of extreme events is generally considered acceptably reliable for periods up to three times the data period, ie. a 30 year data record should provide a reasonable estimate of a 100 year return period extreme. As the wave height increases, its potential for damage increases largely in proportion to the square of the height. Accordingly, estimates of the probability of extreme values of wave height in a region are an essential input to any comprehensive assessment of coastal processes and related coastal works. The longer the period of data available, the more accurate will be the prediction of extreme conditions.

This report presents an overview of various extreme wave height analysis techniques from the literature which are presently in common use throughout the world and makes specific recommendations as to the methods which could be used for the analysis of Departmental wave data. The report does not address the associated question of estimating the periods of extreme waves at this time.

Background to Extreme Value Statistical Analysis

Although some earlier classical references exist, Fisher and Tippet (1928) are generally acknowledged as having established the basic theory of *extreme value statistical analysis*. The theory expounds that if data are independent and identically-distributed (defined later) then there are three, <u>and only three</u>, asymptotic limiting distributions for their maxima (Muir and Shaarawi (1986)) regardless of the parent population (eg. normal, Rayleigh, Lognormal etc). These are traditionally termed the Fisher-Tippet distributions Types I, II and III and are said to represent the "zone of attraction" for the specific underlying distributions. Jenkinson (1955) later represented these three separate extreme value (EV) distributions in a generalised manner as the "GEV" and popular usage also includes the use of the terms EV1, EV2 and EV3. All of the above refer to the same distribution forms although there are some differences in nomenclature between, say, the traditional statistical literature, hydrology literature and meteorological or engineering literature. The FT-I distribution, also sometimes termed the *double-exponential*, was popularised by Gumbel (1958) in regard to flood estimation and often now bears his name by default. Likewise the FT-II is sometimes referred to as the Frechet and the FT-III is sometimes (strictly incorrectly) referred to as the Weibull because of its almost identical functional similarity to Weibull (1939). Note that the EV1 and EV2 are two parameter distributions whilst the EV3 has three parameters.

Notwithstanding the above theoretical treatment of extremes, Galambos (1978) presents the proof (as referenced by Carter and Challinor (1981)) that other asymptotic extreme distributions can exist if the earlier data pre-conditions are not met but they will not necessarily be those of Fisher-Tippet. For example, the GEV assumptions imply regularly sampled data such as annual maxima (Muir and Shaarawi (1986)) which are often not practically available in sufficient number. (It should be noted that "annual", an essentially arbitrary timescale, is only relevant because it encompasses a full seasonal variability and is therefore likely to satisfy the independent and equally-distributed criteria.) At this stage the problem can tend to become confusing and this has resulted in a number of other preferred distributions throughout the literature.

In coastal and ocean engineering practice it will be shown that the only additional distributions generally considered for these types of analyses are the Lognormal, the Weibull and a variation sometimes called the Poisson-Gumbel. Likewise, the Frechet and EV3 are rarely used in practice, many years of empirical experimentation with wave data having lead the fraternity towards those distributions which generally appear to best fit the natural data series - the Gumbel and Weibull types. That said, the most broadly held consensus would then be that there is no strictly theoretical argument for preferring one distribution over another and selection is invariably subjective even though based on various objectively-posed criteria. Some of the issues arising from these facts are presented in the development which follows. Table 1 summarises the various candidate distributions commonly in use for extreme value analysis.

Distribution Type	No. of Parameters	Other Names
FT-I	2	EV1 Gumbel
FT-II	2	EV2 Frechet
FT-III	3	EV3 "Weibull"
Lognormal	2	
Weibull	3	
Poisson-Gumbel	3	

Table 1: Commonly Used Extreme Value (EV) Distributions

The Essential Issues

The following development broadly follows Mathiesen *et al* (1994) but also draws from many other references and also personal experience.

The sequence of events is normally as follows:

- Data selection and sampling
- Model selection
- Goodness of fit and confidence limit estimation
- Establishment of design conditions

Data Selection

Selection of the data set is a key element in the estimation of extreme values. A consistency in approach is required which must be ultimately underpinned by a philosophical statement about the problem to be solved.

The time variation of wave height at any location is a non-stationary process due to the natural progression of storm systems, normally on a seasonal basis, which might influence height, direction, persistence etc. In shallow waters, wave height may be limited by breaking. All of these factors can potentially interfere with the necessary assumptions for extreme value analysis, which are discussed below.

1. Data Independence

Samples must be independently occurring such that only the one peak condition from a single storm event can be used. The "traditional" method of only selecting *annual maximum* series is one philosophical level implying independence (neglecting the effects of longer term weather cycles such as El Nino etc). At the other extreme, use of a *total data* set of hourly sampled data produces many related near-peak events which are more likely to follow the Rayleigh distribution rather than an asymptotic limit assumed for EV analysis. It also follows that data sets cannot be combined unless they are independent samples from the same population. Hence, data from two nearby wave buoys would normally be separately analysed or else a composite non-overlapping (gap filled) data set formed. Data from two sparsely located buoys could arguably be combined to produce a "longer" record if they independently sample events from the same ocean basin, for example, and the goal is to derive a statistic applicable to the basin as a whole.

The inclusion of all independent but not necessarily annual maxima data in the analysis is called use of a *partial duration* series. Selection of the independent data set can be assisted by computing the auto-correlation for the full data time series using a number of time lags, eg. 6, 12, 18, 24h. Selecting a time interval for "independence" can then be based on the shortest lag having a relatively low correlation level (eg. < 0.5). In practice, this is likely to vary in the range of 2 to 5 days, depending on the particular climatology of the area, as represented by the typical time of passage of high and low pressure cells. Having selected the lag as a suitable data window interval, the approach is normally then to also select a threshold for the maxima to ensure very common wave heights during periods of relative calm will not contaminate and swamp the extreme conditions. Selecting a threshold is also called *data censoring* and when combined with independent selection the overall approach is generally termed the *peak over threshold method.*

Finally, since probability of exceedance is fundamentally related to the passage of time, the average time between sampled events must be retained. This is normally done by simply recording the number of events in the overall period of record.

2. Equally Distributed Data

This requirement relates to the homogeneity of the data sample. Care must be taken that all data are likely to be drawn from the same statistical population. If not, then the data must be segregated into either seasonal or other climatological populations and the analyses then separately posed and solved. It is a straightforward matter to recombine a series of separate distributions later to give an overall, say, annual return period curve for the area.

Typical examples where segregation might be needed would be (summer) tropical cyclones versus (winter) easterly trough lows etc. In the case of tropical cyclones, further discrimination might also be warranted on the basis of track in some circumstances. This can also be extended to include directional wave analyses if a predominant bias occurs in the data set. Where known seasonal influences occur but appear equally balanced there may not be a need to separate the data sets.

The presence of significant gaps in a data set can be validly infilled from another nearby source if the circumstances are appropriate. Failing this, the period of the gap must be accounted for in the analysis by adjusting the total period of record to suit.

It is clear that a high degree of subjective decision making may be needed in many practical situations - hence the need for a consistent analysis philosophy.

3. Data Accuracy

Notwithstanding the above considerations, the general accuracy of the wave data used in the analysis must also be considered. The mixing of spectral and non-spectrally derived estimates, for example, could have a significant impact on the analysis and changing individual buoy characteristics, eg. non-directional to directional, may have subtle effects. Either way it is clearly important to have a good knowledge of the long term variability and accuracy of the data set.

One particular aspect which can also directly affect the extreme wave height data set is due to the temporal sampling standards used. If 6 hourly sampling has been supplanted in recent times by 3 hourly, for example, it would be advisable to investigate the likely transfer function between these two base sampling intervals and perhaps adjust the earlier data. The differences will be due to the local climatology, ie. the peakedness of storm hydrographs. Since the aim of the exercise is to estimate the exceedance of extreme heights it is important to try and ensure that the true peak waves in any storm event have been sampled. For this reason the highest frequency-sampled data set (eg. hourly or less if available) would always be used and in any case, some statement about the temporal sampling used should accompany the extreme value analysis estimates, eg. "the analysis provides an estimate of the 3 hourly sampled significant wave height". At a further, finer analysis scale, base spectral sampling periods may also have some impact on the final peak value selected and this could also be investigated if deemed critical to the particular analysis.

Distribution Selection

The background introduction has already alluded to the absence of a clear consensus on this issue. This is unavoidably related to the fundamental problem in statistics, namely uncertainty about the "true" form of the distribution. The FT-I (EV1 or Gumbel) and the Weibull have been shown to provide the most consistent matching with field data for waves. This may be as a result of a natural "law" or simply an artefact of the particular data sets which have figured in prominent analyses, a result of wave buoy inaccuracies or sample aliasing etc.

The forms of the two most popular EV distributions are as follows:

The FT-I, EV1 or Gumbel:

$$F(x) = \exp(-\exp(-(x-a)/b)) \qquad -\infty \le x \le \infty$$

and the Weibull:

$$F(x) = 1 - \exp(-((x-a)/b)^k) \qquad a < x \le \infty$$

where $F(x)$ = $\text{Prob}[H \geq H']$

H = sampled wave height

H' = a specific wave height

and the parameters to be fitted are:

a = the *location*

b = the *scale*

k = the *shape*

(sometimes c is used rather than k)

and the so-called *reduced variates* useful in data plotting and analysis are given by:

Gumbel: $-\ln[-\ln(F(x))]$

Weibull: $(-\ln[1-F(x)])^k$

It is not surprising that the Weibull (with three parameters) will almost always produce a better data match than the Gumbel (with two parameters). However, the Weibull is more "unstable" because of its third parameter (the *shape*). Mathematically, one must assume *a priori* a value for either the *shape* or the *location* in order to fit the curve to the data or indeed to even plot the data. For this reason, many popular methods (eg. Petrauskas and Aagaard (1971) and Goda (1988)) choose a selection of *shape* factors which have been found to generally cover the range of experience for wave height analyses. Typically, a Weibull *shape* near 1.3 will closely match a Gumbel curve for the same data set. However, while the data fits may appear near identical, the extrapolated extreme wave heights at long return periods may be significantly different.

Care should also be taken that if an analysis is undertaken over a finite spatial region having an expected variation in climatology (often a consideration for ocean basin hindcasts), then philosophically one would expect a gradual variation in extreme value parameters. This is more likely to occur with the two parameter Gumbel curve than the Weibull and it may be wise to clamp the Weibull *k* value to obtain a consistency and continuity of predictions.

A further theoretical complication arises when a particular data set is censored by applying a threshold. In this case, if the philosophical stance is that the data which has been removed *do* belong to the distribution which is to be fitted then an allowance should be made for their omission. This necessitates using the truncated forms of each distribution (refer Mathiesen (1994) for example). If however the philosophical stance is that the removed data *do not* form a part of the population to be fitted then they can be ignored, subject to adjustment of the inter-arrival parameter (see later), and the complete form of the distribution used instead. Clearly there is a high degree of subjectivity in this situation. It also begs the question of why data would be removed if it is actually deemed to be a part of the population. This again reflects the fundamental uncertainty of which is the underlying distribution and the search for the "best fit". (A personal view here is that the mathematical exactness of the truncated forms is swamped by the ultimate subjectivity of the selection of the data set.)

Fitting Methods

There are essentially three main methods of fitting statistical data to any distribution. These are listed below in order of commonly accepted complexity, sophistication and accuracy:

1. Method of Moments
2. Least Squares
3. Maximum Likelihood

The method of moments involves the matching of the mean and standard deviation of the sample data to that of the posed distribution. It can be quite adequate in situations where the data is (fortuitously) well matched to the distribution. Strictly, this method cannot be used with censored data sets.

The least squares method is the traditional minimisation of the square of the offsets of a line from a series of points. The method therefore requires a series of independent probability points upon which to minimise the offset to the line. This necessitates assigning a probability to each data item to form a pair of datum (eg. {H, probability of H}). This is a circular argument since ultimately the probability of H is the aim of the data fitting exercise. Nevertheless, ranking of the data sample is undertaken and slices of probability are allocated to the ranked position of the sample. How the allocation of probability is done will clearly influence the ultimate position of the line and debate over the "best" way of doing this continues to rage. It is generally, but not universally, accepted that an unbiased estimate of the so-called *plotting position* is required. Accordingly, the Weibull (1939) method (as originally used by Gumbel (1958)) is generally discredited and many studies have shown its high bias effects. The method of Gringorten (1963) is accepted as the preferred plotting position formula for the Gumbel (FT-I, EV1) while Petrauskas and Aagaard (1971) and Goda (1988) produced slightly different "unbiased" formula for the Weibull distribution. (The reason why there could be two different formula for the same distribution is because each is a simple approximation to a complex function which requires integration and so subjective opinion plays a part - again.) Once the data pairs are selected the method is straightforward in its application of linear regression in the *reduced variate* space.

The method of *maximum likelihood* is generally favoured by statisticians since it does not rely on the plotting position, is asymptotically unbiased and also efficient (it produces low variance estimates). It does this by constructing a mathematical likelihood function which is coupled with the chosen distribution form, and the parameters of the distribution are obtained by locating the maximum of the function. Opponents of this method claim it is biased for small samples (eg. < 20), too complex to solve and not significantly better than the other alternatives. Certainly there are some difficulties with the Weibull distribution where the solution can only be given simultaneously for all three parameters if k>2.0 (Bury(1975)), which is beyond the (empirically) accepted "wave" range of 1 to 2. Accordingly, either *k* or *a* is then fixed to allow the optimisation to proceed on a two parameter basis, *a* most often set to a value just below the smallest sample. The reasons for avoiding the method still seem somewhat dismissive given that it is undoubtedly objective and the necessary numerical methods have been widely available for the past twenty years. It is also acknowledged that maximum likelihood is the most robust method, being much less sensitive to outliers in the data than the least squares approach. Unfortunately, in order to visualise the data fit, the use of plotting positions is still required and comparisons of the points (especially the outliers) and the fit are inevitable. Ironically, the least squares fit is then in the enviable position of normally showing the "better" visual fit to the data.

Goodness of Fit Tests

Since the choice of distribution is almost always somewhat subjective, the question remains as to whether it was a "good" choice. The usual approach is either to adopt a particular distribution as a standard (eg. Gumbel) and compare the fit against some criteria, or, to select a range of distributions to fit the data and pick the "best" from that set.

The use of Chi-squared or correlation coefficient testing is normally considered too insensitive for this task.

Visual comparisons are essential in aiding comprehension of the fit and noting the presence of outliers but not always quantitatively useful for discriminating goodness-of-fit between one or more possible distributions.

The quantile-quantile (Q-Q) plot is often recommended whereby the computed and observed wave heights are cross-plotted with the expectation of a straight line occurring. This method also allows the overplotting of different distributions which is not normally possible for the Weibull because k appears in the reduced variate term.

Beyond Q-Q is the realm of distribution-free tests such as the Kolmogorov-Smirnoff and Anderson-Darling (both favouring the middle of the distribution) and the Cramer-Von Mises (said to favour the tail or low probability region). All these tests perform best for large data sets.

Monte Carlo methods are also often used for determining goodness of fit (eg. Petrauskas and Aagaard (1971), Goda (1988)). The technique is to randomly generate a large number of synthetic data sets based on the fitted distribution and compare, say, the mean squared deviation (MSD) of these points chosen from the distribution in each case with the original data set MSD. If the original data MSD is shown to be better than 50% of the randomly chosen samples then it is accepted as a “good” fit. Ranking of alternate distributions is then similarly applied.

In recent years *The Bootstrap* method (Effron (1979)) has become more widely used and can be thought of as a variation on the direct *Monte Carlo* method. *The Bootstrap* deals with the actual sample data set rather than the assigned form of distribution and can therefore claim to be non-parametric.

Confidence Limits

From a practical design perspective it is always desirable to have some indication as to the range of values a particular return period wave height estimate might take, on the assumption that the stated value represents a mean position. While the alternative question is “how wide is the return period estimate of a specific wave height likely to be”, this rarely seems to be considered but is equally valid.

The confidence limit surrounding an estimate is normally based on the assumption that the chosen distribution type is indeed the true distribution type, ie. as derived from the goodness of fit test. The problem then to be considered, given that the data record is not infinitely long, is the accuracy in estimation of the distribution parameters themselves. Put another way, if the data set were much longer (such as being based on another 10 years of records) is it likely that the fitted distribution parameters would be very different? If the distribution parameters might indeed be different then this translates into an uncertainty surrounding any particular estimate, such as the 100 yr return period wave height.

Monte Carlo and now Bootstrap methods are again favoured in this situation since numerical experiments can be performed to generate synthetic data sets which mimic the original set. Normally, based on such trials, tabulated statistics can be produced which are a function of the sample size and the sample moments (eg. mean, standard deviation and sometimes skewness etc). Lawless (1974) presents detailed arguments on the subject of confidence limits for both the Gumbel and Weibull types.

Return Period and Encounter Probability

As most engineers are aware, probability of extreme events is often conveniently given as a *Return Period* expressed in years. The return period is simply the *average* (ie. the mean) time in years between equalling or exceeding a particular event level. Some prefer the term *Average Recurrence Interval* or ARI, but the meaning is the same. Use of the *Return Period* is variously misunderstood, however, with the public sometimes assuming that a given event will not occur again until a fixed period of time has elapsed. There appears no easy way to overcome this confusion except to always express probability of exceedance on an annual basis, but this leads to dealing with quite low values which are easily confused and misquoted anyway.

It should be noted that there is nothing special about the time units chosen for the return period but that years are a convention which suits the purpose of implying low likelihood. In the case of using partial duration series the essential return period is actually the average inter-arrival time of the chosen events, which for a random process is the inverse of the *Poisson* parameter λ - the average number of events, say, per year. This event-time-interval must therefore be accounted for in the calculation of the return period when expressed as years, ie.

$$F(\, x(T_r)\,) = 1 / (\lambda\, T_r)$$

where T_r is the return period in years.

Remembering that T_r is simply the average period, this implies that the same event magnitude may sometimes occur more or less frequently in any particular similar length of time. This is due to the natural variability of random sampling from a population. In fact, the N year event magnitude has a 63% chance of being equalled or exceeded in any N year period. This derives from the following formula due to Borgman (1963) which relates the return period R to an encounter interval L , the so-called *encounter probability* P in that interval being:

$$P(\, x(T_r),\, L\,) = 1 - \exp\, (\, -L/\, T_r)$$

or, via the often used but approximate form

$$P(\, x(T_r),\, L\,) = 1 - (1 - 1/\, T_r)^L$$

This particular formula is often preferred over the return period alone because it relates the risk of an event to a potential lifetime of exposure, which can vary depending on the use of the data. For example, the design of temporary works during construction of a port could be based on a criteria that the project will accept a 10% chance that a given wave height could be exceeded within a six month period. This means that the equivalent return period being considered is actually close to the 5 year value. Similarly, the 500 year return period event might be considered as a criteria for the lifetime of the permanent port structures since it has about a 10% chance of being exceeded in any 50 year period of time. Figure 1 presents a graphical view of the relationship between encounter probability, lifetime and return period.

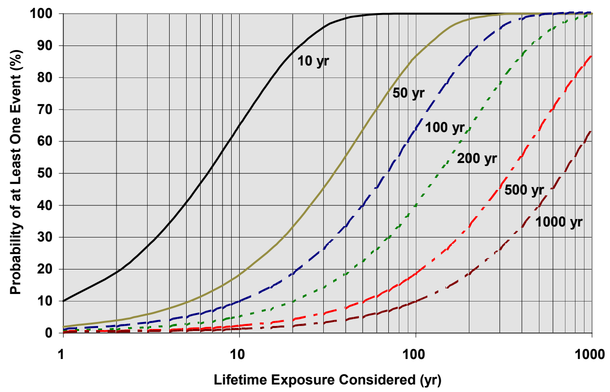

Figure 1: Encounter Probability versus Return Period

A Brief Review of Some Published Methods

Like all evolving sciences, the literature is peppered with different approaches and varying theoretical arguments. Whilst the references presented here are not exhaustive they represent a good cross-section of the work done which specifically relates to the estimation of extreme wave heights for engineering purposes. In an attempt to add an historical understanding to the major developmental themes, Figure 2 ascribes a rough dependency to the studies summarised here in the form of common symbol shapes. A particular symbol shape signifies that in the time progression, the author has drawn either directly or indirectly from earlier work having the same symbol shape. (This is an attempt to avoid the rather complicated visual linkages which might occur using lines and arrows.)

The absence of a final single position in this development should rightly be interpreted as an evolution which is still continuing, albeit on a narrower scale. The following summaries omit the essentially statistical works (as indicated in Figure 2) which underpin the engineering studies, although from time to time relevant points derived from these works are cited.

Phase 1 - 1960's Awareness

The theoretical works by Fisher and Tippet (1928), Weibull (1939), Jenkinson (1955) and Gumbel (1958) were well established by the early 1960's and had found wide application across various disciplines.

Bretschneider (1958) and Borgman (1961) are perhaps the earliest references to the problem by researchers who became prominent in the wave prediction arena. Borgman's early work is based on his student thesis and necessarily theoretical; Borgman (1963) followed up with a more readable entry into the engineering literature and introduces the very practical concept of non-encounter probability. Draper (1963) presents an example of solving the very problem which is the subject of this report. He used a *total data* approach and showed that a Lognormal distribution fitted his data well. The total data approach using the Lognormal distribution for estimating *extremes* is now not encouraged, mainly on the theoretical grounds of data independence and the asymptotic limits of the Lognormal actually being FT-I. Its accepted use continues however in the area of operational wave exceedance and persistence.

As an aside, Gringorten (1963) at this time had identified the need for unbiased probability plotting rules for the EV1 (Gumbel) distribution. This fact has not been universally agreed or taken up but will be shown to be a recurring theme in the least-squares method of distribution fitting.

Phase 2 - 1970's Practical Needs

During this time the offshore oil and gas industry spurred analyses for extreme design criteria in many parts of the world. Thom (1971, 1973) was amongst the first to try to rationalise a global view of extreme winds and waves.

Petrauskas and Aagaard (1971) of Chevron Oil presented one of the most enduring and arguably advanced analysis methods for extreme waves which has probably underpinned the majority of offshore installations in the past 20 years. Their method is based on least squares fitting of the data set to eight separate distributions - the Gumbel and seven Weibull, based on a range of *shape* factors. The "best" distribution is then assigned by a statistical test involving *Monte Carlo* sampling and a goodness of fit measure. They also considered confidence limits, utilised the Gringorten (1963) plotting rules as well as developing corresponding unbiased plotting rules for the Weibull distribution, recognised the role of the Poisson inter-arrival parameter and advocated use of encounter probabilities rather than return periods.

Other developments through this period can be attributed to Nolte (1973), addressing the emerging

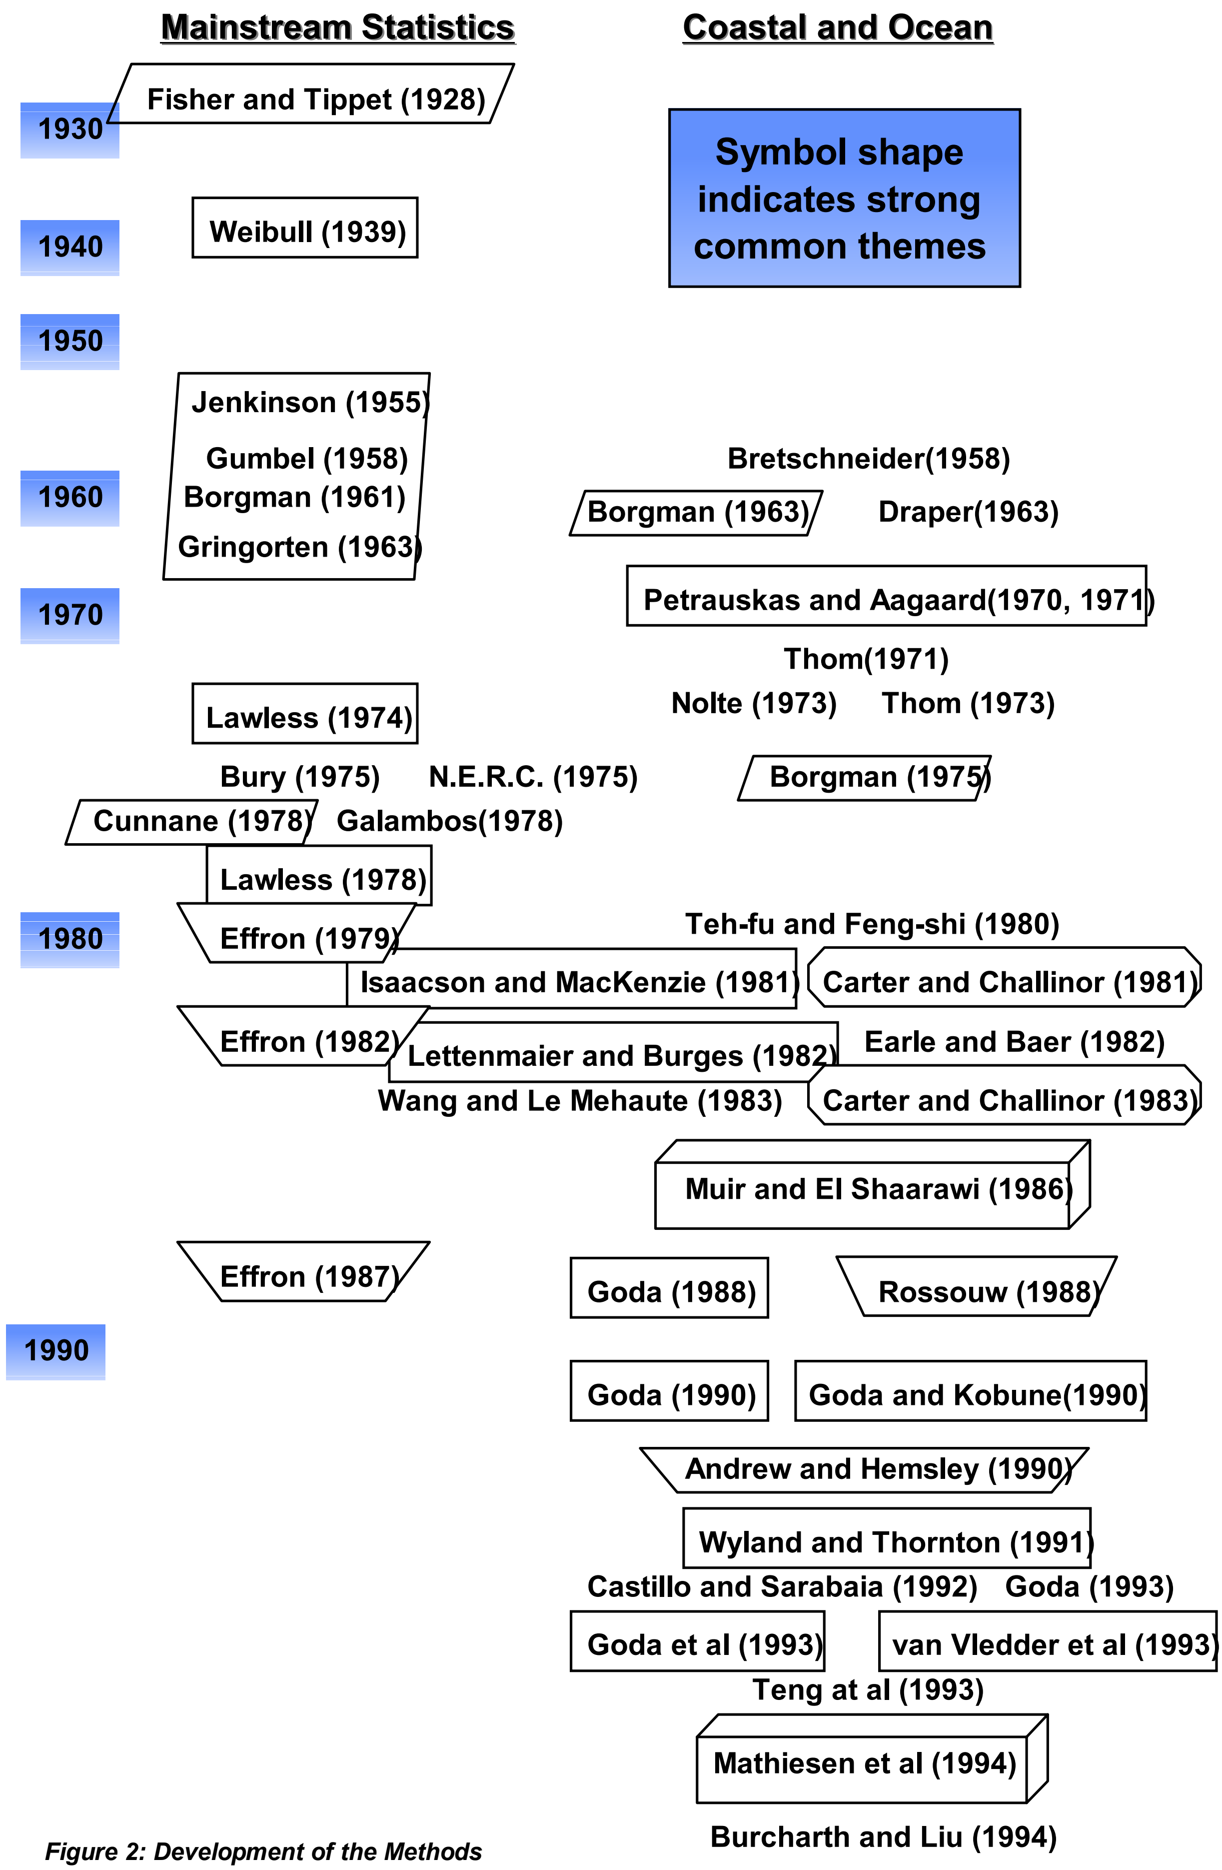

Figure 2: Development of the Methods

Arctic developments and Borgman (1975) now firmly in the engineering context.

The mid-1970's also saw significant developments in associated areas such as the emerging work by Lawless (1974, 1978) and an extensive study on flood frequencies in the British Isles (N.E.R.C. (1975)). From this latter work Cunnane (1978) argues most strenuously the case for unbiased plotting formula.

Phase 3 - 1980's Experimentation and Review

Teh-fu and Feng-Shi (1980) emphasised the difficulties of using annual maxima series in regard to typhoon estimates. They proposed the use of the compound Poisson-Gumbel distribution to allow incorporation of all events above an arbitrary threshold, rather than the Poisson return period scale shift used by others.

Isaacson and MacKenzie (1981) present perhaps the first attempt at an objective review of the subject for general wave analysis applications. They cite Petrauskas and Aagaard (1971) extensively and generally retain the value thereof as well as introducing the earlier work of others. They treat the FT-III in some detail, pointing out the difficulty of plotting data without some *a priori* selection of one of its three parameters, normally the *shape* or the *location*.

In the same year Carter and Challinor (1981) looked in detail at the differences between annual maxima and monthly maxima data series for the EV1. They found the annual maxima approach could underpredict extreme values and that if a monthly series were independent and EV distributed, then the annual maxima could not be EV distributed. They recommended a compounded annual distribution based on the joint probability product of the monthly (or seasonal) maxima. Their paper is the first to endorse use of the maximum likelihood estimator rather than the least squares method and plotting positions. They followed in Carter and Challinor (1983) with a review of the various data fitting methods, concluding like Gringorten (1963) that "Gumbel's method using the Weibull plotting position" was highly biased and should not be used. They found the method of moments almost as good as the maximum likelihood for their (large) data series.

Lettenmaier and Burges (1982) did a similar study which reached the same conclusions, again for the EV1, and criticised continued textbook use of the so-called Weibull plotting formula.

Earle and Baer (1982) however did not heed the call and undertook studies based on the Lognormal distribution and using the Weibull plotting formula. Their simulation work designed to address the question of uncertainties is judged by Goda (1988) as inadvertent proof of the Weibull plotting bias. Others criticised their use of the Lognormal as an inappropriate choice for *extreme* value analysis.

Throughout this period Effron (1979, 1982, 1987) was proposing a relatively bold statistical sampling theory to be known as *The Bootstrap*, later to be taken up by others and applied to this area.

Wang and Le Mehaute (1983) examine the issue of independence of data samples and develop relationships to show its influence on the extrapolation to design heights. Curiously they determine that their 6 hourly data is apparently independent. They also present a technique for estimating how the length of the data sample directly affects the confidence interval of, say, the 100 year return period.

Muir and El Shaarawi (1986) present a very detailed review which also addresses general uncertainty in measured and analysed wave heights prior to any extreme value extrapolation. They cover aspects such as distribution type, data selection (including censoring), goodness of fit and confidence limits. They recommend use of the maximum likelihood estimator with (usually) censored data.

Goda (1988) enters the field with an extension of the Petrauskas and Aagaard (1971) method. He reduces the candidate Weibull distributions from seven to four and develops slightly modified plotting position formula. Also using *Monte Carlo* methods, he examines the potential bias in estimates when the true distribution in unknown and develops confidence limits based on empirical factors. This method has been implemented by Leenknecht et al (1992) in the US Army Corps of Engineers Coastal Engineering Research Centre software package ACES - **A**utomated **C**oastal **E**ngineering **S**ystem.

Rossouw (1988) presents an analysis of eight years of wave buoy data from South Africa. He discusses use of lag correlation to derive a suitable sample of maxima and chooses to group the data in a seasonal way. Curiously he found his data was insensitive to independence issues when using the method of moments for fitting. He provides a useful description of the differences between random variations of a sample and confidence limits of its parameters and is the first to advocate use of *The Bootstrap* for deciding between distributions.

Phase 4 - 1990's Debate Continues ...

Andrew and Hemsley (1990) also apply *The Bootstrap* for this purpose as well as providing a useful review of the general extreme analysis approach. They warn of the ability of the Weibull and EV1 to provide equally likely fits to data but to predict quite different extreme height at long return periods. They use a data splitting technique to separately drive *The Bootstrap* method but use of the (discredited) Weibull plotting formula may impact their results.

Goda and Kobune (1990) reject *The Bootstrap* as being ineffective in discriminating between candidate distributions. In this paper they also considered the EV2 and developed a plotting formula for its use.

Goda (1990) is essentially the same methodology as his preceding publications but with a wider treatment of some related aspects.

Wyland and Thornton (1991) show how three different hindcast data sets respond to the various types of extreme value analyses (mainly based on Isaacson and MacKenzie (1981) and Muir and El Shaarawi (1986)). They criticise the MSD error criteria formulation of Petrauskas and Aagaard (1971) and also Goda (1988). In their Weibull analyses they chose *a priori* location values rather than shapes, which would appear to offer some benefit.

Statisticians Castillo and Sarabaia (1992) entered the field, sweeping away the burdens of the past and giving specific advice on the type of EV distributions which should be used. They advocated restricting analyses to the extreme tail of the data set only by using a weighted least squares approach. Goda (1993) then responded in a vigorous exchange which challenged their basic understanding of the problem.

In 1990, the Section on Maritime Hydraulics of the International Association for Hydraulic Research (IAHR) had organised a Working Group on Extreme Wave Statistics with the aim of reaching a mutual understanding of the merits and demerits of the methods used in extreme statistical analysis. Representatives were invited from eight separate institutions to cooperatively develop a recommended procedure. The first results of the Working Group are presented in the Proceedings of the Second International Conference on Ocean Wave Measurement and Analysis as Goda et al (1993). In this paper, different methods of fitting an example Weibull (k=1.4) data set are described. The trial concluded maximum likelihood was best for uncensored data and least squares better for censored data (in the absence of a technique for maximum likelihood which correctly treats censored data). While individual estimates of the 100 year wave height varied with a standard deviation of between 6% to 12%, the overall bias was only 3%.

The companion paper by van Vledder et al (1993) allowed each of the eight groups to use their preferred method to determine the extreme value distribution of two sets of data. The 100 year wave height estimates from all groups were within 10% and the 90% confidence bands were within 10% of the 100 year wave height. The paper comments specifically on data censoring, pointing out that as the threshold for inclusion was decreased, the 100 year wave height prediction increased.

In the same publication, Teng et al (1993) use 13 years of wave buoy data and fit a Lognormal distribution (contrary to popular opinion). They also modify the plotting formula of Gringorten (1963) against the recommendation of Goda et al (1993).

Mathiesen et al (1994) presents the final recommendation of the IAHR working group. They conclude (or rather concur) that the partial duration (points over threshold) approach is preferred and that the Weibull distribution is superior. No recommended fitting method is proposed. Various goodness of fit tests are advocated, especially the Q-Q plots, and use of confidence limits derived from *Monte Carlo* modelling are suggested. All in all this paper is rather tame - there are clearly still some strongly held beliefs between the various institutions involved.

The final reference cited here is that of Burcharth and Liu (1994). They advocate EV1 and Weibull as the candidate distributions and claim least squares is better than maximum likelihood when the true distribution is unknown, based on the data set they used. Interestingly they maintain the Weibull plotting position is the preferred one amongst all the options available!

Conclusions and Recommended Method of Analysis

The “best” method of extreme value analysis is yet to be fully agreed amongst a fairly wide range of differing views, each affected by disciplinary, organisational, cultural and personal factors. The previous section highlights many of these different approaches. There is a degree of “baggage” evident in the literature which relates to the early lack of computational power for performing data fits such as the maximum likelihood and also undertaking *Monte Carlo* style experiments. There would seem to be much less reason now for not adopting such methods and *The Bootstrap* is an emerging technique which will probably eventually take root. The existing methods are heavily influenced by the subjectivity of institutional approaches - this is evident in the Mathiesen (1994) summary paper which unfortunately fails to draw more specific recommendations on this topic. My conclusion here is that the “best” methodology is not necessarily generally available at this time but will gradually emerge.

With this in mind, the presently recommended approach for the analysis of Departmental wave data would be that of Goda (1988) as implemented by Leenknecht et al (1992) in the CERC ACES software system (Version 1.07). The software is available in the Department, the necessary analysis procedures are well documented and a range of output options are available. All that remains is a consistent approach to data set sampling based on the guidelines presented earlier. Appendix A provides some guidance on the data sampling methods which could be developed.

It is possible that ACES will be updated in time to more robust methods but in the interim it will provide a consistent and recoverable standard for extreme wave height analysis which can also be used for other extreme data sets such as storm tides with the appropriate interpretation being applied. Appendix A includes an example data analysis using ACES and explains the necessary steps in detail.

References

Andrew, M.E. Smith, O.P. and McKee, J.M. (1985). Extremal Analysis of Hindcast and Measured Wind and Wave Data at Kodiak, Alaska. Dept of the Army. Tech Report CERC-85-4. June.

Andrew, M.E. and Hemsley, J.M. (1990). Resampling Approach to Extreme Wave-Height Analysis. Jnl Waterway, Port , Coastal and Ocean Engin. ASCE. Vol 116 No. 4. Jul/Aug. pp 444 - 458.

Borgman, L. (1961). The Frequency Distribution of Near Extremes. Jnl Geophysical Research. Vol 66 No. 10. Oct. pp 3295 - 3307.

Borgman, L. (1963). Risk Criteria. Jnl Waterway, Port, Coastal and Ocean Div. ASCE. Vol 89 No. WW3. Aug. pp 1 - 35.

Borgman, L. (1975). Extremal Statistics in Ocean Engineering. Proc. Civil Engin in the Oceans, III. ASCE. pp 117 - 133.

Bretschneider, C.L. (1958). Selection of Design Wave for Offshore Structures. Jnl Waterways and Harbors Div. ASCE. Vol 84 No. WW2. Mar. pp 1 - 37.

Burcharth, H.F. and Liu, Z. (1994). On the Extreme Wave Height Analysis. Proc Int Conf Hydro-Technical Engin for Port and Harbor Constr. Yokosuka. Oct. pp 123-142.

Bury, K. (1975). Statistical Models in Applied Science. John Wiley and Sons. 625pp.

Carter, D.J.T. and Challenor, P.G. (1981). Estimating Return Values of Environmental Parameters. Qrtly Jnl Royal Meteorological Society. Vol 107. pp 259 - 266.

Carter, D.J.T. and Challenor, P.G. (1983). Methods of Fitting the Fisher-Tippet Type 1 Extreme Value Distribution. Jnl of Ocean Engineering. Vol 10 No 3. pp 191 -199.

Castillo, E. and Sarabia, J.M. (1992). Engineering Analysis of Extreme Value data: Selection of Models. Jnl Waterway, Port, Coastal and Ocean Engin. ASCE. Vol 118 No. 2. Mar/Apr. pp 129 - 146. (See also discussion by Goda, Y. (1993) Jul/Aug)

Cunnane, C. (1978). Unbiased Plotting Positions - A Review. Journal of Hydrology. Vol 37. pp 205 - 222.

Draper, L. (1963). Derivation of a Design Wave From Instrumental Records of Sea Waves. Proc. Institute of Civil Engineers. Vol 26. pp 291 - 304.

Earle, M.D. and Baer, L. (1982). Effects of Uncertainties on Extreme Wave Heights. Jnl Waterway, Port, Coastal and Ocean Div. ASCE. Vol 108 No. WW4. Nov. pp 456 - 478. (See also discussion by Guedes Soares, C.)

Effron, B. (1979). Bootstrap Method - Another Look at the Jackknife. The Annual of Statistics. Vol 7 No. 1. pp 1 - 26.

Effron, B. (1982). The Jacknife, the Bootstrap and Other Resampling Plans. Soc for Industrial and Applied Mathematics. Philadelphia.

Effron, B. (1987). Better Bootstrap Confidence Intervals. Jnl American Statistical Assoc. Vol 82 No. 397. pp 171 - 200.

Fisher, R.A. and Tippet, L.H.C. (1928). Limiting Forms of the Frequency Distributions of the Largest or Smallest Member of a Sample. Proc. Cambridge Philosophical Society. Vol 24. pp 180 - 190.

Galambos, J. (1978). The Asymptotic Theory of Extreme Order Statistics. J. Wiley and Sons. New York.

Goda, Y. (1988). On the Methodology of Selecting Design Wave Height. Proc. 21st Int. Conf. on Coastal Engin. ASCE. pp 899 - 913.

Goda, Y. and Kobune, K. (1990). Distribution Function Fitting for Storm Wave Data. Proc. 22nd Int. Conf. Coastal Engin. ASCE. pp 18 - 31.

Goda, Y. (1990). Distribution of Sea State Parameters and Data Fitting. In Herbich, J.B. (Ed). Handbook of Coastal and Ocean Engineering. Gulf Pub Coy. Houston. Chap 7.

Goda, Y. Hawkes, P. Mansard, E. Martin, M.J. Mathiesen, M. Peltier, E. Thompson, E. and Van Vledder, G. (1993). Intercomparison of Extremal Wave Analysis Methods Using Numerically Simulated Data. Proc. 2nd Int. Symp. Ocean Wave Measurement and Analysis. ASCE. July. pp 963 - 977.

Gringorten, I.I. (1963). A Plotting Rule for Extreme Probability Paper. Jnl Geophysical Research. Vol 68 No.3. pp 813 - 814.

Gumbel, E.J. (198). Statistics of Extremes. Columbia Univ Press. New York.

Isaacson, M.Q. and MacKenzie, N.G. (1981). Long-Term Distributions of Ocean Waves: A Review. Jnl Waterway, Port, Coastal and Ocean Div. ASCE. Vol 107 No. WW2. May. pp 93 - 109.

Jenkinson, A.F. (1955). The Frequency Distribution of the Annual Maximum (or Minimum) Values of Meteorological Elements. Qrtly Jnl Royal Meteorological Society. Vol 81. pp 158 - 171.

Lawless, J.F. (1974). Approximations to Confidence Intervals for Parameters in the Extreme Value and Weibull Distributions. Biometrika. Vol 61. pp 123 - 129.

Lawless, J.F. (1978). Confidence Interval Estimation for the Weibull and Extreme Value Distributions.

Technometrics. Vol 20 No. 4. Nov. pp 355 - 364 plus discussion.

Leenknecht, D.A. Szuwalski, A. and Sherlock, A.R. (1992). Automated Coastal Engineering System - User's Guide and Technical Reference. Waterways Experiment Station. US Army Corps of Engineers. Coastal Engineering Research Center. Dept of the Army. NTIS. Sept. 1992.

Lettenmaier, D.P. and Burges, S.J. (1982). Gumbel's Extreme Value I Distribution: A New Look. Jnl of Hydraulics Div. ASCE. Vol 108 No. HY4. Apr. pp 502 - 514.

Mathiesen, M. Goda, Y. Hawkes, P.J. Mansard, E. Martin, M.J. Peltier, E. Thompson, E.F. and van Vledder, G. (1994). Recommended Practice for Extreme Wave Analysis. Jnl of Hydraulic Research. IAHR. Vol 32 No. 6. pp 803 - 814.

Muir, L.R. and El-Shaarawi, A.H. (1986). On the Calculation of Extreme Wave Heights: A Review. Journal of Ocean Engineering. Vol 13 No. 1. pp 93 - 118.

N.E.R.C. (1975). Flood Studies Report Vol II: Meteorological Studies. Natural Environment Research Council. London.

Nolte, K.G. (1973). Statistical Methods for Determining Extreme Sea States. Proc. 2nd Int. Conf. Port and Ocean Engin. Under Arctic Conditions. pp 705 - 742.

Petrauskas, C. and Aagaard, P.M. (1971). Extrapolation of Historical Storm Data for Estimating Design Wave Heights. Society of Petroleum Engin. Trans. Vol 25. pp 23 - 37.

Rossouw, J. (1988). Design Waves and Their Probability Density Functions. Proc. 21st Int. Conf. Coastal Engin. ASCE. pp 822 - 834.

Teh-fu, L. and Feng-shi, M. (1980). Prediction of Extreme Wave Heights and Wind Velocities. Jnl of Waterway, Port, Coastal and Ocean Div. ASCE. Vol 106 No. WW4. Nov. pp 469 - 479.

Teng, C. Timpe, G.L. Palao, I.M. and Brown, D.A. (1993). Design Waves and Wave Spectra for Engineering Applications. Proc. 2nd Int. Symp. Ocean Wave Measurement and Analysis. ASCE. July. pp 993 - 1007.

Thom, H. C. S. (1971). Asymptotic Extreme Value Distributions of Wave Heights in the Open Ocean. Journal of Marine Research. Vol 29. pp 19 - 27.

Thom, H. C. S. (1973). Extreme Wave Height Distributions Over Oceans. Jnl Waterways, Harbors and Coastal Engin Div. ASCE. Vol 99 No. WW3. Aug. pp 355 - 374.

van Vledder, G. Goda, Y. Hawkes, P. Mansard, E. Martin, M.J. Mathiesen, M. Peltier and E. Thompson, E. (1993). Case Studies of Extreme Wave Analysis: A Comparative Analysis. Proc. 2nd Int. Symp. Ocean Wave Measurement and Analysis. ASCE. July. pp 978 - 992.

Wang, S. and Le Mehaute, B. (1983). Duration of Measurements and Long Term Wave Statistics. Jnl of Waterway, Port, Coastal and Ocean Div. ASCE. Vol 109 No. 2. May. pp 236 - 249.

Weibull, W. (1939). A Statistical Theory of Strength of Materials. Ing. Vet. Ak. Handl. Stockholm.

Wyland, R.M. and Thornton, E.B. (1991). Extremal Wave Statistics Using Three Hindcasts. Jnl of Waterway, Port, Coastal and Ocean Engin. ASCE. Vol 117 No. 1. Jan/Feb. pp 60 - 74.

Appendix A - Suggested Analysis Procedure

This section presents some suggestions for data sampling and provides an example of the use of the ACES (Version 1.07) analysis package for extreme wave height analysis. At the outset it is worth noting that this package is somewhat cumbersome in regard to its data input/output requirements - “user antagonistic” would be an appropriate description. However, there is a reasonably comprehensive user guide available and this is recommended reading in order to avoid some of the more subtle pitfalls.

It will soon become apparent in the following development that there is not necessarily a step-by-step progression when preparing the recorded wave height data set for extreme value analysis. Many of the decisions which need to be made may be dependent on other decisions and so typically an iterative approach is required. This stems as much from the philosophical problem as to what constitutes an “event” as well as more physical issues such as accuracy of sampling of the true peak values.

<u>Step 1: Data Assembling</u>

The base time series of wave height should first be assembled from either a single site record or an amalgamated (gap filled) record as appropriate, dependent on whether gap filling can be justified from an adjacent site. The justification for gap filling would need to consider aspects such as proximity, exposure and water depth. The highest sampled-frequency data should be used in the first instance but this could be initially reduced to (say) an hourly *peak* data set using windowing if retention of a full data set is prohibitive.

Where data sets which have different base sampling frequencies (eg. six hourly versus hourly etc) are to be combined and considered as a single data set, some consideration should be given to standardising the record. For example, an hourly data set is preferred since it is more likely to sample the true peak, depending on the characteristic storm hydrograph shapes in the region. To convert, say, a six hourly data set into an equivalent one hourly set requires the development of a transfer function between the two sets. This would only be reliable if the available one hourly data record were of comparable or longer duration than the six hourly set. By comparing the characteristic errors between the six hourly and one hourly discretely sampled points from the one hour data set alone, it would be possible to develop a regression line relating the potential error in the “true” (one hour) **peak** value when using only six hourly sampled data. This could then be applied to the six hourly recorded data set to provide a “synthetic” one hourly peak data set equivalent. The success of this technique would depend on whether or not similar weather conditions were experienced during the two periods when the sampling frequency was different. To be effective, this approach also needs to be done within the context of individual events rather than all the data taken together, since it is the adjusted **peak** value of the event which is ultimately sought. On this basis it may need to be deferred until the results of Step 2 are known.

Consideration at this stage should also be given to possible separation of the data set to ensure equally distributed populations. For example, a simple seasonal split may be indicated. In other situations the populations may be mixed throughout the year due to tropical and extra-tropical influences at various times and the separation may need to be undertaken on the basis of a more detailed assessment of weather charts, wind records or wave direction (where available). Separation could also be achieved as part of the event sampling process itself (refer Step 2) on the basis of specifically identified event *parameters*. For example, wave hydrograph shape parameters (long versus short duration, sharply versus mildly peaked etc) may indicate appropriate population splits.

<u>Step 2: Event Sampling</u>

Initially, a base threshold of wave height for consideration is normally selected. This is logically connected to the event population itself although *a priori* it may not be obvious which value to use. The rule of thumb is to settle on a threshold value which will at least trigger something like the “expected” number of events each year. The “expected” number might be loosely based on a meteorological overview, for example, or be derived from operational exceedance or persistance information.

The use of auto-correlation analyses using a series of time lags may prove a useful initial tool in this process. The minimum time interval between successive events should be somewhat longer than the time lag for which the auto-correlation function is 0.3 to 0.5 (Mathiesen et al (1994)). A period of several days is the norm but this will depend on the particular wave climatology of the area.

Alternatively, assuming the above approaches do not lead to an obvious means of identifying the data of interest, it may be useful to attempt a parameterisation of the time series to identify “events” with certain characteristics. For example, and by way of suggestion only, the data could first be filtered to retain, say, peak three hourly values. This would remove a degree of diurnal variability and noise (note that there is no particular need to retain accurate timing for the data peaks at this stage so long as the filtering period is much less than the inter-event interval). The smoothed (filtered but with peaks retained) data set could then be analysed akin to a traditional “zero crossing” analysis or a discrete event variant on a persistance analysis to identify the characteristics of the time series. Properties of interest might include the statistics of half-height duration of peaks, ie. the period of time that wave height exceeds half of the peak height. From such statistics a picture should emerge of the classes of events as a function of magnitude, duration and perhaps shape.

The above method, if successful, would directly produce the sample of peak events of interest. However, if a simpler "data window" is selected based on time lags or other means then the series must still be sampled to obtain the extreme data set. On this basis the time series is conditionally sampled so that the largest wave height (above the selected threshold) is selected within each data window. This essentially means moving the window each time to the time of the next occurrence of a wave height which exceeds the threshold value, selecting the highest wave in the window and then repeating the steps above.

Preparing a series of plots showing the event count yield from these procedures based on a range of assumptions can be of considerable benefit. For example, plotting the number of detected events versus the window interval (eg. the time lag) or versus the selected height threshold can be useful.

It should be remembered also that the wave period associated with each peak wave height might need to be retained as part of a steepness assessment. The unfiltered (raw) time series would then need to be referenced to locate the relevant data of interest.

Step 3: Analysis

The ACES system is relatively straightforward in terms of analysing the final data set. The "Extremal Significant Wave Height Analysis" option is contained within the so-called *Single Case* "Wave Prediction" module. (Note that the technique is not limited alone to Hs but could also be applied to Hmax or other quantities such as storm tide levels. The output labelling however cannot be changed and so data would need to be additionally plotted and presented.)

Care should be taken to ensure ACES has been correctly setup via the INSTALL program if hardcopy graphics output is required. Also, on entering ACES, it requires specification of the display units (eg. metric) and the print device. It also permits explicit naming of the default output files, which in this case are the "TRACE.OUT" file which will contain any on-line edits made to the input data file, and the "PLOTDAT1.OUT" file which will contain the details of the data analysis and can be used for subsequent external plotting and presentation. No other files are used.

ACES assumes initially manual data entry but the format of the file is readily derived and should be able to be externally generated to suit.

The input parameters are as follows:

1. units of the data (eg. meters)
2. the *declared* integer number of events in the data period (NT)
3. the data period in decimal years (K)
4. the site water depth - a simple check is made on extrapolated heights and a warning issued
5. a title for the data set
6. the data set itself, being N values (maximum up to NT values)

ACES calculates the Poisson parameter LAMBDA and the censoring parameter NU based on this information and these are displayed on the output for reference purposes. LAMBDA is simply NT/K but the software seems to override the reported value of NT to be be not less than K. There should be no theoretical reason why the time period between events should be limited to a maximum of one year. In any event this overriding of NT appears to not affect the true calculation (probably a "bug").

NU is simply N/NT indicating the proportion of events being analysed via the truncated distribution form. Hence, if NT = N then the data set at this point is considered *"uncensored*". If fewer than NT values are specified then the data set is considered *"censored"* by the user, ie. not all of the available values are to be included in the best fit analysis. If censored, then the truncated form of the particular distribution is invoked rather than the unlimited form.

Figure A-1 presents the data input screen for the case of an example set comprising 50 extreme Hs values (Hs_TEST.IN) from a 10 year record in 20m depth, based on a 2.6m threshold value. A further screen is provided to select from a range of confidence limits (95% is the default and is used here). The higher the confidence limit selected, the wider the indicated spread of values is likely to be.

After analysis, ACES presents a series of summary screens showing the fits produced for each of the candidate distributions, ie. FT-I and Weibull with k = 0.75,1.0,1.4 and 2.0.

The first screen is given as Figure A-2 which reflects the base parameter values of N, NT, NU and LAMBDA and shows the predicted Hs values for return periods out to 100 years. Goodness of fit statistics are also shown for each distribution. Note that ACES, based on Goda (1988), does not utilise the types of non-parametric goodness of fit or Q-Q style tests described earlier. Two statistics are used here, the *correlation parameter* and the *sum square of residuals*. The best fit should indicate a correlation value closest to 1.0 and the smallest value of the residual. In this case the Weibull k=0.75 produces the best fit.

The second screen presents the confidence limits about the mean values, given in Figure A-3.

Finally, the third screen gives a reference table relating the return period to the encounter probability, shown in Figure A-4

A plotting option is available on-line whereby a particular data fit or all data fits can be selected and some control over colours and windowing etc is possible. The graph may be dumped to a plotting device or alternately copied onto the clipboard. Figure A-5 presents the best fit data for this example where *k*=0.75. The graph scale can be log or linear only, ie. specific Weibull or Gumbel scales are not available.

The output file (default PLOTDAT1.OUT) has all the necessary information to enable a plot to be externally generated, including the derived *a*, *b* and *k* parameters for each distribution. Additionally, the Goodness of Fit statistics can be simply plotted to aid the interpretation of the analysis as shown in Figures A-5 and A-6. In this case there is a clear preference to select the *k*=0.75 Weibull option using both of these statistics, which is also easily supported by visual inspection of the graphs for the other *k* options.

If the data for a region has been split into different populations (eg. seasonal) for the purpose of fitting separate distributions, they can be recombined to produce a total distribution for the site. Assuming independence, the joint probability of the combined sets is simply the multiplication of the probabilities of each predicted wave height level.

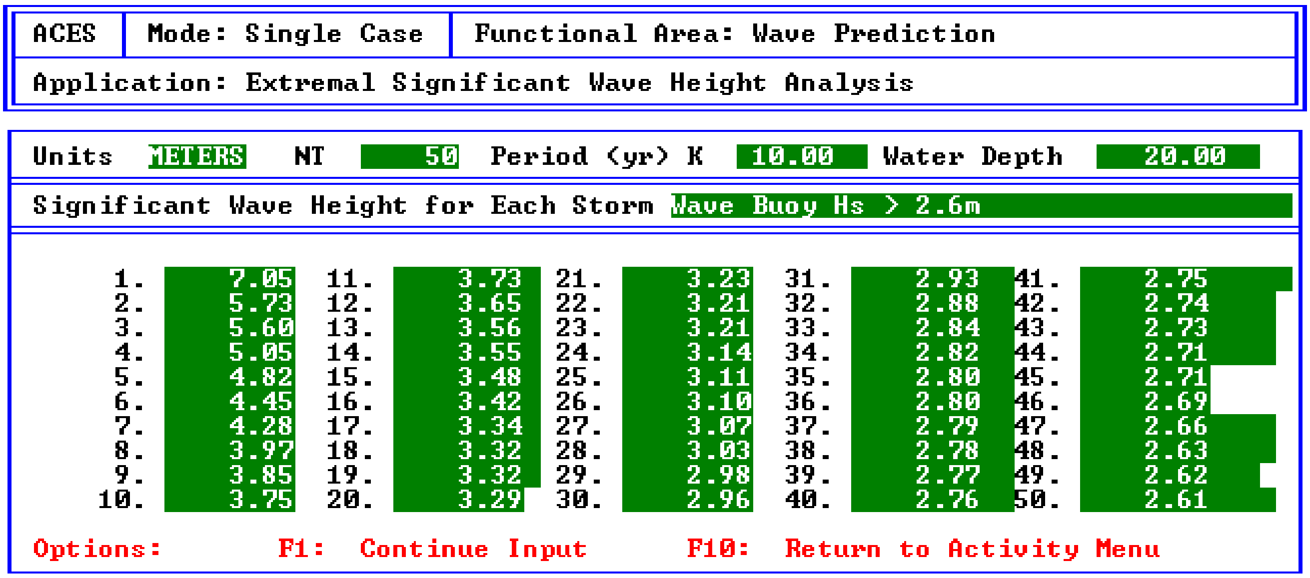

Figure A- 1 Input Data Set Screen

ACES	Mode: Single Case	Functional Area: Wave Prediction
Application: Extremal Significant Wave Height Analysis		

N = 50 NU = 1.00 NT = 50 K = 10.0 yr LAMBDA = 5.00	FT-I	Weibull Distribution			
		k=0.75	k=1.00	k=1.40	k=2.00
Correlation Sum Square of Residuals	0.9514 0.3283	0.9956 0.0426	0.9893 0.1654	0.9606 0.3197	0.9205 0.4441
Return Period (Yr)	Hs (m)	Hs (m)	Hs (m)	Hs (m)	Hs (m)
2	4.55	4.49	4.60	4.61	4.55
5	5.21	5.50	5.45	5.28	5.06
10	5.70	6.34	6.10	5.74	5.40
25	6.34	7.53	6.95	6.32	5.80
50	6.82	8.47	7.60	6.74	6.08
100	7.30	9.46	8.24	7.15	6.34

Options: F1: Next Screen F10: Return to Activity Menu

WARNING: RETURN PERIODS > 30 Yrs May Not Be Meaningful.

Figure A- 2 Return Period Screen

ACES	Mode: Single Case	Functional Area: Wave Prediction
Application: Extremal Significant Wave Height Analysis		

95% Confidence Interval (Lower Bound - Upper Bound) UNITS (m)

Return Period (Yr)	FT-I	Weibull Distribution			
		k=0.75	k=1.00	k=1.40	k=2.00
5	4.5 - 5.9	3.9 - 7.1	4.4 - 6.5	4.5 - 6.0	4.5 - 5.6
10	4.8 - 6.5	4.3 - 8.4	4.8 - 7.4	4.9 - 6.6	4.8 - 6.0
25	5.3 - 7.4	4.8 - 10.3	5.3 - 8.6	5.3 - 7.4	5.1 - 6.5
50	5.6 - 8.0	5.2 - 11.8	5.7 - 9.5	5.6 - 7.9	5.3 - 6.8
100	6.0 - 8.6	5.6 - 13.3	6.1 - 10.4	5.9 - 8.4	5.5 - 7.2

Options: F1: Next Screen F2: Previous Screen F10: Return

Figure A- 3 Confidence Limits Screen

ACES	Mode: Single Case	Functional Area: Wave Prediction
Application: Extremal Significant Wave Height Analysis		

Percent Chance for Significant Height Equaling or Exceeding Return Period Hs

Return Period (Yr)	Period of Concern (Yr)					
	2	5	10	25	50	100
2	75	97	100	100	100	100
5	36	67	89	100	100	100
10	19	41	65	93	99	100
25	8	18	34	64	87	98
50	4	10	18	40	64	87
100	2	5	10	22	39	63

Options: F2: Previous Screen
F10: Return to Activity Menu

Figure A- 4 Encounter Probability Table

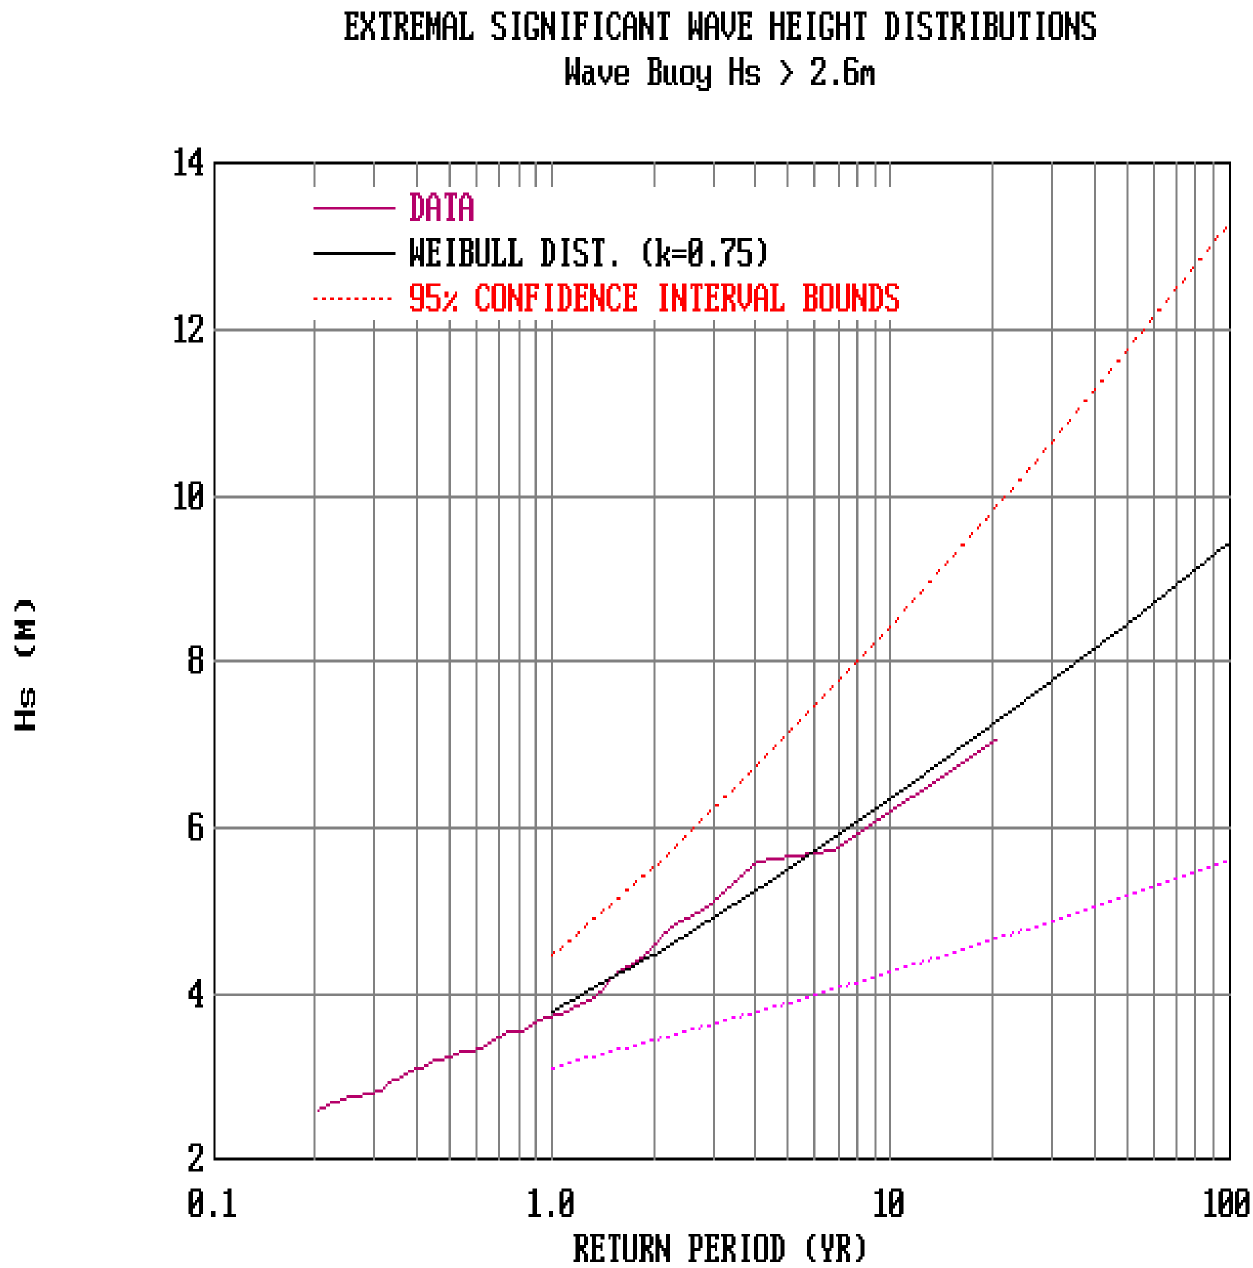

Figure A- 5 Graph of "Best Fit" Weibull k=0.75

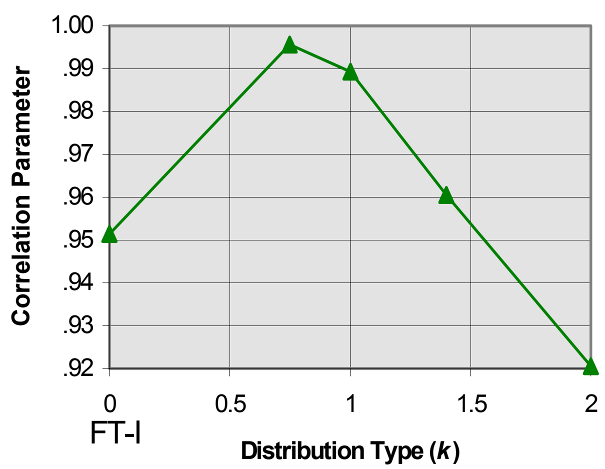

Figure A- 6 Correlation Parameter

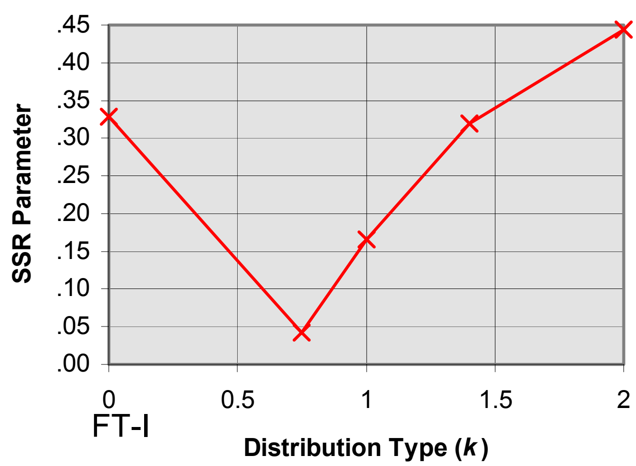

Figure A- 7 Sum Square of Residuals