EXTREME STILL WATER LEVELS

Sofia Caires[1], Ferdinand Diermanse[1], Douwe Dillingh[2] and Reimer de Graaff[1]

[1]WL|Delft Hydraulics, PO Box 177, 2600 MH Delft, The Netherlands.
Corresponding email: Sofia.Caires@wldelft.nl, Tel: +31-15-2858819.
[2]National Institute for Coastal and Marine Management (RIKZ), The Netherlands

1. INTRODUCTION

Still water level (SWL) is the level that the sea surface (at a given point and time) would assume in the absence of wind waves. SWLs are influenced by astronomical and meteorological effects. The estimation of extremes of SWL, required in metocean studies and used in the design of coastal structures, is not straightforward. One of the main problems it faces is the inhomogeneity, sparsity and scarcity of the data. Moreover, it is not always clear which modelling approach is most appropriate for estimating the extremes.

The following approaches are currently used:

1. Extreme value analysis of the SWLs.

2. Estimation of extreme water levels from the convolution of the extreme value distribution of the surge (or that of a non-synchronous difference between SWL and tide) with the empirical distribution of tidal levels. Compared to 1., this is thought to make better use of the data, and of the sometimes complete tidal information.

3. Estimation of extreme surge levels from extreme weather conditions (winds and atmospheric pressures) and computation of pessimistic or conservative SWL estimates by adding the Highest Astronomical Tide to them.

The purpose of this study is to assess approaches 1. and 2. and to provide guidelines as to which should be used in a given situation. In each of them, two different extreme value analysis methods will be considered: the peaks-over-threshold and annual maxima methods. Using the results of our analyses with Approach 2., we shall also provide indications about the tidal level that should be used in Approach 3.

Besides a decision about which approach to take, a decision must also be made about which offset/surge to be used in Approach 2; this depends on the location of the data.

In this study we have been fortunate to have had a long-term timeseries of still water level measurements made available to us by the Dutch Ministry of Transport, Public Works and Water Management. They are the measurements of the gauge located at Hoek van Holland, The Netherlands (see Figure 1 in Section 3.1) and extend from 1887 to 2006. This long and well documented dataset is ideal for a study as that proposed here since it allows for reliable statistics and conclusions to be drawn from the data.

2. DESCRIPTION OF THE STATISTICAL METHODOLOGY USED

2.1 Introduction

In the framework of approaches 1. and 2. outlined in the Introduction, different paths of analysis can be followed. For example, in both cases three different methods of extreme value analysis may be used: the annual maxima, the peaks-over-threshold or the r-largest method. The second and the third methods both make use of more observations than the first, but since the peaks-over-threshold method is generally thought to be somewhat superior we shall not consider the r-largest method in this study (see Coles, 2001). Subsections 2.2 and 2.3 describe briefly the annual maxima and peaks-over-threshold methods.

For a given method of extreme value analysis, Approach 2 consists of

1. separating the astronomic tide from the water level observations,
2. analysing the astronomic tide and the residual (storm surge) separately,
3. computing the empirical distribution of the astronomic tide,
4. carrying out a univariate extreme value analysis of the storm surge data, and
5. computing the convolution of the empirical distribution of the astronomic tide with the extreme value distribution estimated from the storm surge data.

As stated in the Introduction, in this study we will work with data for which the tide and the residual signals have already been separated. Harmonic analysis of the still water levels is a problem in itself and will not be considered here.

Section 2.4 briefly describes the computation of the convolution. The estimation of model parameters and the computation of confidence intervals are based on the same methods in all the extreme value analyses; they are described in Section 2.5.

2.2 Annual Maxima/GEV method

In order to explain the basic ideas, let us define $M_n = \max\{X_1,\ldots,X_n\}$, where $X_1, X_2 \ldots$ is a sequence of independent random variables having a common distribution function *F*. In its simplest form, the *extremal types theorem* states the following: If there exist sequences of constants $\{\sigma_n > 0\}$ and $\{\mu_n\}$ such that $\mathrm{P}\{(M_n - \mu_n)/\sigma_n \le z\} \to G(z)$ as $n \to \infty$, where *G* is a non-degenerate cumulative distribution function, then *G* must be a generalized extreme value (GEV) distribution, which is given by

$$G(z) = \exp\left\{-\left[1+\xi\left(\frac{z-\mu}{\sigma}\right)\right]^{-1/\xi}\right\},$$

where *z* take values in three different sets according to the sign of ξ: $z > \mu - \sigma/\xi$ if $\xi > 0$ (the domain of *z* has a lower limit), $z < \mu - \sigma/\xi$ if $\xi < 0$ (the domain of *z* has an upper limit), and $-\infty < z < \infty$ if $\xi = 0$.

In other words, if the distribution function of (a normalization of) the maximum value in a random sample of size *n* converges to a distribution function as *n* tends to infinity, then that distribution function must be a GEV distribution. Moreover, this and other results of extreme value theory *hold true even under general dependence conditions* (Coles, 2001).

The GEV parameters μ, σ and ξ are called the location, scale, and shape parameters and satisfy $-\infty < \mu < \infty$, $\sigma > 0$ and $-\infty < \xi < \infty$. For $\xi = 0$ the GEV is the Gumbel distribution, for $\xi > 0$ it is the Fréchet distribution, and for $\xi < 0$ it is the Weibull distribution. For $\xi > 0$ the tail of the GEV is "heavier" (i.e., decreases more slowly) than the tail of the Gumbel distribution, and for $\xi < 0$ it is "lighter" (decreases more quickly and actually reaches 0) than that of the Gumbel distribution. The GEV is said to have a type II tail for $\xi > 0$ and a type III tail for $\xi < 0$. The tail of the Gumbel distribution is called a type I tail. See the book of Coles (2001) for more information.

The extremal types theorem gives rise to the *annual maxima* (AM) method of modelling extremes, in which the GEV distribution is fitted to a sample of block maxima.

One of the main applications of extreme value analysis is the estimation of the once per m year (1/m yr) return value. The 1/m yr return value based on the AM method/GEV distribution, z_m, is given by

$$z_m = \begin{cases} \mu - \dfrac{\sigma}{\xi}\left(1 - \left\{-\log\left(1 - \dfrac{1}{\mathrm{m}}\right)\right\}^{-\xi}\right), & \text{for } \xi \neq 0 \\ \mu - \sigma \log\left\{-\log\left(1 - \dfrac{1}{\mathrm{m}}\right)\right\}, & \text{for } \xi = 0. \end{cases}$$

2.2 Peaks-over-threshold/GPD method

The sample sizes of annual maxima data are usually small, so that model estimates, especially return values, have large uncertainties. This has motivated the development of more sophisticated methods that enable the modelling of more data than just block maxima. These methods are based on two well-known characterizations of extreme value distributions: one based on exceedances of a threshold, and the other based on the behaviour of the *r* largest, for small values of *r*, observations within a block.

The approach based on the exceedances of a high threshold, hereafter referred to as the POT (Peaks Over Threshold) method, consists of fitting the generalized Pareto distribution (GPD) to the peaks of 'clustered' excesses over a threshold, the excesses being the observations in a cluster (of successive exceedances) minus the threshold, and calculating return values by taking into account the rate of occurrence of clusters (see Pickands, 1971 and 1975, and Davidson and Smith, 1990). Under very general conditions this procedure ensures that the data can have only three possible, albeit asymptotic, distributions (the three forms of the GPD) and, moreover, that observations belonging to different peak clusters are (approximately) independent. In the POT method, the peak excesses over a high threshold *u* of a time series are assumed to occur in time according to a Poisson process with rate λ_u and to be independently distributed with a GPD, whose distribution function is given by

$$F_u(y) = 1 - \left(1 + \xi \frac{y}{\tilde{\sigma}}\right)^{-1/\xi},$$

where $0 < y < \infty$, $\tilde{\sigma} > 0$ and $-\infty < \xi < \infty$. The two parameters of the GPD are called scale ($\tilde{\sigma}$) and shape (ξ) parameters. For $\xi = 0$ the GPD is the exponential distribution with mean $\tilde{\sigma}$, for $\xi > 0$ it is the Pareto distribution, and for $\xi < 0$ it is a special case of the beta distribution. As for the GEV, the GPD is said to have a type II tail for $\xi > 0$ and a type III tail for $\xi < 0$. The tail of the exponential distribution is a type I tail.

The 1/m yr return value based on a POT/GPD analysis, z_m, is given by

$$z_m = \begin{cases} u - \dfrac{\tilde{\sigma}}{\xi}\{(\lambda_u \mathrm{m})^{\xi} - 1\}, & \text{for } \xi \neq 0 \\ u - \tilde{\sigma}\log(\lambda_u \mathrm{m}), & \text{for } \xi = 0. \end{cases}$$

Just as block maxima have the GEV as their approximate distribution, the threshold excesses have a corresponding approximate distribution within the GPD. Moreover, the parameters of the GPD of threshold excesses are uniquely determined by those of the associated GEV distribution of block maxima. In particular, the shape parameter is the same, and the scale parameters of the two distributions are related by $\tilde{\sigma} = \sigma + \xi(u - \mu)$.

The choice of threshold (analogous to the choice of block size in the block maxima approach) represents a trade off between bias and variance: too low a threshold is likely to violate the asymptotic basis of the model, leading to bias; too high a threshold will generate fewer excesses with which to estimate the model, leading to high variance.

An important property of the POT/GPD approach is the threshold stability property: if a GPD is a reasonable model for excesses of a threshold u_0, then for a higher threshold u a GPD should also apply; the two GPD's have identical shape parameter and their scale parameters are related by $\tilde{\sigma}_u = \tilde{\sigma}_{u_0} + \xi(u - u_0)$, which can be reparameterized as

$$\sigma^* = \tilde{\sigma}_u + \xi u$$

Consequently, if u_0 is a valid threshold for excesses to follow the GPD then estimates of both σ^* and ξ, hence the quantile estimate itself, should remain nearly constant above u_0. This property of the GPD can be used to find the minimum threshold at which a GPD model applies to the data.

2.4 Convolution of astronomical tides and extreme surge levels

Extreme value analyses based on the SWL alone are considered by some to be wasteful of data. In order to take advantage of the sometimes fully available tidal information, extreme values of SWLs can also be estimated from the convolution of the distribution of the residual (the synchronous or not offset/surge) extremes and the empirical distribution of the tidal levels. More precisely, given that the SWL is the sum of the residual and the tide and that these variables can, under certain conditions, be assumed independent, another approach for obtaining the extreme value distribution of the SWL is to estimate the distribution function of 'large values' of SWL by the convolution integral

$$F(z) = \int G_r(z - x) f(x) dx$$

where G_r is the distribution function of 'large values' of the residual (either the GPD or the GEV) and f is the (in principle fully known) density function of the tide levels.

If G_r is the GPD, the 1/m yr return value of SWL can be computed from F by finding z_m such that $1-F(z)=1/(\lambda m)$, where λ has been defined in Section 2.2. If G_r is the GEV, the 1/m yr return value is computed by solving the same equation with λ replaced by 1.

2.5 Parameter estimation and confidence intervals

There are several methods available for the estimation of the parameters of extreme value distributions. Most of them, for instance the methods of moments and of probability weighted moments, give explicit expressions for the parameter estimates. The maximum likelihood (ML) method tends to be the preferred estimation method since it is quite general and more flexible than other methods, especially when the number of parameters is increased as for instance when extending the extreme value approach to account for non-stationarity.

There are big uncertainties with estimates pertaining to extreme values which need to be quantified. When obtaining ML estimates, the variances of the estimates can be obtained from the expected information matrix or from the observed information matrix. An alternative, and usually more accurate, method is the profile likelihood method (Coles, 2001, p. 57), which is based on the deviance function and yields asymmetric confidence intervals. Other alternatives are based on bootstrap procedures with adjustments as suggested by Coles and Simiu (2003). According to a recent study on the coverage rate of confidence intervals of extreme value estimates based on various methods, WL (2007), the adjusted percentile bootstrap method turns out to produce the best confidence intervals from of the point of view of coverage rates. For this reason this will be the method employed in our analyses.

In the case of return values computed from the convolution integral, the bootstrap is applied by resampling from the sample of the residuals only (as the tide is regarded as known).

3. DESCRIPTION OF THE ANALYSIS OF THE DATA

3.1 Introduction

In compliance with the Flood Defences Act of The Netherlands, the primary coastal structures must be checked every five years for the required level of protection on the basis of the Hydraulic Boundary Conditions and the Safety Assessment Regulation. These Hydraulic Boundary Conditions must be derived anew every five years and established by the Minister of Transport, Public Works and Water Management. Extreme still water levels are one of the components of these hydraulic boundary conditions. They are defined using the rich dataset of still water level measurements along the Netherlands coast, which is maintained due to the high safety standard of the country that are the result of the tragic 1953 flooding of The Netherlands on which more than 18 hundred people died. Some of the available timeseries contain more than 100 years of measurements and the corresponding tidal levels are also available. For this study we have been provided with the still water level and high tide data from the Hoek van Holland measuring location; see Figure 1.

The measurements go from 1887 to 2006. This extensive dataset is ideal for a study as the one proposed here since it allows for reliable statistics and conclusions to be drawn from the data. The measurements up to 1985 have been extensively analysed by Dillingh et al. (1993). This makes the dataset even more appealing since most of the peculiarities of the data are known.

In many countries, the design criteria of coastal structures require that they should withstand the 1/100 years return values of the loads. Given that the data chosen for this study are from The Netherlands, where the design criteria for sea dikes vary from 1/2000 years to 1/10000 years, we shall be estimating the 1/100, 1/1000 and 1/10000 return values of SWLs. Even with such a long dataset, the estimation of the 1/10000 years return value is somewhat ambitious.

Figure 1 – Google earth aerial view of the Hoek van Holland tide gauge location.

3.2 Pre-processing of the data

Choice of the surge variable

In shallow waters, the wind set-up can be rather high and cause a significant increase of the propagation velocity of the tidal wave. Furthermore, the currents caused by the wind will also influence the propagation of the tidal wave. Consequently, the instantaneous offset between the SWL and the tidal levels will also include effects of the interaction between the tidal wave and the atmospheric factors and cannot always be considered independent from the co-occurring tidal height. This phenomenon is relevant in the North Sea, where the tidal range is rather high and shallow regions are common. And indeed, as Dillingh et al (1993) report, it is relevant for the Hoek van Holland data considered here. For this reason, Dillingh et al (1993) have decided to statistically analyse the skew High Water offset (the offset between the SWL peak and the tidal high water, independently of a time difference between the two) instead of the vertical, synchronous offset; see Figure 2. Consequently, the skew High Water offset will also be the surge variable considered in this study.

Choice of the population

The data consist of the high SWL (the SWL peak at high tide), the corresponding high tide, and the high water offset from August 1887 until December 2006. From now on we refer to these data as

SWL, tide and surge, respectively. The tide in Hoek van Holland is semi-diurnal with an average time interval of about 12h and 25min between high waters, which implies that the data set contains on average two observations per day.

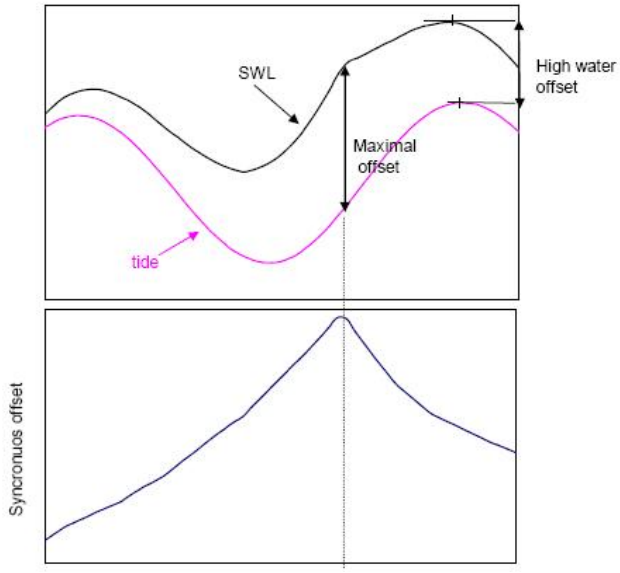

Figure 2 –Schematic representation of high water offset and synchronous offset.

In order to carry out the extreme value analysis the data should be homogeneous; seasonal and other variations should therefore be filtered out. Following the study of Dillingh et al. (1993), where the homogeneity of the data was thoroughly assessed, we shall only consider data from the long winter season of October to March. The data to be used in extreme value analyses should also be independent. Dillingh et al. (1993) report that the average duration of North Sea surge storms of about 2.41 days. Therefore, in this study, SWL peaks (the POT sample) above a threshold are collected from the original time series in such a way that they can be modelled as independent observations. This is, as usual, done by a process of *declustering* in which only the peak (highest) observations in clusters of successive exceedances of a specified threshold are retained and, of these, only those which in some sense are sufficiently far apart (so that they belong to more or less 'independent storms') are considered. Specifically, cluster maxima at a distance of less than 60 hours apart were treated as belonging to the same cluster (storm/extreme event).

Removal of trends

As already noticed by Dillingh et al. (1993), the series of yearly means of the high SWLs (October to March, 1887 until 2006) shows a trend due to sea level variations caused by global warming, dredging, coastal works and/or morphological changes. This trend, which is of about 0.26 cm/year, was removed from the data, adjusting the SWLs to the levels of the 2006 long winter season (October 2005 until March 2006).

3.3 Analysis of the results

POT/GPD analysis of the SWL data

We have used the threshold stability property mentioned in Section 2 to choose the most appropriate threshold for selecting a sample of peak excesses and fitting the GPD to it. More precisely, we have looked for threshold values around which the estimate of the shape parameter and σ^* seem to be stable before becoming rather variable due to reduction of the sample size. Figure 3 shows the estimates of the shape parameter, of σ* and of the 1/10000 years return value as functions of the threshold, obtained with the SWL (October to March 1887/88 – 2005/06) data. The threshold that we have chosen is marked by a vertical line. The return value plot of the corresponding GPD fit is shown in Figure 4 and the model parameter estimates are presented in Table 1.

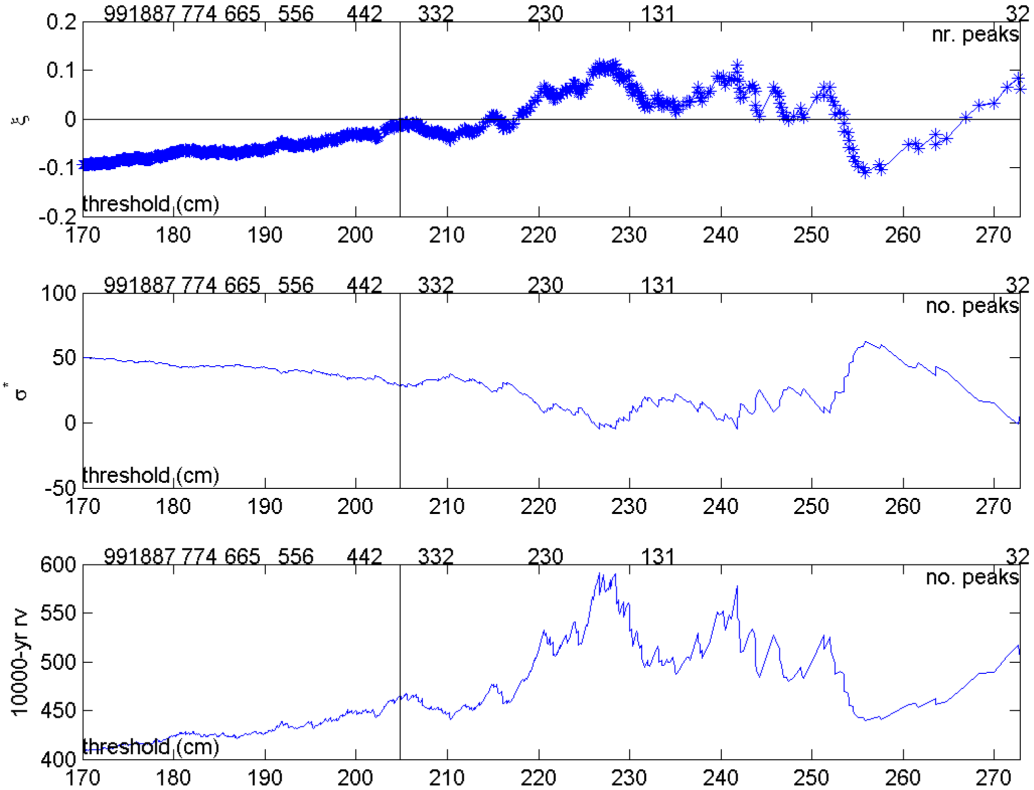

Figure 3 – Variation of the estimates of ξ, σ* and 1/10000 years SWL return values of the GPD model with the threshold used to collect the SWL POT sample.

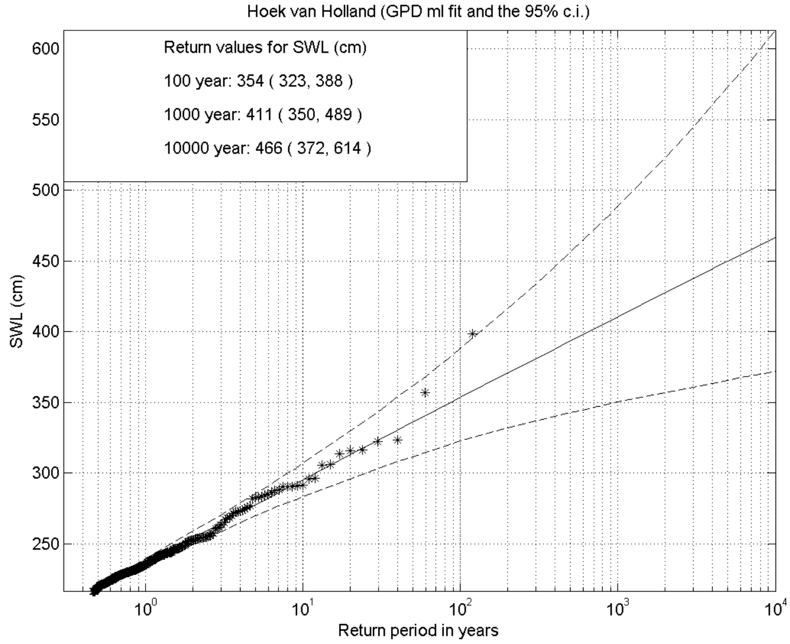

Figure 4 – Return value plot of the GPD model fitted to the SWL data obtained with the ML method (solid black line) and associated 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.

The return value plot suggests that the GPD model is appropriate for the data. Noteworthy is the position of the highest peak in the data, which corresponds to the SWL reached in the tragic storm of 1 February 1953. According to the plot, the return period of this event is much longer than the 119-year period covered by the data.

As can be seen in Table 3, the estimate of the shape parameter is close to zero, which suggests that the data have a type I tail.

We should note that at the beginning of our analysis we have followed Dillingh et al. (1993) in selecting only SWL peaks for which the skew offset was at least 30cm. However, this constraint proved to be unnecessary as it does not affect the estimates.

Sample size	u	$\hat{\xi}$	$\hat{\sigma}$
393	205	-0.01 (-0.13, 0.09)	26 (23, 30)

Table 1 Parameter estimates and associated 95% confidence intervals of the GPD model fitted to the SWL POT data.

AM/GEV analysis of the SWL data

Figure 5 shows the return value plot of the GEV fit to the 'annual' (long winter season) maxima of the SWL data. Table 2 gives the corresponding parameter estimates. The return value estimates for the 1/100, 1/1000 and 1/10000 years are given in Figure 5. Comparing these with the estimates obtained with the GPD model (Figure 4), one can conclude that they are compatible and rather close. From Table 1 and Table 2 one can also see that the estimates of the shape parameter are rather similar, supporting the validity of both models. As was to be expected (due to the difference in sample sizes), the confidence intervals provided by the GEV model are wider than those of the GPD.

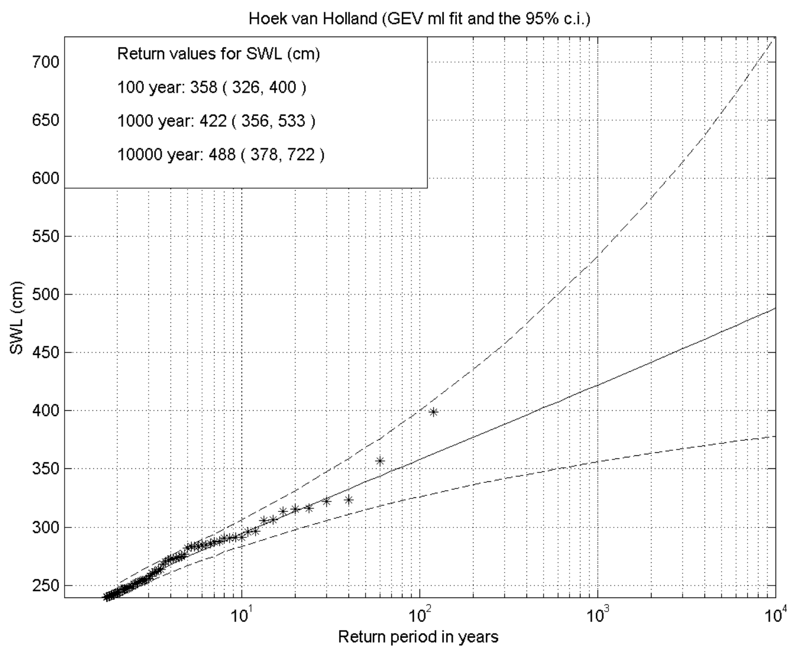

Figure 5 – Return value plot of the GEV fitted to the SWL data obtained with the ML method (solid black line) and associated 95% confidence intervals (dashed black lines). The AM data are represented by the asterisks.

Sample size	$\hat{\xi}$	$\hat{\sigma}$	$\hat{\mu}$
119	0.01 (-0.14, 0.16)	26 (22, 30)	235 (230, 241)

Table 2 Parameter estimates and associated 95% confidence intervals of the GEV model fitted to the SWL AM data.

Convolution of astronomical tides and peak surge heights

As in the POT/GPD analysis of the SWL data, we have used the threshold stability property to choose the most appropriate threshold for selecting a sample of surge peak excesses and fitting the GPD to it. Figure 6 shows the estimates of the shape parameter, of σ* and of the 1/10000 return value as functions of the threshold obtained with the surge (October to March 1887/88 – 2005/06) data. The thresholds that we have chosen are marked by vertical lines. The return value plot of the corresponding GPD fit is presented in Figure 7, and the model parameter estimates in Table 3.

From Table 3 and Table 1, we see that the estimate of the shape parameter of the GPD is the same with surge and SWL data.. Comparing the GPD fits of the surge and SWL data (figures 4 and 7), one can say that the fit of the surge data is a bit poorer. As before, the highest peak of the surge data corresponds to the 1953 storm, which is again associated with a return period longer than that spanned by the measurements.

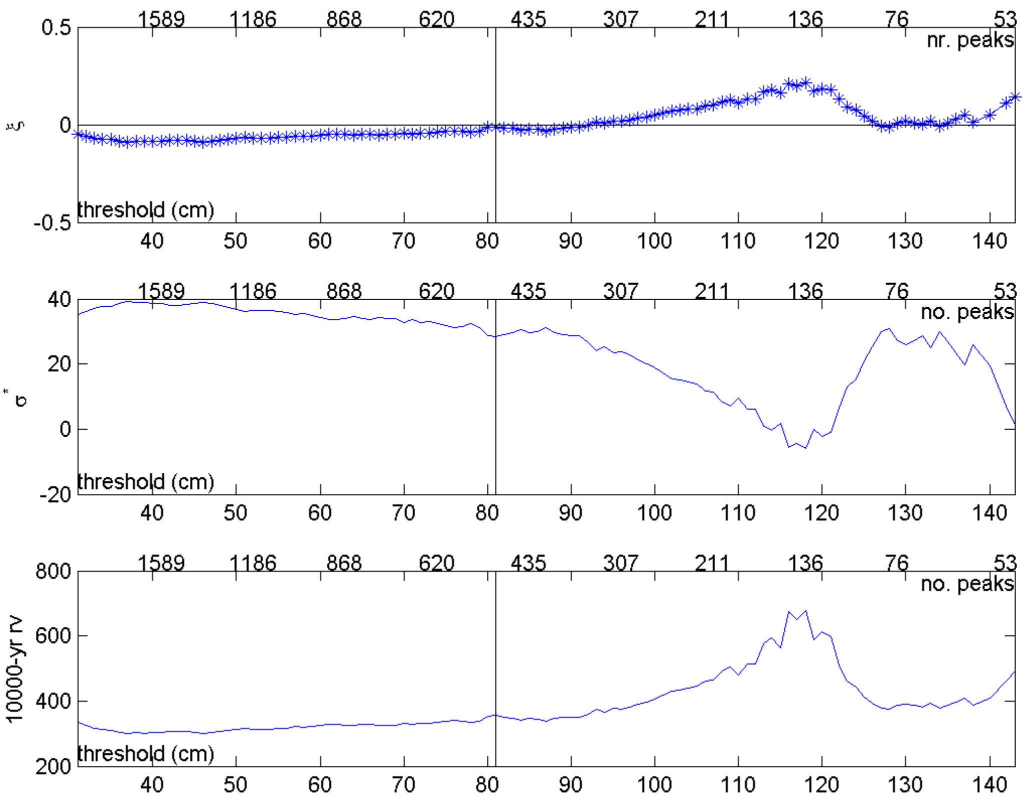

Figure 6 – Variation of the estimates of ξ, σ* and 1/10000 years surge return values of the GPD model with the threshold used to collect the surge POT sample.

Figure 8 shows the return value plot obtained from the convolution of the empirical distribution of the astronomical tide and the GPD fit just described. The fit seems appropriate and the return value estimates are compatible with those obtained in Approach 1.

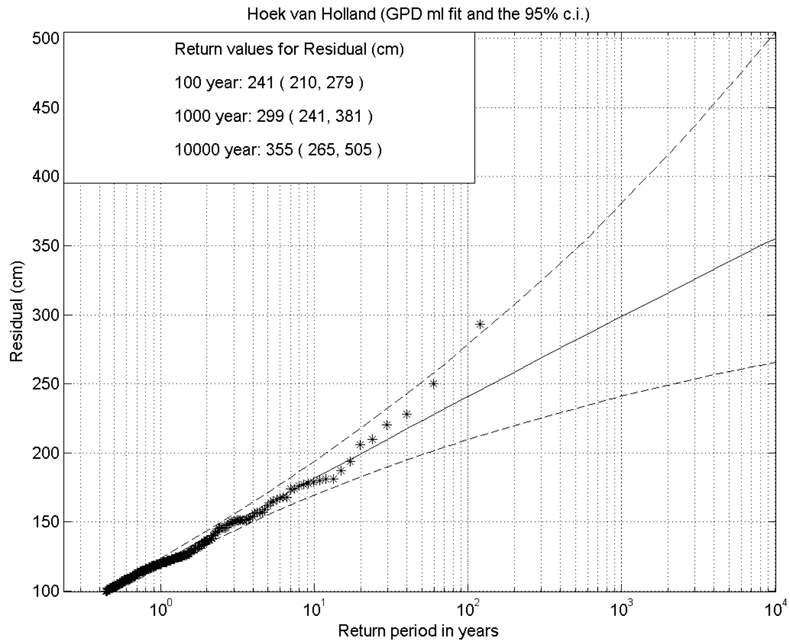

Figure 7 – Return value plot of the GPD fitted to the surge data obtained with the ML method (solid black line) and associated 95% confidence intervals (dashed black lines). The POT data are represented by the asterisks.

Sample size	u	$\hat{\xi}$	$\hat{\tilde{\sigma}}$
506	81	-0.01 (-0.11, 0.08)	27 (24, 31)

Table 3 Parameter estimates and associated 95% confidence intervals of the GPD model fitted to the surge POT data.

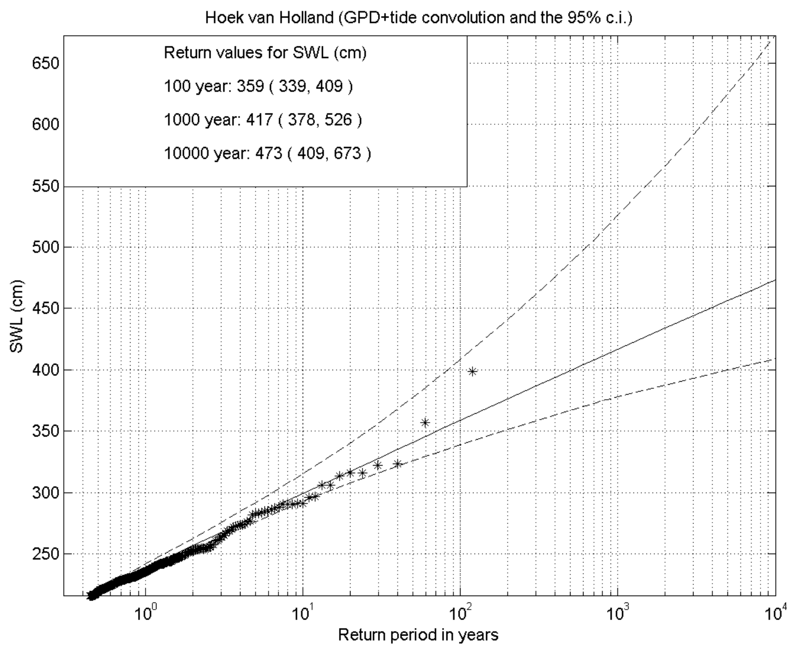

Figure 8 – Return value plot of the convolution of the astronomical tide and the surge levels (solid black line) and associated 95% confidence intervals (dashed black lines). The corresponding SWL data are represented by the asterisks.

Convolution of astronomical tide and AM surge heights

Figure 9 shows the return value plot of the GEV fit to the annual maxima of the surge data and the surge 1/100, 1/1000 and 1/10000 years return value estimates. Table 4 gives the corresponding parameter estimates. Comparing these with the estimates obtained with the GPD model (Figure 7), one can conclude that they are compatible and rather close. From Table 3 and Table 4 one can also see that the estimates of the shape parameter are rather similar.

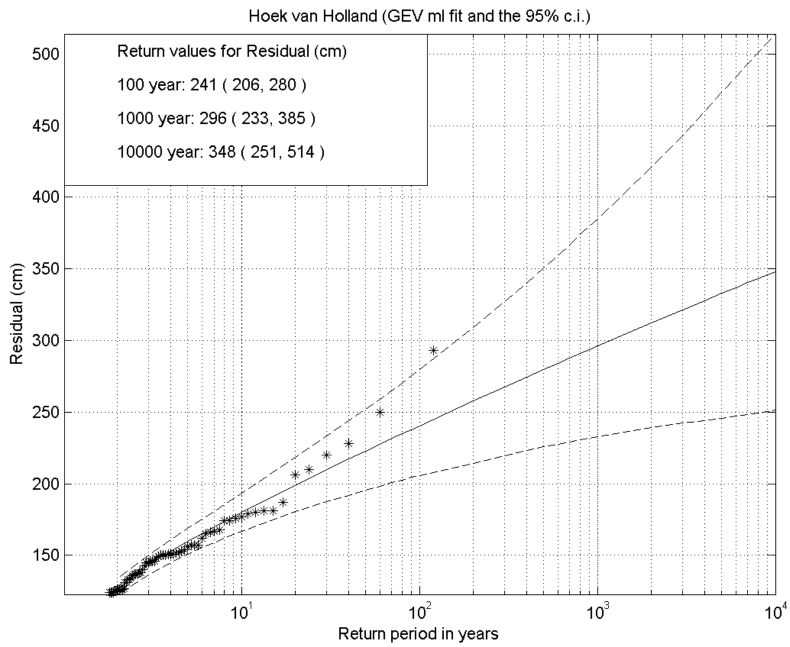

Figure 9 – Return value plot of the GEV fitted to the surge data obtained with the ML method (solid black line) and associated 95% confidence intervals (dashed black line). The surge data are represented by the asterisks.

Sample size	$\hat{\xi}$	$\hat{\sigma}$	$\hat{\mu}$
119	-0.03 (-0.17, 0.09)	28 (24, 33)	118 (113, 123)

Table 4 Parameter estimates and associated 95% confidence intervals of the GEV model fitted to the surge AM data.

Figure 10 shows the return value plot obtained from the convolution of the empirical distribution of the astronomical tide and the GEV fit just described. The fit seems appropriate and the return value estimates are compatible with those obtained from the SWL data using the other models.

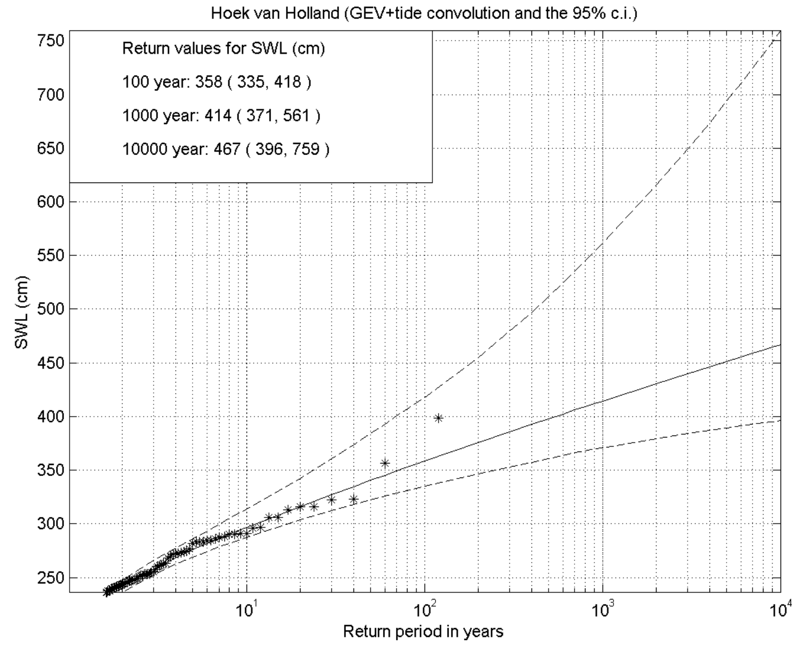

Figure 10 – Return value plot of the convolution of the astronomical tide and the AM surge levels (solid black line) and associated 95% confidence intervals (dashed black lines). The corresponding SWL AM data are represented by the asterisks.

4. DISCUSSION AND RECOMMENDATIONS

The goal of this study is to assess approaches 1. and 2. to estimating extremes of SWLs, and to provide guidelines as to which should be used in a given situation. The approaches have been applied to the Hoek van Holland data. We shall now discuss the results obtained and give recommendations.

To facilitate the discussion, the results have been gathered in Figure 11 and in Tables 5 and 6.

Figure 11 shows the return value plots of the various estimates and associated 95% confidence intervals. The figure shows a striking agreement between the return value point estimates provided by the different methods. The compatibility between the estimates can be further seen in the relative differences between the SWL return value estimates obtained from the different analyses and those of the POT/GPD analysis of the SWL given in Table 5. Differences between estimates are always less than 5%. Given the amplitude of the 95% confidence intervals associated with the estimates (cf. Table 6), these differences do not appear significant.

It can thus be concluded that, given a large enough and reliable dataset, the estimates of return values from the different approaches and variants are compatible.

From the amplitude of the 95% confidence intervals in relation to the associated point estimate, given as a percentage in Table 6, the following conclusions can be drawn:

- The wider confidence intervals are those for estimates using the AM/GEV model. This is to be expected and should be more noticeable with datasets smaller than the one considered here.

- The use of Approach 2. does not result in shorter confidence intervals. This goes against the idea that not using the known tidal information is data wasteful. It could be explained from the fact that the uncertainty in the estimates is due to the rare extreme events and not the well determined tide.

- As expected, the relative amplitude of the confidence intervals increases with the return period. This sets limitations to the actual use in practice of the 'more extreme' return value estimates computed from a dataset with a given length.

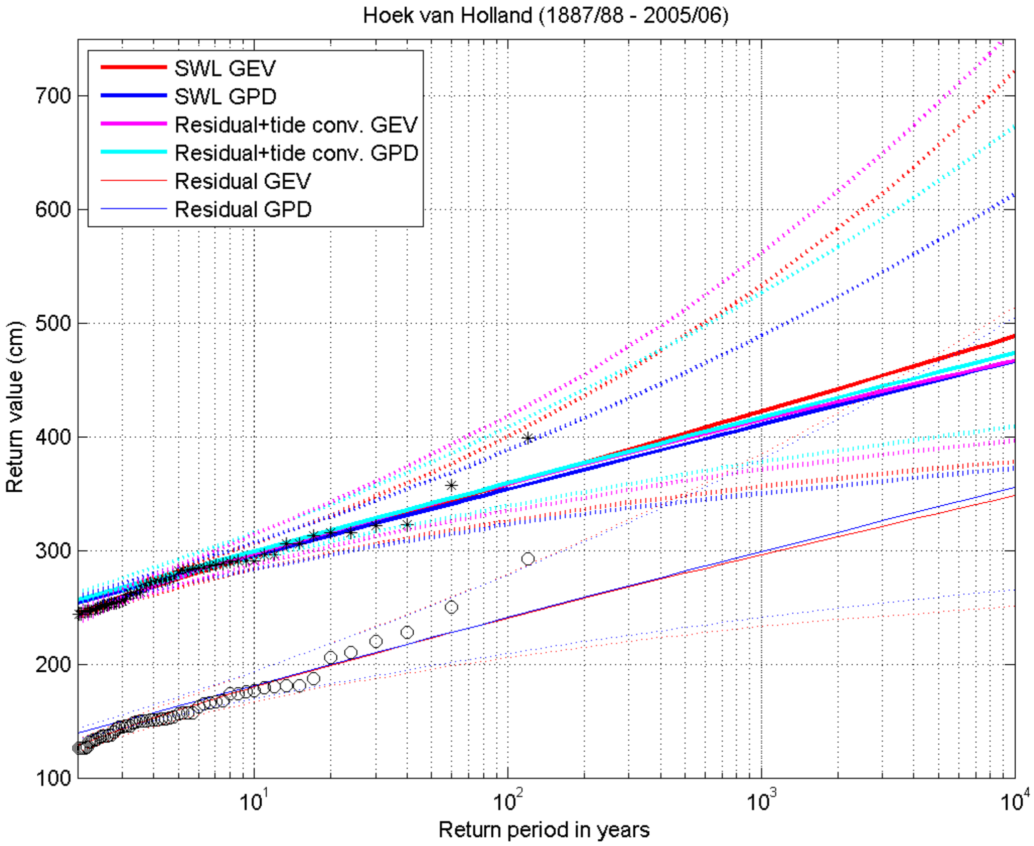

Figure 11 – Return value plot constructed with the results from all the analyses, including 95% confidence bands as dashed lines. The SWL data are represented by the asterisks and the surge data by the circles.

Return period	Convolution Residual POT/GPD	SWL AM/GPD	Convolution Residual AM/GPD
1/100 year	1.50	1.25	1.37
1/1000 years	1.54	2.83	0.95
1/10000 years	1.48	4.63	0.07

Table 5 Relative differences between the SWL return values estimates obtained from the different analyses and those of the POT/GPD analysis of the SWL. The values are percentages of the POT/GPD SWL return value point estimates.

For the return periods longer than 10 years, the return value estimates of the convolution are equal to those of the associated surge analysis plus a constant. Both in the POT and GPD models of Approach 2. this constant is 118 cm. The average of the high tide data for the period October to March 1887/88 – 2005/06 is 111 cm, and the maximum is 157 cm. The Mean High Water Spring at this site is 130cm. This indicates that adding the Highest Astronomical Tide to the water levels associated with the extreme weather conditions of Approach 3. is a too conservative procedure. A value between Mean High Water and Mean High Water Spring should be used.

	Return period	SWL (cm)	rel. amp. of c.i. (%)
SWL POT/GPD	1/100 years 1/1000 years 1/10000 years	354 (323, 388) 411 (350, 489) 467 (372, 614)	19 34 52
Residual POT/GPD	1/100 years 1/1000 years 1/10000 years	241 (210, 279) 299 (241, 381) 355 (265, 505)	29 47 67
Convolution Residual POT/GPD	1/100 years 1/1000 years 1/10000 years	359 (339, 409) 417 (378, 526) 473 (409, 673)	19 35 56
SWL AM/GEV	1/100 years 1/1000 years 1/10000 years	358 (326, 400) 422 (356, 533) 488 (378, 722)	21 42 70
Residual AM/GEV	1/100 years 1/1000 years 1/10000 years	241 (206, 280) 296 (233, 385) 348 (251, 514)	31 51 75
Convolution Residual AM/GEV	1/100 years 1/1000 years 1/10000 years	358 (335, 418) 414 (371, 561) 467 (396, 759)	23 46 78

Table 5 Return value estimates, associated 95% confidence intervals and relative amplitude of the confidence intervals.

The following recommendations arise from this study:

- The POT/GPD approach is generally preferable to the AM/GEV approach since the estimates of the latter have greater variability, even with long datasets.

- Approach 2. does not seem to be superior, in terms of reduction of uncertainty of estimates, to Approach 1. It is therefore preferable to use Approach 1. since this is simpler and/or does not require the determination of the tidal signal. In the case of the POT/GPD approach, this of course assumes that the threshold has been taken high enough so as to exclude peaks with no surge component.

- The choice of the offset to be used in Approach 2. should take into consideration the characteristics of the basin under study. For the North Sea, the basin of the example used here, the instantaneous offset between the astronomical tide and the SWL should not be used since the two may be correlated. For other basins, such as for instance the Mediterranean Sea, where water depths are rather high and the wind set-up less important, the instantaneous offset can in principle be used.

- In Approach 3., the tidal level that should be added to the water level associated with extreme weather conditions should be somewhere between Mean High Waters and Mean High Water Spring.

ACKNOWLEDGEMENTS

This research was funded by the Dutch Ministry of Economic Affairs. We are thankful to the Dutch Ministry of Transport, Public Works and Water Management for the Hoek van Holland water levels data.

References

Coles, S., 2001: An Introduction to the Statistical Modelling of Extreme Values. Springer Texts in Statistics, Springer-Verlag: London.

Coles, S., and E. Simiu, 2003: Estimating uncertainty in the extreme value analysis of data generated by a hurricane simulation model. J. Engrg. Mech., 129 (11), 1288-1294.

Davison, A. C., and R. L. Smith, 1990: Models for exceedances over high thresholds (with discussion). J. Roy. Stat. Soc., 52B, 393–442.

Dillingh, D., L. de Haan, R. Helmers, G.P. Können and J. van Malde, 1993: De basispeilen langs de Nederlandse kust; statistisch onderzoek (In Dutch). Rijkswaterstaat, Dienst Getijdenwateren /RIKZ, Report DGW-93.023.

Pickands, J., 1971: The two-dimensional Poisson process and extremal processes. Journal of Applied Probability, 8, 745-756.

Pickands, J., 1975: Statistical inference using extreme order statistics. Annals of Statistics, 3, 119-131.

WL, 2007: Extreme wave statistics. Confidence intervals. WL | Delft Hydraulics Report H4803.30. Author: S. Caires.