

# Extreme offshore wave statistics in the North Sea

K. Roscoe, S. Caires, F. Diermanse & J. Groeneweg  
*Deltares, The Netherlands*

## Abstract

As part of the flood risk assessment in The Netherlands, extreme hydraulic boundary conditions (HBC) at coastal water defenses are calculated to assess the safety and sufficiency of the defenses. For this purpose, extreme offshore values of wave height and wave period with a frequency of occurrence of 1/10,000 years need to be estimated. The estimation is carried out using measurements at nine offshore locations in the North Sea, with a record length of 30 years.

To estimate extreme wave heights and wave periods, the Generalized Pareto Distribution (GPD), which follows from extreme value theory, was fitted to peaks-over-threshold (POT) at each of the nine measurement stations. The fitted GPD parameters and consequently the extrapolated extreme values are particularly sensitive to the choice of threshold. A reproducible method was employed for threshold selection based on the anticipated behavior of the GPD shape parameter as the threshold approaches the domain of attraction. A modified regional frequency analysis (MRFA) was subsequently carried out to spatially smooth the GPD shape parameters. Final GPD fits were made using the shape-parameter output of the MRFA to produce final extrapolated extreme value estimates. The extreme values were compared with those resulting from an operationally-used distribution function in The Netherlands known as the conditional Weibull distribution.

*Keywords: extreme value theory, flood risk, threshold selection, generalized Pareto distribution, extreme wave parameters.*



# 1 Introduction

## 1.1 Background

The Flood Defences Act of The Netherlands requires that the primary flood defences be assessed every five years. The defences must be able to withstand extreme conditions, with return periods along the coast varying per location from 1/4,000 years to 1/10,000 years. The assessment is carried out anew every five years to allow for the incorporation of new knowledge, such as improved models, improved methodologies, extension of measured time series, and updated model input such as changes in bathymetry. The focus in this study was on coastal defences. The assessment of these defences requires the estimation of extreme hydraulic boundary conditions (HBC), such as water level, wind speed and wave parameters at the defence structure. In the current study, the focus was on wave heights and wave periods. For the calculation of HBC at the defence structures, univariate extreme statistics are first calculated using measurements collected at offshore stations in the North Sea. These statistics are then input into probabilistic and physical models such as the model Hydra-K (Stijnen et al. [1]) and SWAN (Booij et al. [2]) which translate these statistics into loads at the defence structures.

## 1.2 Objectives

The univariate extreme statistics at offshore locations are currently calculated using a probability distribution referred to as the conditional Weibull distribution. This distribution is not the asymptotic distribution predicted by extreme value theory (Coles [3]). The objective of the current work was to estimate the univariate statistics of wave height and wave period using the generalized Pareto distribution, which follows from extreme value theory, and to compare these results with those obtained using the distribution currently in use.

# 2 Extreme value theory

Extreme value theory provides an analogue to the central limit theorem for the extreme values in a sample. According to the central limit theorem, the mean of a large number of random variables, irrespective of the distribution of each variable, is distributed approximately according to a Gaussian distribution. According to extreme value theory, the extreme values in a large sample have an approximate distribution that is independent of the distribution of each variable. Extreme value theory can be applied to annual block maxima or to exceedances above a high threshold. The former are described by the general extreme value (GEV) distribution; the latter by the generalized Pareto distribution (GPD). For shorter time series, it is advantageous to opt for the modelling of peak excesses with the GPD (Caires [4]).



## 2.1 Peaks-over-threshold method

The approach based on the exceedances of a high threshold, hereafter referred to as the *Peaks-over-Threshold* (POT) method, consists of fitting the GPD to the peaks of clustered *excesses* over a threshold, the excesses being the observations in a cluster minus the threshold, and calculating return values by taking into account the rate of occurrence of clusters (Pickands [5, 6] and Davison and Smith [7]). Under very general conditions this procedure ensures that the data can have only three possible, albeit asymptotic, distributions (the three forms of the GPD given below) 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  $\lambda_u$  and to be independently distributed with a GPD, whose distribution function is given by

$$F_u(y) = \begin{cases} 1 - \left(1 - \xi \frac{y}{\sigma}\right)^{1/\xi}, & \text{for } \xi \neq 0 \\ 1 - \exp\left(-\frac{y}{\sigma}\right), & \text{for } \xi = 0, \end{cases} \quad (1)$$

where  $y > 0$ ,  $\sigma > 0$  and  $1 - \xi y/\sigma > 0$ . The two parameters of the GPD are called the scale ( $\sigma$ ) and shape ( $\xi$ ) parameters. Note that the sign convention of the shape parameter can differ; that is, the shape parameter is sometimes considered the negative of what is shown in eqn (1). When  $\xi = 0$ , the GPD is said to have a Type I tail and amounts to the exponential distribution with mean  $\sigma$ ; when  $\xi < 0$ , it has a Type II tail and it is the Pareto distribution; and when  $\xi > 0$ , it has a Type III tail and it is a special case of the beta distribution. If  $\xi > 0$ , the support of the GPD has an upper bound,  $\sigma/\xi$ , which is called the *upper end-point* of the GPD and is to be thought of as the upper-limit of the excesses, the upper limit of the variable of interest being then  $u + \sigma/\xi$ .

The  $m$ -year return value based on a POT/GPD analysis,  $z_m$ , is given by

$$z_m = \begin{cases} u - \frac{\sigma}{\xi} \{(\lambda_u m)^{-\xi} - 1\}, & \text{for } \xi \neq 0 \\ u + \sigma \ln(\lambda_u m), & \text{for } \xi = 0. \end{cases} \quad (1)$$

The choice of the threshold in the POT 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 correct model for excesses above a threshold  $u_0$ , then for

a higher threshold  $u$  a GPD should also apply; the two GPDs have identical shape parameter and related scale parameters (Coles [3]).

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. However, in ordinary extreme value analyses, like the ones we are concerned with in this study, the flexibility provided by the ML method is not necessary, and for the range of tails typically found with wave data (not too heavy-tailed distributions) and for small to moderate sample sizes the method of probability-weighted moments (PWM) performs better (Hosking et al. [8] and Hosking and Wallis [9]). For this reason, PWM was used for parameter estimation in the current study.

### 3 Operationally-applied distribution

For the estimation of HBC in 2006, a similar methodology to the one described in this paper was applied, but with the use of what is known as the conditional Weibull distribution (CWD) for the extrapolation of extreme wave parameters; the equation of the CWD is presented in eqn (3).

$$F(x) = 1 - \exp[-(\omega/\sigma)^\alpha - (x/\sigma)^\alpha], \text{ for } x \geq \omega, n \quad (2)$$

where  $\omega$  is the location parameter (threshold),  $\sigma$  is the scale parameter, and  $\alpha$  is the shape parameter.

The choice of the CWD was based on historic usage and followed from recommendations by the former Dutch National Institute for Coastal and Marine Management (RIKZ). The CWD is principally supported for what is referred to as its robustness. That means that for changes in the threshold and variations in shape parameter, the changes in the resulting statistics are not terribly sensitive, especially compared with the GPD.

The genesis of the CWD is based on the assumption that the following Weibull distribution represents the parent distribution of the wave parameters.

$$F_w(x) = 1 - \exp\left[-\left((x - \nu)/\sigma\right)^\alpha\right], \quad x > \nu, \quad (3)$$

where  $\nu$  is the location parameter,  $\sigma$  is the scale parameter, and  $\alpha$  is the shape parameter. This form of the Weibull distribution is one of the forms of the GEV for minima, and falls into a Type I tail type, meaning that there is no upper limit to the exceedances. The CWD is the distribution of the exceedances from a Weibull distribution above the threshold  $\omega$ .



There are a couple of criticisms about the use of this model for the estimation of extreme statistics. First, the arguments in favour of the CWD were addressed and questioned in a report by Caires [10], in which the robustness is criticized as a managerial advantage, but not advantageous if the produced statistics are incorrect. That is, the choice of a robust model is one of bias over variance. Further, the assumption that the (non-truncated) wave parameters follow a single distribution is disputed (Ferreira and Guedes Soares [11]). Additionally, the Weibull distribution is not a universally appropriate model for describing entire sets of wave data; for example, the lognormal distribution is commonly used. Lastly, wave parameters are naturally physically limited; this limitation is not represented by a Type I tail model, which has no upper limit.

## 4 Initial extreme value analysis

### 4.1 Data collection for POT analysis

The measurements that were used in the current study are based on 30 years of 3-hourly wave data at nine offshore locations in the North Sea (see Figure 1).

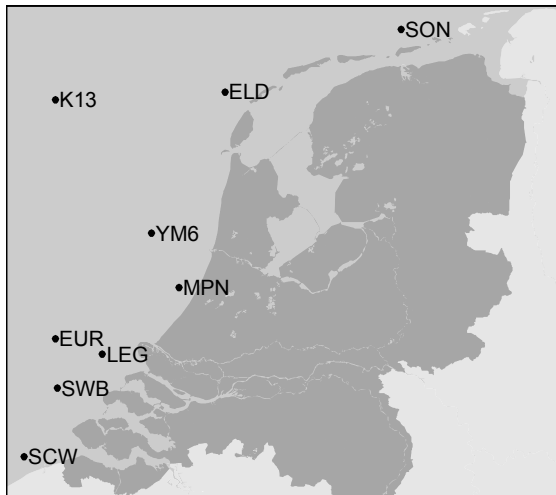


Figure 1: Dutch offshore measurement stations in the North Sea.

The data were continuously collected in records of 10 minutes length. Every hour, two 10-minute records were combined to estimate the wave parameters. The wave parameters considered in this analysis are the significant wave height ( $H_{m0}$ ), the mean wave period ( $T_{m-1,0}$ ) and the peak wave period ( $T_p$ ). In the current approach, storm peaks above a threshold (the POT sample) are selected from the 3-hourly time series in such a way that the sample extracted from the original time series can be modelled as independent observations. With wave and similar data this is usually done by a process of de-clustering in which only the

peak (highest) observations in clusters exceeding a specified threshold are retained and, of these, only those which are sufficiently far apart in time (so that they belong to more or less independent storms) are considered. Specifically, in the current study, cluster maxima at a distance of less than 48 hours apart were treated as belonging to the same cluster (storm). Following, a filter was applied to compensate for over- and underestimation due to sampling variability (Forristall et al. [12]).

#### 4.2 Choice of threshold

The threshold stability property mentioned in Section 2.1 was used to choose the most appropriate threshold for selecting a sample of peak excesses to which the GPD was fitted. More precisely, threshold values were sought around which the estimate of the shape parameter showed the least variation.

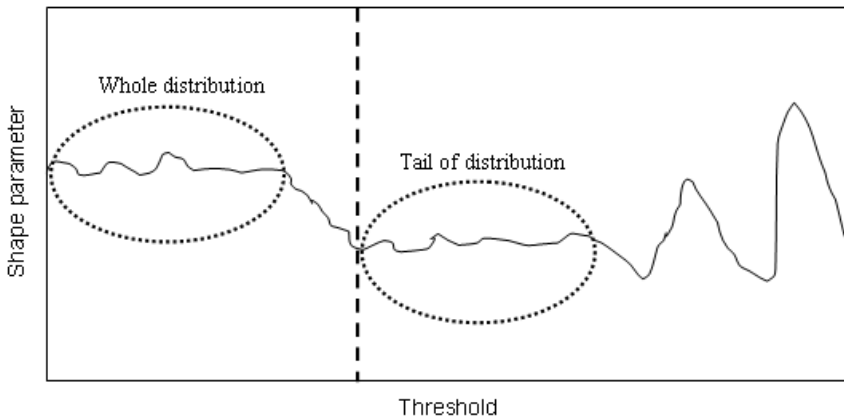


Figure 2: Schematic showing the behaviour of the shape parameter with increasing threshold. The vertical line shows the optimal choice.

To reduce subjectivity and retain reproducibility in this process the choice of the threshold was automated such that the selected threshold was the lowest that resulted in a shape parameter in the stable region (see Figure 2). The stable region refers to the range of thresholds that result in a low variance of shape parameter for changes in threshold. Note that theoretically there will be two stable regions: one that represents the fit to the full set of storm peaks, and a second that represents the fit to the extreme of the storm peaks. The latter is referred to as the domain of attraction and represents the stable region for the optimal threshold; to avoid the selection of thresholds in the first stable region, a constraint was applied that thresholds resulting in more than ten storms per year were not considered.

Once the threshold is selected, fitted GPD parameters for that threshold define the initial extreme value statistics at the given station. That is, wave heights and wave periods for any return period of interest can be calculated.

## 5 Regional frequency analysis

The initial extreme value analysis derived GPD fits for each station independently. Without the application of a regional frequency analysis (RFA), specific outliers at a given location can have too much influence on the statistics for that region. Application of RFA uses data at all stations to estimate the GPD shape parameter; in doing so, it helps reduce the effect of randomness that is inherent in short data sets; RFA is therefore often referred to as ‘trading space for time’ (Hosking and Wallis [13]). Because the stations are not from a homogeneous region, a modified regional frequency analysis (MRFA) was applied which takes into account the local depth and fetch length (Den Heijer et al. [14]).




Figure 3: Effect of MRFA on shape parameters.

Figure 3 shows that the MRFA results in more spatially consistent patterns of shape parameter. Open icons represent  $\xi$  values prior to the MRFA; closed icons represent  $\xi$  values following MRFA. A lower shape parameter translates to a higher return value, which is expected in the north, where fetch lengths are much longer. The MRFA returns new shape parameters, for each station and wave parameter, that will be used in the final extreme value analysis.

## 6 Final extreme value analysis

The final extreme value analysis refers to the final fitting of the GPD. Following the MRFA, fixed shape parameters are assigned to each of the wave parameters.



Prior to the MRFA, a threshold was selected for which there was an associated GPD shape parameter (input to MRFA). Following the MRFA, new thresholds must be selected that are associated with the MRFA shape parameters. This is done by investigating which thresholds result in a set of POT that leads to the MRFA shape parameter; among the possible choices, the threshold which is nearest to the original (pre-MRFA) threshold is selected.

Once the threshold parameter is selected, the scale parameter can be adjusted to fit the GPD with the MRFA shape parameter to the peak excesses above the selected threshold. The resulting fit produces the final statistics.

## 7 Results

The results of the initial extreme value analysis are presented in Table 1. Shown are the 1/10,000-year return values for the three wave parameters and in parentheses the lower and upper bounds on the return values, calculated using bootstrap with 1,000 samples. The uncertainties are considerably large, principally due to the flexibility in the GPD parameters; specifically that all three tail types can be represented by the GPD shape parameter.

Table 1: 1/10,000-year return values and uncertainty: GPD, pre-MRFA.

Station ID	$H_{m0}$ (m)	$T_{m-1,0}$ (s)	$T_p$ (s)
SON	10.62 (7.8, 16.14)	15.82 (12.66, 23.58)	20.17 (15.52, 33.38)
ELD	8.47 (7.23, 10.99)	14.37 (11.65, 21.2)	16.74 (13.48, 25.31)
K13	9.59 (7.77, 13.45)	16.83 (13.21, 26.23)	18.74 (14.58, 28.02)
YM6	9.34 (7.17, 14.7)	13.08 (11.4, 16.92)	19.27 (14.54, 32.57)
MPN	7.54 (6.23, 10.02)	12.31 (10.19, 20.98)	14.45 (12.35, 20)
EUR	7.35 (6.18, 9.95)	10.87 (9.46, 14.39)	14.16 (11.41, 21.09)
LEG	7.89 (6.49, 10.50)	9.98 (9.2, 11.84)	11.78 (10.66, 14.04)
SWB	6.36 (5.49, 8.01)	9.58 (8.9, 11.18)	11.13 (10.19, 13.31)
SCW	4.75 (4.28, 5.63)	10.11 (9.01, 12.61)	13.58 (10.71, 23.52)

Table 2: 1/10,000-year return values: GPD.

Station ID	$H_{m0}$ (m)	$T_{m-1,0}$ (s)	$T_p$ (s)
SON	10.12	16.02	20.30
ELD	9.86	14.15	17.72
K13	9.52	15.19	17.89
YM6	8.46	13.84	17.03
MPN	7.20	12.29	15.48
EUR	7.40	10.79	12.98
LEG	7.16	10.24	12.45
SWB	6.20	9.91	11.77
SCW	5.02	9.91	12.00





The return values resulting from the final extreme value analysis (i.e. following the MRFA) are presented in Table 2. The differences between the initial and final analyses are sometimes considerable, with maximum differences of 1.4m, 1.6s, and 2.2s for wave parameters  $H_{m0}$ ,  $T_{m-1,0}$ , and  $T_p$ , respectively. These differences are all within the uncertainty in the initial estimates. Table 3 presents the final extreme value analysis using the CWD as well as the difference between the CWD and GPD return values. The CWD consistently returns higher return values than the GPD. This is because the CWD is a distribution with a fixed Type I tail, while the GPD shape parameter tends to a Type III tail for wave data. The implications of this are discussed in the next section.

Table 3: 1/10,000-year return values: CWD and difference between CWD and GPD.

Station ID	10 <sup>-4</sup> Return Values CWD			Change in 10 <sup>-4</sup> Return Values [CWD-GPD]		
	H <sub>m0</sub> (m)	T <sub>m-1,0</sub> (s)	T <sub>p</sub> (s)	H <sub>m0</sub> (m)	T <sub>m-1,0</sub> (s)	T <sub>p</sub> (s)
SON	11.71	17.36	21.63	1.59	1.34	1.33
ELD	11.05	15.05	18.57	1.19	0.90	0.85
K13	10.67	15.31	18.30	1.15	0.12	0.41
YM6	9.34	14.53	17.39	0.88	0.69	0.36
MPN	7.81	12.91	15.96	0.61	0.62	0.48
EUR	7.77	11.30	13.14	0.37	0.51	0.16
LEG	7.72	11.18	13.39	0.56	0.94	0.94
SWB	6.77	10.72	12.56	0.57	0.81	0.79
SCW	5.46	10.50	12.20	0.44	0.59	0.20

## 8 Discussion

The GPD is theoretically the correct distribution for extrapolating POT, regardless of the parent distribution of the peaks. However, in practice the GPD can be an unattractive choice from a management perspective.

The GPD is very flexible; it allows for all three tail types, letting the data themselves determine the type. This can be advantageous because in the case of wave parameters, they are physically limited. That is, they cannot bypass certain physical maxima due to physical constraints, such as depth. The CWD, which has a fixed Type I tail, does not have the flexibility to respond to this physical information in the data, while the GPD does. The flexibility is less advantageous from a management perspective. The choice of threshold can have a substantial effect on the return values; this can be perceived as unsettling because the choice of threshold is not an exact science and relies on a degree of subjectivity. Figure 4 shows the impact that this subjective choice can have on the return value. The variation is less extreme in the case of the CWD, largely because the tail type is fixed.



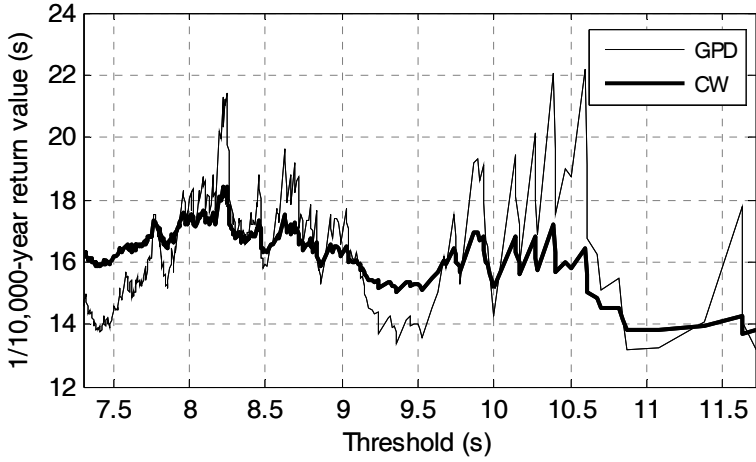


Figure 4: Sensitivity of 1/10,000-year return value to threshold (example for wave parameter  $T_{m-1,0}$  at station K13).

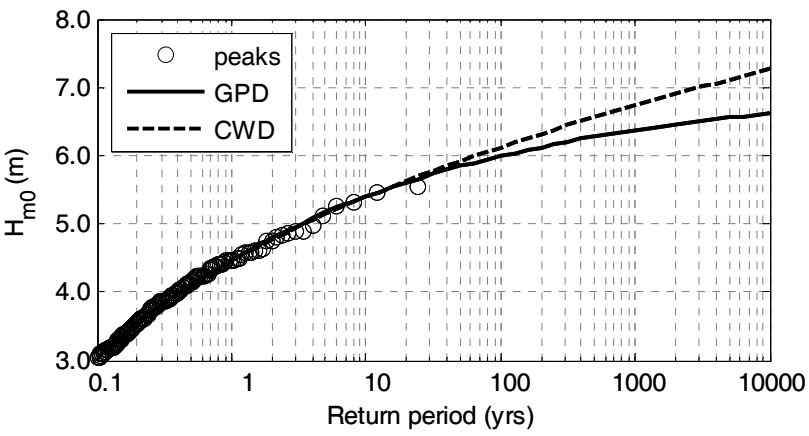


Figure 5: GPD and CWD fits to the same set of POT (example for wave parameter  $H_{m0}$  at station SWB).

The CWD returns consistently higher return values than the GPD. This is again explained by the fixed Type I tail of the CWD. While the Type III tail is understood to better represent the trend of wave data, the large degree of uncertainty means that the estimated curvature and the estimated maximum may be underestimated. The CWD is less likely to underestimate the return values due to its Type I tail, which means that while there is uncertainty around the CWD estimate, it is a more conservative estimate. Figure 5 shows a

representative example of the difference in the GPD and CWD fits for the same set of POT.

## 9 Conclusions and recommendations

The GPD was used to estimate extreme values of wave height and wave period at nine offshore measurement stations in the North Sea. The GPD shape parameters were subsequently spatially smoothed using a modified regional frequency analysis. The GPD was then refitted to a set of POT using the MRFA shape parameter. The results of the analysis were compared with those of conditional Weibull, which is currently operationally applied in The Netherlands.

The GPD consistently returned lower return values, principally due to the flexibility in shape parameter compared with the fixed Type I tail of the CWD. The fitted GPD parameters and consequently the extrapolated extreme values were particularly sensitive to the choice of threshold. An automated method was employed for threshold selection based on the anticipated behavior of the GPD shape parameter as the threshold approaches the domain of attraction. However, for shorter time series this selection method involves much uncertainty, and while reproducible, still involves subjective choices and is not an exact science. The combined sensitivity of the GPD to threshold and the comparatively low extreme value estimates make the application of GPD as carried out in the current study an unattractive choice for the case of short time series.

Because of how critical the choice of threshold is, further research into the selection of the correct threshold is recommended and will be carried out in 2010. Specifically, formal threshold selection techniques discussed in Einmahl et al. [15] will be applied to generated time series to test the applicability, especially in the case of shorter time series.

## Acknowledgement

The presented work is part of the WTI project commissioned by Rijkswaterstaat – Centre for Water Management in The Netherlands.

## References

- [1] Stijnen, J.W., Thonus, B.I., Diermanse, F.L.M., & Geerse, C.P.M., *Hydra-K – Functional documentation, version 3.5 (in Dutch)*, WL | Delft Hydraulics/HKV Report PR1314 for RIKZ, 2007.
- [2] Booij, N., Ris, R.C. & Holthuijsen, L.H., A third-generation wave model for coastal regions, Part 1: Model description and validation, *J. Geophys. Res.*, 104 (C4), 7649– 7666, 1999.
- [3] Coles, S., *An Introduction to the Statistical Modelling of Extreme Values*. Springer Texts in Statistics, Springer-Verlag, London, 2001.
- [4] Caires, S., *A comparative simulation study of the annual maxima and the peaks-over-threshold methods*, Deltares report 1200264-002, 2009



- [5] Pickands, J., The two-dimensional Poisson process and extremal processes. *Journal of Applied Probability*, 8, 745-756, 1971.
- [6] Pickands, J., Statistical inference using extreme order statistics. *Annals of Statistics*, 3, 119-131, 1975.
- [7] Davison, A. C. & Smith, R. L., Models for exceedances over high thresholds. *J. Roy. Stat. Soc.*, 52B, 393-442, 1990.
- [8] Hosking, J. R. M., Wallis, J. R. & Wood, E.F., Estimation of the generalized extreme-value distribution by the method of probability-weighted moments. *Technometrics*, 27, 251-261, 1985.
- [9] Hosking, J. R. M. & Wallis, J. R., Parameter and quantile estimation for the Generalized Pareto Distribution. *Technometrics*, 29, 339-349, 1987.
- [10] Cairns, S., *Extreme wave statistics: methodology and application to North Sea wave data*. WL | Delft Hydraulics Report H4803.30, February 2007.
- [11] Ferreira, J. A. & Guedes Soares, C., Modelling distributions of significant wave height. *Coast. Eng.*, 40, 361-374, 2000.
- [12] Forristall, G.Z., Heideman, J.C., Legget, I.M., Roskam, B. & Vanderschuren, L., Effect of sampling variability on hindcast and measured wave height. *J. Wtrwy., Port, Coast., and Oc. Engrg.*, 122 (5), 216-225, 1996.
- [13] Hosking, J.R.M. & Wallis, J.R., *Regional Frequency Analysis: An approach based on L-Moments*. Cambridge University Press, Cambridge, UK, 1997.
- [14] Den Heijer, F., Diermanse, F.L.M. & van Gelder, P.H.A.J.M., Extreme wave statistics using regional frequency analysis, *paper for the ISSH symposium* (International Symposium on Stochastic Hydraulics), 2005.
- [15] Einmahl, J.H.J., Li, J. & Liu, R.Y., Thresholding events of extreme in simultaneous monitoring of multiple risks. *Journal of the American Statistical Association*, 104 (487): 982-992, 2009.

