Application of a multivariate extreme value approach to system flood risk analysis

D Wyncoll [1] and B Gouldby[1]

[1] HR Wallingford, Howbery Park, Wallingford, Oxfordshire, OX10 8BA

Published in the Journal of Flood Risk Management

Abstract

Effective management of flooding requires models that are capable of quantifying flood risk. Quantification of flood risk involves both the quantification of probabilities of flooding and the associated consequences. Modern flood risk models account for the probabilities of extreme hydraulic loading events and also include a probabilistic representation of the performance of flood defence infrastructure and its associated reliability. The spatial and temporal variability of flood events makes probabilistic representation of the hydraulic loading conditions on the flood defences complex. In the system method used widely within England and Wales, simplifying assumptions relating to the spatial dependence of flood events are made. This paper utilises relatively recent advances in multivariate extreme value methods to refine this method. The paper demonstrates how the refinements can improve estimates of flood consequences for single events and how this can be captured within a risk analysis that incorporates the likelihood of defence failures.

Keywords

Flood risk analysis, Spatial dependence, Multivariate extremes, Systems modelling, Flood defence failure, Reliability analysis

Introduction

Flooding is a global problem with extreme flood events having major consequences recorded at different locations relatively frequently. It is long been recognised that risk management approaches offer a number of benefits over more traditional deterministic design event based approaches, USACE (1996), Sayers *et al.* (2002), for example. A primary component of successful flood risk management is the ability to quantify flood risk and to simulate the risk reduction that results from introducing different mitigation measures.

Flood risk is generally regarded as function of probability and consequence and quantifying flood risk can be complex. Extreme flood events at any particular spatial location occur, by definition, rarely and there is limited data with which to verify models. The temporal variability of flooding can also be complex. Some flash flood events last for a short period of time, a day, for example, whilst others can last for many weeks. Flooding can be localised in terms of its impact or can effect large spatial scales. Within the UK for example, the flooding of the village of Boscastle, that took place over a day Roca-Collel and Davison (2010), can be contrasted with the summer floods of 2007 Marsh and Hannaford (2007), that lasted for a period of six weeks and effected most of England and parts of Wales. Flooding can arise from different sources; pluvial, fluvial, waves and surges for example. These different sources can often interact to exacerbate the impact of the flood event.

As risk is a function of probability, establishing the joint probability of different combinations of flood inducing events arising has been the subject of much research. Examples include extreme, fluvial flows at confluences, estuarine water levels, astronomical tides and metrological surges, waves and sea levels, see for example, Tawn and Vassie (1989), Acreman (1994), Bruun and Tawn (1998), Hawkes *et al.* (2002), Hawkes (2008). These approaches have, in the past, tended to have been restricted in terms of the number of variables that have been considered and the spatial extents that are covered. The reasons for this primarily relate to constraints placed by the statistical models that have been used. In particular, traditional multivariate extreme-value statistical models have constraints relating to the handling of the dependence structure. These constraints make extensions to large spatial scales, where the degree of dependence between extremes of a given variable can cover a wide range, impractical to implement. Recent developments in multivariate extreme value methods, Heffernan and Tawn (2004), have however, removed some of these constraints and opened opportunities for improving flood risk analysis methods.

Risk analysis models of flooding systems that incorporate the likelihood of flood defence failure are now recognised as effective means of supporting the flood risk management process. These models have been applied at national and regional scales for supporting decisions relating to long term climate change adaptation, strategic planning and asset management prioritisation purposes, see USACE (1996), Hall *et al.* (2003), Evans *et al.* (2006), Apel *et al.* (2004), Gouldby *et al.* (2008) and Vorogushyn *et al.* (2010), for example. Often, when aggregating flood risks over large spatial scales, an assumption of independence between extreme events occurring at different spatial locations is made. Keef *et al.* (2009) and Lamb *et al.* (2010) note that this simplifying assumption can lead to inaccuracies when seeking to define single event impacts at large spatial scales.

This paper describes the application of the Heffernan and Tawn (2004), multivariate extreme value method, updated by Keef *et al.* (2009), to generate boundary conditions for a system flood risk analysis method that has been widely used in England and Wales for national and regional flood risk analysis, Gouldby *et al.* (2008). The application enables the assumptions regarding spatial dependence to be explored in the context of a systems based flood risk analysis. The new system has been applied on a case study in the North West of England, with the results described below.

Background to system flood risk analysis

System risk models that are currently applied in practice typically define risk through consideration of the aleatory uncertainty associated with the random nature of extreme flood events and the epistemic uncertainty associated with the potential for flood defence infrastructure to suffer failure. Whilst there are many other sources of uncertainty, approaches to quantify these, Merz and Thieken (2009), Gouldby *et al.* (2010), for example, are not commonly applied in practice, primarily due to the computational burden associated with the implementation of the methods. The primary components of the risk models are: hydraulic loads, described by extreme value distributions, flood defence infrastructure, defined by fragility curves, flood inundation simulation representation and functions that relate the simulated floods to consequences. The flood defences are typically defined as discrete lengths with a specific fragility curve prescribed for each length. The models that are applied in practice within England and Wales assume the performance of each of the defence lengths is independent and hence the risk, expressed in terms of the Expected Annual Damage (EAD), is given by:

$$EAD = \int \sum_{i=1}^{2^n} P(d_i|X) f_X(X) g(d_i, X) dX \tag{1}$$

where n is the number of defence lengths, f_X is the joint density of hydraulic loads over the defence lengths, d_i is the defence system state (a vector that comprises a representation of the state as failed or undamaged of each defence in the system) and g is a function that relates the defence system state and hydraulic loads to the economic damage to property.

The derivation of the joint density of the hydraulic loads can be complex to define over large spatial areas and a simplifying assumption is therefore made in current practice. The hydraulic loading conditions are assumed to be fully dependent in terms of recurrence interval (return period) within a flood area. This enables the integration of the joint density of the hydraulic loads over the consequence function to be undertaken in terms of a single likelihood of hydraulic load using a simple integration procedure:

$$EAD \approx \sum_{i=1}^{N} P\left(\frac{x_{i-1} + x_i}{2} < X \leq \frac{x_i + x_{i+1}}{2}\right) \overline{g}(x_i) \tag{2}$$

where $\overline{g}(x_i)$ is the expected economic damage for the hydraulic load x_i. Flood areas are typically less than 10 km^2 and Eqn. 2 is evaluated independently for each flood area. To assess the risk at spatial scales larger than a single flood area, the results for each flood area are aggregated. It is of note however, that important information relating to the nature of flood events can be lost when considering only the EAD. This approach does not, for example, enable assessment of the probability and consequences associated with a single flood event that occurs at spatial scales larger than a flood area, when there is partial dependence between the hydraulic loading conditions. Analysis of single event consequences at large spatial scales (ie bigger than a single flood area) can be of interest and importance for the insurance industry, where single event loss damages are of interest and emergency planners, for example, where large-scale evacuation plans are developed.

Within the analysis described here, the simplifying assumption of full dependence between flood areas is replaced with a multivariate extreme value model. This model is described in more detail below. The use of the multivariate approach enables information on the probability and consequences associated with single events to be captured more accurately, whilst also considering the likelihood of the flood defences failing.

Method

The primary aim of the application of the multivariate approach is to solve Eqn. 1 with an explicit consideration of the dependence, in the extremes, of the hydraulic loading conditions (the vector X) at different spatial locations. It is of note that the dependence in the extremes is likely to vary significantly between the sites. Often, but not always, sites close to each other will exhibit high dependence with one another, sites further a-field will often be weakly dependent. The statistical model needs to be sufficiently flexible to accurately capture the range of dependence between the sites.

The steps in the analysis are

1. Fit statistical model to joint extremes of hydraulic loads at multiple locations within the study area.

2. Simulate extreme hydraulic loading events from the fitted model that form the hydraulic boundary condition element of the system risk model.
3. For each simulated hydraulic loading event, undertake multiple realisations of flood defence system states and associated inundation extents and economic impacts.
4. Integrate the results of Eqn. 3 to determine the risk expressed as EAD.

Statistical model for hydraulic loads

Of the range of multivariate extreme value models that are available, the approach of Heffernan and Tawn (2004) offers most flexibility in capturing the range of dependence exhibited by flood related variables. Other approaches, Coles and Tawn (1991) and Ledford and Tawn (1996) for example, have restrictive assumptions relating to the dependence among multiple variables. The Heffernan and Tawn (2004) approach is therefore the model of choice for this study.

Let $\boldsymbol{X_t} = (X_1, \ldots, X_j)_t$ be a time series of flood related hydraulic loading conditions at a collection of locations within the area of interest. Before dependencies between locations are considered, the extreme loads at each location are studied marginally. For this the standard peaks-over-threshold approach of Davison and Smith (1990), is used: cluster maxima are identified from the time series and the excesses above a suitably high threshold are fitted to the generalised Pareto distribution (GPD). This defines a probability model for large values of the variable X_i:

$$P\{X_i > x \mid X_i > u_i\} = \left[1 + \xi_i \frac{(x - u_i)}{\beta_i}\right]_+^{-1/\xi_i} \quad \text{for } x > u_i, \qquad (3)$$

where $\beta_i > 0$, $\xi_i \in \Re$ are the GPD parameters and $[y]_+ = \max(y, 0)$. The threshold u_i is chosen to be just large enough to ensure a stable estimate for the shape parameter ξ_i for all larger thresholds.

In order to separate the marginal characteristics from the dependence analysis, it is usual to standardise the data to common margins using the probability integral transform. This is often referred to as the copula approach. The Heffernan and Tawn (2004) model uses standard Gumbel marginal distributions which are obtained by setting $Y_i = -\log(-\log[\hat{F}_i(X_i)])$ where $\hat{F}_i$ is an estimate of the cumulative distribution function for X_i. For this the GPD fit above the threshold is combined with the empirical distribution $\tilde{F}_i$ of the X_i values to give the following semi-parametric function (first used by Coles and Tawn (1991)):

$$\hat{F}_i(x) = \begin{cases} \tilde{F}_i(x) & x \le u_i, \\ 1 - (1 - \tilde{F}_i(u_i))\left[1 + \xi_i \dfrac{(x - u_i)}{\beta_i}\right]_+^{-1/\xi_i} & x > u_i. \end{cases} \qquad (4)$$

The transformed multivariate time series $\boldsymbol{Y_t} = (Y_1, \ldots, Y_d)_t$ retains the dependence structure of the original data but satisfies $P\{Y_i > y\} = \exp(-\exp[-y])$ for each different location *i*. The primary aspect of the Heffernan and Tawn (2004) approach is to model the dependence between extreme values of Y_i and typical values of the remaining variables. The analysis is repeated for each site *i* so that extreme values of all variables are considered.

Let $\boldsymbol{Y}_{-i}$ denote the vector of all variables Y_j excluding Y_i. The Heffernan and Tawn (2004) approach is typically applied using the multivariate non-linear regression model

$$\boldsymbol{Y}_{-i} \mid Y_i = \boldsymbol{a}\, Y_i + Y_i^{\boldsymbol{b}}\, \boldsymbol{Z} \qquad \text{for } Y_i > v, \tag{5}$$

where v is a high threshold on Y_i, $\boldsymbol{a} \in [0, 1]$ and $\boldsymbol{b} < 1$ are vectors of parameters and $\boldsymbol{Z}$ is a vector of residuals. Vector arithmetic should be interpreted component-wise so that each Y_j is modelled as a function of Y_i using parameters $a_{j|i}$ and $b_{j|i}$ and residual $Z_{j|i}$.

The regression parameters $a_{j|i}$ and $b_{j|i}$ are estimated using maximum likelihood under the temporary assumption that $Z_{j|i}$ follows a Normal Distribution with unknown mean and variance. This fit uses all pairs (Y_i, Y_j) corresponding to cluster maxima of $Y_i > v$ to be consistent with the marginal GPD fits made to cluster maxima of X_i. Heffernan and Tawn (2004) show that asymptotically $Y_i > v$ is statistically independent of the residual $Z_{j|i}$. The threshold v is therefore chosen to be just large enough for this condition to hold. Once all parameter estimates have been found a non-parametric estimate of the joint distribution of $\boldsymbol{Z}$ is constructed from the empirical distribution of the sample residuals.

The above description assumes variables Y_i and Y_j occur concurrently so is not appropriate for modelling temporally dependent data with extreme events that are lagged between gauges. Keef *et al.* (2009) overcome this deficiency by fitting the conditional model of $Y_{j,t+\tau} \mid Y_{i,t}$ for a selection of lags τ for each gauge $j \neq i$. This allows the model to be used to simulate new events over a range of lags so that the largest values in each time window need not occur concurrently.

Statistical simulation of extreme hydraulic loading events

The fitted model provides parameter estimates $a_{j|i,\tau}$ and $b_{j|i,\tau}$ for all locations i, $j \neq i$ and for a range of lags τ. Additionally, for each location i, an empirical sample of joint residuals $\boldsymbol{Z}$ is available, each of which provides values $Z_{j|i,\tau}$ for each $j \neq i$ and lag τ. These can be used to simulate a large number of dependent peak events $\boldsymbol{Y}$, each of which consists of a single peak value for every location in the network with associated lags between peaks. These are then transformed to give samples of spatially dependent peak flow events $\boldsymbol{X}$ on the original scale.

To simulate an event, a conditioning site, i must be selected. The value Y_i is sampled above the threshold v and the remaining Y_j values are sampled from the fitted model for $\boldsymbol{Y}_{-i,\tau} \mid Y_i$. In order to control the proportion of events where each site is most extreme, the value Y_i is constrained to be largest by rejection sampling.

The simulation consists of repeating the following steps, after selecting a conditioning site i:

1. Sample a value Y_i from the standard Gumbel distribution conditioned to exceed v.
2. Independently select one of the joint residuals $\boldsymbol{Z}$ for site i.
3. Calculate $Y_{j,\tau} = a_{j|i,\tau} Y_i + Y_i^{b_{j|i,\tau}} Z_{j|i,\tau}$ for all $j \neq i$ and lags τ.
4. For each j, set Y_j to be the maximum $Y_{j,\tau}$ over selected lags τ.
5. The joint sample $\boldsymbol{Y}$ is rejected unless Y_i is maximum over all gauges.

These simulated events have Gumbel marginal distributions which can be transformed to the original scales using the probability integral transform, inverting the transform applied to the original data using GPD fits and the empirical distribution.

To select the proportion of events where each gauge is most extreme, Lamb R *et al.* (2010) propose simulation from the fitted model without rejection and this approach has been adopted here. For each gauge

i, a large sample of samples is used to estimate $P\{Y_i \geq Y_j \, \forall \, j \neq i \mid Y_i > v\}$, the probability that Y_i is largest when i is the conditioning gauge. Since all simulated events have at least one threshold exceedance, this can be used to estimate the proportion of all events where Y_i is largest and hence variable X_i is most extreme.

The output from this analysis is a simulated set of hydraulic loading events. Each event comprises a peak flood variable at each location. An example of fluvial flow events simulated from the statistical model (ignoring lags and rejection) and compared to the observed data is given on Gumbel scales in Figure 1. The black dots show the empirical data for sites i and j whilst the overlain blue dots show those generated from the multivariate model with conditioning site i above the threshold. It is apparent that the dependence in the extremes for site i within the simulated data reproduces well that within the observations. Further examples are available within Heffernan and Tawn (2004) and for flood related variables within Lamb *et al.* (2010), for example.

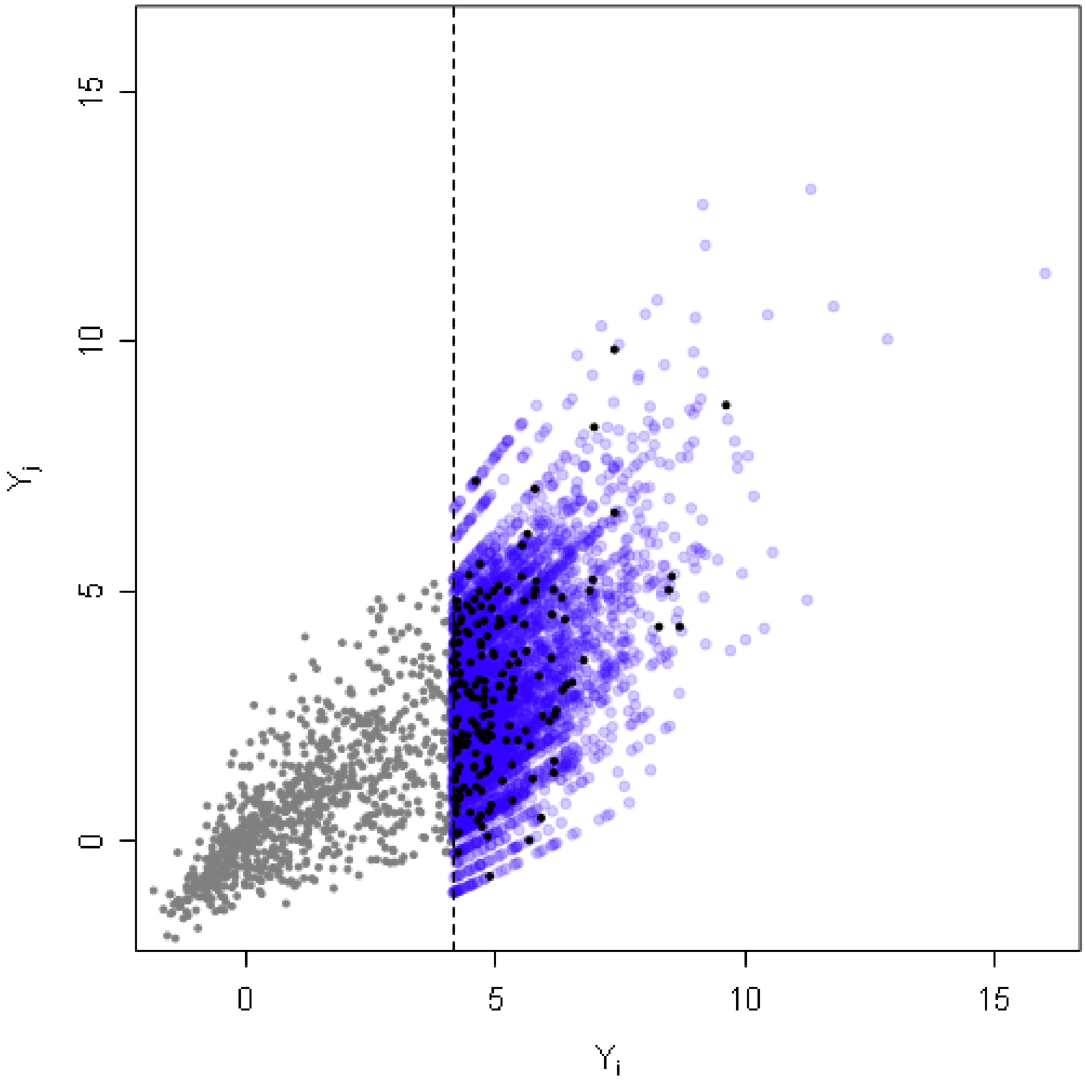

Figure 1: Comparison of observed peak flow events (black) against samples from the Heffernan and Tawn (2004) model (blue) for conditioning site i plotted on the Gumbel scale.

Defence system state reliability and floodplain analysis

Each flood event that has been simulated from the statistical model forms the boundary conditions for the flood defence system reliability analysis and subsequent inundation and economic consequence evaluation. The line of flood defences that forms the boundary between the river or coast and the floodplain area is discretised into sections that have different performance characteristics (ie type of defence, elevation and condition). Fragility curves are used to define the performance of each defence section, USACE (1996), Simm *et al.* (2009), Schultz *et al.* (2010), Vorogushyn *et al.* (2009), for example, and the performance of each defence section is assumed to be independent from any other. It is assumed that each defence can exist in two possible states, failed (ie structural failure or breached), or undamaged. During a flood event, water can enter the floodplain through any particular defence if it is a breached defence or if it remains undamaged but is overtopped. As there are a finite number of defence sections and only two possible system states for each section, there are a finite number of defence system states (combinations of failed and undamaged defences). It is, however, computationally impractical to run an inundation simulation for all of the different defence system states and, for each hydraulic loading event, a Monte Carlo simulation of the defence system states is therefore undertaken.

For each simulated defence system state, floodplain discharge volumes are calculated for each flood defence. This water is then propagated across the floodplain using a computationally efficient spreading algorithm Lhomme *et al.* (2008). Flood depths are then determined and their associated economic damages estimated with reference to widely applied functions, Penning-Rowsell *et al.* (2005).

Within any particular discrete location within the floodplain, the probability of exceeding any particular flood depth (h), conditional on the specified hydraulic loading level (x) is therefore:

$$P(H > h \mid x) \approx \left({}^{m_h}\!/_{m}\right) \qquad (6)$$

where m is the total number of defence system state Monte Carlo realisations and m_h is the number that exceed h, under the specified loading level x. Similarly, the probability of exceeding a specific economic damage level (c) is:

$$P(C > c \mid x) \approx \left({}^{m_c}\!/_{m}\right) \qquad (7)$$

Combining the Monte Carlo realisations for the hydraulic loading with the defence system states, it then follows that the unconditional annual probability of exceeding any particular level of economic damage is:

$$P(C > c) \approx \frac{(l_c / l)}{n_y} \qquad (8)$$

where l_c is the number of Monte Carlo realisations that exceed c, l is the total number of Monte Carlo realisations and n_y is the number of years of data that has been simulated. The risk, expressed in the usual terms of EAD is:

$$E(C) \approx \frac{\sum_{i=1}^{l} c_i}{n_y} \qquad (9)$$

where c_i is the economic damage arising on the ith realisation of the hydraulic loading conditions. The number of realisations required can be controlled in a number of different ways, for example, through specification of convergence criteria on the quantity of interest.

Case study application

Study area description

The case study covers a catchment in the northwest of England and includes the city of Carlisle (Figure 2). Carlisle has a history of flooding with significant events having been recorded in 1963, 1968, 1979, 1980, 1984 and most recently in January 2005, when nearly 3000 homes were flooded, 3 people died and the resulting flood losses were estimated as £400m, Geographical Association (2009). The floodplain protection varies significantly across the study area and includes high ground in rural areas, embankments and a range of heavily engineered vertical structures within urban areas, Carlisle for example. The study area has been separated into 10016 flood areas and is consistent with those used in the National flood risk assessment of England, Environment Agency (2009).

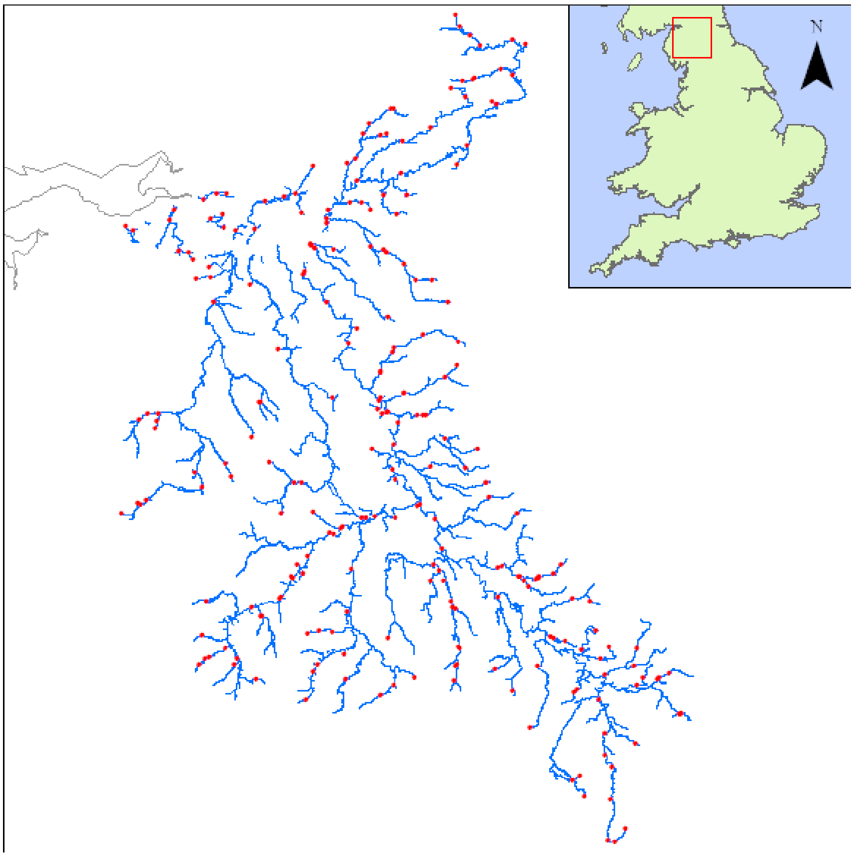

Figure 2: Maps showing the study area for the case study with the representative nodes spanning the river network.

Hydraulic load analysis

Synthetic flow data form the basis of the analysis and these stem from a series of linked models. Time series rainfall, on a 25km spatial grid and an hourly time interval provide the driving data for a gridded hydrological model (Grid-to-Grid, or "G2G", Bell (2007). These rainfall data, together with estimates of potential evaporation (PE) have been generated from the UK Met. Office Hadley Centre global climate model (GCM), HadCM3, dynamically downscaled to 25km using the HadRM3 regional model, Jones R *et al.* (2004). The G2G Model operates on a 1km grid covering the whole of the UK, and, using a gridded representation of the kinematic wave equation, simulates surface and sub-surface flows. A full description of the G2G model is provided in Bell (2007).

The G2G Model provides 27 years of time series flows on a 1km UK grid along the river network. The gridded data were first interpolated to each of 4119 nodes that span the river network by selecting the largest value of each surrounding grid point at every time step. Rather than applying the statistical analysis to every node, variables that had a correlation over 0.99 with a remaining node were removed to leave a representative subset of 284 sites.

The marginal extremes at each of these selected nodes have been analysed separately. Peak flow clusters were identified using the runs method of Smith and Weissman (1994) before the cluster maxima above a high threshold were fitted to GPD. After transformation to Gumbel marginal distributions, cross-correlations were analysed using node pairs lagged by up to ±6 days. In the majority of cases, the largest correlation was obtained by a lag within ±2 days. Hourly lags of up to ±3 days were therefore used to fit the multivariate statistical model to every pair of nodes in the subset, adopting a similar approach to that of Lamb *et al.* (2010).

Peak flows were then paired with local maxima at every other node so that extreme events that had at least one threshold exceedance could be identified. 709 such events were counted giving an average of 26 extreme events per year. With this, 1000 years worth of hydraulic loading conditions were simulated from the fitted model on the Gumbel scale. Each sample comprised a peak flow level for all 284 selected nodes and had at least one threshold exceedance. Rejection sampling was used to control the proportion of events that are most extreme for each site.

After transformation back to the original scale, linear relationships identified between nodes with large correlations were used to extend the peak flow samples to all 4119 river nodes. These were then converted to water levels using the Conveyance Estimation System, McGahey C *et al.* (2009) and mapped to all 34006 discrete river defence sections in the catchment.

Figure 3 shows a comparison of hydraulic loading data (peak fluvial flows) simulated from the statistical model for a subset of stations. The figure shows the range of dependence that is present across the study area and highlights the flexibility required of the multivariate model to capture this. Perhaps unsurprisingly, sites in close proximity exhibit a high degree of dependence whilst others, further a-field are only partially dependent on one another.

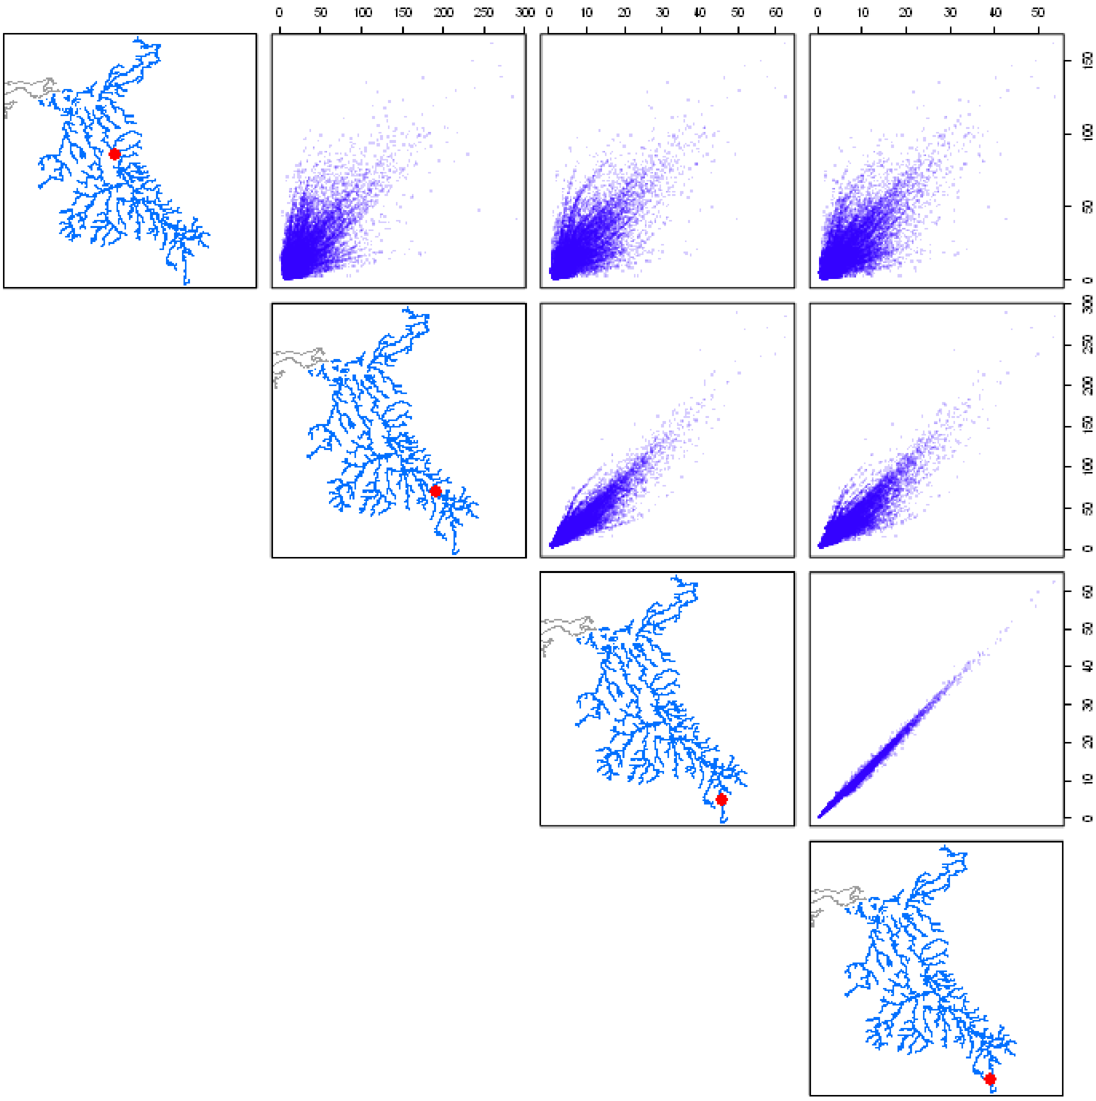

Figure 3: Comparison of peak flow events sampled from the statistical model four sites on the river network.

Figure 4 shows the variation in return period of the hydraulic loading levels across the study area for a subset of single realisations from the multivariate model. The spatial variation in the intensity of the flood events is apparent. Some events spatially diverse and extreme flows arise across the study area, whereas other events are very much localised with extremes arising in relatively small areas over the catchment. This highlights the difficulties that have arisen in the past to describe flood events in terms of the severity of their loading conditions - “The event was a 1 in 100 year (flow) event”, for example. It is apparent that statements about flood events, expressed in terms of the extremes of the hydraulic loading conditions, are often only valid for small spatial areas. Attempting to express widespread flood events, like those that affected England Wales in the summer of 2007, in terms of return period of the hydraulic loading conditions, as often seems to be a requirement, is clearly a frivolous exercise and something that is probably best avoided.

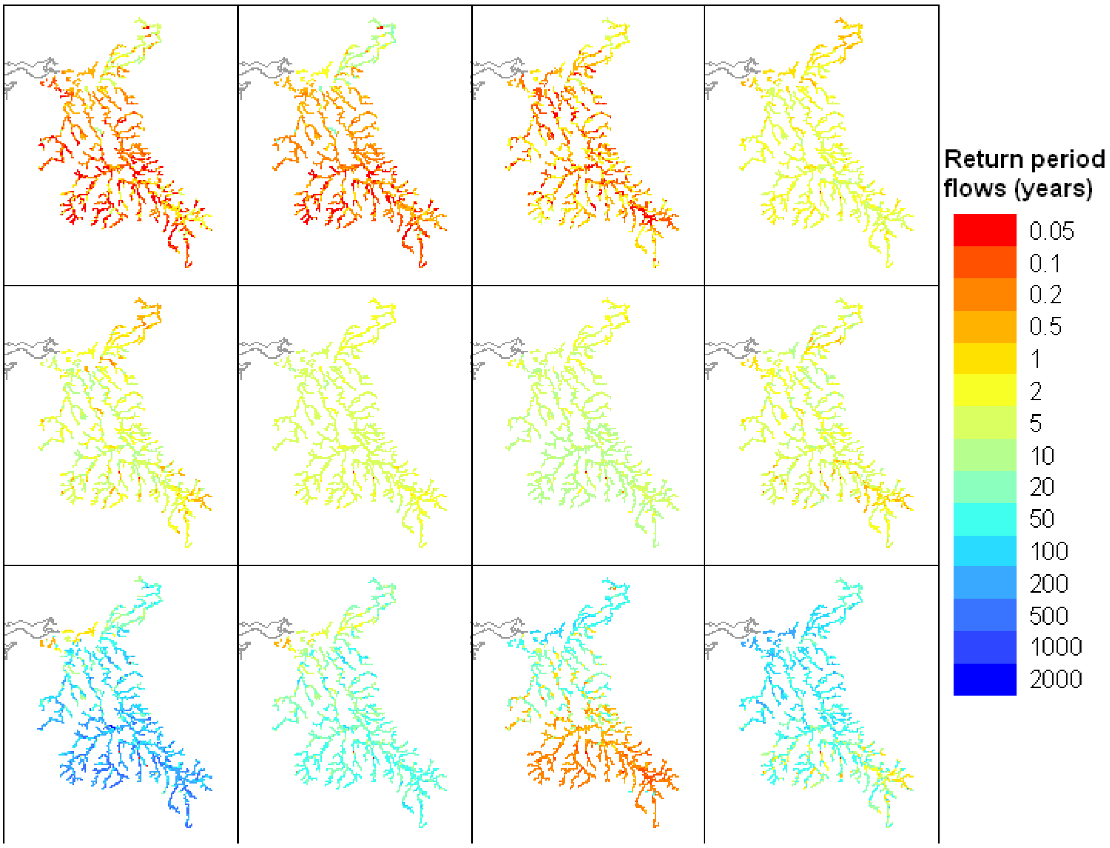

Figure 4: Variation in return period water levels across the study area for a subset of sampled hydraulic loading events.

Lewis et al. (2011) draw attention to the validity of making an assumption of complete dependence of hydraulic loading level (equal return period) within a flood area with system based flood risk analysis models, for example Hall et al. (2003), Gouldby et al. (2008). This assumption is made to enable a simplified integration of the joint density of the hydraulic loads over the consequence function (Eqn. 2). Their analysis was conducted within the context of coastal systems and it is possible to explore the validity of this assumption further here, within the context of this fluvial system. Figure 5 shows the variation of return period within one of the largest flood areas for the same events shown in Figure 4. The flood area comprises 30 defences, the maximum and minimum return periods and average return period for each of the hydraulic loading events are shown in Table 1. It is apparent that even within a flood area there can be significant variation in the return period of the water level and the assumption of equal likelihood can potentially introduce significant bias in some cases. This reinforces the observations on coastal systems of Lewis et al. (2011).

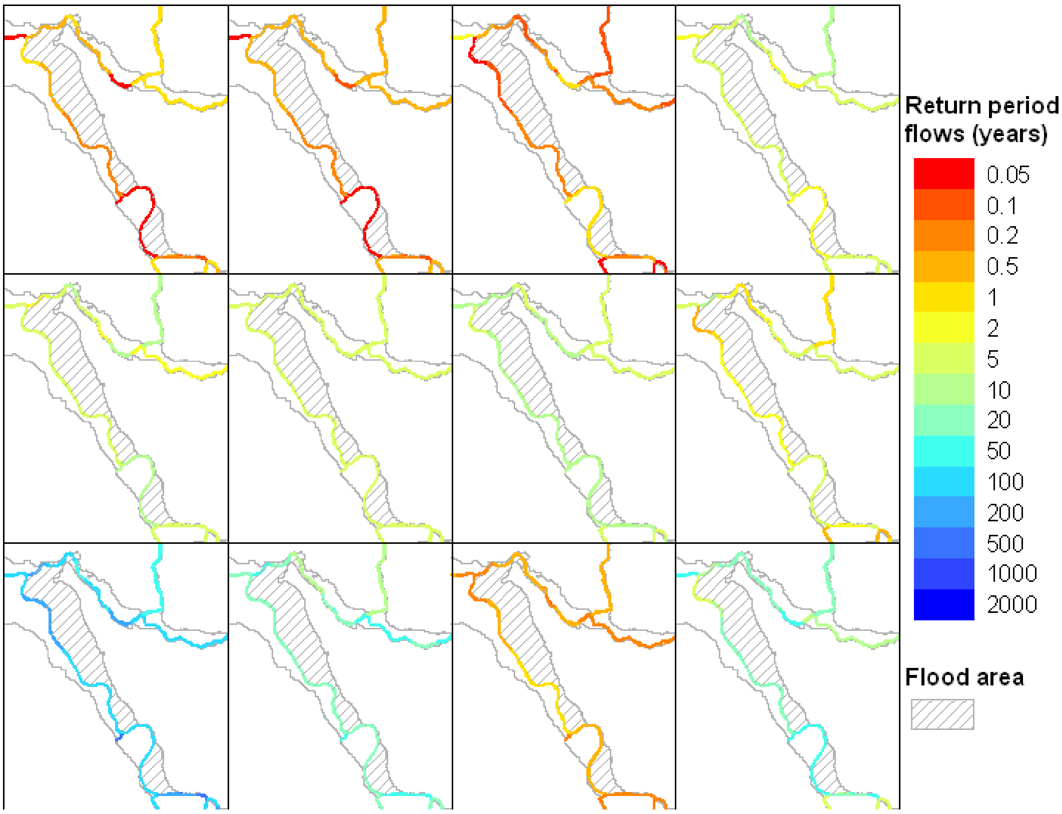

Figure 5: Variation in return period water levels within a particular flood area.

Economic damages and risk estimation

Within the current National Flood Risk Assessment of England, Environment Agency (2009), the EAD is derived at the level of a flood area and then aggregated to obtain a regional or national value of risk. The method does not enable estimates of return period damages, or flood hazard characteristics of depth and velocity, for single events that are greater than a flood area. The introduction of the multivariate method does however, enable this. For each realisation of the multivariate model of the hydraulic loads, a Monte Carlo simulation of the defence system states and associated inundation and consequence analysis, within each flood area, has been undertaken. The resulting risk, in terms of EAD, over the whole study area is £19.7 million.

The distribution of damage aggregated over the whole study area is shown in Figure 6. A comparison has also been made under the assumption of full dependence and complete independence of hydraulic load, on a flood area basis. It is important to note that whilst the EAD for the three cases in Figure 6 does not change, the distribution of damages is quite different. It is this information that can be of importance for the insurance industry, for example, where estimates of likelihood of losses for single flood events are often required.

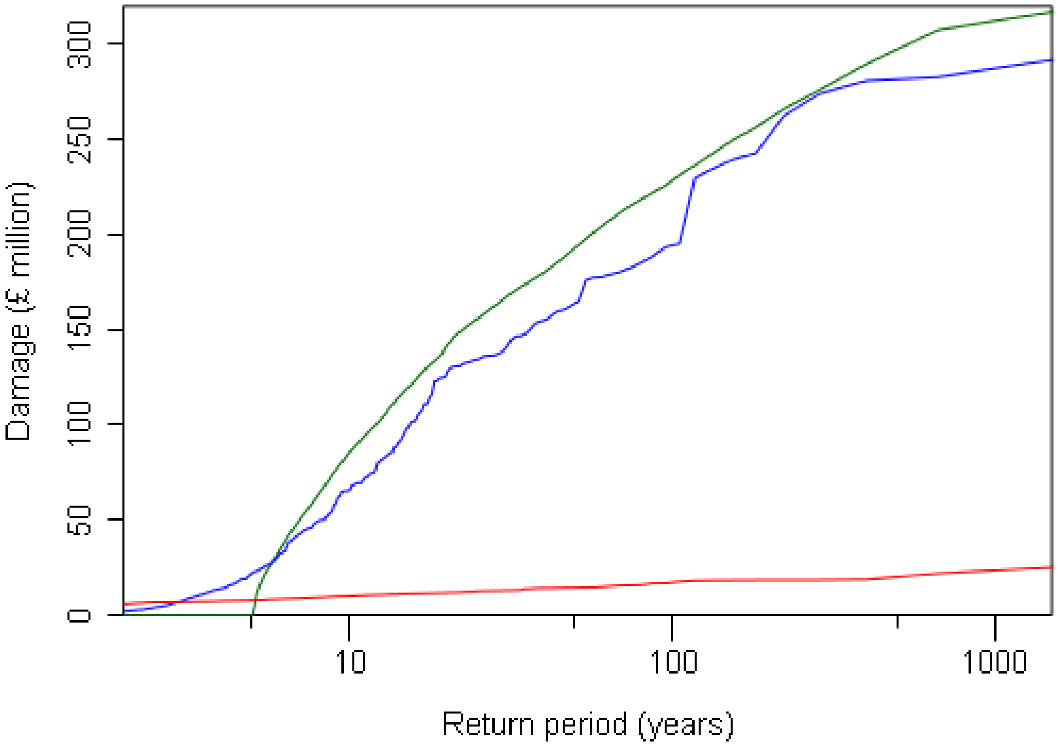

Figure 6: Return period total damage for the study area for the fully dependent (green), partially dependent using the statistical model (blue) and independent (red) cases.

In this particular study area the results show the modelled dependence of economic damages to be close to full dependence. It is important to note that this finding is only applicable for this study area and other study areas will differ. For example, Lamb *et al.* (2010) found the degree of dependence was closer to independence for a study area in the northeast of England.

Conclusions

Flooding is a global problem. Effective flood risk management decision making is underpinned by quantitative system models of flood risk. The models that are applied in current practice within England and Wales incorporate a wide range of hydraulic loading conditions, a representation of the likelihood of failure of the flood defence assets and a computationally efficient inundation model. The current models do however, utilise a simplifying assumption to facilitate the evaluation of the flood risk. The hydraulic loading events within a flood area are assumed to be fully dependent in terms of likelihood. This assumption is made to facilitate the calculation of risk in a computationally efficient manner and to overcome the complexities relating to the dependence structure within the extremes of the hydraulic loading conditions.

An artefact of this assumption is the inability of the current method to provide information on the likelihood of damages for single flood events at spatial scales larger than a flood area. Within the current method, risk, expressed as EAD is calculated on a flood area basis and aggregated to obtain risk at larger spatial scales. Information on the probability of exceeding a specified economic damage threshold for a single event is not currently available. This information can be important for the insurance industry and emergency planners, for example.

To overcome these deficiencies within the current method, a multivariate extreme value method has been applied to determine the joint probability density of the hydraulic loading conditions. A Monte Carlo simulation of the fitted statistical model enables the hydraulic boundary conditions across the flood defence system to be represented more realistically. The resulting method retains the system characteristics of the current methodology and its ability to reflect the performance of the flood defence infrastructure. This is implemented through a secondary Monte Carlo simulation of the defence system states, conditional on the output of the multivariate model.

This development has enabled the exploration of the assumption of full dependence of hydraulic loading likelihood within a flood area. The analysis shows that in some cases, for a single flood event the return period can vary significantly, even within a flood area. This is likely to introduce a bias in estimates of flood risk that utilise this assumption. This concurs with the findings of recent research relating to coastal hydraulic loading events.

The implementation of the multivariate method enables the probability of single flood event damages to be quantified. This refinement can potentially offer insights into the spatial characteristics of single flood events that are not currently possible with the existing system. It also highlights the limitations of using the EAD. EAD is a relatively simplistic measure of flood risk and analysis of additional information comprised within the full distribution can be of importance.

Acknowledgements

This work was supported by the Natural Environment Research Council, Grant Ref. NE/E002420/1 - through the Flood Risk for Extreme Events (FREE) NERC directed mode programme and HR Wallingford´s internal research programme. The authors are grateful to Dr Vicky Bell and Dr Nick Reynard of the Risk Modelling centre at CEH Wallingford, for the support provided with regard to their Grid2Grid Model.

References

Acreman, M. C. (1994). "Assessing the Joint Probability of Fluvial and Tidal Floods in the River Roding." Water and Environment Journal **8**(5): 490-496.

Apel, H., Thieken, A., Merz, B. and Bloschl, G. (2004). "Flood risk assessment and associated uncertainty." Natural Hazards and Earth System Sciences **4**: 295 - 308.

Bell, V., Kay A, Jones R, Moore R (2007). "Development of a high resolution grid-based river flow model for use with regional climate model output." Hydrology and Earth System Sciences **11**(1): 532 - 549.

Bruun, J. T. and Tawn, J. A. (1998). "Comparison of approaches for estimating the probability of coastal flooding." Journal of the Royal Statistical Society: Series C (Applied Statistics) **47**(3): 405-423.

Coles, S. G. and Tawn, J. A. (1991). "Modelling Extreme Multivariate Events." Journal of the Royal Statistical Society. Series B (Methodological) **53**(2): 377-392.

Davison, A. C. and Smith, R. L. (1990). "Models for exceedances over high thresholds." Journal of the Royal Statistical Society, Series B (Statistical Methodology) **52**(3): 393-442.

Environment Agency (2009). Flooding in England: A National Assessment of Flood Risk. Bristol, UK, Environment Agency

Evans, E., Hall, J., Penning-Roswell, E., Sayers, P., Thorne, C. and Watkinson, A. (2006). "Future flood risk management in the UK." Water Management **159**(WM1): 53 - 61.

Geographical Association (2009). Managing flood risk - Carlisle Case Study. http://www.geography.org.uk/resources/flooding/carlisle.

Gouldby, B., Sayers, P., Mulet-Marti, J., Hassan, M. and Benwell, D. (2008). "A methodology for regional-scale flood risk assessment." Proceedings of the Institution of Civil Engineers-Water Management **161**(3): 169-182.

Gouldby, B. P., Sayers, P. B., Panzeri, M. C. and Lanyon, J. E. (2010). "Development and application of efficient methods for the forward propagation of epistemic uncertainty and sensitivity analysis within complex broad-scale flood risk system models." Canadian Journal of Civil Engineering **37**(7): 955-967.

Hall, J., Dawson, R., Sayers, P., Rosu, C., Chatterton, J. and Deakin, R. (2003). "A methodology for national-scale flood risk assessment." Proceedings of the Institution of Civil Engineers, Water and Maritime Engineering **156**(3): 235 - 247.

Hawkes, P. J. (2008). "Joint probability analysis for estimation of extremes." Journal of Hydraulic Research **46**(sup2): 246-256.

Hawkes, P. J., Gouldby, B. P., Tawn, J. A. and Owen, M. W. (2002). "The joint probability of waves and water levels in coastal engineering design." Journal of Hydraulic Research **40**(3): 241-251.

Heffernan, J. E. and Tawn, J. A. (2004). "A conditional approach for multivariate extreme values (with discussion)." Journal of the Royal Statistical Society: Series B (Statistical Methodology) **66**(3): 497-546.

Jones R, Noguer N, Hassell D, Hudson D, Wilson S, Jenkins G and Mitchell J (2004). Generating high resolution climate change scenarios using PRECIS. M. O. H. Centre. Exeter**:** 40.

Keef, C., Tawn, J. and Svensson, C. (2009). "Spatial risk assessment for extreme river flows." Journal of the Royal Statistical Society: Series C (Applied Statistics) **58**(5): 601-618.

Lamb R, Keef C, Tawn J, Laeger S, Meadowcroft I, Surendran S, Dunning P and Batstone C (2010). "A new method to assess the risk of local and widespread flooding on rivers and coasts." Journal of Flood Risk Management **3**(4): 323-336.

Lamb, R., Keef, C., Tawn, J., Laeger, S., Meadowcroft, I., Surendran, S., Dunning, P. and Batstone, C. (2010). "A new method to assess the risk of local and widespread flooding on rivers and coasts." Journal of Flood Risk Management **3**(4): 323-336.

Ledford, A. W. and Tawn, J. A. (1996). "Statistics for near independence in multivariate extreme values." Biometrika **83**(1): 169-187.

Lewis, M., Horsburgh, K., Bates, P. and Smith, R. (2011). "Quantifying the Uncertainty in Future Coastal Flood Risk Estimates for the U.K." Journal of Coastal Research: 870-881.

Lhomme, J., Sayers, P., Gouldby, B., Samuels, P., Wills, M. and Mulet-Marti, J. (2008). Recent development and application of a rapid flood spreading method. Flood Risk 2008, Oxford.

Marsh, T. and Hannaford, J. (2007). The summer 2007 floods in England and Wales – a hydrological appraisal. C. f. E. a. Hydrology. Wallingford

McGahey C, Samuels PG and Knight DW (2009). "Advice, methods and tools for estimating channel roughness." Proceedings of the Institution of Civil Engineers, Water Management **162**(6): 353–62.

Merz, B. and Thieken, A. H. (2009). "Flood risk curves and uncertainty bounds." Natural Hazards **51**(3): 437-458.

Penning-Rowsell, E., Johnson, C., Tunstall, S., Tapsell, S., Morris, J., Chatterton, J. and Green, C. (2005). The Benefits of Flood and Coastal Risk Management: A manual of assessment techniques, Middlesex University Flood Hazard Research Centre.

Roca-Collel, M. and Davison, M. (2010). "Two dimensional model analysis of flash-flood processes: application to the Boscastle event." Journal of flood risk Management **3**(1): 63-71.

Sayers, P., Hall, J. and Meadowcroft, I. (2002). "Towards risk-based flood hazard management in the UK." Civil Engineering **150**: 36 - 42.

Schultz, M., Gouldby, B., Simm, J. and Wibowo, J. (2010). Beyond the Factor of Saftey: Developing Fragility Curves to Characterize System Reliability, Water Resources Infrastructure Program. USACE, Washington, D.C. ERDC SR-10-1.

Simm, J., Gouldby, B., Sayers, P., Flikweert, J., Wersching, S. and Bramley, M. (2009). Representing fragility of flood and coastal defences: Getting into the detail. Flood Risk Management: Research and Practice. P. Samuels, S. Huntington, W. Allsop and J. Harrop, Taylor and Francis Group, London, UK.

Smith, R. L. and Weissman, I. (1994). "Estimating the extremal index." Journal of the Royal Statistical Society, Series B (Methodological) **56**(3): 515-528.

Tawn, J. and Vassie, I. (1989). " Extreme sea levels: the joint probabilities method revisited and revised." Proc Inst Civ Eng(87): 429–442.

USACE (1996). Risk-based Analysis for Flood Damage Reduction studies. Engineer Manual EM 1110-2-1619, U.S. Army Corps of Engineers, Washington, D.C.

Vorogushyn, S., Merz, B. and Apel, H. (2009). "Development of dike fragility curves for piping and micro-instability breach mechanisms." Nat. Hazards Earth Syst. Sci. **9**(4): 1383-1401.

Vorogushyn, S., Merz, B., Lindenschmidt, K. E. and Apel, H. (2010). "A new methodology for flood hazard assessment considering dike breaches." Water Resources Research **46**: 17

Tables

Table 1: Minimum, mean and maximum return period flows over the 30 defences in the highlighted flood area for each of the events in Figure 5. The events in the figure are numbered in rows beginning in the top-left.

Event	Return period (years)		
Number	Minimum	Mean	Maximum
1	<1	<1	1
2	<1	<1	<1
3	<1	<1	<1
4	2	5	8
5	1	5	8
6	3	5	9
7	6	8	11
8	1	3	14
9	65	145	400
10	8	22	57
11	<1	<1	1
12	3	25	65