# Efficient joint probability analysis of flood risk

Feifei Zheng, Michael Leonard and Seth Westra

**Feifei Zheng** (corresponding author)
**Michael Leonard**
**Seth Westra**
School of Civil,
Environmental and Mining Engineering,
University of Adelaide,
Adelaide 5005,
SA, Australia
E-mail: feifei.zheng@adelaide.edu.au

## ABSTRACT

Flood attributes such as the water level may depend on multiple forcing variables that arise from common meteorological conditions. To correctly estimate flood risk in these situations, it is necessary to account for the joint probability distribution of all the relevant forcing variables. An example of a joint probability approach is the design variable method, which focuses on the extremes of the forcing variables, and approximates the hydraulic response to forcing variables with a water level table. In practice, however, application of the design variable method is limited, even for the bivariate case, partly because of the high computational cost of the hydrologic/hydraulic simulations. We develop methods to minimise the computational cost and assess the appropriate extent and resolution of the water level table in a bivariate context. Flood risk is then evaluated as a bivariate integral, which we implement as an equivalent line integral. The line integral is two orders of magnitude quicker and therefore beneficial to settings that require multiple evaluations of the flood risk (e.g., optimisation studies or uncertainty analyses). The proposed method is illustrated using a coastal case study in which floods are caused by extreme rainfall and storm tide. An open-source R package has been developed to facilitate the uptake of joint probability methods among researchers and practitioners.

**Key words** | design variable method, extremal dependence, flood risk, joint probability analysis, R package

## INTRODUCTION

Numerous studies have demonstrated that extreme water levels may be attributed to multiple forcing variables. Examples include high ocean levels caused by concurrent waves and storm tides (Hawkes *et al.* 2002), coastal floods involving both elevated storm tides and high river flows (Zheng *et al.* 2013; Kew *et al.* 2013), and fluvial floods resulting from extreme rainfall occurring on an already wet catchment (Pathiraja *et al.* 2012). Statistical dependence between the forcing variables is common due to the significant spatial and temporal dependences that exist in the climate system (Leonard *et al.* 2014). Methods that incorporate these dependent forcing variables into flood risk estimation are often referred to as joint probability analyses (Hawkes 2008).

There are three general methods that encapsulate the joint probability of forcing variables in a statistically rigorous manner: *the response variable method*, *the continuous simulation method* and *the design variable method* (Hawkes 2008; Zheng *et al.* 2014). Each method differs in its assumptions, data requirements and computational complexity. For example, the response variable method directly estimates flood risk as a univariate frequency analysis of the response variable observations at the location of interest (e.g., flood level or flood volume). In this method, the impact of the dependence between the forcing variables on flood risk has been implicitly accounted for by the response variable. Despite being relatively simple, application of the response variable method is limited in practice because recorded data are rarely available at the location of interest.

The continuous simulation method obtains the distribution of the response variable (whether extreme or not) by means of simulation, followed by a univariate frequency

analysis (Cai *et al.* 2008). Similar to the response variable method, this approach does not require complex statistical modelling of extremes, but it does require long and continuous sequences of the forcing variables (such as rainfall and storm tide) that are often not available. It can also be computationally demanding to predict the response variable with hydrologic/hydraulic models. Given that computational budgets are typically limited in practice, the continuous simulation method can be costly to implement.

The design variable method significantly differs from the response variable and continuous simulation methods in that it requires explicit modelling of the extremal dependence of the forcing variables (and associated joint probability density), followed by an integral of the joint density to estimate the flooding risk (Coles & Tawn 1994). Since this method focuses on the extremes of the response variable, it is appreciably less computationally demanding compared to the continuous simulation method that requires simulation of both extreme and non-extreme situations. Widespread implementation of the design variable method is still limited, however, due to difficulty in quantifying the extremal dependence between forcing variables, and computational challenges associated with implementation. This study focuses on the latter.

To address the computational challenges of joint probability analyses, a number of probabilistic computation methods have been developed. For example, the first-order reliability method and the second-order reliability method were developed to approximate the failure region for efficient probability calculation (Zhao & Ono 1999). Monte Carlo importance sampling, directional sampling and subset simulation methods generated samples in the region of interest rather than the entire space (Schueremans *et al.* 2011), and adaptive response surface techniques were proposed to minimise the number of model runs (Nguyen *et al.* 2009). In contrast to these approaches, the design variable method uses a response surface to approximate the behaviour of the hydraulic model at a fixed set of boundary conditions and estimate the failure probability using a numeric integration (Coles & Tawn 1994). To improve its implementation efficiency in a bivariate context, this study proposes a multi-step approach as well as a method for transforming the bivariate integral to an equivalent line integral.

Improved efficiency is especially important whenever aspects of a design require repeated evaluation. Examples include: systems having many locations, return periods or durations of interest; optimisation studies (Nicklow & Hellman 2004); and uncertainty analyses (Fu & Kapelan 2013). Two key issues for enhancing the efficiency of the design variable method are: (i) the construction of a water level table (a table of water levels from different combinations of two forcing variables) that uses the minimum number of hydraulic model evaluations required for the design; and (ii) the estimation of exceedance probabilities for a given water level. We address the first issue by using a multi-step approach combined with techniques that assess the adequacy of the table's extent and resolution, and the second issue by converting the exceedance probability from a bivariate integral to an equivalent univariate line integral. Flood risk is normally defined as the probability of the hazard times the consequences. This study considers the probability of the hazard only, measured in terms of water levels.

## RATIONALE FOR JOINT PROBABILITY ANALYSES

Univariate estimation methods have been used previously for many flood risk estimation problems, whereby the frequency of a single forcing variable (e.g., extreme rainfall or storm tide) is assumed to be equivalent to the frequency of the corresponding water level. It is noted that the water level is used to represent the flood risk in this study, although the total flood volume or flood velocity can also be used to measure the flood risk (Svensson & Jones 2004). The exceedance probability $\Pr(H \geq h)$ can be defined as

$$\Pr(H \geq h) = \int_{x=x_0}^{\infty} f(x)dx, \quad h = B(x_0) \tag{1}$$

where $B(x)$ is a function relating the water level to a single forcing variable $X$ (e.g., rainfall or storm tide); $x_0$ is the value of the forcing variable that causes the water level $h$; and $f(x)$ is a density function of $X$ at extreme levels. To obtain $\Pr(H \geq h)$, one needs to first estimate the corresponding $x_0$ that causes the water level $h$, then estimate the exceedance probability of $x_0$ $\Pr(X \geq x_0)$ and assign

this to $\Pr(H \geq h)$, i.e., $\Pr(H \geq h) = \Pr(X \geq x_0)$. Typically, an annual maximum, $r$-largest or a peak-over-threshold method is used to obtain the tail distribution of $X$ (Coles 2001).

The estimation procedure becomes complicated in a multivariate setting since the exceedance probability of any given forcing process is no longer equivalent to the exceedance probability of the water level, i.e., $\Pr(X \geq x_0) \neq \Pr(H \geq h)$. Considering the design variable $h$ (also often referred to as a limit state) forced by two variables $X$ and $Y$, Coles & Tawn (1994) defined a 'failure region' $A_h$ as

$$A_h = \{(x, y) \in \mathfrak{R}^2 : B(x, y) > h\} \tag{2}$$

where $B(x, y)$ is a 'boundary function' that maps the two-dimensional space $(x, y)$ of the forcing variables to the one-dimensional response variable $h$. In flood estimation, hydraulic models are typically used to obtain a water level $h = B(x, y)$ in terms of boundary conditions such as rainfall and storm tide $(x, y)$. The failure region $A_h$ can be interpreted as the set of values of the forcing variables $(x, y)$ that cause water levels greater than a specified design water level $h$. The corresponding exceedance probability $\Pr(H \geq h)$ is given as

$$\Pr(H \geq h) = \iint_{A_h} f(x, y) \tag{3}$$

where $f(x, y) = \partial F(x, y) / \partial x \partial y$ is the joint density function of the two variables $X$ and $Y$, and $F(x, y)$ is their corresponding joint cumulative distribution function.

Figure 1 illustrates the difference between the univariate method (top panel) and the design variable method as an example of a joint probability method (bottom panel) for a hypothetical scenario in which floods are caused by two forcing variables $X$ and $Y$. In the top panel, the grey shaded region represents the exceedance probability $\Pr(H \geq h)$ where $h$ (the dashed line) is determined by a single forcing variable (e.g., rainfall or storm tides). The grey shaded region in the bottom panel illustrates the exceedance probability for the region $A_h$ where $h$ (the solid line) depends on both forcing variables. The probability $\Pr(H \geq h)$ is evaluated as the integral of the joint density $f(x, y)$ (contours, lower panel) across the whole failure region $A_h$ (Equation

(3)) and thus allows for the statistical dependence between the two forcing variables.

To enable comparison, the failure region determined by the univariate method (the grey shaded region in the top panel) is also shown in the bivariate setting (the bottom panel) by the grey shaded region to the right side of the dashed vertical line. It is clearly seen that the univariate failure region is smaller than the $A_h$ indicated by the joint probability method, suggesting that the univariate method will underestimate the flood risk in this case.

## A MULTI-STEP METHOD FOR EFFICIENT EVALUATION OF JOINT PROBABILITY

The proposed method comprises five steps: (1) perform a preliminary analysis, (2) specify the bivariate probability distribution of the forcing variables, (3) determine the failure region $A_h$ for the given water level $h$, (4) integrate the joint density in the failure region $A_h$ to obtain the flood risk estimate $\Pr(H \geq h)$ and (5) check the bounds on the estimates due to the water level table approximation. Where the bounds are too large for the given design context, further iteration of Steps 3–5 is required.

### Step 1: Perform a preliminary analysis

Joint probability analyses require additional computational effort when compared to conventional univariate methods for flood risk estimation and this is especially the case whenever gridded (two-dimensional) hydraulic models are used. Therefore, a preliminary analysis is proposed, which aims to determine the potential improvement in water level estimates to be gained by undertaking a joint probability analysis.

The default (Case 0) for the preliminary analysis is the conservative assumption of complete dependence between the respective distributions of the boundary values (e.g., rainfall and tide). Thus a water level with $Y$ year average recurrence interval (ARI) is determined from a $Y$ year ARI rainfall event (from which streamflow is derived) coinciding with a $Y$ year ARI storm tide event at the river's lower reaches. To evaluate the complete independence case, two alternative situations are modelled: Case 1, in which a $Y$
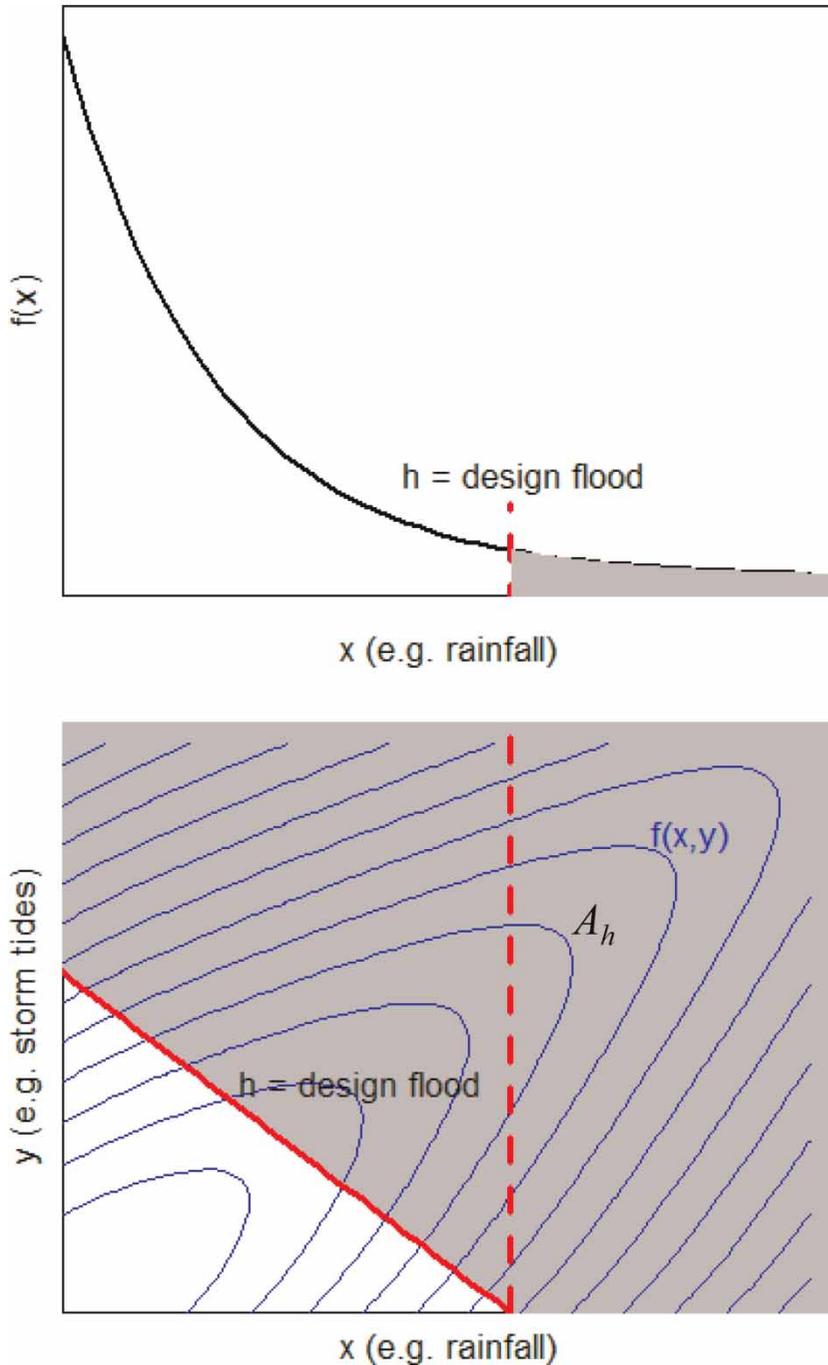
**Figure 1** | The grey shaded region (top panel) illustrates $\Pr(H \geq h)$ obtained from a univariate analysis, which corresponds to the shaded region to the right of the red dashed line in the bottom panel. The grey shaded region (bottom panel) shows the failure region $A_h$ and $\Pr(H \geq h)$ in the bivariate setting by integrating the $f(x, y)$ (blue contours). Please refer to the online version of this paper to see this figure in colour: http://www.iwaponline.com/jh/toc.htm.

year ARI storm tide event occurs in isolation (i.e., with no rainfall); and Case 2, in which a $Y$ year ARI rainfall event occurs in isolation (i.e., with the mean sea level (MSL)). This results in a total of three runs of the hydraulic model

before proceeding with more analyses for a particular ARI of interest for the water level.

The three runs are compared to determine whether the differences between complete dependence and complete

independence of the forcing variables are large enough for the given design context to warrant the additional modelling effort. The tolerance is defined by the user and it is a trade-off between the benefit of having a more accurate assessment versus the cost of implementing the joint probability analysis. In other words, the additional computational cost of the joint probability analysis should be at least proportional to the benefit of the additional precision. This trade-off may vary according to different locations and design problems. If the differences are small, the conservative assumption of complete dependence might serve as a suitable and expedient estimate for the location of interest and hence Steps 2–5 given below can be skipped, otherwise the entire multi-step method (Steps 2–5) needs to be performed for estimating flood risk. The same approach can be repeated for multiple ARIs.

The preliminary analysis can be especially beneficial when multiple locations are considered as a single hydrodynamic simulation can often provide water levels for all locations. The impacts of the joint probability of the forcing variables to flood risk (e.g., rainfall dominating, surge dominating or transition zones) can be simultaneously obtained for all these locations without additional computational overheads compared to considering a single location.

## Step 2: Specify the joint cumulative distribution function $F(x, y)$

To proceed with a joint probability analysis it is necessary to specify the cumulative distribution function $F(x, y)$ between the two forcing variables. There are many statistical models available for modelling bivariate extremes (Kotz & Nadarajah 2000), and we use the bivariate logistic model as it only requires one parameter to model the extremal dependence and shows similar performance to more complex alternatives (Zheng *et al.* 2013). The $F(x,y)$ of the logistic model is given as (Tawn 1988)

$$F(x, y) = \exp\{-(\tilde{x}^{-1/\alpha} + \tilde{y}^{-1/\alpha})^{\alpha}\}, \tilde{x} > 0, \tilde{y} > 0, 0 < \alpha \leq 1 \tag{4}$$

where $\tilde{x}$ and $\tilde{y}$ are standard Fréchet-transformed values of original observations $x$ and $y$. The parameter $\alpha$ describes the dependence strength between the two forcing variables $X$ and $Y$, such as extreme rainfall and storm tides, with

$\alpha = 0$ and $\alpha = 1$, respectively, representing complete dependence and independence.

Two methods are available for estimating the dependence parameter $\alpha$: the bivariate threshold-excess method (Resnick 1987) and the point process method (de Haan 1984). The former method is selected in this study due to its approximately unbiased estimation, as demonstrated by Zheng *et al.* (2014). In the bivariate threshold-excess method, the marginal extremes above a high threshold $u$ are modelled using the generalised Pareto distribution (GPD) (Thibaud *et al.* 2013)

$$G(\bar{z}) = 1 - (1 + \frac{\xi\bar{z}}{\sigma})^{-1/\xi}, \quad \bar{z} = z - u \tag{5}$$

defined on $\{\bar{z}: \bar{z} > 0$ and $(1 + \xi\bar{z}/\sigma) > 0\}$, where $z$ represents one of the original margins (either $x$ or $y$); $\sigma > 0$ and $-\infty < \xi < \infty$ are scale and shape parameters, respectively. If $\xi = 0$, $G(\bar{z}) = 1 - \exp(-\frac{\bar{z}}{\sigma})$. The maximum likelihood method is employed to estimate GPD parameters in this study as it is a widely applied and reliable method (Coles 2001). Alternative approaches include the method of moments and the method of probability-weighted moments (Kjeldsen & Jones 2004).

An appropriate threshold $u$ needs to be determined for each margin and is critical for the use of the threshold-excess method. This is because a value of $u$ that is too low may produce biased estimates due to an invalid asymptotic justification, and a value that is too high will result in a high variance due to the limited sample size (see Coles (2001) for details). The marginal thresholds were selected based on diagnostic plots such as mean residual life plots and scale and shape parameter plots (Coles 2001; Zheng *et al.* 2013).

As shown in Equation (4), the use of the logistic model requires the margins to follow a standard Fréchet distribution, which can be obtained from the original observations using

$$\tilde{z} = \begin{cases} -\left(\log\left\{1 - \hat{\varsigma}_{u_z}\left[1 + \frac{\hat{\xi}_z(z - u_z)}{\hat{\sigma}_z}\right]^{-1/\hat{\xi}_z}\right\}\right)^{-1} & z > u_z, \hat{\xi}_z \neq 0 \\ -\left(\log\left\{1 - \hat{\varsigma}_{u_z}\exp\left(-\frac{z - u_z}{\hat{\sigma}_z}\right)\right\}\right)^{-1} & z > u_z, \hat{\xi}_z = 0 \end{cases} \tag{6}$$

where $\tilde{z}$ is the standard Fréchet value corresponding to the $z$ in the original scale; $u_z$ is an appropriately high threshold for the margin $z$ ($z = x$ or $y$); $\hat{\varsigma}_{u_z} = \Pr\{Z > u_z\}$, which can be empirically determined by the proportion of the data points above the threshold $u_z$, and $\hat{\sigma}_z$ and $\hat{\xi}_z$ are estimated parameters of the GPD in Equation (5).

After transforming both marginal extremes $(X, Y) \rightarrow (\tilde{X}, \tilde{Y})$, the censored likelihood method described in Coles & Tawn (1991) is typically used to estimate $\hat{\alpha}$, with the log-likelihood function $\ell$ given as

$$\ell(\alpha; (x_1, y_1), ..., (x_n, y_n)) = \sum_{i=1}^{n} \log\{\phi(\alpha; (x_i, y_i))\} \qquad (7)$$

where $(x_i, y_i)$ is the $i$th observed joint event and $n$ is the total number of the observed joint events. Equation (7) implies that all joint events are included in the likelihood to estimate the $\hat{\alpha}$, although non-extreme pairs have censored contributions. The $\phi(\cdot)$ is defined as

$$\phi(\alpha; (x, y)) = \begin{cases} \partial^2_{xy} F(\alpha; (x, y)) & x > u_x, y > u_y \\ \partial_x F(\alpha; (x, u_y)) & x > u_x, y \le u_y \\ \partial_y F(\alpha; (u_x, y)) & x \le u_x, y > u_y \\ F(\alpha; (u_x, u_y)) & x \le u_x, y \le u_y \end{cases} \qquad (8)$$

where $\partial_x$, $\partial_y$ and $\partial^2_{xy}$ represent derivatives with respect to $x$ and $y$. $F(\theta; (x, y))$ is the logistic distribution function given in Equation (4). As shown in Equation (8), for the joint exceedances $(x > u_x, y > u_y)$, $\partial^2_{xy} F(\theta; (x, y))$ provides the appropriate likelihood component, while the observations that lie below the threshold provide only a censored contribution to the likelihood (Coles & Tawn 1991). The dependence parameter $\alpha$ is estimated by maximising the log-likelihood function $\ell$ in Equation (7).

## Step 3: Approximate the hydraulic response

Many different combinations of $x$ and $y$ can result in an identical water level $h$. In the design variable method, the failure region $A_h$ for a given water level $h$ is interpolated from a water level table. This table is obtained from hydraulic modelling and represents an approximation of the hydraulic response since it is only evaluated for certain pairs of boundary values across a finite range of ARIs (focused on extreme events). An example of the water level table is given later for the case study.

## Step 4: Estimate the exceedance probability of water levels

Having determined the joint distribution $F(x, y)$ and the failure region $A_h$, it is possible to estimate $\Pr(H \ge h)$. Despite the method usually being expressed in terms of the joint density function $f(x, y)$ (Equation (3)), it is more computationally efficient to use the joint cumulative distribution function $F(x, y)$. This is because $F(x,y)$ is already an integral of $f(x,y)$ and this can be exploited to reduce the bivariate integral to a one-dimensional integration problem. In other words, the bivariate integration of the failure region in Equation (3) can be converted to a univariate line integral along the boundary function $B(x,y) = h$. This is implemented numerically as

$$\Pr(H \ge h) = 1 - \sum_{j=1}^{m-1} \{F(x_j, y_j) - F(x_j, y_{j+1})\} \qquad (9)$$

defined as $\{(x_j,y_j): B(x_j,y_j) = h, y_{j+1} = y_j - \Delta y\}$, where $m$ is the number of points $(x_j,y_j)$ along the boundary line $h$ (the red solid line in Figure 1) and $\Delta y \ge 0$. By taking the finite difference of $F(x,y)$ in only one dimension, the probability of being less than $x_i$ is obtained for an increment of width $\Delta y$. Moving along the boundary function line (the solid line in the bottom panel of Figure 1), the probability increments are obtained for all corresponding pairs $(x_j,y_j)$, and the non-exceedance probability $1\text{-}\Pr(H \ge h)$ is the sum over all increments.

## Step 5: Check the impact of hydraulic approximations on the water level estimate

The values in the water level table are obtained using hydrologic/hydraulic models and obtaining larger tables requires greater computational and modelling effort. To minimise the modelling effort while maximising precision it is important to determine the influence of the selected boundaries and resolution of the table. The methods used here are diagnostic tools to evaluate the adequacy of any given water

level table (no matter how coarse it may be), allowing it to be refined subsequently with additional runs if necessary.

One issue relates to the specification of the upper boundary of the water level table when it does not capture the entirety of the boundary function $B(x, y)$. To address this issue, we compute the return probabilities of two alternate cases. A water level of $h = 2$ m is taken from the case study results to illustrate the two cases as shown in Figure 2 (solid black lines). The first case (left panel) projects the $h = 2$ m line horizontally back to the $y$-axis at the storm tide intersection ARI = 500 years (dotted black line). This caps the largest storm tide at ARI = 500 years, which will result in an overestimate of the return probability (the failure region $A_h$ is shown by the grey shaded region in the left panel of Figure 2). The second case extends the black solid line of $h = 2$ m vertically to infinity (black dotted line in the right panel), i.e., assumes the largest water level produced by the storm tides is $h = 2$ m, even if the storm tide ARI $\rightarrow \infty$. For this case, the failure region $A_h$ is given by the shaded region to the right side of the black solid and dotted lines, and the return probability is underestimated.

Another issue relates to the specification of the lower boundary of the water level table. The lower boundary of the water level table should give the case of 'rainfall only' flooding when the tide is not extreme, and 'tide only' flooding when there is no rain. The sensitivity of the water level estimates to these cases can be tested by removing them and using the lowest ARI tide/rain values in their place.

A third issue relates to whether the resolution of the water level table leads to sufficient precision in the boundary function interpolation for a given water level. Variability in the specification of the water level table leads to potential variation in the interpolated boundary function as well as variation in the associated estimated exceedance probability. The limited resolution of the water level table represents a trade-off between the precision in the water level estimates and the computational cost of running hydrologic/hydraulic models. Quantifying the sensitivity of the results to the resolution of the water level table will allow for judicious selection of additional model runs used to refine the table. We implement a simple method based on adjacent values from the water level table as upper and lower bounds on the interpolated boundary function.

Based on the results in Step 5, one can determine whether or not the water level table needs to be further increased in extent and resolution. For example, if the upper boundary of the water level table causes significant variations in estimates, further large ARIs of the forcing variables are required to construct the water level table; if the lower boundary (no-rainfall or MSL) results in a large difference in water level estimates, this indicates that the inclusion of the lower boundary is critical, offering guidance for building water level tables. Similarly, a higher-resolution water level table is needed when the estimation differences between upper and lower bounds of the interpolated boundary function are substantial.
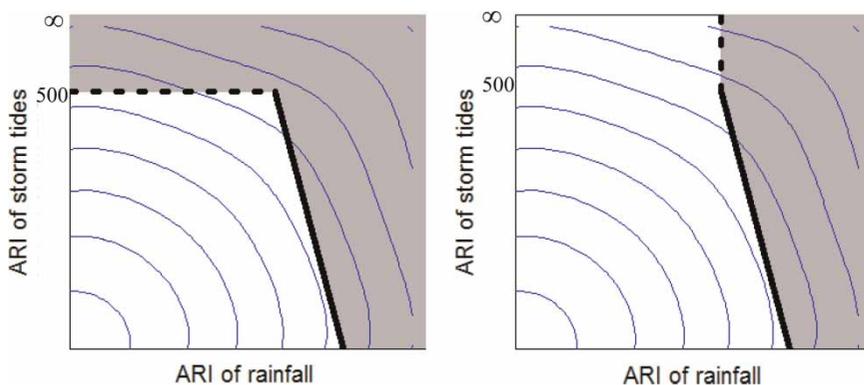


**Figure 2** │ Illustration of the method to address the upper boundary issue of the water level table using $h = 2$ m (the black solid line) from the case study (see Figure 3). The black dotted lines in the left and right panels represent the extensions of the design water level of $h = 2$ m to enable its return probability estimation. The grey shaded region is the failure region $A_h$. Blue contours are joint density $f(x, y)$. Please refer to the online version of this paper to see this figure in colour: http://www.iwaponline.com/jh/toc.htm.

## CASE STUDY RESULTS AND DISCUSSION

The proposed method is demonstrated for a drainage network of a small coastal catchment located in Perth, Western Australia. This drainage network discharges into the Fremantle region near the outlet of the Swan River to the Indian Ocean. Water levels in the drainage system at the catchment outlet are influenced by the concurrent rainfall on the catchment and tide levels. Flows in the urban drainage network are modelled using the Infoworks CS package (Innovyze Inc. 2011).

The design event method is employed to determine flows according to Australian design guidelines (ARR 1987), using rainfall parameters and initial-loss/continuing-loss parameters for the Perth region as specified by the guidelines. The design method accounts for antecedent moisture conditions by means of the loss parameters (Rahman et al. 2002). The method was repeated for several storm-burst durations and a 1-hour duration was ultimately selected as the critical duration as this duration produced the maximum water levels for all ARIs. Only the results for the 1-hour duration are presented in this paper. The dependence analysis was conducted for the cumulative 1-hour rainfall paired with the maximum concurrent storm tide (Zheng et al. 2013). It is noted that the 'water level' terminology was used for the case study, and floods only occur when water overflows from the drainage network.

### Preliminary analysis

Table 1 shows the results of the preliminary analysis (Step 1 of the proposed method), considering ARI = 2, 20 and 100 years of the marginal variables (rainfall and storm

**Table 1** | Water levels of different combinations of rainfall (1 hour duration) and storm tide in terms of ARI (years). Only the highlighted cells have to be determined

| | | Rainfall events (ARIs) | | | |
| --- | --- | --- | --- | --- | --- |
| | | No-rainfall | 2 | 20 | 100 |
| Storm tide events (ARIs) | Mean sea level | | 1.16 | 1.76 | 2.24 |
| | 2 | 1.05 | 1.67 | | |
| | 20 | 1.25 | | 2.29 | |
| | 100 | 1.35 | | | 2.83 |

tides). The 'no rainfall' and 'MSL' cases are represented in grey shades along each margin of the water level table while the diagonal terms are the cases of complete dependence. It can be seen that the diagonal values are considerably higher than either margin for all ARIs considered. This indicates that the interaction of the two variables has a significant influence on the water level at this study location and that a joint probability analysis is necessary.

### Estimation of dependence

Having determined the necessity for a joint probability analysis, it is necessary to estimate the dependence between the two forcing variables (rainfall and storm tides). Using the method outlined in Step 2, Equations (4)–(8) are used to estimate the dependence parameter $\hat{\alpha}$ between the extreme rainfall and storm tides. A comprehensive analysis of dependence has been documented in Zheng et al. (2013) and hence we do not present the details in the current study. We estimated the dependence parameter to be $\hat{\alpha} = 0.95$ using a large number of rainfall-storm tide data sets from the case study catchment (Zheng et al. 2013). The fitted joint density function $f(x, y)$ for $\hat{\alpha} = 0.95$ is shown by the blue contours in Figure 3.

### Approximation of hydraulic response

In Step 3 of the method, a water level table has to be constructed to specify the boundary of the failure region $A_h$ as necessary for the density integral above a given water level $h$. Table 2 shows a relatively higher-resolution water level table for a range of ARIs (10 tide ARIs × 10 rainfall ARIs gives 100 model runs). The temporal patterns of the extreme rainfall (1 hour duration) events in Table 2 with different ARIs were taken from Australian Rainfall and Runoff guidance (ARR 1987).

The interpolated contours from Table 2 are shown in Figure 3 (solid red lines). The overlay of the probability density (blue contours) with the boundary function is a useful visual aid for conceptualising the mechanics of Equation (3): the exceedance probability $\Pr(H \geq h)$ is the summation of the probability density function $f(x, y)$ (blue contours) for all combinations of $x$ and $y$ that exceed the specified water

**Table 2** │ A water level table with water levels of different combinations of rainfall and storm tide in terms of ARI (years)

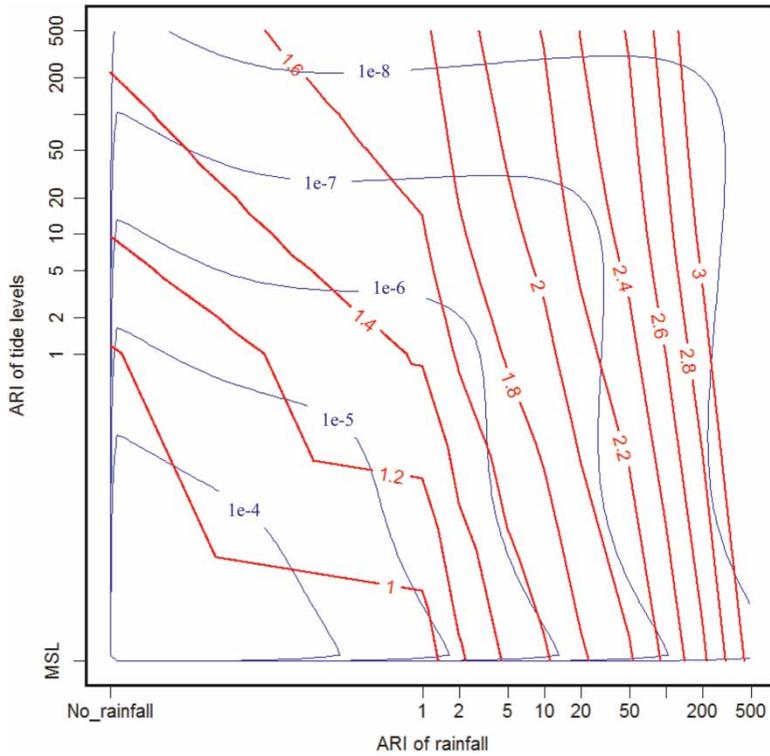| | | Rainfall (ARI) | | | | | | | | | |
| --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- |
| | | No rain | 1 | 2 | 5 | 10 | 20 | 50 | 100 | 200 | 500 |
| Storm tide (ARI) | MSL | 0.79 | 0.88 | 1.16 | 1.43 | 1.57 | 1.76 | 1.97 | 2.24 | 2.55 | 3.07 |
| | 1 | 0.98 | 1.42 | 1.63 | 1.82 | 1.94 | 2.13 | 2.34 | 2.60 | 2.92 | 3.44 |
| | 2 | 1.05 | 1.48 | 1.67 | 1.86 | 1.98 | 2.17 | 2.38 | 2.64 | 2.96 | 3.48 |
| | 5 | 1.14 | 1.54 | 1.73 | 1.92 | 2.03 | 2.22 | 2.43 | 2.70 | 3.02 | 3.54 |
| | 10 | 1.20 | 1.58 | 1.77 | 1.95 | 2.07 | 2.26 | 2.47 | 2.74 | 3.05 | 3.58 |
| | 20 | 1.25 | 1.62 | 1.81 | 1.98 | 2.10 | 2.29 | 2.50 | 2.77 | 3.09 | 3.61 |
| | 50 | 1.31 | 1.66 | 1.85 | 2.02 | 2.13 | 2.33 | 2.54 | 2.81 | 3.12 | 3.65 |
| | 100 | 1.35 | 1.69 | 1.87 | 2.04 | 2.16 | 2.35 | 2.56 | 2.83 | 3.15 | 3.67 |
| | 200 | 1.39 | 1.72 | 1.90 | 2.07 | 2.18 | 2.38 | 2.59 | 2.86 | 3.18 | 3.70 |
| | 500 | 1.44 | 1.76 | 1.93 | 2.10 | 2.21 | 2.41 | 2.62 | 2.89 | 3.21 | 3.73 |



**Figure 3** │ The density of the logistic model $f(x, y)$ with $\hat{\alpha} = 0.95$ (blue contours), and the water levels $h$ for various combinations of rainfall and storm tide events (red lines). The rainfall and storm tide events are given in terms of ARIs. Please refer to the online version of this paper to see this figure in colour: http://www.iwaponline.com/jh/toc.htm.

level $h$. In the format of Figure 3 it is straightforward to visualise that the exceedance probability depends on the shape of the density function via the estimated parameter $\hat{\alpha}$, so that the more (less) strongly dependent the two variables are, the more (less) density occurs along the increasing diagonal, and accordingly, the exceedance probability is greater (lower).

It is also straightforward to see that for some contours the exceedance probability is influenced mostly by one forcing variable (such as the 3 m near-vertical contour) whereas for other contours it is influenced by combinations of both variables (such as the 1 m contour). This suggests that very high water levels of this case study are dominated by the rainfall while more frequent flood events are

increasingly influenced by the joint dependence between extreme rainfall and storm tides.

## Water level estimation

The water level estimates for the case study are presented in Figure 4 (solid black line). These results required 1.32 seconds to evaluate 100 water levels ranging from 1 to 3 m using R3.0.1 on a 3.4 GHz Pentium PC. We also implemented the direct integral method specified in Equation (3) through nested use of the standard 1-D *integrate* function from the R 'stats' package. To obtain the results (to similar accuracy) given in Figure 4, the bivariate integral method of Equation (3) required 131 seconds for the same range and number of water levels, which is approximately two orders of magnitude slower than the proposed method specified in Equation (9).

Although the difference between the proposed integration and the direct integral methods is on the order of seconds (contrasted with hydraulic models running on the order of days), the cumulative efficiency gain is significant given that many water levels, locations and duration scales may be considered. Furthermore, many applications involve



**Figure 4** | Estimated water levels for $\hat{\alpha} = 0.95$ (black solid line) against different annual recurrence intervals (ARIs). The black dashed line indicates the lower boundary of estimates. $\hat{\alpha} = 0$ (red dotted line) and $\hat{\alpha} = 1$ (blue dot-dashed line) represent complete dependence and independence, respectively. Please refer to the online version of this paper to see this figure in colour: http://www.iwaponline.com/jh/toc.htm.

numerous iterations of the estimation procedure, such as when optimising the design of flood infrastructure and when performing uncertainty analyses. For example, 1,000 samples of marginal and dependence parameters were generated from a posterior distribution and were used to conduct a parametric uncertainty at the location corresponding to Table 1. The analysis, at this one location, required 20 minutes using Equation (9) as compared to 1.5 days using Equation (3).

For comparison with the estimated dependence case ($\hat{\alpha} = 0.95$), Figure 4 shows the cases where the rainfall and storm tide are completely dependent ($\hat{\alpha} = 0$, red dotted line) and completely independent ($\hat{\alpha} = 1$, blue dot dash line). There are substantial differences between the estimated ARIs of the independence and complete dependence cases for all given water levels. For example, for the water level of 2.0 m, the estimated ARI with $\hat{\alpha} = 0.95$ is approximately 18 years, the ARI with $\hat{\alpha} = 1$ is 40 years (assuming complete independence), and the ARI with $\hat{\alpha} = 0$ is approximately 8 years (assuming complete dependence).

Since no recorded water level data are available for this case study catchment, the verification on the estimated water levels from the design variable method cannot be made. However, the accuracy of the design variable method has been demonstrated by a number of previous studies, such as Coles & Tawn (1994) and Bruun & Tawn (1998). The focus of this study is on the efficiency improvement of the bivariate integration as well as the proposal to effectively construct water level tables for flood risk estimation though the design variable method.

## Upper boundary of the water level table

Using the approach outlined in Figure 2, the two approaches for extrapolating the water level table above the highest estimated values were adopted, and which by design will either over- or under-estimate the flood risk. These are shown in Figure 4 (the solid black and dashed lines). For the given ARIs, the differences in the water levels from the two cases are relatively minor, suggesting that the provided upper boundaries in Table 2 with ARI = 500 for each margin are sufficient for estimating water levels with ARI up to 100 years. If this difference were greater, it would be necessary to expand the water level table using larger ARIs of the rainfall and
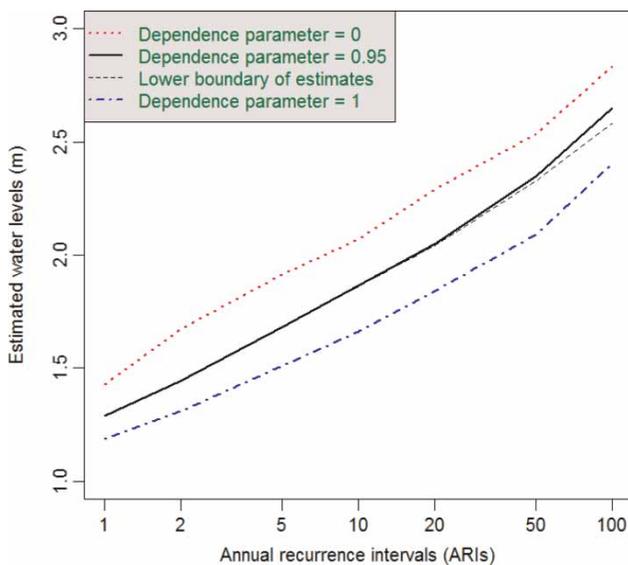
storm tides (e.g., ARI = 1,000 years). The proposed specification method can be used as a diagnostic to verify whether the given upper boundary is sufficient or not for estimating the specified water level range.

## Lower boundary of the water level table

To consider the influence of the lower boundary, a basic analysis of the sensitivity to variations of Table 2 was conducted. The results are summarised in Figure 5 where the solid black line is the same as that in Figure 4. The green dashed and red dotted lines represent the cases where the lower boundaries have been removed (i.e., the removal of the no-rainfall column and the 'MSL' row in Table 2, respectively). The results show that the estimates of the case study are sensitive to the lower boundary: the green dashed line shows that estimates of low water levels are sensitive to excluding the no-rainfall case, but that all estimates are sensitive to the removal of the MSL case as shown by the red dotted line. This indicates that it is important to include the low boundaries of no-rainfall and mean sea level in the water level table.

The significant impact of the lower boundary is expected, because in practice, rainfall or surge events will frequently occur separately as a result of their weak
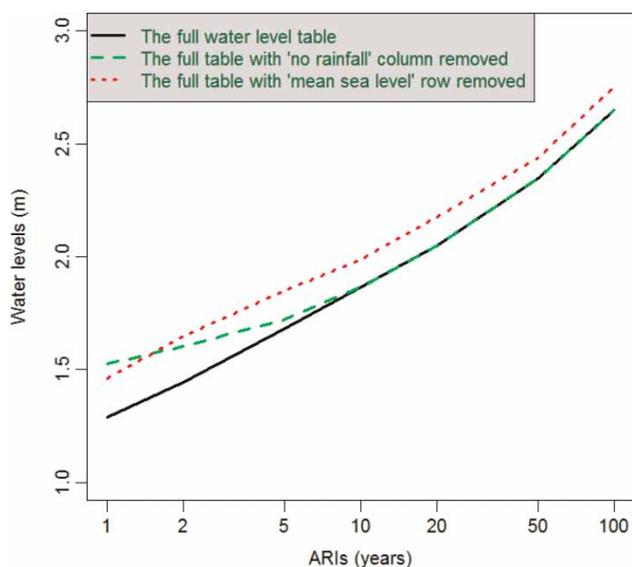


**Figure 5** | Estimated water level estimates using variously sized water level table. The black line uses Table 2 in entirety. The green dashed and red dotted lines remove the no-rainfall column and the mean sea level (MSL) row, respectively. Please refer to the online version of this paper to see this figure in colour: http://www.iwaponline.com/jh/toc.htm.

dependence (Zheng *et al.* 2013). Therefore the removal of the no-rainfall (or MSL) case indicates that a significant proportion of the extreme surge (or rainfall) events are ignored, resulting in substantial increases of water level estimates using Equation (9). As shown in Figure 3, the rainfall has large impact on water levels across all ARIs, while the storm surge only influences the water levels with low ARIs for the case study. This explains the significant effects of the MSL case to water level estimates across all ARIs (red dotted line in Figure 5) as well as the limited impact of the no-rainfall case on water levels with high ARIs (green dashed line).

## Resolution of the water level table

The method to assess the resolution of the water level table is illustrated in Figure 6 using two water levels $h = 1.52$ and $h = 2.54$ m (solid red lines) – the same as in Figure 3. All values in the water level table (Table 2) are shown in Figure 6 to facilitate the interpretation of the proposed method.

From Figure 6, the respective boundary functions for $h = 1.52$ and $h = 2.54$ m are shown as the linearly interpolated solid lines. The interpolation bounds are obtained from the bordering values of Table 2 for the two separate cases. The lower bound case uses all the values in the water level table that are immediately lower than the specified water level $h$, resulting in a boundary function that has a larger failure region and consequently a larger exceedance probability (red dotted lines in Figure 6). The upper bound case takes all the values in the water level table that are immediately larger than the water level $h$, producing a boundary function with a smaller failure region and hence a smaller exceedance probability (red dashed lines in Figure 6).

Figure 7 shows the lower and upper bounds for the estimated water levels with ARI ranging from 1 to 100 years using two different sized water level tables. The red dashed lines are based on the $10 \times 10$ water level table presented as Table 2 whereas the blue dotted lines are from a coarser $5 \times 5$ water level table (MSL/no-rainfall, 2, 10, 50 and 500 ARIs). The range between the bounds is a measure of the coarseness of the water level table in its approximation of the hydraulic response and the improvement in the bounds on the estimate that arises from using a $10 \times 10$ table over a $5 \times 5$ table can clearly be seen. Note that the $10 \times 10$

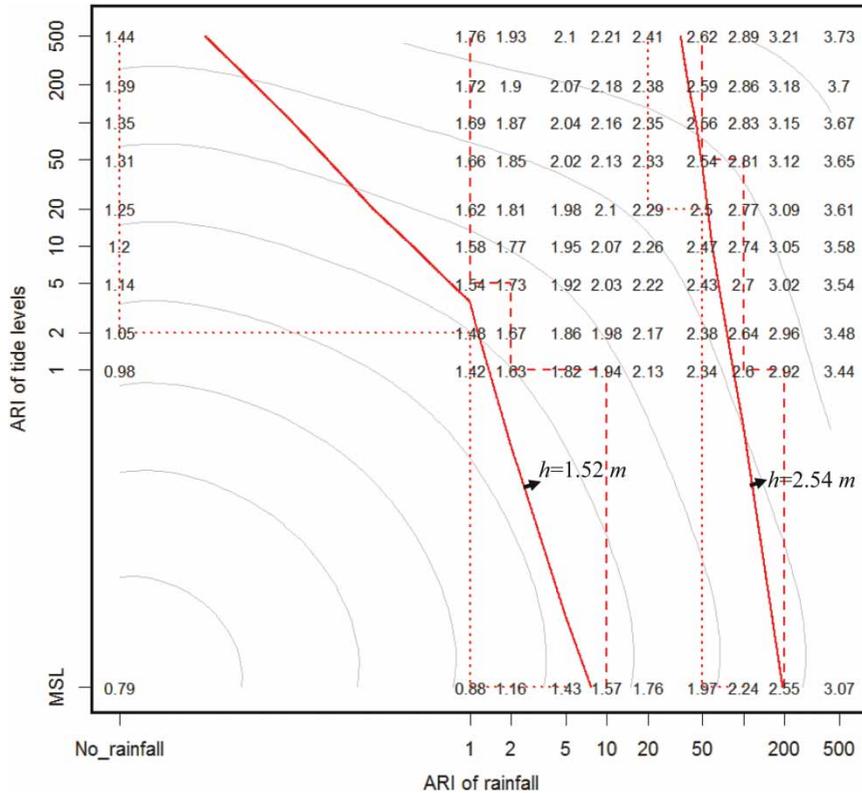**Figure 6** │ Illustration of the proposed method to quantify interpolation bounds for different water levels ($h = 1.52$ m, 2.54 m) based on the water level table (Table 2). The dotted and dashed red lines are respectively the lower and upper bounds of the interpolated water level (solid red lines). The grey contours are the joint density. Please refer to the online version of this paper to see this figure in colour: http://www.iwaponline.com/jh/toc.htm.
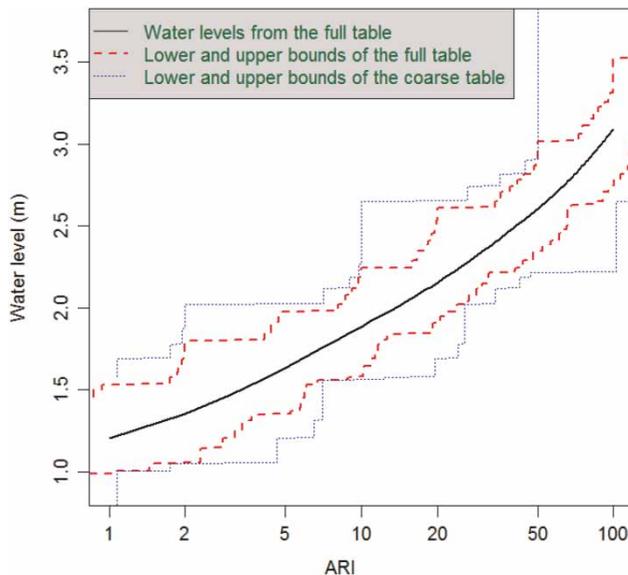


**Figure 7** │ Estimated water levels and the bounds for $\hat{\alpha} = 0.95$ (black line) against different ARIs. The red dashed lines are based on the full table (Table 2). The blue dotted lines use a coarse water level table (MSL/no-rainfall, 2, 10, 50 and 500 ARIs). Please refer to the online version of this paper to see this figure in colour: http://www.iwaponline.com/jh/toc.htm.

table was the initial implementation in this case study, so the $5 \times 5$ comparison was not strictly used as a diagnostic.

To reduce the bound range, it is possible to perform additional runs of the hydraulic model and refine the water level table. The amount of refinement would depend on whether the bounds are acceptable for a specific water level of interest. If they are deemed too large, then it may be necessary to increase the resolution in the rows or columns of the table immediately above/below the water level of interest. For example, if one is interested to reduce the bound range of $h = 2.54$ m as shown in Figure 7 for the $10 \times 10$ table, additional model runs should target the region between the dotted and dashed lines associated with $h = 2.54$ m as illustrated in Figure 6.

While the ranges of the respective bounds in Figure 7 are large at places, it is important to emphasise that the lines represent the utmost extreme cases of departure from the interpolated 'best' line. In other words, they should not be interpreted in a statistical manner (as with confidence

limits). Furthermore, unless there is a specific 'higher-resolution' feature in the hydraulic model that might induce a large shift in the water level, it is reasonable to expect subsequent refinements of the water level table to give interpolations that lie toward the centre of the bounded range than to the outer extent.

The method demonstrated in Figure 7 has a diagnostic ability since it allows the sensitivity of the range of water levels to be evaluated for any given degree of resolution of the water level table. Along with the diagnostic checks for the extent of the water level table (Figures 4 and 5) it is possible to assess the adequacy of the water level table as an approximation of the hydraulic response and any impact it may have on the results. This method allows for development of the water level table in a manner that minimises the number of evaluations of computationally expensive hydraulic models.

## CONCLUSIONS

The dependence of floods on multiple forcing variables complicates the calculation of flood exceedance probabilities, because a single water level can be caused by many different combinations of the forcing variables. The current study implements an efficient joint probability method (the design variable method) by using a water level table combined with a univariate line integral. For studies involving many repeated design scenarios (locations, nodes, parameters) the cumulative saving of this integration method can be substantial. Additionally, a range of techniques are implemented to further minimise the computational effort in the use of the design variable method. These include: (i) a preliminary analysis to determine the necessity of a joint probability analysis at the location of interest; and (ii) diagnostic tools to evaluate the adequacy of the water level table in terms of its extent and resolution.

The proposed method has been demonstrated using a coastal flooding case study affected by both rainfall and storm tides. Results for this case study show that although the dependence between extreme rainfall and storm tides appears to be weak, it has significant implications for flood risk. This suggests that it is critical to incorporate the dependence between extreme rainfall and storm tides to correctly estimate the coastal flood risk.

While the proposed method has been illustrated here using extreme rainfall and extreme storm tides, it can also be used with other hydrological variables, such as waves and tides in coastal regions, rainfall and downstream river heights in urban drainage systems, or flows from two tributaries in large-scale catchments. An open-source R package ('jointPm' http://cran.r-project.org) has been developed so that this method can be accessible to a wide audience. Important future focuses along this research line include: (i) extension of the proposed method to multivariate cases of more than two forcing variables; and (ii) incorporation of uncertainty in the selection of rainfall durations and temporal patterns.

## ACKNOWLEDGEMENTS

## REFERENCES

Australian Rainfall and Runoff 1987 A Guide to Flood Estimation. Institution of Engineers, Barton, ACT.

Bruun, J. T. & Tawn, J. A. 1998 Comparison of approaches for estimating the probability of coastal flooding. J. R. Stat. Soc. C Appl. 47 (3), 405–423.

Cai, Y., Gouldby, B., Hawkes, P. & Dunning, P. 2008 Statistical simulation of flood variables: incorporating short-term sequencing. J. Flood Risk Manag. 1 (1), 3–12.

Coles, S. G. 2001 An Introduction to Statistical Modelling of Extreme Values. Springer, London.

Coles, S. G. & Tawn, J. A. 1991 Modelling extreme multivariate events. J. R. Stat. Soc. B Meth. 53 (2), 377–392.

Coles, S. G. & Tawn, J. A. 1994 Statistical methods for multivariate extremes: An application to structural design. J. R. Stat. Soc. C Appl. 43 (1), 1–48.

de Haan, L. 1984 A spectral representation for max-stable processes. Ann. Probab. 12 (4), 1194–1204.

Fu, G. & Kapelan, Z. 2013 Flood analysis of urban drainage systems: probabilistic dependence structure of rainfall characteristics and fuzzy model parameters. J. Hydroinformat. 15 (3), 687–699.

Hawkes, P. J. 2008 Joint probability analysis for estimation of extremes. J. Hydraul. Res. 46 (suppl 2), 246–256.

Hawkes, P. J., Gouldby, B. P., Tawn, J. A. & Owen, M. W. 2002 The joint probability of waves and water levels in coastal engineering design. J. Hydraul. Res. 40 (3), 241–251.

Innovyze Inc 2011 Infoworks CS Technical Review, Innovyze Inc., Wallingford, UK, http://www.innovyze.com/products/infoworks_cs/infoworks_cs_technical_review.pdf.

Kew, S. F., Selten, F. M., Lenderink, G. & Hazeleger, W. 2013 The simultaneous occurrence of surge and discharge extremes for the Rhine delta. Nat. Hazard Earth Syst. Sci. 13 (8), 2017–2029.

Kjeldsen, T. R. & Jones, D. A. 2004 Sampling variance of flood quantiles from the generalised logistic distribution estimated using the method of L-moments. Hydrol. Earth Syst. Sci. 8 (2), 183–190.

Kotz, S. & Nadarajah, S. 2000 Extreme Value Distributions Theory and Applications, Imperial College Press, London.

Leonard, M., Westra, S., Phatak, A., Lambert, M., van den Hurk, B., McInnes, K., Risbey, J., Schuster, S., Jakob, D. & Stafford-Smith, M. 2014 A compound event framework for understanding extreme impacts. Clim. Change 5 (1), 113–128.

Nicklow, J. W. & Hellman, A. P. 2004 Optimal design of storm water inlets for highway drainage. J. Hydroinformat. 6 (4), 245–257.

Nguyen, X. S., Sellier, A., Duprat, F. & Pons, G. 2009 Adaptive response surface method based on a double weighted regression technique. Probab. Eng. Mech. 24 (2), 135–143.

Pathiraja, S., Westra, S. & Sharma, A. 2012 Why continuous simulation? The role of antecedent moisture in design flood estimation. Water Resour. Res. 48 (6), W06534.

Rahman, A., Weinmann, P. E., Hoang, T. M. T. & Laurenson, E. M. 2002 Monte Carlo simulation of flood frequency curves from rainfall. J. Hydrol. 256 (3–4), 192–210.

Resnick, S. I. 1987 Extreme Values, Regular Variation, and Point Processes. Springer, Berlin, Heidelberg, New York.

Schueremans, L., Smars, P. & Van Gemert, D. 2011 Safety of arches – a probabilistic approach, 9th Canadian Masonry Conference edition, June, New Brunswick, Canada. Available from http://smars.yuntech.edu.tw/papers/brunswick.pdf.

Svensson, C. & Jones, D. A. 2004 Dependence between sea surge, river flow and precipitation in south and west Britain. Hydrol. Earth Syst. Sci. 8 (5), 973–992.

Tawn, J. A. 1988 Bivariate extreme value theory: models and estimation. Biometrika 75 (3), 397–415.

Thibaud, E., Mutzner, R. & Davison, A. C. 2013 Threshold modeling of extreme spatial rainfall. Water Resour. Res. 49 (8), 4633–4644.

Zhao, Y. & Ono, T. 1999 A general procedure for first/second-order reliability method (FORM/SORM). Struct. Saftey 21, 95–112.

Zheng, F., Westra, S. & Sisson, S. A. 2013 Quantifying the dependence between extreme rainfall and storm surge in the coastal zone. J. Hydrol. 505, 172–187.

Zheng, F., Westra, S., Leonard, M. & Sisson, S. A. 2014 Modeling dependence between extreme rainfall and storm surge to estimate coastal flooding risk. Water Resour. Res. 50 (3), 2050–2071.