

Modelling the probable maximum precipitation for large catchments

Melanie E. Roberts¹

Mark McGuinness²

Xing Ju Lee³

Bill Whiten⁴

(Received 23 November 2015; revised 4 October 2016)

Contents

1 Problem description	M206
2 PMP and the GTSRM	M211
3 Statistical method	M212
3.1 Exponential distribution with an offset	M214
4 Volumetric examination of extreme weather events in the GTSM zone	M216

DOI:10.21914/anziamj.v57i0.10274, © Austral. Mathematical Soc. 2016. Published October 13, 2016, as part of the Proceedings of the 2015 Mathematics and Statistics in Industry Study Group. ISSN 1445-8810. (Print two pages per sheet of paper.) Copies of this article must not be made otherwise available on the internet; instead link directly to the DOI for this article. Record comments on this article via

<http://journal.austms.org.au/ojs/index.php/ANZIAMJ/comment/add/10274/0>

Table 1: acronyms

AEP	annual exceedance probability
ALARP	as low as reasonably practicable
BCa	bias-corrected and accelerated
BFD	Burdekin Falls Dam
CI	confidence interval
EO	exponential distribution with an offset
GTSM	generalised tropical storm method
GTSRM	generalised tropical storm revised method
GP	generalised Pareto
PAR	population at risk
PMP	probable maximum precipitation
RMS	root mean square
TL	tera-litre, one million mega-litres

4.1 Statistical analysis of storm volume data M218

5 Identifying the Annual Exceedance Probability of extreme dam events **M221**

5.1 Statistical analysis of flood level data M222

5.2 Statistical analysis of flood flow transformed data M228

6 Conclusion **M232**

7 Acknowledgements **M234**

1 Problem description

The Burdekin Falls Dam (BFD) in North Queensland provides bulk non-potable water, predominantly to agriculture, local towns and the city of

Townsville. At capacity, the dam holds 1.86 teralitres (one teralitre, TL, is one million megalitres) and the lake behind it covers an area of 220 km². The dam discharge is an overflow weir 504 m wide capable of a flow that is 15 m above the weir lip. This weir routinely spills for several months each year. The dam catchment covers an area of 114 770 km², making this the largest dammed catchment in Queensland. The catchment is subject to extreme rainfall events originating from both monsoonal and tropical cyclone events, and falls within the Generalised Tropical Storm Method (GTSM) zone [11].

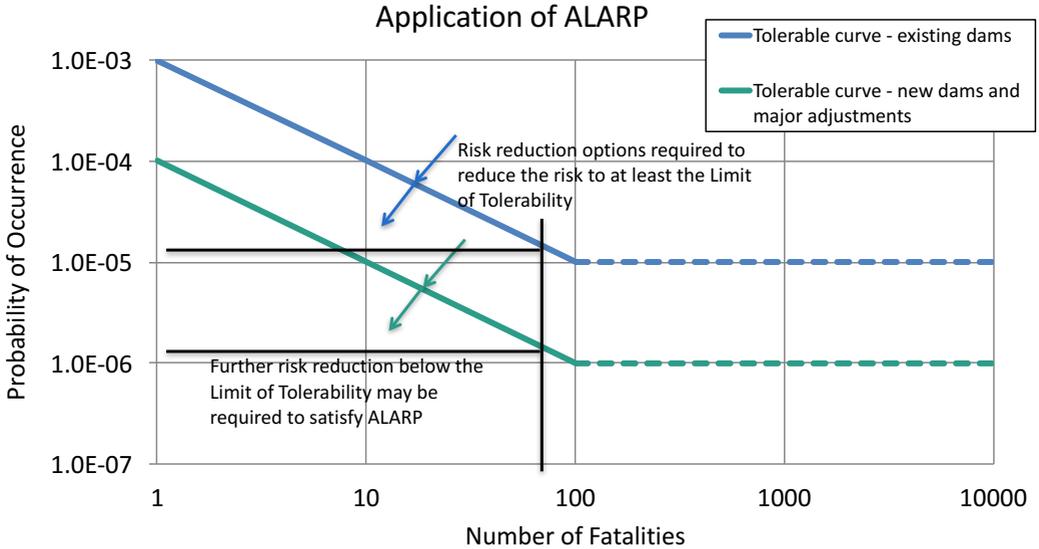
SunWater operate the BFD, and are considering increasing the stored capacity. Any structural changes to the dam are required to comply with recently revised Dam Safety Guidelines [7, 8, 10]. Existing dams in Queensland must, in any case, comply with these guidelines by 2035.

A review of the revised guidelines as they apply to the BFD by the dam operators, SunWater, identified a difficulty associated with the large catchment size.

A first step in the guidelines is to determine the population at risk, in the event of a flood that causes dam failure, over and above the population at risk (PAR) if the dam was not there. SunWater assess the PAR value for BFD to be 87. Then Table 2 of the guidelines, reproduced and annotated with the PAR of 87 in Figure 1, indicates that under the principal of a risk that is as low as reasonably practicable (ALARP), the dam must be designed for an event with an annual exceedance probability (AEP) of 10^{-6} or smaller, that is, approximately a one in a million year event. This also falls at the limit of the range of events considered in Australian Rainfall and Runoff [2], which contains details about how to calculate the design rainfall given the AEP, as illustrated in Figure 2, placing it at the most extreme end of extreme events, and well beyond the credible limit of extrapolation. Extreme value probabilities are complicated by the sparsity of comparable events in the historical record, and thus standard measures of extrapolation are not suitable, as opposed to moderate or large events where interpolation or standard extrapolation methods can be employed, even for a large catchment area.

A dam is usually designed by deciding what rainfall or flood event has the

Figure 1: The required risk (annual exceedance probability) of the designed maximum rainfall, for existing and for new dams [9, Fig. 1]. The values for a PAR of 87 for the Burdekin Dam are indicated by the black lines.

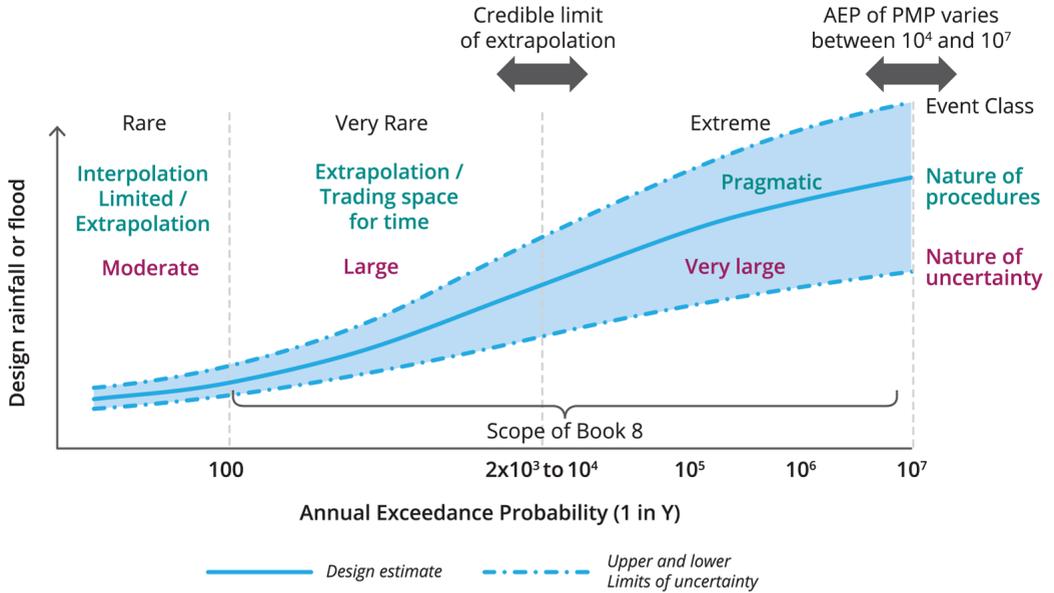


desired AEP, and designing for that as an extreme event. The AEP is related to a series of design figures determined for each dam known as the probable maximum precipitation (PMP). The AEP of PMP is determined as a log-log function of the catchment area in Australian Rainfall and Runoff [2, Figure 6] and is reproduced here in Figure 3.

This assumed relationship is defined in the relevant guidelines [2, 8] only for catchments up to 100 000 km² in area, with the BFD catchment falling outside of this range at 114 770 km². Our figure extrapolates to larger areas, assuming linear behaviour in the log-log plot, so that we can show where the BFD falls (dashed lines in Figure 3).

However, the AEP of the PMP that is obtained in this way is 10^{-4} , which is a higher risk than the required ALARP value of 10^{-6} , so interpolation to the required value is not possible. Extrapolation beyond the PMP is required, to

Figure 2: A conceptual graph of the design rainfall or flood versus the AEP [3, Figure 8.1.1]. The AEP for the Burdekin Dam is at the rightmost limit at one in 10^6 years.

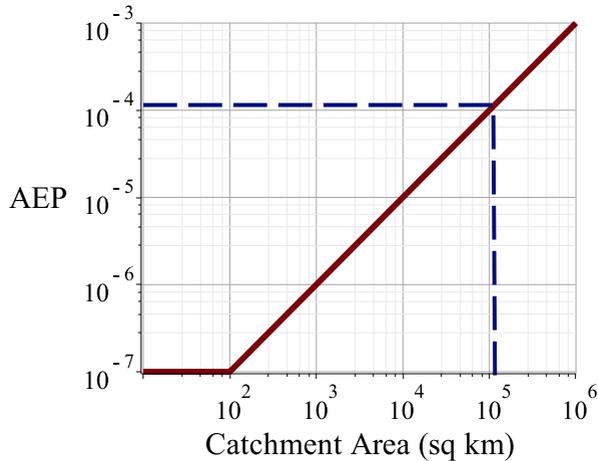


more extreme events with a smaller risk, and there is at this time no basis for such extrapolation in existing regulations. Consequently, it is not possible, under the usual rules, for the BFD to comply with the requirements as the large dam catchment area places the BFD just outside the usual regulatory structure.

The authors of Australian Rainfall and Runoff acknowledge, in regard to assigning an AEP to the PMP, that “there is no conceptually sound, defensible basis upon which to make recommendations for design practice” [2, p.17, section 3.5.1].

The somewhat arbitrary nature of the assignment of risk to dams in this category, and the resultant conclusions for large dam catchments, motivates SunWater to seek a more rigorous understanding of risk and severe weather

Figure 3: The AEP of PMP plotted against catchment area with extrapolation to larger catchments (solid line), adapted from the Guidelines on acceptable flood capacity for water dams [8]. The dashed line shows where the Burdekin Dam catchment falls.



events in the region in order to better understand the risks to be managed in the dam's construction and operation.

The problem, as posed to the Study Group, is to either work on better ways to assign an AEP to the PMP, or to examine the estimation method for the PMP, incorporating local characteristics.

2 Probable Maximum Precipitation events and the Generalised Tropical Storm Revised Method

The probable maximum precipitation (PMP) refers to the greatest depth of precipitation that is meteorologically possible over a defined geography in a given period of time at a particular time of the year, without allowance for long term climate trends [11]. The annual exceedance probability (AEP) of the PMP is theoretical, and acknowledges that the method of determining PMP is itself subject to risk. A range of methods are used to estimate the PMP [1] and its AEP. Within Australia, the Bureau of Meteorology endorses two methods for estimating the PMP depending on the location under consideration, and for the Burdekin Falls Dam (BFD) the appropriate method is the Generalised Tropical Storm Revised Method (GTSRM) for Estimating Probable Maximum Precipitation as specified by Walland et al. [11] and related publications.

This article does not concern itself with the methodology for arriving at a PMP. Since PMP events are both spatially and temporally dependent, it is standard practice to consider a series of events relating to storms of various duration, where spatial dependencies are incorporated through consideration of topographic enhancement and the area of the catchment.

When assessing the impact of a PMP event on the dam, spatio-temporal factors are taken into account, and a rainfall that varies across the catchment is modelled in detail, with a different spatial distribution used for different storm durations, based on historical observations and measurements of large storm events in North Queensland over twenty years. For the BFD, the most critical PMP event is the 72 hour 570 mm event, as this results in the maximum flow over the weir of any of the PMP scenarios stipulated for consideration. Hereafter referred to as the design PMP, the 72 hour 570 mm PMP event is the focus of this article.

The dam catchment has three main sub-catchments, one in the north, one

to the west, and one in the south. The northern sub-catchment provides the greatest concern for flood risk at the dam: the region is associated with a fast response; with flows reaching the weir on the order of a day; and the north-east of the catchment having a small area where topographic enhancement is relevant. The non-uniform response times for the sub-catchments is acknowledged in runoff routing simulations that have been conducted by SunWater, summing the catchments to obtain the resulting flows at the dam, and also simulating the resulting river flows beyond the dam to the sea. In this way, modelling accounts for the possibility that rainfalls are temporally aligned such that the flood peaks of the sub-catchments coincide at the dam lake.

One of the challenges facing the Study Group was to interpret the regulations for dam construction in Queensland, and the problem summary above reflects our interpretation of those regulations. The main issue is that the PMP, which usually serves as a worst event with a smaller risk or AEP than is to be designed for, cannot serve as a worst event for BFD as the regulations stand. Any redesign needs to involve calculations of risk (AEP) that is much less probable than that of the PMP event.

A major focus of the group was to explore a range of measures of risk for extreme weather events. Following a description of the statistical methods in Section 3, the catalogue of extreme weather events used to develop the PMP profiles is examined in Section 4, indicating that the identified PMP may not accurately reflect a theoretical worst case scenario. An alternative measure of risk for the BFD, namely that associated with the height of the dam lake, is explored in Section 5. The article concludes with a discussion in Section 6.

3 Statistical method

The choice of statistical method used in the analyses of Sections 4 and 5 is motivated by the focus of the analyses on extreme data above a certain high threshold value, or exceedance data.

If the distribution of the data \mathbf{x} is known, then it is straightforward to specify the distribution for the exceedance data \mathbf{y} above a threshold \mathbf{u} ($\mathbf{y} = \mathbf{x} - \mathbf{u}$) using the conditional probability

$$\Pr(\mathbf{X} > \mathbf{u} + \mathbf{y} \mid \mathbf{X} > \mathbf{u}) = \frac{1 - F(\mathbf{u} + \mathbf{y})}{1 - F(\mathbf{u})}, \quad \mathbf{y} > 0, \quad (1)$$

where $F(\mathbf{x})$ is the distribution function of the data.

Often, as is the case in the analyses to follow, the distribution of \mathbf{X} is unknown, thus specification of the exceedance data distribution relies instead on an approximation by its limiting distribution, which is the generalised Pareto distribution for a wide range of data [5].

The generalised Pareto (GP) distribution is characterised by its shape parameter ξ , scale parameter σ , and threshold \mathbf{u} , with cumulative distribution function

$$F(\mathbf{y}) = 1 - \left(1 + \frac{\xi \mathbf{y}}{\sigma}\right)^{-1/\xi}. \quad (2)$$

Estimation of the GP distribution parameters involves a two-step procedure. First, the threshold \mathbf{u} value is ascertained. Then, fixing \mathbf{u} , the shape and scale parameters are estimated by maximising the likelihood function $L(\cdot)$ (product of the data points probability [12]) of the GP distribution for the particular data set (satisfying the fixed threshold condition). The threshold is determined here by fitting the GP distribution to a range of threshold values and monitoring the stability of the other two parameter estimates. The smallest candidate threshold value where the other two parameters stabilise is taken to be the estimated threshold value. An alternative method of determining the threshold value is through using the mean residual life plot.

The goodness-of-fit of the fitted GP distribution is then assessed using four common visual diagnostics comparing the distribution function (probability plot), quantiles (quantile plot), return levels (return level plot), and distribution (density plot) of the fitted distribution with the empirical estimates obtained directly from the data set.

The return levels \mathbf{y}_p are obtained by inverting $F(\mathbf{y}_p) = 1 - p$ for a specified ‘return period’ or ‘repeat time’ $1/p$. The return levels are also commonly used for model inference in such applications. The confidence intervals associated with the return level estimates are approximated by the return level’s profile log-likelihood as explicit representation of the confidence limits are not available.

The profile log-likelihood ($l_p(\mathbf{y}_p)$) is the log-likelihood function $l(\mathbf{y}_p, \theta) = \log L(\mathbf{y}_p, \theta)$ expressed as a function of solely the parameter of interest (\mathbf{y}_p here). The explicit dependence of these functions on the data is omitted for clarity. The other parameters, denoted generically by θ , are set to the corresponding values that maximises the likelihood, given a value of \mathbf{y}_p , that is

$$l_p(\mathbf{y}_p) = \max_{\theta} l(\mathbf{y}_p, \theta).$$

The confidence interval for \mathbf{y}_p is then given by

$$2 [l_p(\hat{\mathbf{y}}_p) - l_p(\mathbf{y}_p)] \sim \chi_1^2$$

where $\hat{\mathbf{y}}_p$ is the estimate of \mathbf{y}_p , and χ_1^2 is the χ^2 distribution with one degree of freedom. Denoting the $(1 - \alpha)$ quantile of the χ_1^2 distribution as $\chi_{1,1-\alpha}^2$, the $100(1 - \alpha)\%$ confidence interval is formed by the two values of \mathbf{y}_p where $l_p(\mathbf{y}_p)$ is $\frac{1}{2}\chi_{1,1-\alpha}^2$ less than $l_p(\hat{\mathbf{y}}_p)$.

3.1 Exponential distribution with an offset

A special case of the GP distribution is considered in Section 5, namely the exponential distribution with an offset (EO distribution), which arises as the limit of the GP distribution as $\xi \rightarrow 0$. The EO distribution is an exponential distribution with mean σ whose values are shifted by \mathbf{u} ($\mathbf{y} = \mathbf{x} - \mathbf{u}$). The cumulative distribution function for the EO distribution is

$$F(\mathbf{y}) = 1 - \exp(-\mathbf{y}/\sigma). \quad (3)$$

The use of the EO distribution is driven by the poor fit of the GP distribution to the small data set in Section 5. The EO distribution can be viewed as a simplification of the GP distribution in that there are only two parameters to estimate (since ξ is fixed at 0). The trade-off for this simplification is that the return level for the EO distribution only increase linearly on the log scale whereas the GP distribution allows greater flexibility in the return level behaviour.

The small data set in Section 5 also motivates modifications to both the estimation procedure and goodness-of-fit check, which are more robust.

The EO distribution is fitted to the data by non-linear least squares with squared errors on the probability scale normalised by the variance from order statistics, thus minimising the sum of squares

$$\sum_{i=1}^n \frac{[i/(n+1) - (1 - e^{-(x(i)-u)/\sigma})]^2}{i(n+1-i)(n+1)^2/(n+2)}, \quad (4)$$

where $x(i)$ is the ordered data i , and n is the length of the data set. The numerator of the fractions here are the square of the difference between the expected probability and that given by the distribution, whereas the denominator is the variance of the corresponding order statistic.

The goodness-of-fit for the EO distribution is evaluated using three variants of the Kolmogorov–Smirnov test. These three variants assess three measures of discrepancies in probabilities between the empirical and fitted distributions; maximum difference in probability, maximum weighted difference in probability where the weights are the standard deviation of the order statistics, and root mean square (RMS) of the weighted differences in probability (that is, the square root of expression (4)).

The distributions of the test statistics are generated by a parametric bootstrap using samples from the exponential distribution, and the statistics obtained from the data are compared to determine the adequacy of the fit.

The confidence intervals for the return levels are estimated using bias-corrected and accelerated (BCa) bootstrap [6].

4 Volumetric examination of extreme weather events in the GTSM zone

A volumetric analysis of the 122 storms in the Catalogue of Significant Rainfall Occurrences of Tropical Origin Over Australia [4] is conducted as a first step to understanding the risk (or return time) associated with precipitation events.

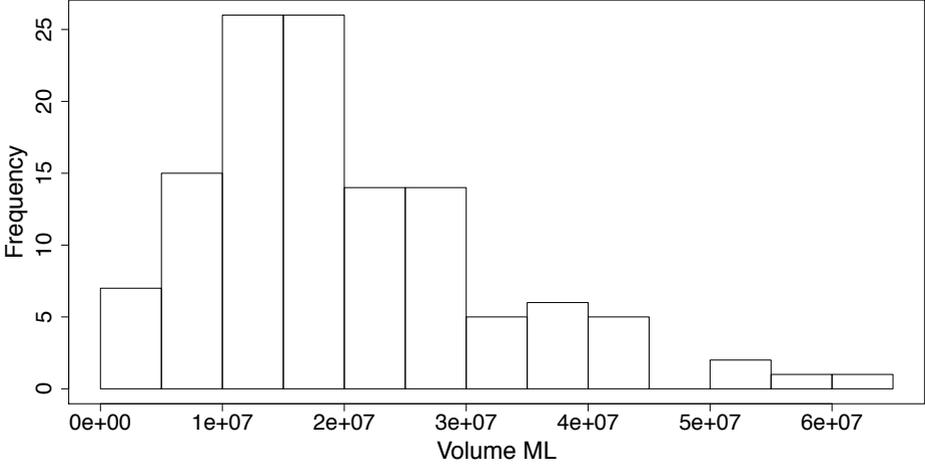
Beesley et al. [4, Table 3] provide the mean rainfall depth over each standard area for storms in the catalogue. The standard areas reported range from 100 to 150 000 km². The maximum volume of rain that could have fallen over a Burdekin Falls Dam (BFD) catchment area rounded to 115 000 km² is calculated for each of the storms. Analysis indicates that the greatest volume observed over this catchment size is 70.35 TL, from the storm 1974JAN23-6, which was centred in the district of Fielding in far North Queensland, approximately 520 km west of Townsville. Under the prescribed assumptions, the volume of water associated with this storm exceeds that corresponding to the design probable maximum precipitation (PMP) event, namely 65.5 TL. This assumes a spatially uniform distribution within each areal polygon, which is unlikely, and has the potential to result in the impact on the dam being under-estimated.

Approximating the storms as concentric circles, we fit the rainfall data from Beesley et al. [4, Table 3] to an intensity profile of the form

$$f(\mathbf{r}, \mathbf{t}) = \mathbf{a} \mathbf{t} \mathbf{e}^{-\mathbf{b} \mathbf{r}}, \quad (5)$$

where \mathbf{r} is the circle radius [km], \mathbf{t} is the storm duration [days], and with the parameters \mathbf{a} [km/day] representing a constant rainfall rate, and \mathbf{b} [1/km] the spatial scale of the storm. The parameters \mathbf{a} and \mathbf{b} are determined under

Figure 4: Histogram of volumetric measure of rainfall (ML) for a catchment of size 115 000 km².



the concentric circle assumption, the average rainfall depth is

$$\bar{d}(R) = \frac{2\pi \int_0^R f(r, T) r \, dr}{\pi R^2} = \frac{2aT}{b^2 R^2} [1 - (1 + bR)e^{-bR}]. \quad (6)$$

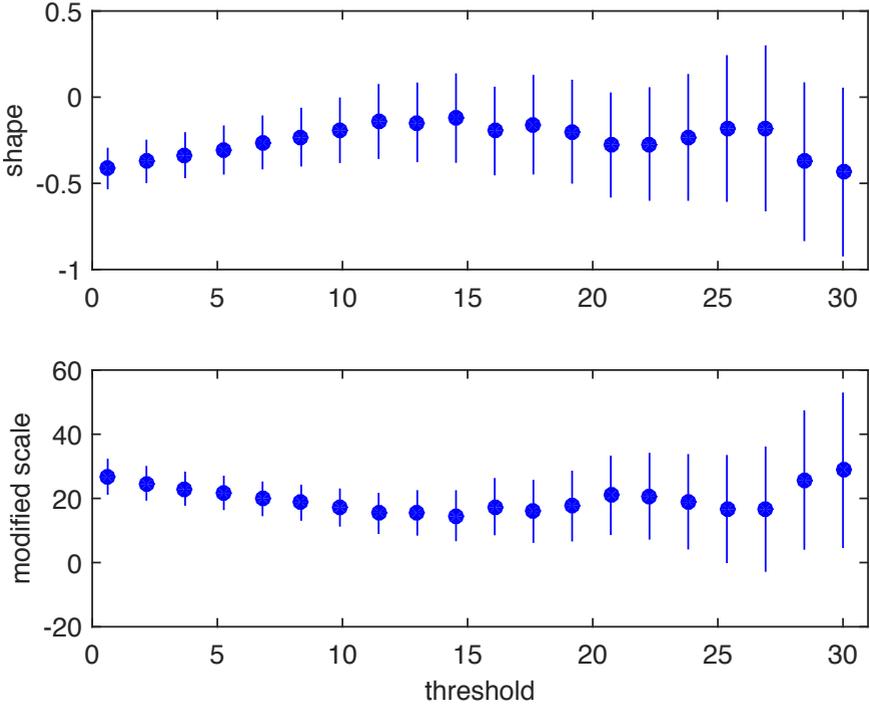
The total volume of rain in a catchment of size R_* due to the storm is therefore

$$V = \frac{2\pi aT}{b^2} [1 - (1 + bR_*)e^{-bR_*}], \quad (7)$$

where $R_* = \sqrt{A_*/\pi}$ with $A_* = 115\,000 \text{ km}^2$ for the BFD catchment.

Applying (7) to the storms surveyed by Beesley et al. [4] provides a distribution of extreme rainfall events by volume for catchments of the size of the BFD, which is shown in Figure 4. This model results in a maximum volumetric rainfall of 60.5 TL, which is of a similar magnitude, albeit less than, the design PMP event. This is also less than the rainfall observed in storm 1974JAN23-6 due to smoothing of the distribution.

Figure 5: Threshold parameter search for the storm volume analysis.

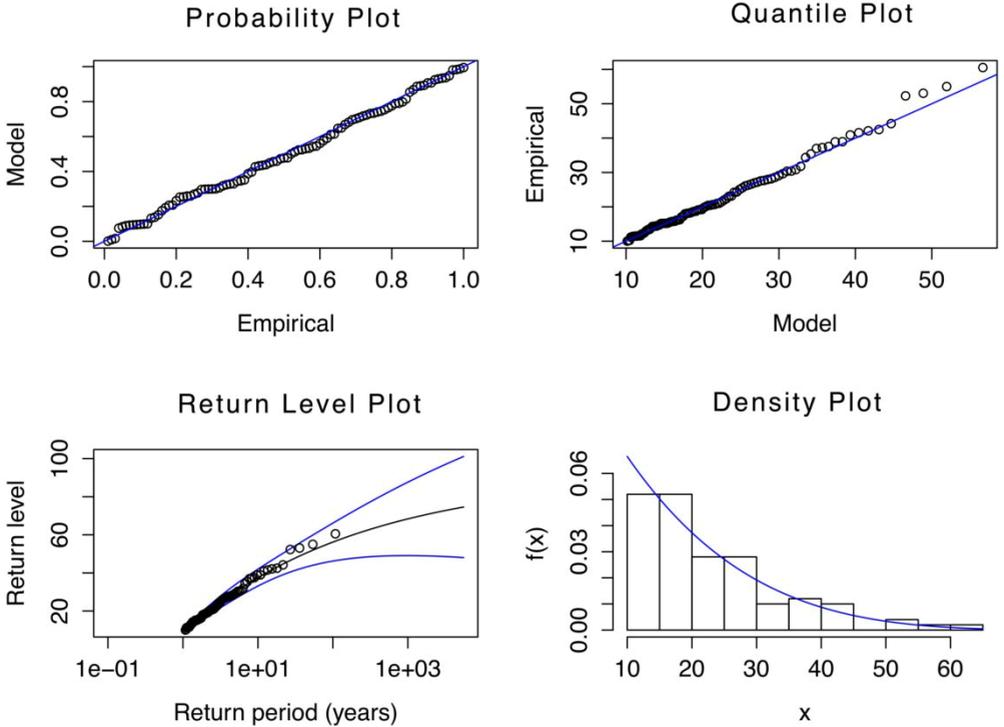


4.1 Statistical analysis of storm volume data

The estimated rainfall data for the 115 000 km² area for each of the 122 storms, shown in Figure 4, are then fitted to a generalised Pareto (GP) distribution. From Figure 5, the GP threshold value is estimated to be 10 as the shape and scale parameter estimates stabilised (excluding statistical error) for threshold values of 10 and higher.

Fixing the threshold value at 10, the estimated shape and scale parameters (with standard errors in parentheses) are -0.185 (0.096) and 15.048 (2.064) respectively. The fitted GP model did not exhibit any obvious lack of fit from the model fit diagnostics (Figure 6).

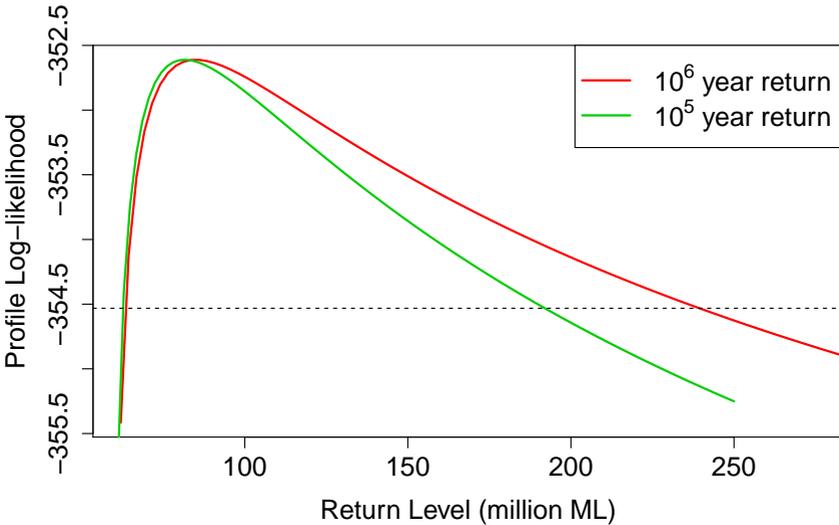
Figure 6: Goodness-of-fit plots for the fitted GP distribution of rainfall in TL over 115 000 km² with threshold set to 10.



The one in a million year event estimate from the fitted distribution is 84.9 TL with a 95% confidence interval (CI) of [63.4, 239.2] (obtained from the profile log-likelihood of Figure 7). The estimate is in excess of the design PMP event which, from this analysis, corresponds to approximately a 1-in-529 year event.

The storm catalogue of extreme events is relevant to the Generalised Tropical Storm Region, and thus the estimated return levels are for extreme events throughout this region. That is, the above analysis identifies a one in a million year event occurring in the Generalised Tropical Storm Region, which may

Figure 7: Return level profile log-likelihoods for the 10^6 year and 10^5 year return period event estimates. The dashed horizontal line denotes the maximum profile log-likelihood value (-352.6) less $\chi_{1,0.95}^2$ to obtain the 95% confidence limits for the estimates. The maximum profile log-likelihood values were the same for both cases.



or may not impact on the BFD. The probability of a severe event impacting on the BFD is yet to be accounted for. Approximately one in ten events in the catalogue impacted the BFD catchment, therefore there is approximately a probability of 0.1 of an extreme event impacting the catchment, which suggests that the appropriate one in a million year event corresponds to the repeat time of 10^5 years, namely a storm volume of 81.5 TL (95% CI: [62.7, 191.9] TL obtained from profile log-likelihood of Figure 7).

The impact of this on lake levels can only be calculated with careful and detailed flow analyses of the entire catchment, and also depends on the temporal and spatial distribution of the rainfall, and the bathymetry at

5 *Identifying the Annual Exceedance Probability of extreme dam events* M221

the lake. This analysis was not conducted, instead a rough indication for comparison is obtained as the PMP for BFD includes a 570 mm total rainfall event that occurs over a total of 120 hours. Converting 81.5 TL to mm of rain evenly distributed over an area of the BFD catchment gives a total rainfall of about 710 mm.

The above analysis indicates that caution must be employed in the use of the design PMP event to represent a theoretical worst case scenario for dam safety planning. From a strictly volumetric perspective, it is evident that storms of a similar magnitude to the design PMP event have been recorded [11]. Further, the design PMP event corresponds to a 1-in-529 year event, with the one in a million year event being approximately 40% more severe. Care must be taken with drawing conclusions from such a limited dataset, with the estimates obtained dependent on the statistical assumptions made and the choice of extreme events.

Given the underlying assumptions to the above model for extreme events, both in terms of the spatio-temporal distribution of rainfall and the statistical assumptions of the model, the identified extreme events and return period may be directly employed for dam safety planning. Nevertheless, this analysis demonstrates that the risk associated with the design PMP event identified may be unacceptably high, and provides a preliminary estimate of the magnitude of the event that corresponds to the prescribed risk ceiling.

In the next section, an alternative measure of risk is explored through the analysis of observations of weather impact at the dam.

5 Identifying the Annual Exceedance Probability of extreme dam events

Although the references to risk for dam safety relate primarily to the annual exceedance probability (AEP) of probable maximum precipitation (PMP), the

guidelines are designed to ensure an acceptable level of risk to downstream communities. It follows that the principle concern is the ability of the dam to safely withstand severe weather events. As discussed in Section 4, the nominal most severe PMP event relevant to the BFD, a 72 hour 570 mm rainfall, is not sufficiently improbable to satisfy community expectations.

A key design criteria for dam safety at Burdekin Falls Dam (BFD) is the maximum height of water above the weir lip before overtopping the adjacent dam walls. This height is currently 15 m. Overtopping the dam walls carries a risk of scouring away the earthworks and bypassing the existing weir with a wall rupture, or undermining the weir and toppling it, resulting in a flood event downstream.

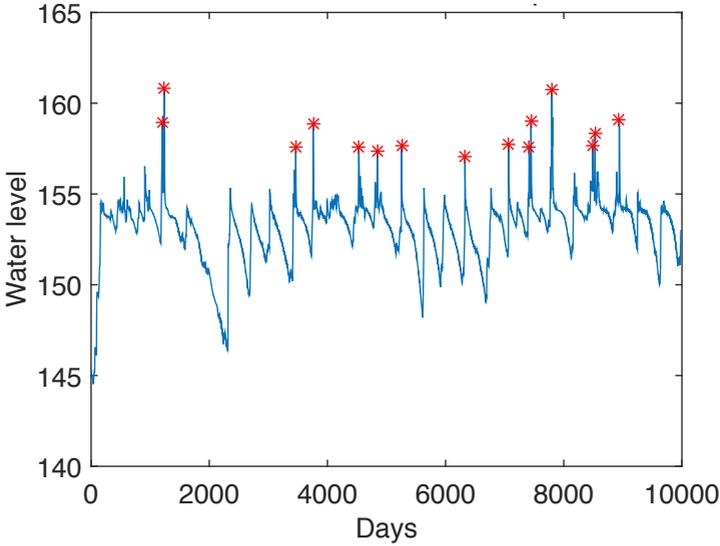
Water heights in Lake Dalrymple, the body of water retained by the dam, have been recorded at various timescales, usually at least daily, over the 27 years of the dam's operation, although with the introduction of automatic recording, the recording frequency has reduced to three-hourly (Figure 8).

These records of water height over the spillway are utilised to identify an exceedance probability, expressed as a return time, for severe events, and hence identify the required dam wall height to satisfy the nominal one in a million year event criteria required for BFD as discussed in Section 1. This degree of extrapolation, even to 10^4 years, introduces a large degree of uncertainty into the estimate, with further uncertainty due to the choice of model. The inherent uncertainty in such estimates must be fully accounted for in designing dam safety.

5.1 **Statistical analysis of flood level data**

Due to the presence of the spillway, it is rather natural to consider measurements above the height of the spillway to have a different interpretation compared with measurements below spillway. This motivates the use of the GP distribution of daily maximum with a minimum threshold value set to

Figure 8: Recorded water level at BFD from construction in 1987 to 2015. The fifteen extreme flood peaks are identified by red asterisks. On this scale, the weir lip is at a height of 154.2 m with a maximum design water height of 169.2 m.



be the height of the spillway. The GP distribution estimation for a range of threshold values are shown in Figure 9. The maximum likelihood estimations for threshold values above 158.5 were unreliable as they converged to a boundary point on the parameter space and were omitted.

From Figure 9, the estimates do not appear to stabilise across the range of threshold values. There is therefore no clear choice for a suitable threshold value based on these estimates. Hence, the threshold value was instead set to the height of the spillway (154.2 m) to enable the analysis to move forward, noting that the model fit may be poor. The goodness-of-fit for this fitted GP model is first assessed by comparing both the probabilities and quantile plots (as shown in Figure 10). The strong deviation from linearity in the quantile plot indicates that the model fit is poor and any extrapolation (as required)

Figure 9: Maximum likelihood estimates and their confidence intervals for the shape (top) and modified scale (bottom) parameters of the generalised Pareto distribution for threshold values from 154 to 158.5 in increments of 0.25.

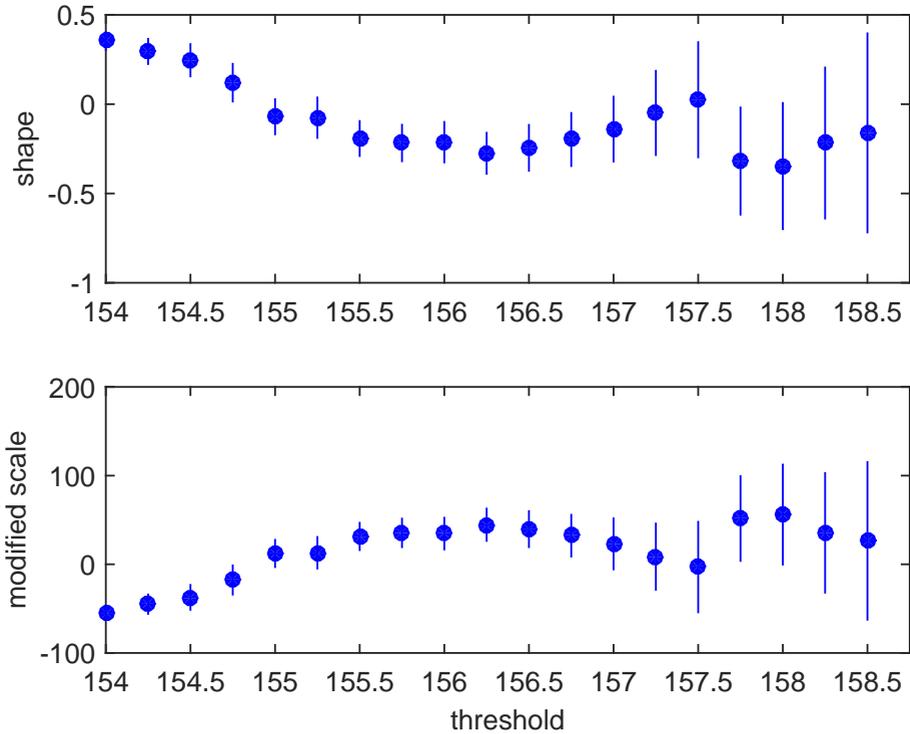
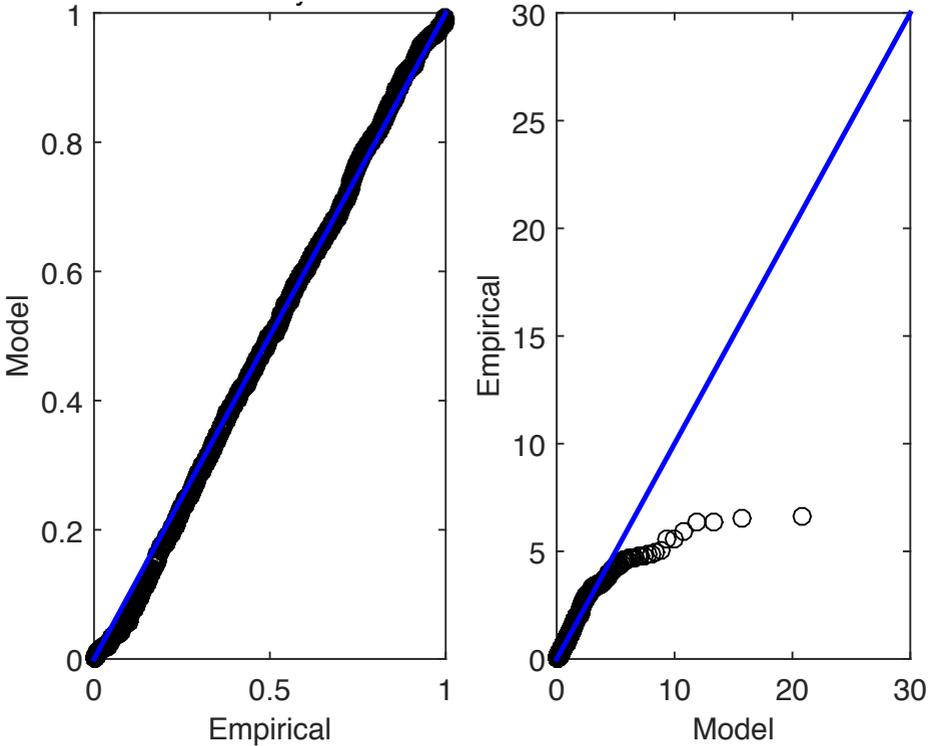


Figure 10: Probability (left) and quantile (right) plots of the fitted GP model for the flood level data with the threshold value set to be the height of the spillway.



using the model would be suspect.

Due to the difficulty in ascertaining an appropriate threshold value and the poor model fit for the GP model with a threshold value selected from the application, we instead resort to a simplification of the model, namely the EO distribution.

Another concern is that even taking only daily maxima, the data could very well still exhibit dependence from previous time points. Specifically for the application, it is possible for an extreme rainfall event to affect the dam water

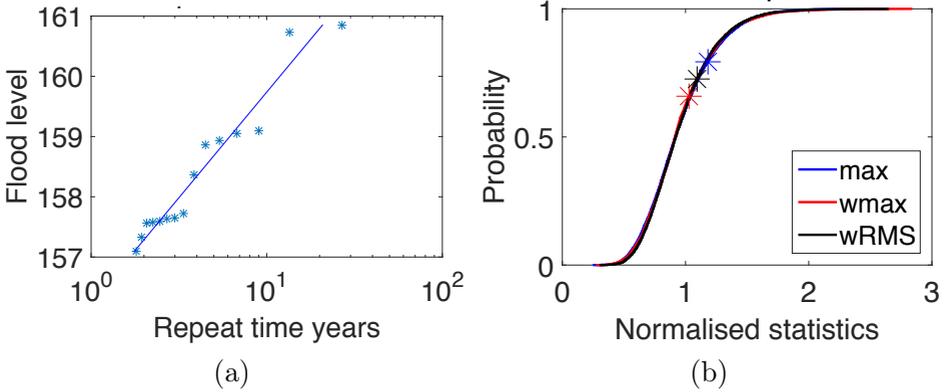
Table 2: Peak flood levels (metres above weir) for BFD, as identified in Figure 8.

Date	Level
1991-01-15	4.73
1991-02-21	6.65
1997-03-26	3.40
1998-01-13	4.67
2000-02-27	3.36
2001-01-04	3.13
2002-02-19	3.45
2005-01-27	2.89
2007-02-05	3.53
2008-01-18	3.37
2008-02-19	4.85
2009-02-06	6.53
2010-12-29	3.43
2011-02-06	4.17
2012-03-23	4.90

levels for multiple days given the time it takes for the rain to flow downstream to the dam as well as potentially rainfall events that last more than one day. Given the relative ease of fitting the EO distribution, we instead use only a selection of extreme flood events from the data.

Fifteen extreme flood events were identified from the recorded history of dam lake heights, as indicated by the red asterisk in Figure 8 and listed in Table 2. While some of these events may appear co-incident, these peaks are at least 30 days apart, and are therefore considered to be independent events. As our interest here is in the peak flow, as opposed to the total flow volume during the event, we consider only the observed peak flow corresponding to the extreme flood event, which are hereafter referred to as the flood peaks.

Figure 11: (a) Fit of EO distribution to the flood peaks with the probability shown as a repeat time in years; and (b) the goodness of fit test.



The available records indicate a maximum flood of 6.65 m (above the spillway), with the current dam design featuring a weir depth of 15 m. Using the formula for weir flow the relative volume flow for the depth of 15 m is $(15/6.65)^{3/2} = 3.4$, so the weir is capable of handling a flood more than three times the maximum observed in its 27 years of operation.

An EO distribution is fitted to the flood peaks to estimate the maximum expected flood over prolonged periods. The fitted distribution is shown in Figure 11a with the probability converted to a repeat time in years, which is the reciprocal of the probabilities.

Comparing the statistics from the flood peaks with the simulated distributions in the goodness of fit tests gives probabilities (standard measures) of 0.80, 0.67 and 0.73 for the three tests mentioned above respectively, as shown in Figure 11b. These goodness-of-fit probabilities are in the expected range, and therefore provide no evidence of a lack of fit. In the worst case, 20% of randomly generated samples fit worse than the flood peak data.

While the EO distribution is a good fit to the data, this does not preclude other distributions giving a similar fit, and such alternative distributions will

give different probabilities.

The probability distribution obtained is extrapolated to estimate the expected maximum flood height in a given number of years; such figures are used in design calculations. Extrapolating the EO distribution from Figure 11a out to a one million year time horizon indicates that the current dam wall relates to a one in 5 000 year event, while the one in a million year event corresponds to 177 m. This is about 7 m above the current design height for water at the spillway of the BFD.

Figure 12 shows the extrapolation of the fitted distribution to estimate the average time period in which a flood level will reach the top of the dam (5 000 years with 95% CI [283, $5.86 \cdot 10^5$]) and the estimated maximum flood level for a million years (177.3 m with 95% CI [169.4, 187.5]). The CIs associated with both these estimates are large.

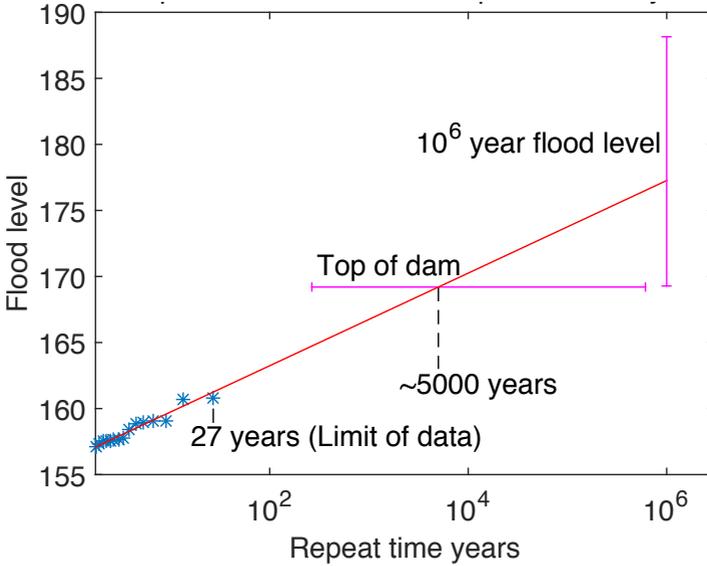
The repeat times are averages and for this application it is more useful to look at the distributions. Estimates of the distributions of the return time for a given flood level or the estimated maximum flood for a given time period are obtained by bootstrapping, by random selection from the data set with repeats. Figure 13a shows the bootstrap distribution of the time to a flood reaching the top of the dam. Figure 13b shows the bootstrap distribution of flood levels for one hundred, ten thousand, and a million years.

These distributions do not give reliable values for extreme probabilities (e.g., reaching the top of the dam in 100 years). For these, a distribution could be fitted to the bootstrap values and then extrapolated. Of course, values obtained would still be subject to variation due to limited data and model assumptions.

5.2 Statistical analysis of flood flow transformed data

As an alternative model assumption, the flood volume flow, rather than the flood level, is fitted to an EO distribution. The flood levels are converted to flow volumes (m^3/s) using the rectangular weir formula that gives the flow

Figure 12: Extrapolation of the probability distribution for flood peak heights out to a one in a million year event. The current dam wall height, indicated, corresponds to a 5 000 year return time. 95% confidence intervals are shown for repeat time to top of dam, and maximum flood level over 10^6 years



proportional to the $3/2$ power of the level

$$\text{flow} = 1.02 \cdot 504 \cdot (\text{level})^{1.5}. \quad (8)$$

Figure 14a shows the fit of an EO distribution to the data, and Figure 14b evaluates the model's goodness of fit using the same three test statistics used in Section 5.1. The probabilities for these are 0.88, 0.72 and 0.75, slightly higher than for the flood levels, but still in an acceptable range, with in the worst case 12% of randomly generated samples giving a worse fit than the flow data.

Again the distribution is extrapolated to give the repeat time for a flood to reach the top of the dam (240 000 years with 95% CI [1991, $8.578 \cdot 10^8$]), and

Figure 13: Bootstrapped probabilities of (a) the flood level reaching the top of the dam, and (b) probable flood level for time periods of 100, 10^4 and 10^6 years, obtained using the flood peak heights data.

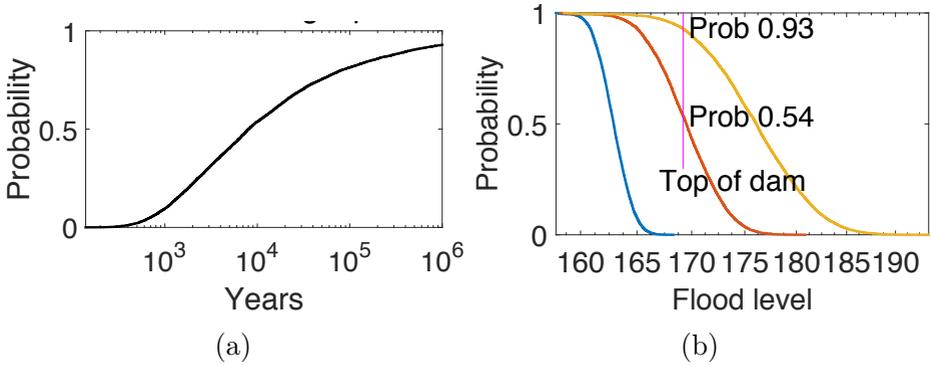


Figure 14: (a) Fit of EO distribution to the peak flood flow, and (b) goodness of fit test.

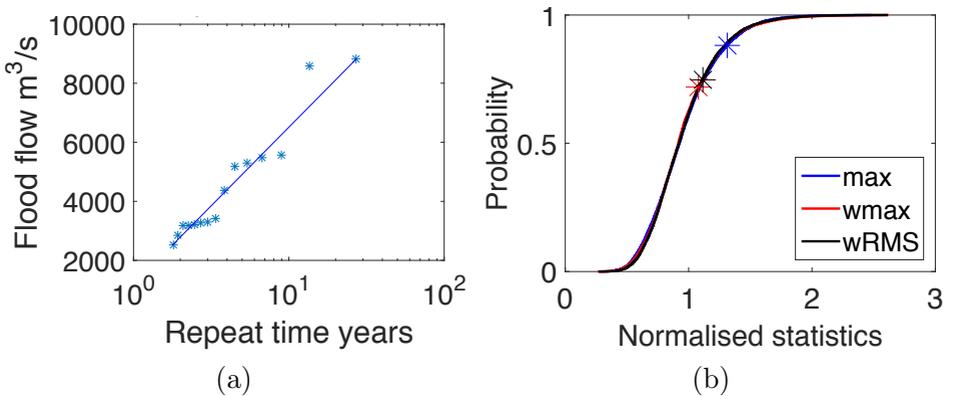
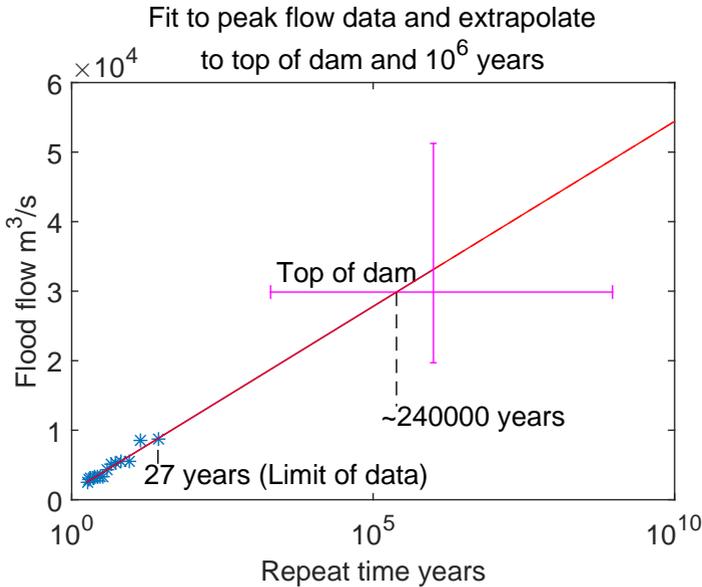


Figure 15: Flood flow repeat times for extended time periods. 95% confidence intervals are shown for repeat time to top of dam, and maximum flood level over 10^6 years

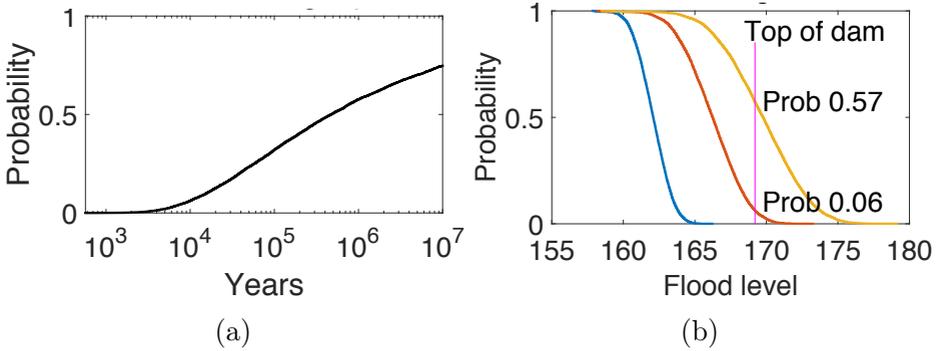


the flow and thus level (170.2 m with 95% CI [165.6, 175.7]) for a repeat time of a million years (Figure 15). The wide CIs associated with both estimates highlights the lack of reliability of the estimates due to the small data set used.

Figure 16a gives the bootstrap estimated time for a flood to reach the top of the dam, and Figure 16b gives the bootstrap distribution for maximum flood heights for one hundred, ten thousand, and a million years.

Figures 14 to 16, calculated using the volume flow, are compared with Figures 11 to 13, which used the level above the weir. This comparison shows the difference underlying assumptions make. Whereas both analyses produce estimates with wide CIs due to the small data set, there are still some notable differences. In particular, the flood level analysis produced a one in a million year flood estimate level where the lower limit of the 95% CI was just above

Figure 16: Bootstrapped probabilities of (a) the flood level reaching the top of the dam, and (b) probable flood level for time periods of 100, 10^4 and 10^6 years, obtained using the flood flow data.



the current height of the dam. The associated estimate from the flow analysis indicates that the current dam height was well within the 95% CI of the one in a million year estimate. However, the flow data also had a poorer fit to the EO distribution compared with the flood level data. Hence, more weight should be assigned to the flood level analysis instead and the inference that the current dam height is not in compliance with the new guidelines.

6 Conclusion

An examination of the recently revised regulations for designing dams with a large catchment area in Queensland reveals that the usual design flood is too likely, that is having too short a return time, to comply with the requirement to design for a one in a million year event. That is, a probable maximum precipitation (PMP) based on storms observed over one hundred years in Northern Queensland is not improbable enough to qualify, and further analysis is required to calculate what rainfall or flood would correspond to the required annual exceedance probability for the Burdekin Falls Dam (BFD).

A number of analyses, preliminary in nature, attempted to identify a one in a million year design rainfall event for the BFD by extrapolating beyond events observed in the past 107 years in the Generalised Tropical Storm Zone of northern and central Australia, and 27 years in Lake Dalrymple.

A review of the extreme storm catalogue reveals that one storm of a similar size to the current design maximum for BFD has already occurred elsewhere in Northern Queensland. Extrapolation of the storm catalogue identifies a one in a million year event to be in excess of the rainfall volume of the current design maximum. This is suggestive of a higher design requirement than currently in place. That extrapolation suggests that rainfalls at this risk level might be double the currently designed PMP total rainfall event, and points to the need to try to push the very detailed PMP storm information gathered to bigger, lower risk storms events, then running the required computer simulations of flows in catchments and tributaries and rivers, down to the Lake, to predict the required design spillway maximal height.

Historical extreme heights of the Lake above the spillway lip have also been analysed and extrapolated. These suggest a compliant design height above the spillway that is 7 m above the current design height of 15 m. They also suggest that as presently designed, the dam can handle a one in fifty thousand years flood event. A similar analysis of extreme flood flow volumes (rather than heights) over the spillway suggest a compliant design height that is just 1.1 m above the current design height.

To predict a one in a million year event from twenty-seven years of data giving fifteen data values is a heroic extrapolation. The problem is that estimations of maximum flood levels are necessary to design safe dams and to validate dam safety. The estimates are statistical, thus the results must be a probability distribution. This distribution allows a 100 or 1000 year flood to occur in the next year, but with a low probability. A thousand year flood can occur any time within the thousand years, and may occur zero or several times in the thousand years. In addition to the range given by the probability distribution, the estimates obtained also depend on the statistical

assumptions made and to some extent on the methods of analysis used.

Changes in the bathymetry as the lake level rises above currently observed ranges are not included in this analysis, except an implicit assumption that they are comparable to the changes in the current range of operation. Changes in climate are another possibility that could alter the estimates made. The occurrence of more extreme events is a commonly acknowledged consequence of global warming.

7 Acknowledgements

We are grateful to the industry representatives James Stuart and Rob Ayre for their support and enduring enthusiasm. We also thank the many people who contributed to this project, including Agah D. Garnadi, Bob Anderssen, Brenda Vo, Glenn Fulford, Julia Piantadosi, Laura Karantgis, Tony Pettitt, Tom Witelski and Tony Gibb. Mark McGuinness acknowledges financial support from Science Foundation Ireland Grant SFI/12/IA/1683 through MACSI, University of Limerick, Ireland. Xing Ju Lee receives PhD scholarship funding from the Centre of Research Excellence in Reducing Healthcare Associated Infections (NHMRC Grant 1030103) and is also affiliated with the ARC Centre of Excellence for Mathematical & Statistical Frontiers.

References

- [1] Ashgo, K., Kuczera, G., Green, J., Weinmann, E., & Laurenson, E. (2000) Estimation of Extreme Rainfall Exceedance Probabilities: Nondimensional Stochastic Storm Transposition, Proceedings in Hydro 2000: Interactive Hydrology. Barton, A.C.T.: Institution of Engineers, Australia, 2000: 60–65. National conference publication. <https://search.informit.com.au/documentSummary;dn=255173836804935;res=IELENG M211>

- [2] Australian Rainfall and Runoff, A guide to flood Estimation, Book V1, Estimation of large to extreme floods. Engineers Australia, Vol One, reprinted 2001. [M207](#), [M208](#), [M209](#)
- [3] Nathan, Rory and Weinmann, Erwin, Australian Rainfall and Runoff, Book VIII, Very to extreme flood estimation, Engineers Australia, Nov 2015, downloaded 23 May 2016. <http://arr.ga.gov.au/arr-guideline/books-and-chapters> [M209](#)
- [4] Beesley, C. A., Meighen, J., & Xuereb, K. C. (2004) Catalogue of Significant Rainfall Occurrences of Tropical Origin Over Australia, Hydrology Report Series, No. 9, Hydrology Unit, Commonwealth Bureau of Meteorology, Melbourne, Australia. <http://www.bom.gov.au/water/designRainfalls/document/HRS9.pdf> [M216](#), [M217](#)
- [5] Coles, S. (2001) An introduction to statistical modeling of extreme values. Springer London. [M213](#)
- [6] Efron, B. & Tibshirani, R. J. (1994) An introduction to the bootstrap. CRC press. [M216](#)
- [7] Guidelines on Acceptable Flood Capacities for Dams, Queensland Department of Natural Resources and Water, Feb 2007. <http://trove.nla.gov.au/version/45001274> [M207](#)
- [8] Guidelines on Acceptable Flood Capacity for Water Dams, Queensland Department of Energy and Water Supply, January, 2013. http://www.floodcommission.qld.gov.au/__data/assets/file/0020/7193/Barton_Maher__Annex_BM3a.pdf [M207](#), [M208](#), [M210](#)
- [9] Guidelines on Acceptable Flood Capacity for Water Dams, Queensland Department of Energy and Water Supply, August, 2016. https://www.dews.qld.gov.au/__data/assets/pdf_file/0003/78834/acceptable-flood-capacity-dams.pdf [M208](#)
- [10] Queensland Dam Safety Management Guidelines, Queensland Department of Natural Resources and Mines, February 2002.

https://www.dews.qld.gov.au/__data/assets/pdf_file/0007/78838/qnrm02013.pdf M207

- [11] Walland, D., Meighen, J., Xuereb, K., Beesley, C., & Hoang, T. (2003) Revision of the Generalised Tropical Storm Method for Estimating Probable Maximum Precipitation, Hydrology Report Series, No. 8, Hydrology Unit, Commonwealth Bureau of Meteorology, Melbourne, Australia. M207, M211, M221
- [12] Wasserman, L. A. 2005. All of statistics: a concise course in statistical inference, Springer, New York, ISBN 0-387-40272-1 M213

Author addresses

1. **Melanie E. Roberts**, IBM Research - Australia
<mailto:melanie.roberts@au1.ibm.com>
orcid:0000-0003-4027-9651
2. **Mark McGuinness**, Victoria University of Wellington, New Zealand
<mailto:Mark.McGuinness@vuw.ac.nz>
orcid:0000-0003-1860-6177
3. **Xing Ju Lee**, Queensland University of Technology, Australia
<mailto:xj.lee@qut.edu.au>
orcid:0000-0002-8803-2810
4. **Bill Whiten**, 4 Magnet Cl, Riverhills 4074, Australia
<mailto:billwhiten@tpg.com.au>
orcid:0000-0002-9778-3632