Swan and Helena Rivers Flood Study

EASTERN METROPOLITAN REGIONAL COUNCIL

Review of HARC Draft Report by Associate Professor Rory Nathan

EMRC Ref 2015/108

UoMC Ref 2015-252

1 April 2016

THE UNIVERSITY OF MELBOURNE

Contents

1. Introduction ... 1

1.1 Background ... 1

1.2 Qualifications of Reviewer ... 1

1.3 Conduct of Review ... 2

1.4 Report outline ... 2

2. Review Comments ... 3

2.1 Overall Approach ... 3

2.2 Available Data ... 3

2.3 Frequency Analyses ... 4

2.4 Model Configuration ... 5

2.5 Calibration ... 5

2.6 Design Inputs ... 6

2.7 Verification ... 7

2.8 Design estimates ... 8

3. Conclusions ... 9

1. Introduction

1.1 Background

This report summarises a Technical Peer Review of the investigation undertaken by Hydrologic and Risk Consulting (HARC) into the Swan and Helena Rivers Flood Study. This review, and the investigation undertaken by HARC, were both commissioned by the Eastern Metropolitan Regional Council (EMRC). The member councils of EMRC comprise the City of Swan, the City of Belmont and the City of Bayswater, and these investigations were undertaken in partnership with the Department of Water (DoW) and the State Emergency Management Committee, with additional financial support provided by the Commonwealth Government Disaster Resilience Program.

The study undertaken by HARC aims to provide a better understanding of the full range of flood behaviour in the Swan and Helena Rivers. The Helena River has a total catchment area of around 1,655 km^2, 90% of which lies above Mundaring Weir. It is located in the lower reaches of the Swan River catchment, which is the outfall for a number of upstream catchments totalling more than 120,000 km^2 in area. The investigation undertaken by HARC involves consideration of the local flood history, available collected flood data, and the development of hydrologic model/s calibrated and verified, where possible, against significant historic flood events and extended, where appropriate, to determine the full range of flood behaviour. The predominant focus of their investigations was on the determination of the 10%, 5%, 2%, 1%, 0.2% and 0.05% Annual Exceedance Probability event design flows and Probable Maximum Flood estimates for key locations within the Swan and Helena rivers study area.

This Technical Peer Review was undertaken by Rory Nathan, who is Associate Professor of Hydrology and Water Resources at the University of Melbourne. The review was contracted through UoM Commercial Ltd on behalf of The University of Melbourne. This review commenced in October 2015 (around six weeks after HARC commenced their investigations) and was completed in February 2016.

1.2 Qualifications of Reviewer

Associate Professor Rory Nathan has over thirty years' consulting and research experience in various organisations in Australia and overseas. He has particular expertise in the estimation of flood risk and is lead author of the national guidelines on the estimation of extreme floods. He has received a number of national and international awards for his research publications. Engineers Australia has formally recognised his significant impact on design practice: in 2000 he was awarded national "Civil Engineer of the Year" and in 2009 he was listed as one of the nation's "most influential engineers".

He has been contracted to provide many reviews similar to this assignment over the past twenty years. He was contracted by Queensland Floods Commission of Inquiry to provide expert review of factors relevant to the magnitude of the January 2011 flood, and this included a specific focus on the need to develop and apply Monte Carlo methods for the Brisbane River catchment. In Western Australia he was a member of panel undertaking risk review of the Dam Safety Program for WA Water and the South-West Irrigation Dams (2002 and 2007), which involved an assessment of the flood studies undertaken for all catchments upstream of the dams. Over the past 15 years he has also reviewed catchment flood studies commissioned (or undertaken by) South East Queensland Water Board, Snowy Hydro, Hydro Tasmania, Sydney Catchment Authority, WA Water, and Melbourne Water.

Internationally, Associate Professor Nathan has been contracted by a number of agencies in North America to advise on the development and application of stochastic (Monte Carlo) methods to estimate flood risk, which has included the US Bureau of Reclamation and the US Army Corps of Engineers. In the past four years he has been contracted by the US nuclear industry (through the Electric Power Research Institute) and the international hydropower industry (through the Dams Safety Industry Group of the Centre for Energy Advancement through Technological Innovation) to provide advice on advanced flood estimation techniques.

1.3 Conduct of Review

The review was undertaken progressively throughout the study, focussing on milestone stages as agreed with HARC and EMRC at commencement. The milestone reviews covered salient aspects related to the proposed approach, focusing on the development and calibration of the coarse and contributing areas of the catchment, the characterisation of design inputs, the statistical analysis of historical evidence, and the derivation of design floods.

Formal review comments were provided in four stages:

i) Proposed methodology and design rainfalls, based on:
  - the report prepared by HARC: “Swan and Helena Rivers Flood Study – Hydrology – Data Collection and Project method. Draft A”, dated 16/9/2015
  - Teleconference meeting notes prepared by EMRC “Contract Inception Meeting”, dated 21/10/2015
  - the file note prepared by HARC: “Design rainfall estimates for the Swan Avon River catchment “ (undated), issued 25/11/2015

ii) Configuration of the coarse hydrologic model, based on:
  - the file note prepared by HARC: “RORB model verification for the Swan Avon River catchment “, dated 21/12/2015
  - Discussion at a meeting held at HARC offices (with DoW in attendance by telephone) on 23/12/2015

iii) Calibration of the coarse hydrologic model, based on:
  - the file note prepared by HARC: “RORB model calibration “, dated 20th January 2016
  - Discussion at a teleconference with HARC and DoW on 22nd January 2016

iv) The whole investigation, based on:
  - the draft report prepared by HARC: “Swan and Helena Rivers Flood Study: Hydrology“, dated 10/2/2016
  - Discussion at a video conference with HARC, EMRC and DoW on 18/2/2016
  - Supplementary material provided by HARC on 29/02/2016, and further on 10/03/2016
  - The final report (Version 1) prepared by HARC: “Swan and Helena Rivers Flood Study: Hydrology“, dated 30/3/2016

This report includes selected comments from the above reviews where they remain relevant to the material included in the final report prepared by HARC in March 2016.

1.4 Report outline

A summary of the salient technical comments made on different aspects of the study are provided in Section 2 of this report. These comments assume detailed understanding of the hydrological issues involved and are suited to readers with appropriate technical training. The overall comments on the study are provided in Section 3, and these are suited to a more general audience.

2. Review Comments

2.1 Overall Approach

The proposed approach is based on the use of a flood event based on storage routing techniques, which is recommended Australian practice. The parameters of the model are determined in the following two ways:

- Routing parameters are derived from matching the flood response of gauged historic events
- Loss parameters are derived through reconciliation of information obtained from matching historic calibration events with flood quantiles obtained from flood frequency analysis.

While the available gauged record is somewhat limited, information is also available on flows (obtained indirectly from hydraulic simulation of flood levels), and this provides key evidence with which to assess the reasonableness of the derived design floods.

For design purposes the model is implemented within a Monte Carlo framework, and this is well suited to considering the joint probabilities involved in the simulation of the required design floods.

There are two particular aspects to the study which have required careful consideration, namely:

- the total catchment area is 124,000 km^2, which is beyond the limits considered in the derivation of design rainfalls, and larger than the catchments used in the development of empirical evidence (and knowledge base of user experience) relevant to the estimation of floodplain risks; and,
- the uncertainty regarding the contribution of the ~100000 km^2 catchment area upstream of the Yenyenning and Cowcowing Lakes, and Lake Ninan, and the likely dependence of this contribution on flood severity.

It is notable that the adopted approach uses gridded rainfall products to supplement the available historic and design rainfalls, and that the characterisation of flood response in the upper catchment was informed by the undertaking of hydraulic analysis and topographic analysis. Both these considerations add considerable value to the analyses undertaken.

2.2 Available Data

The study involved the collation and review of historic rainfall data, recorded and derived streamflow maxima, flood hydrographs for major flood events, and information on topography. The collation and review of this data is appropriate for the scale and focus of the study, and on the basis of the information provided it is unlikely that further analyses would yield material benefits.

That said, it is worth making mention of three sources of information that could be feasibly subjected to further analysis:

- Estimates of peak flood magnitudes that occurred prior to the systematic gauged record were obtained from prior studies based on hydraulic analysis of historic flood levels. These studies were not considered in this review, and it is not known to what extent they could be improved upon. The resources required to undertake such studies could easily eclipse the effort spent on the current investigation, and as such would represent a material change in scope. As discussed in Section 2.3, the historic flood estimates have a significant influence on the design estimates and further analysis may increase the defensibility of these estimates.

- Additional analyses could be undertaken of the available topographic and spatial data to further inform the nature of the natural storage discharge features in the upper reaches of the catchment. Given the current level of reconciliation between rainfall-based estimates and flood frequency quantiles, there would appear to be little benefit in further analysis at this stage, However, if future analysis of historic flood information indicated a material shift in the inferred flood risks, then this analysis might provide further useful information on the behaviour of the upper catchment relative to the flashier response of the "contributing" catchment.

- The (AWAP) gridded rainfall information could be further analysed to provide information on the "space-time" characteristics of rainfalls across the whole catchment. Alternatively, such information could be used to investigate the correlation structure of rainfalls over the "contributing" and "upstream" areas of the catchment, where the outflows from each could be combined in a joint probability framework. As with the previous point, it would only make sense to consider this if a future assessment of historic flood risk indicated that a more complex modelling approach was required.

2.3 Frequency Analyses

The quantiles derived from frequency analyses provide the best source of information on flood risk that is largely independent of the rainfall-based estimates. This is crucially important information to consider, and the nature of the statistical analyses undertaken (using FLIKE) is based on current best practice.

It is clear that the assessed flood risk depends heavily on the estimates of the historic floods that occurred prior to the commencement of systematic gauging. Indeed, inclusion of the historic data effectively doubles the magnitude of the 1% AEP event. The fact that the two frequency curves based on gauged-only and all-historic data are statistically so different – ie the upper and lower confidence limits of each fitted distribution hardly overlap – indicates that inferences based on these analyses need to be dealt with cautiously.

It is stated in the report (Section 4.2.2.2) that the change in slope of the flood frequency curve provides strong evidence for non-stationarity in the data. While it is agreed that the pre-gauged historic data represents a markedly different flood regime, the evidence for this is in the number of flood exceedances in the early part of the record, and not as stated in the slope of the flood frequency curve. Specifically, the salient concern is that 9 of the top 10 floods all occurred in the first half of the historic period; that is, the derived pre-gauged information suggests that the first half of the record yielded markedly higher floods than found in the second half of the historic period. This is unlikely to be the result of land-use change as de-forestation would tend to increase flood peaks not reduce them. The most likely causes for this shift are either: a) systematic bias in the derived flood peaks based on hydraulic inference, or b) a shift in climate regime. It is not possible with the information available to form a view on which of these causes is the most likely. Certainly a reduction in flood peaks seems consistent with the well-recognised reduction in rainfall and streamflow yields observed in this region, though it should be noted that a reduction in average yields cannot necessarily be equated to a reduction in the magnitude and frequency of extreme floods.

This sensitivity of inferred flood risk to this historic information highlights the importance of relying on estimates from prior hydraulic studies. As flagged in the previous section, any revision of the estimates of the top few historic events may alter the inferred flood risk, which may in turn impact on the degree of reconciliation with rainfall-based estimates. The nature of the cause attributed to this shift has a bearing on the weight that should be given to the flood frequency quantiles in the reconciliation process. If further analysis indicated that there is good confidence in some of the historic peaks and not in others, then this would give weight to adopting a higher 1% AEP than that inferred from the systematic record. Conversely if it was accepted that the cause was due to a shift in climate, then it could be argued that the early part of the record is not relevant to flood risk over the planning horizon, and thus more weight should be given to flood quantiles derived solely from the gauged record.

Further comment on the appropriateness of giving weight to the historic events is provided in Section 2.7.

2.4 Model Configuration

The adoption of RORB is an appropriate model for the purpose, particularly given its ability to explicitly accommodate the joint probability of major flood producing factors.

The layout and configuration of the whole-catchment RORB model appears reasonable; the number of sub-areas is unusually large, though this is perhaps justified to support the comparative analysis undertaken using 2D-hydraulic model. There are also a large number of inter-station areas, but given the size of the catchment and the availability of gauged data to help justify the differences in runoff response, this is justified.

Overall, the focus given to the characterisation of natural storages in the model beyond that attributed to normal overland and reach storages is commendable. The hydraulic analyses undertaken provide a useful means for assessing the appropriateness of the hydrologic routing functions. There is considerable uncertainty about how the flood runoff-response varies spatially, and the approach adopted represents a physical basis for the treatment of selected river reaches (in a manner not easily accommodated by varying routing parameters in a spatially lumped fashion).

The use of the digital elevation data to define the storage-discharge functions between Yenyenning Lakes and Lakes Hinds and Ninan, and the hydraulic reasoning used to define the outflow relationships, provides a reasonable basis to cater for the effects of these natural storages. It is possible that the potential for depression storage is greater than that analysed at the three selected locations, but the adopted approach should provide a conservatively low estimate of this. It would be feasible to spend more effort on analysing the topographic data with additional hydraulic modelling, but given how well the model simulates the nature and degree of flood attenuation in the catchment (see Section 2.6) this effort would seem unwarranted. Nevertheless, if additional information did become available that implies a markedly different flood response then this is a facet of the model that would be worth re-visiting.

Care needs to be taken when introducing additional loss and delay functions as these can be adjusted in an ad-hoc manner to accommodate shortcomings that are better dealt with elsewhere in the model structure. However, it is noted that the adopted functions were implemented in an identical manner for all four calibration events, and this gives an acceptable degree of confidence for the magnitude of events considered. Application of the model to events much larger than those considered in calibration will introduce additional uncertainty, though in reality this issue is common to all aspects of model parameterisation.

2.5 Calibration

There are two largely independent sets of parameters that need calibration, namely those dealing with catchment routing and those with losses. The routing parameters represent the degree of attenuation associated with catchment flood response, and this is reflected in the selection of K_c and m parameters, and in the characterisation of “special storages” to mimic localised features (as discussed in the preceding section). The loss parameters relate to the proportion of incident rainfall that does not immediately appear as surface runoff, and this is governed by the selection of initial and subsequent loss rates. Routing parameters are best identified through calibration to individual historic events, and loss parameters are best identified by reconciling local and regional loss estimates with those required to match the location and shape of the flood frequency curve.

It is considered that the ability of the model to reproduce the timing and shape of the historical flood hydrographs, with only minor variation of routing parameters between events, is remarkably good. A lower standard of calibration could easily be excused given the large heterogeneous nature of the catchment and the paucity of pluviograph information. The use of AWAP data captures well the broad nature of the rainfall gradients involved, but this is a necessary not sufficient reason to achieve such a good standard of calibration. The variation of flood response characteristics (as best indicated by the “unadjusted c0.8 values” in Table 6.26) is quite marked, but this would seem justified on the basis of differences in channel morphometry and catchment characteristics. Overall, it is considered that fits to historic flood evens provide good confidence in the manner in which the gauged flood response has been represented in the model.

While the standard of calibration achieved is very encouraging, it is noted that the magnitude of the largest calibration event has an exceedance probability that is somewhere between only 5% and 2%. Simulation of the 1946 event provides additional comfort that the model is able to simulate an event with an AEP closer to 1%, but given the difference in the quality of information this comparison is of less significance.

The historic events provide good evidence for adoption of the proportional loss model, and there is reasonable consistency in the variation of the loss parameters across the different spatial units and between the events considered.

2.6 Design Inputs

The preparation of design inputs is largely consistent with current best practice as documented in the draft Australian Rainfall & Runoff guidelines, however it is worth making mention of some aspects that needed special consideration:

- It has recently become apparent that the Bureau of Meteorology (BoM) are considering adjusting their IFD2013 estimates for AEPs rarer than 5%AEP. This need has arisen from their subsequent analyses of rarer rainfalls to replace the annual CRC-FORGE estimates, which have to date been based on progressive State-based analyses over a 15 year period. The BoM have yet to finalise their thinking on this, but visual inspection of some preliminary diagnostic information would suggest that the IFD2013 estimates for the Avon R catchment may be around on average 5% to 10% too low. Finalised revised estimates may not be available till the middle of 2016, and when these become available it would be prudent to check to what extent any changes might alter the results.

- The new areal reduction factors derived by Podger et al (2015) are based on a more comprehensive data set than that available to the previous CRC-FORGE study, and these are thus considered more appropriate. However, it is noted that these estimates have only been derived for events up to an AEP of 1%, and thus their application to rarer rainfalls introduces additional uncertainty.

- Use of AWAP data to derive design rainfalls for the whole catchment is appropriate given the large area involved. There is some potential for the 24 hour point rainfall to be biased low due to the restricted sampling of daily rainfalls in comparison with 24-hour periods. If present, any bias would decrease quickly with storm duration and is unlikely to be relevant for storms of 72 hours and longer. This is not a trivial issue to investigate because of the spatial scales involved (the correction factors required are different to that used for point data), and the effort required to defensibly identify appropriate adjustment factors would require considerable research effort that would be expected to lie outside the scope of this project. Overall it is considered that the approach adopted provide more accurate estimates than would be obtained by extrapolation of more traditional design information. The approach used to smooth the distribution parameters across duration to avoid inconsistencies in the tails of the distributions is also appropriate.

- Point rainfall growth factors were used to extend 1%AEP areal rainfall values out to 0.05%. This is potentially overly-pragmatic as it is based on the assumption that the skews of the areal rainfalls are the same as point rainfalls. The change in slope of the rainfall frequency curves at 1% AEP (evident in as a “kink” in the slope of the curves Figure 8-10) suggests that the skews of the point and areal distributions are different. While more effort could be spent on resolving this by the analysis of extrapolated fitted GEV distributions, the likely magnitude of any adjustments would be small compared to other sources of uncertainty. Accordingly, it is considered that the curves as derived are fit for purpose.

- The approach taken to specify initial starting levels in the natural lake systems is appropriately pragmatic. The importance of this for Lakes Hinds and Ninan would appear minimal thus of little significance. The implications for the loss functions in the upper reaches of Yilgarn River are of more significance, but it is agreed that the there is little merit in exploring this further as long as the model is able to yield results consistent with the flood frequency analysis using similar assumptions as adopted for the calibration events.

The adopted Monte Carlo framework is well suited to the consideration of design factors of importance to this study. Further complexity could be added, in particular seasonality could be considered more fully and historical information could have been used to develop “space-time” patterns for use with the design rainfalls (or else the contributions from the upper catchment could be parameterised as a stochastic input in a manner that varies with storm severity). However, any consideration of such factors would introduce further uncertainties which would require considerable effort to resolve; given that the model as configured could be validated using independent information relevant to the extremes of interest, it is considered that there is no justification for introducing additional complexity.

2.7 Verification

The final justification for the adopted loss model is provided by the nature of the agreement between rainfall-based and flood frequency quantiles, as summarised in Figure 9-1 of the HARC report. This is a most important check, as this distils the outcome of two largely independent approaches. The model configuration and parameterisation adopted for this comparison, along with the adopted design inputs, are appropriate for the comparison.

There are two key points to note in these results. Firstly, it is apparent that similar peak flow results can be obtained with appropriate loss rates from either the lower or complete catchment. This represents the commonly encountered problem of “equifinality” in modelling, where similar outputs can be obtained by the adoption of markedly different assumptions; unfortunately, it is not possible to justify one set of results over another without obtaining further evidence on one or other of the dominant processes.

Secondly, the rainfall-based results lie midway between the flood frequency quantiles derived from gauged data alone, and those based on consideration of the largest three historic peaks; they lie significantly below the quantiles suggested by inclusion of all 9 historic floods that occurred prior to the start of the gauged record. This outcome is somewhat encouraging, as the “best estimate” based on rainfall-based modelling equates to a prudently conservative interpretation of the flood frequency estimates derived using the gauged record. That is, it gives some weight to the possibility that the latter historic period is more relevant to climate conditions expected in the future, but allows for the possibility for increases in flood risk if the catchment returns to wetter conditions.

These results highlight the importance for the need to incorporate uncertainty into future decision-making. It would be inappropriate to adopt some flood quantile, say the 1%AEP event, as a fixed basis for planning without giving due regard to the inherent uncertainty involved in its estimation. Future decisions on floodplain management need to recognise that it is possible that a given planning level might be exceeded, either because of the (aleatory) uncertainty associated with the occurrence of extreme events, or the (epistemic) uncertainty associated with a non-stationary climate and the data upon which these estimates have been derived.

There is potential for the modelling approach to consider additional complexity. For example, more information could be obtained on the distribution of antecedent lake levels in the upper reaches of the catchment, or on the seasonal variation of some key flood producing factors. However, given the uncertainty involved in the available evidence to validate such modelling, it is considered that the adopted complexity is reasonably commensurate with the available information, and there is no compelling justification to introduce further sophistication.

2.8 Design estimates

Given the foregoing it is considered that the design estimates as reported are appropriate. The results are slightly lower than derived in previous studies, though it is noted the current study has the benefit of an additional 30 years of record.

While adoption of a proportional loss model best suits the estimation of flows over the notional range of historical floods, it is not well suited to estimation of the Probable Maximum Flood. The approach adopted to constrain the uncertainty of this parameterisation is sensible, and it provides a pragmatic means to defining this upper limiting value. The level of agreement between this estimate and world maxima provides a reassuring sanity check.

3. Conclusions

The study has considered estimates obtained from two largely independent methods. One method, based on the statistical analysis of historic and gauged maxima, is subject to the intrinsic uncertainties associated with estimating peak flows from recorded levels. The other method, based on the use of rainfalls in conjunction with a flood model, is subject to the uncertainties associated with model parameterisation and the paucity of relevant information, particularly regarding the contribution of the upper reaches of the catchment.

This review has considered the conceptual approach involved in application of the two methods. It is considered that the approach followed is consistent with best practice as defined in the draft (2015) Australian Rainfall and Runoff national guidelines. The review did not include numerical checks of any data, inputs, or of model specification.

The defensibility of the final estimates rests largely on the level of agreement achieved between the two (largely independent) methods, and on the fact that the study has made good use of the available information. The explicit consideration of joint variability in the flood producing factors is a strength of the rainfall-based method, and the incorporation of historic peaks that occurred prior to the gauged record has highlighted the considerable inherent uncertainty in the nature of the available evidence. It is considered that the design flood estimates represent a prudently conservative interpretation of the flood frequency estimates derived using the gauged record. That is, the derived results give some weight to the possibility that the latter historic period is more relevant to the climate conditions expected in the future, but they allow for the possibility for increases in flood risk if the catchment returns to wetter conditions.

It would be feasible to undertake further work to reduce the residual uncertainties in the estimates. The areas of most benefit would likely involve consideration of the relative uncertainties in estimation of the historic flood peaks, and in the physical and antecedent factors that control runoff from the upper reaches of the catchment. That said, the scope of the investigations undertaken and the level of effort expended represents a high standard of hydrologic practice. It is considered that the adopted approach is commensurate with the available information, and there is no compelling justification to introduce further sophistication.

Overall it is concluded that the study undertaken by HARC provides a defensible set of flood estimates that are suitable for design and planning purposes, as long as due consideration is given to the inherent uncertainty involved in their derivation.