Trinity River Basin Dam Safety Analysis with HEC-WAT

Adams, Lea, Chief – Water Resource Systems Division, USACE – Hydrologic Engineering Center, Davis, CA, lea.g.adams@usace.army.mil
Lehman, William, Senior Economist, USACE – Hydrologic Engineering Center, Davis, CA, william.p.lehman@usace.army.mil

Introduction

The Trinity River Basin spans 18,000 square miles in Texas. The basin contains both USACE and non-USACE lakes, with approximately 24% of the watershed area managed by USACE reservoirs. Although the non-USACE lakes do not store floodwaters, they do have an effect on operation of the USACE lakes. Five USACE dams in the Trinity River Basin were classified with a Dam Safety Action Classification of 2 (Urgent) or 3 (High Priority) in the past few years: Grapevine Dam, Benbrook Dam, Lewisville Dam, Ray Roberts, and Joe Pool. HEC-WAT, in conjunction with other HEC software tools, was applied to evaluate extreme hydrologic loading conditions at the dams of interest.

Study Analysis Overview

The USACE Dam Safety Program embraces a risk-informed decision-making process when prioritizing investment decisions. Risk is comprised of both the likelihood of an event occurring and the consequences of that event. USACE identified hydrologic risk drivers at each dam, including overtopping, wave overwash erosion and surface slides. Given the proximity of the dams to the Dallas-Fort Worth metropolitan area downstream, the consequences of a failure at one or more of the dams in the Trinity system could be catastrophic. Estimates range from 1,000 to 10,000 lives lost. This study was intended to evaluate the likelihood portion of the risk equation, with development of stage-frequency hydrologic loading curves at the five USACE dams.

HEC-WAT was well-suited to evaluate the Trinity River watershed as an integrated reservoir system, considering systems operations and coincident flows, and to explicitly consider uncertainty in the analysis as required by USACE policy. Regional precipitation and temperature data were generated using a weather generator program developed specifically for this study. The precipitation was transformed to flow using HEC-HMS and routed through the system reservoirs with HEC-ResSim, all within the HEC-WAT framework.

Model Setup

Trinity Weather Generator

A custom weather generation tool was developed in Python for the Trinity study, and was based on an earlier weather generator created for the USACE evaluation of Herbert Hoover Dike in Florida. The weather generation approach was selected to incorporate a regional precipitation-frequency analysis within a well-developed hydrologic context. Continuous precipitation and temperature inputs were generated from historical data at a regional scale that considered the various central Texas storm types: local storms, mesoscale storms with embedded convection, mid-latitude cyclones and tropical storm remnants. The computational burden of evaluating enough storm events to reasonably assess extreme hydrologic conditions was initially a concern, and the study team chose a 4-month simulation window instead of an annual simulation window as a balance between accuracy and efficiency in the modeling effort.

The spring season of March 1 to June 30 was chosen as the critical period for reservoir operations because it has the highest background precipitation and highest average reservoir pool stages. The storm typing approach had two main benefits: it works well for larger regions which are subject to different types of storms and it is very useful when trying to evaluate extremely rare events because it retains the probability of each type of event. The final precipitation data set consisted of continuous one-hour data for the 4-month simulation window. The continuous nature of the data sets was important because storm sequencing affected peak reservoir stages. The study team theorized, and eventually the study results confirmed, that multiple smaller storms in sequence could produce more extreme hydrologic conditions than a single large storm.

Another strategy pursued to reduce the computational burden was stratification of the precipitation and temperature data sets. The team planned to evaluate between 400,000 and 1,000,000 events in HEC-WAT. Stratification allowed more events beyond 1xE-06 to be evaluated without a corresponding increase in the number of model simulations, although it is difficult to estimate the frequency of the rarest events because variance is high for those events.

HEC-HMS

The HEC-HMS model used for the Trinity dam safety analysis was originally developed as part of the Corps Water Management System (CWMS) modeling effort. Because compute times are tightly coupled to model setup and application, a number of model adjustments were made to strike a balance between efficiency and accuracy, given the purpose of this particular study:

- Reduced the number of subbasins from 283 to 43, and truncated the watershed footprint at Dallas.
- Changed the routing method in shorter reaches from Muskingum Cunge to Muskingum.
- Removed all modeling features not needed for this particular study, such as observed flows, forecasting information and other extraneous paired data.
- Removed several water supply reservoirs from the HEC-HMS model and added them into the HEC-ResSim model.

Uncertainty in hydrologic parameters, including times of concentration, storage coefficients, loss rates, and recession rates, were evaluated with HEC-HMS as part of this analysis. A new feature in HEC-HMS 4.3, Markov Chain Monte Carlo (MCMC) optimization, was used to generate uncertainty values for the suite of parameters, and the specified values uncertainty method was applied in the model.

HEC-HMS 4.3 was used in the Trinity study because of availability of two new features: the Hamon temperature-only evapotranspiration method, and the specified values sampling uncertainty analysis method (which used results from the MCMC optimization trials). The specified values sampling method is important because it maintains correlation between model parameters and prevents unrealistic combinations of model parameters. After all the various modeling adjustments and additions were made, the team was able to reduce HEC-HMS runtimes from 90 seconds for a 2.5-month simulation to about 52 seconds for the final 4-month HEC-WAT runs (at a 1-hour timestep).

HEC-ResSim

As with the HEC-HMS model, the original HEC-ResSim model came from the CWMS modeling library. The number of inflow locations in the HEC-ResSim model was reduced to match handoff locations within the simplified HEC-HMS model, and a series of model adjustments were made:

- Fixed broken references in reservoir operating rules.
- Removed redundant operating rules, as well as rules that were unnecessary during flood operations.
- Updated model to correctly account for overtopping flow rates at each dam when experiencing extreme flows.
- Replaced downstream reservoir control rules with maximum release rules.

Uncertainty sampling in HEC-ResSim was limited to sampling starting reservoir pool elevations. The effect of sampled starting pool elevations on seasonal maximum pool elevations was minimal, however, given the long simulation periods.

HEC-ResSim 3.3 beta was used for the Trinity study. The study team was able to reduce HEC-ResSim runtimes from about 6 minutes to 18 seconds for the final 4-month HEC-WAT runs (at a 3-hour timestep).

HEC-WAT

An HEC-WAT model of the Trinity watershed, using HEC-WAT 1.1 beta, was constructed to facilitate modeling hundreds of thousands of hydrologic events. As noted above, precipitation and temperature data were generated externally to HEC-WAT using a study-specific weather generator. This data was imported into HEC-WAT via a plugin called the Stochastic Data Importer, which allows HEC-DSS data sets to be read as inputs to HEC-HMS and HEC-ResSim. The Model Linking Editor in HEC-WAT was used to manage the flow of information between the input data sets and the two study models.

Hydrologic events were organized in the HEC-WAT model by lifecycles and realizations. A lifecycle is the smallest grouping of hydrologic events within HEC-WAT, and is typically 20 to 100 events in length. Lifecycles are then grouped into larger realizations. This structure is typically used to organize the two types of uncertainty sampling within HEC-WAT: natural variability within lifecycles and knowledge uncertainty within realizations. For the Trinity study, the lifecycle-realization structure was further leveraged to organize the compute structure on a distributed compute network, and manage the record keeping for stratification and destratification.

Two custom HEC-WAT plugins were created for this study. The "Trinity" plugin computed critical duration flow and stage averages for various points in the study area, and was run as part of the compute sequence. This added customized study results to the standard set of available results from the HEC-HMS and HEC-ResSim plugins that summarize Monte Carlo outputs during an HEC-WAT uncertainty simulation. The "Merger" plugin combined several post-processing operations. It gathered all HEC-DSS results files from each compute from across a distributed compute network, organized the files with a numbering scheme, merged them into a master HEC-DSS file and then destratified the results. During the destratification process, convergence criteria were applied to each output to determine whether the results had converged.

Validation and Stress Testing

It was important to ensure that the models were functioning properly over the full range of expected hydrologic conditions prior to launching the study production runs, for purposes of both quality and distributed compute robustness. To this end, the HEC-HMS model was calibrated to one observed event; validation to a second observed event was desirable but constrained by study timelines. The HEC-HMS and HEC-ResSim results were then confirmed to be identical in both standalone and inside the WAT. Uncertainty sampling was turned on and the team ensured that sampling was occurring as expected in both HEC-HMS and HEC-ResSim.

Once the model results were validated, several rare precipitation events created by the Trinity weather generator were applied in HEC-WAT to stress test the HEC-HMS and HEC-ResSim model setups for performance under extreme hydrologic loading conditions. A number of adjustments were made to ensure that the models could pass large flows without failing, as problems in either model could cause a large-scale HEC-WAT compute to fail.

Study Analysis

Distributed Computing

After considering contracting options, the study team chose instead to use an in-house distributed compute network for greater flexibility in performing runs. The compute network consisted of 20 desktop machines managed with several virtual machine (VM) software tools, including Hyper-V, PDQ Inventory and PDQ Deploy. This suite of tools allowed remote access for team members located in different parts of the country to manage the VMs, deploy batch processes and initiate HEC-WAT runs.

The final HEC-WAT compute was structured with 50 events per lifecycle, which met the requirements of stratification scheme, and 20 lifecycles per realization. The production distributed compute network consisted of four grids of 20 compute nodes, plus one master node for storage of aggregated results, but the HEC-WAT model was initially stress tested on a smaller grid to ensure that the distributed compute software was functioning as expected. HEC-WAT supports distribution of computes via either lifecycle or realization. Occasional compute failures were expected to happen due to software or hardware problems, and distribution by lifecycle was chosen to minimize the amount of time needed to re-run a simulation if a failure occurred.

The study plan called for evaluation of up to 1,000,000 events, but initial investigations suggested that results could converge between 600,000 and 800,000 events. The production runs were organized into 100,000 event lifts, with 25,000 events distributed to each VM grid. Convergence looked promising after 400,000 events, but the next compute lift revealed a very large outlier and additional lifts were required. The unexpected outlier was related to the issue of several more-frequent storms in sequence creating higher peak flows and reservoir stages than a single less-frequent storm. Ultimately the full complement of 1,000,000 events was modeled, and convergence was observed at most but not all study output locations.

Data management was a concern with a study of this size. Dividing the full simulation into 100,000-event lifts, which were further divided into 25,000-event computes across the four VM grids, helped keep data file sizes manageable. The hourly precipitation input data sets resulted in a file size of 14GB for 25,000 events. The HEC-WAT output data set sizes were about 90MB per 25,000 events, or 3.6GB for the full 1,000,000 events.

Study Results

The final HEC-WAT analysis of 1,000,000 events was completed over the course of 14 days, in ten lifts of 100,000 events per run. Stage-frequency results and uncertainty bounds were developed at each of the five dams (Grapevine Dam, Benbrook Dam, Lewisville Dam, Ray Roberts, and Joe Pool) and at the Dallas Floodway for peak inflow, peak inflow volume for the critical duration (2- or 3-day flow) and total inflow volume for the 4-month simulation window. The stage-frequency curve at Grapevine Dam based on 1,000,000 events evaluated in HEC-WAT is shown in Figure 1.

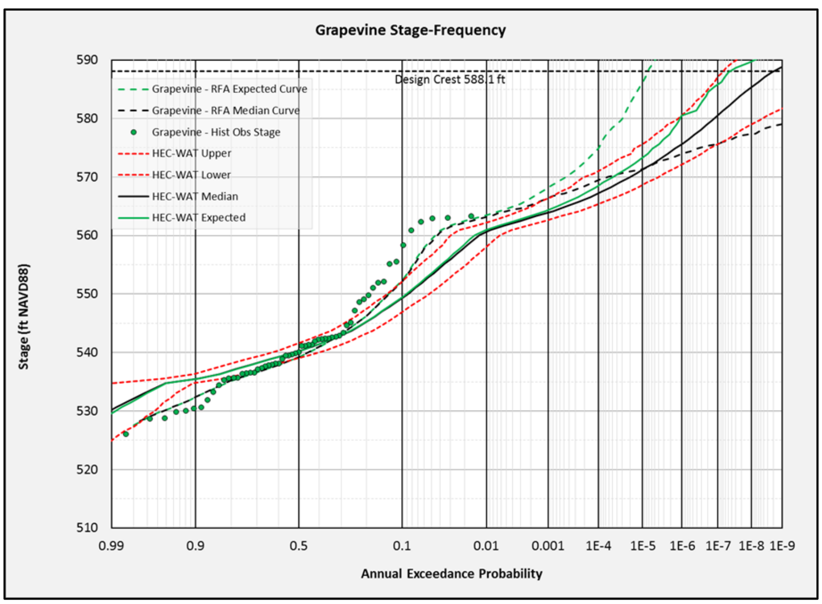

Figure 1. Stage-frequency at Grapevine Dam

Lessons Learned

There were a number of challenges and lessons learned during the Trinity study. The most important items are noted here:

- Questions arose about how to calibrate an HEC-WAT model for an entire watershed. Individual models responsible for one part of the hydrologic process can be calibrated relative to observed events, but it's more difficult to determine where to make adjustments across three models working in concert with their own intrinsic model errors.

- Runtime failures occurred that were initially attributed to the software, since each program was in active development. After investigations didn't reveal any bugs, it was discovered that the source of the failures was actually extremely large precipitation events that were crashing to the HEC-HMS and HEC-ResSim models, despite the initial stress testing. The study setup effort had focused primarily on the five USACE dams, extending volume-elevation curves at those locations, but it was eventually discovered that the curves needed to be extended further at some of the other dams in the basin. The lesson learned was to investigate the simplest, most obvious source of a problem first, before spending time on the more complex possibility of software bugs.
- Stress testing was performed on the stratification bin with the rarest events, but it was later discovered that several of the most extreme floods actually occurred at more frequent events, driven by multiple smaller events in sequence. There wasn't an easy way to identify these critical stress testing events without running a large number of overall events.
- A 2GB file size limitation on file transfers buried in legacy code caused hours of troubleshooting at the beginning of the production runs. This limitation was missed in stress testing because there were more zero-precipitation hours than during the full production runs and the corresponding stress testing HEC-DSS files compressed into files less than the 2GB limit. Large file sizes also required long transfer times to move data onto the compute network.
- The compute logs for HEC-WAT, HEC-HMS and HEC-ResSim were not originally designed to support troubleshooting for a Monte Carlo-style compute. The model logs capture a large amount of detail to help users find problems, which works well for a small numbers of events. However, the information in model logs can become overwhelming when they are filled with detailed information for thousands of events.
- The original work plan consisted of a smaller set of test runs to confirm model validation, followed by one or two full production runs. When results from the large-scale production runs were reviewed, however, the decision was made to adjust methods in the weather generator to better capture the extreme end of the frequency curve. This required generating new precipitation data, often followed by adjustments to the HEC-HMS and HEC-ResSim models, then re-analysis of the production runs. This review-adjustment-re-analyze cycle occurred several times, and took longer than anticipated because evaluation of the extreme end of the frequency curve required producing hundreds of thousands of events.

Acknowledgments

The following USACE staff made significant contributions to the completion of the Trinity River Basin dam safety analysis with HEC-WAT: Allen Avance, Matt Fleming, John Hunter, Greg Karlovits, Joan Klipsch, Chan Modini, Sara O'Connell, and Haden Smith.

References

Karlovits, G. 2018. Synthetic weather simulation for characterization of uncertainty in extension of stage-frequency curves in a system of flood control dams. Poster presented at: American Geophysical Union Fall Meeting, Washington, D.C.

USACE-Risk Management Center. 2018. “Trinity River Basin Flood Hazard Assessment,” RMC-TR-2018-## (draft).