You Are Smarter Than Your Model. Really.

Ellen B. Faulkner, P.E.
Water Resources Engineer, Ayres Associates

Introduction

Dam safety studies often involve the use of numerical models to assess the flooding that could occur from a failure of the dam. These models provide useful information, but also abundant opportunity for frustration, excessive analysis costs, and misleading results. The author's 25 years of working with the HEC-RAS model and other unsteady-flow model applications for dambreak studies have provided numerous lessons on making, recognizing, and correcting model malfunctions and errors. On reflection, most of these problems could have been avoided – or at least detected early - by taking two actions. First, each study should start with the confident application of common sense, plus some simple hydraulic principles that require no computational gymnastics at all. Second, the modeler should refuse to accept results that do not conform with the expectations thus developed.

The necessity of insisting on sensible answers seems painfully obvious. However, in a wilderness of detailed input requirements, miniscule time and distance steps, and dozens of output options, there is a tendency to lose sight of, or even distrust, what is simple and intuitive. Obtaining a technically meaningful result becomes even more difficult when the modeler begins to measure success by simply getting the model to run.

This paper makes reference to the HEC-RAS model, which is presently the most widely used software for analyzing one-dimensional unsteady flow in natural channels in the U.S. However, the focus of the paper is the importance of an independent and critical evaluation of any model's performance. Regardless of the sophistication of the software, every water resources engineer with a class in open channel hydraulics and a pocket calculator has both the tools and the responsibility to build a meaningful, quantitative framework into which the model results should fit.

The Basics Matter.

The biggest difference between a dambreak analysis and the type of calculations that generations of engineering students made by hand (or by cumbersome home-made computer routines) is that the dambreak analysis requires a new calculation for every small increment of time that is being modeled, in some cases over thousands of time increments. Still, in dambreak model analyses there is much to be learned from the classic equations involving steady, uniform flow. For example:

- Normal depth (calculated using Manning's equation) is a good first-cut estimate at the depth for a given discharge at a given section. If the model gives a very different depth, there should be a readily recognizable reason - normally either a drawdown or a backup from downstream. In steady flow, the "drawdown" or "backup" is caused by geometry changes; in dambreak modeling it may also be caused by the position of the flood wave.

- Critical depth – the flow depth below which conditions at a given cross section cannot possibly affect conditions upstream - depends *only* on the discharge and the cross section shape. Factors such as slope, roughness, and what is upstream or downstream do not affect critical depth; they affect where the actual water surface is relative to that critical depth. Remembering this distinction is helpful when trying to troubleshoot through critical or near-critical flow.

- The computed energy grade line is at least as informative as the water surface profile. The energy grade line is always above the water surface profile; if it touches the profile the model has computed a zero velocity for that location. Unlike the water surface profile, which can exhibit brief stretches of adverse slope, the energy grade line *always* decreases in the direction of flow. When it does not, the model is simply not working.

- The weir equation, familiar to all water resources engineers, is what the model uses to compute flow through a dam breach. The model calculations are also affected by factors such as an irregularly shaped breach, the drawdown during the time of formation, and tailwater submergence. Still, a manual weir flow calculation using the starting reservoir elevation and a rectangular approximation of the breach gives a good initial estimate of the peak breach outflow. If the model produces a number that differs significantly from this benchmark, the modeler should insist on understanding why.

Conservation of Mass is Easy.

It is oversimplifying to suggest that the model uses steady-flow equations over and over again to develop the narrative of the flood wave's development and downstream translation. The model's solution technique does use those equations, but there are also time-dependent terms that do not occur in the steady flow equations. The most obvious and influential time-dependent relationship is the conservation of water volume over time. Keeping track of water volumes throughout the model run is one of the easiest ways to understand and gain confidence in the model calculations. For example, the acre-feet listed in the HEC-RAS output hydrograph at a section downstream of a failed dam should be very similar to the acre-feet stored in the reservoir above the breach bottom (after accounting for the background flows not related to the failure). If it is not, the model may have become "stuck" in an iterative calculation and approximate solution that produced a spurious loss or gain in water volume. Another possibility is that the run did not last long enough to allow the entire flood wave to pass through the downstream section.

A reservoir is a particularly easy place to apply conservation principles as a check on the model behavior. When the level is rising, the inflow is greater than the outflow; when the level is falling, the outflow exceeds the inflow; and WHY the difference between inflow and outflow exists should be easily explained by what is known about the spillway's stage-discharge rating, the breach, and the inflow hydrograph.

The charts below are an example from a HEC-RAS run where both the manual breach flow calculation and a volume comparison indicated a problem in the model. In the first chart, the peak breach outflow (about 1,300 cfs more than the background flow) did not match the manual weir equation estimate of 3,700 cfs. Additionally, the volume of the tiny triangle perched on the larger background hydrograph was much smaller than the volume of water that was expected to be released by the breach. An investigation of the model setup revealed that an input error (an incorrectly entered channel distance) was preventing the model from "seeing" all but a small slice of the reservoir. In the second chart, the reservoir cross sections have been corrected; the peak breach outflow of 3,600 cfs roughly matches the manual calculation; and the volume under the breach hydrograph matches the known reservoir volume above the stabilized post-failure elevation of 849 feet.

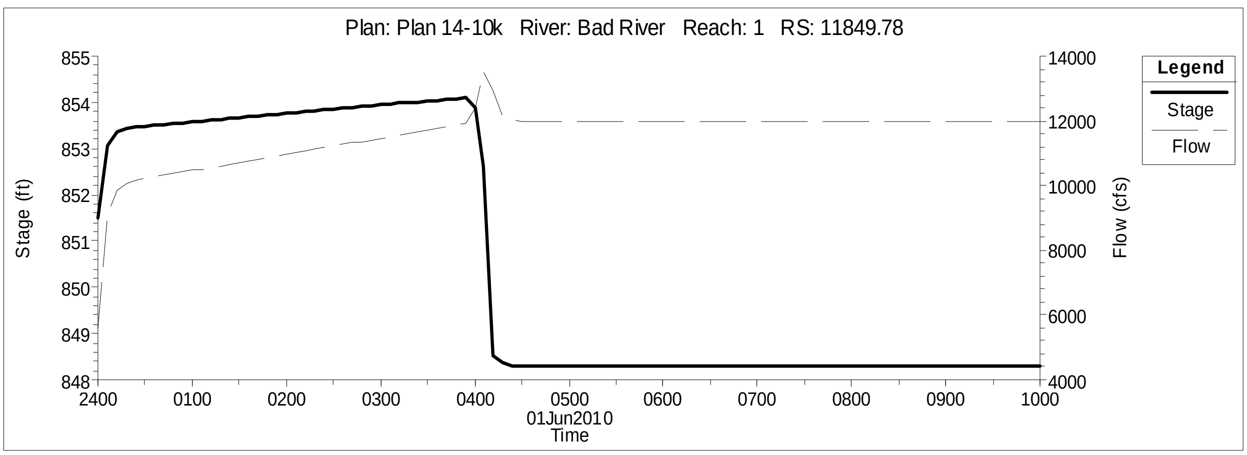

(1) Problem due to incorrect upstream data entry

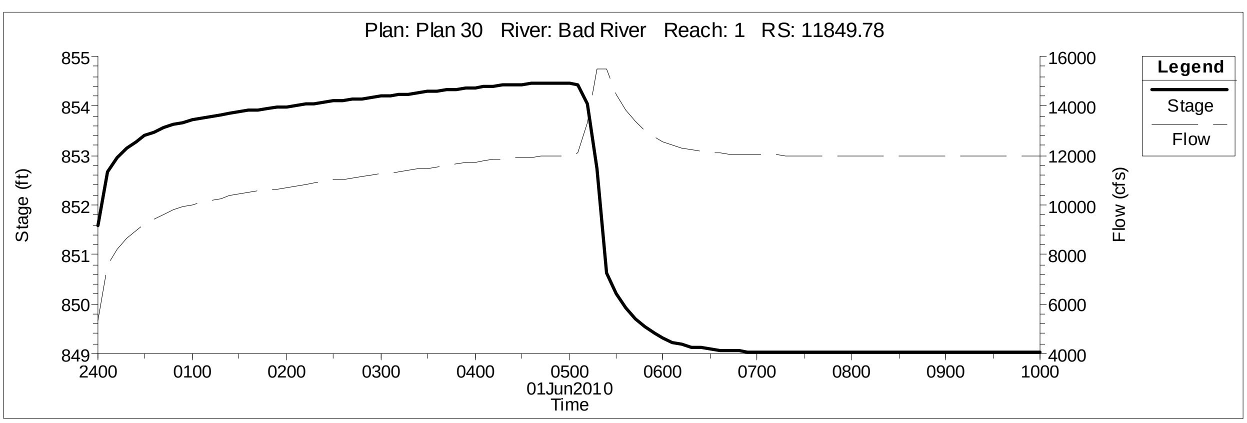

(2) Output matches expectations

Know What to Expect.

There are a number of basic “rules of thumb” and common sense principles that help the modeler frame his expectations. In addition to those discussed above, consider the following:

- As the dambreak flood hydrograph moves downstream, its peak should decrease and its shape should flatten, unless there are other contributions to flow in the channel. The degree to which this happens depends on the initial steepness of the hydrograph and its volume, relative to the storage volume available in the downstream valley.

- Sustained supercritical flow is unlikely in a natural stream, as are very high velocities. Depending on the type of stream, a computed velocity over 10 to 15 feet per second warrants a second look. Any velocity anomalies that are not explainable by channel geometry should also be re-checked.

- The flood wave velocity is not the same as the flow velocity. The front of the flood wave will generally arrive at a downstream location faster than would be estimated using the distance and the computed flow velocity.

- Even for a dambreak situation, irregularities in the hydrograph or the flood profile (such as double peaks, adverse slopes, or sudden changes in water surface slope) are red flags. So is too *much* regularity, as in a hydrograph or water surface profile that suddenly becomes flat.

- Formal calibration data (such as measured river flows and stages) may be lacking, but there is always something that serves as a “reality check.” For example, the normal tailwater elevation at the dam (matching common flow conditions) is usually known or observable. Well-traveled highway bridges are not likely to be overtopped during small floods.   Non-flood flows should stay with the banks. Literature on stream formation tells us the bankfull flow is likely to be have around a 50 percent annual recurrence probability (the “two year flood”); if the model is fitting the 10-percent probability flood inside the banks, there may be a problem.

- Given the widespread availability of aerial imagery, there is no reason not to know how wide the river is at normal flow, where there are rapids or sluggish vegetated reaches, or how much a bridge constricts flow. (This was true even in the days of paper topo maps, but the aerial images are updated much more frequently than the topo maps were.)

No Surprises, Please.

As can be seen in the preceding discussion, before ever constructing a model the user can establish reasonable bounds for expected peak flows, hydrograph volumes, depths, velocities, and wave travel times. But there are exceptions to every rule, and it is perfectly acceptable for a set of model results to defy expectations. What is less acceptable is for the modeler to faithfully report the surprising results without explaining, to her own satisfaction and that of any reviewer, why they occurred. Poring over the software-generated graphics (including profiles, hydrographs, rating curves, velocity distributions, and cross section plots), tracking hydrograph volumes through model time and space, and checking the model sensitivity to particularly worrisome input values will help reconcile the expectations and the results.

Keep The Model Under Control.

The key to having a model you trust is having a model you can understand. The model should be no more complicated than is necessary to develop reliable results at the locations of interest. The photo at the right shows a particularly challenging modeling situation – a narrow gap between bedrock outcrops bordered by thick vegetation or, as the HEC RAS cross section description had it, “The Crack In The Ground.” The location of interest in this particular study was a campground about one-half mile downstream of the Crack. Unfortunately, features like the Crack often prove numerically problematic and may, for example, create an unrealistic wall of water on the upstream side as in the first profile shown below. A solution that appears to work for one flow may produce unrealistic results at a different flow. This justifiably creates a distrust of all the solutions, even the ones that appear plausible.

In the second profile, the problem reach has been smoothed into a trapezoidal chute that is wider than the real topography. Geometric accuracy in that immediate area has been sacrificed for the sake of ensuring that the dam breach hydrograph delivered to the downstream area of interest is realistic.

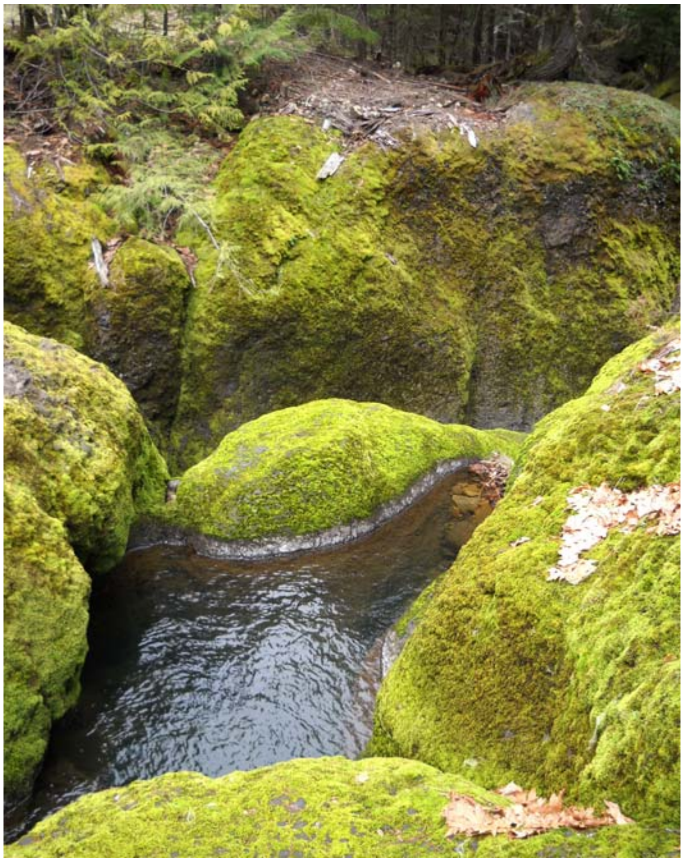

The Crack in the Ground

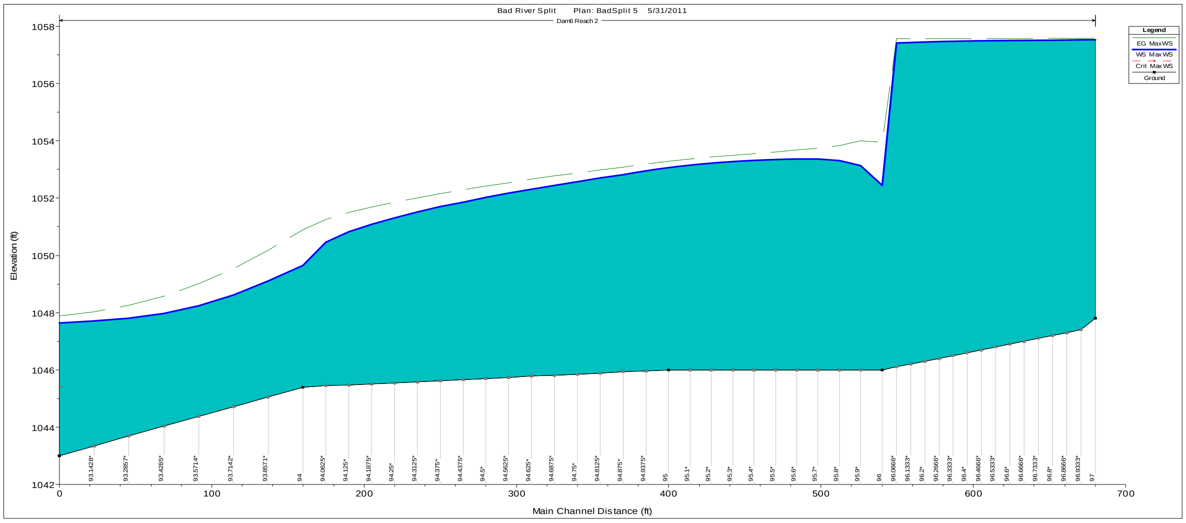

Profile 1: Behaving Badly at the Crack in the Ground.

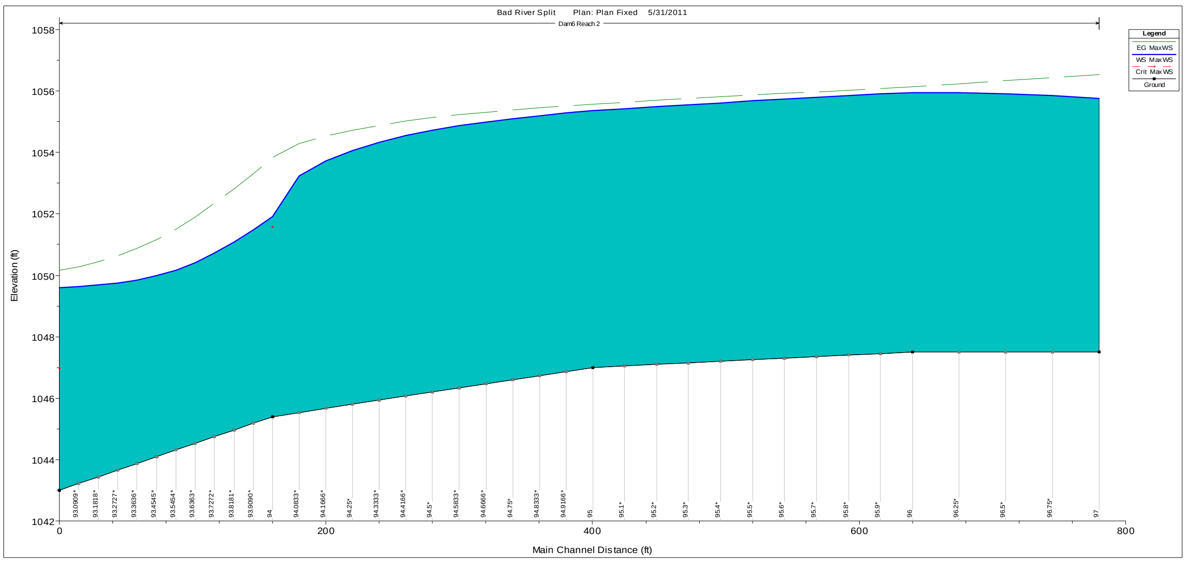

Profile 2: Ignore the Crack, Get Better Downstream Results.

If we are only interested in results at the downstream end of the reach, how do we know that the results corresponding to Profile 2 are more reliable than the first set? Might it not be better to live with a computational blip upstream of the Crack than to misrepresent the earth's surface just to make the model run? In this case, the answer is no. A review of the hydrograph entering and leaving the simplified model reach shows that the hydrograph has a classic, sharp-peaked shape and passes through the reach largely unchanged. In other words, the *worst* effect the simplification could be having is to eliminate any attenuation that might be caused by the Crack.

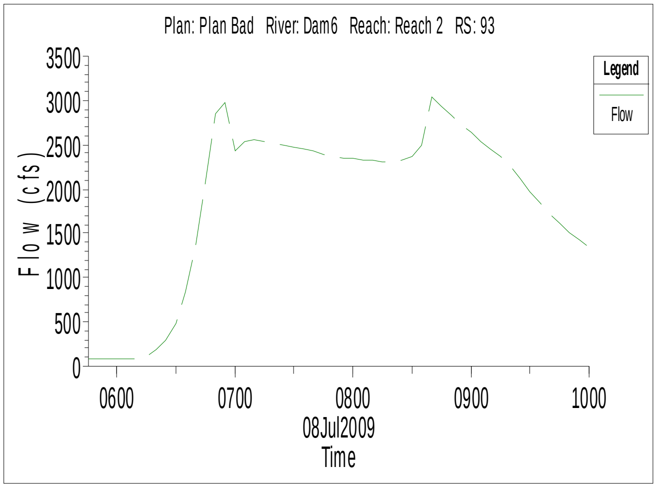

Hydrograph Leaving Model Reach –
Crack is Still in Model (see Profile 1)

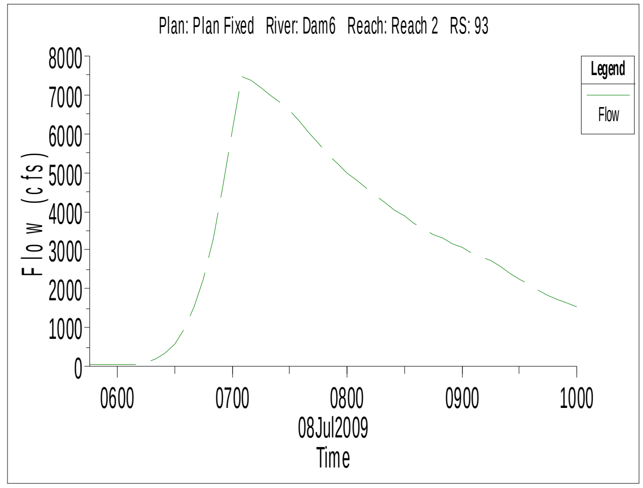

Hydrograph Leaving Model Reach –
Crack Removed From Model (see Profile 2)

In contrast, the hydrograph corresponding to Profile 1 is clearly affected by computational instability. Secondly, the modeled reach is just downstream of a divergent junction, and the effect of creating an artificially heightened water surface at the upstream end was to send far too much water down the other branch, away from the campground. This difference can be seen in both the hydrographs (note the scale difference) and the maximum water surface profiles.

You are the Master of the Initial Conditions.

Running an unsteady flow model can feel like using a windup toy. You set it up the way you want it, then let fly. If it careens off course or crashes into another object, you need to set it up differently. The initial conditions (the hydraulic properties the model calculates at time zero) are a very good starting point for trouble-shooting, because they separate problems related to basic hydraulic geometry from computational issues arising from rapidly changing flow conditions. The initial conditions can be checked by examining the computed profile at time zero, running the model with a short, steady flow hydrograph, or running the steady-state version of the model. If they are not well-behaved or do not meet the modeler's general expectations, there is little point in continuing with the model run. One important thing to check is whether reservoirs and dams are being modeled as expected: is the starting reservoir elevation realistic? Is the outflow what it should be? If the initial reservoir outflow is not the same as the inflow, the reservoir level will rise or fall as soon as the computation begins. Is that the condition you meant to represent?

The initial conditions are one type of boundary condition. Other boundary conditions are also fixed by the user and need to be carefully controlled. The physical locations of the upstream and downstream boundaries establish the model's entire universe. If the reach above the dam is being represented by cross sections (as opposed to a storage area) the upstream cross sections need to contain all of the stored water that could be released by the breach. If there is a feature (such as a lake or a bridge) that affects flow conditions at the downstream end, it needs to be represented either within the model, or as the downstream boundary. The downstream boundary is almost always an approximation. If there is any concern that it might not be a very good one, a sensitivity analysis is in order. For example, with a normal depth boundary condition, the user can check to what degree changing the friction slope affects the model results in the reach of interest.

The Output is Not the Answer.

The take-home messages from all of the preceding text are as follows:

- Skepticism is good.
- What the modeler could figure out with no model is probably generally correct, if imprecise.
- If it looks funny, it is probably wrong – unless it can be explained by a thoughtful review of the physical situation.

Beyond this, the modeler has the responsibility to interpret the model output and report both the output *and* the interpretation. One important function of the interpretation step is to acknowledge the uncertainty in the analysis, recalling that precision and accuracy are not the same thing. This author is virulently opposed to reporting results with excessive and misleading precision, unless there is a regulatory requirement to do so (as in many floodplain applications).

There will always be locations in the model where the results are less reliable than elsewhere. In the example above involving the Crack in the Ground, the model results within the artificially simplified model reach should not be reported without a clear disclaimer about what they represent. If the model output indicates that the water surface elevation immediately upstream of a bridge is slightly lower than that downstream, it would usually be better to report an upstream elevation based on the energy grade line than to insist on reporting an adverse slope through the bridge. Well-reasoned and well-articulated documentation of this decision is what makes the difference between “interpretation” and “fudging.”

In reporting results as well as in developing, running, and checking the model, it is often too easy to let the model call the shots – particularly when it has been a long and frustrating battle to get any results at all. Perhaps the best antidote is to ask oneself if the results, as reported, could be confidently explained to the average homeowner, village board, or 13-year old. If the answer is no, we can do better.