

CFD Project Workflow Guide

Contents

Select a problem with known results for proof-of-concept testing	1
Set up and run a coarse test case	2
Select and calibrate numerical methods	3
Minimize & quantify the mesh dependency.....	4
Quantify the validation uncertainty.....	5
Calibrate mesh-dependent input parameters	6
Model the case of interest.....	7
Report results.....	7

This project workflow guide is not an exhaustive decision-tree for all elements of a project. Rather, it is meant to be a checklist with a middling level of detail. Some possible project elements have been omitted (generating geometry, for example). The Flow Science support staff hopes you find this guide helpful. Please send suggestions to support@flow3d.com.

Select a problem with known results for proof-of-concept testing

Whenever you must model a new class of problems using CFD, it is good practice to check whether you can set up a CFD model that gives expected results. To check, you will need a case that is representative of the problem of interest and that has known results. Once you have tested one or more representative cases for a problem type you can skip this step for new problems of the same type.

Check the online resources for similar cases

Some classes of problems (ogee spillways, for example) have already been thoroughly validated against physical experiments with **FLOW-3D**. You can search to see how other users have modeled the same type of problem:

- www.flow3d.com > Resources > Bibliography > html search (Ctrl-F) for keywords
- www.flow3d.com > Site Search (upper right input box) for keywords
- users.flow3d.com > Conference Proceedings, Validations, Example Simulations, Technical Notes (for cases related to specific physics models), and Training (for tutorials)

Compare the test case to the case of interest

- How were the test case experimental results collected?
- Can experimental uncertainty be quantified for the test case?

- Are the test case results from a reliable source?
- Does the test case include the same physical processes as the case of interest?
- Does the test case have similar physical and time scales as the case of interest?
- Are the dimensionless numbers (Re , Fr , We , etc.) similar between the cases?

A test case may still be valuable even if it doesn't meet all of the above criteria, but there will be relatively less certainty in how to parameterize the case of interest.

Set up and run a coarse test case

Follow the simulation setup guide using the coarsest possible mesh

This setup mesh is purposely kept *as coarse as possible* to make it run *as fast as possible*. After completing the simulation setup guide (on this site) you will have:

- An annotated diagram with all necessary information in a consistent unit system,
- A simulation folder with:
 - all *.stl* files (checked and repaired)
 - a ready-to-run *prepin* file
- A mesh that resolves the thinnest flow sections with 3 or 4 cells,
- A mesh that follows best-possible-practices,
- Rationales for all specified physics models and coefficients.

Run and post-process the coarse case

If deadlines are pressing, only run a portion of the case. Answer the following questions:

- Are the *solver warnings*  acceptable?
- Do the *solver text* and *runtime plots* indicate that the time step sizes, pressure iterations, mean and turbulent kinetic energies, and fluid volume are appropriate?
- Does the model run in a reasonable amount of time?
- Is it possible to get the necessary outputs from post-processing?
- Check *3-D* and *2-D* plots. Do the results look right?
- Check *Probe* and *Text* outputs. Do the results follow the expected trends?

If the answers to the above questions are yes, the simulation is configured adequately. If any answers are no, consider adjusting the simulation setup with the following checklist:

Troubleshoot the coarse case if necessary

- Does the mesh follow best-possible-practices?
- Are geometry files checked and repaired and well-resolved by the mesh?
- Are the boundary conditions appropriate and realistic?
- Are the initial conditions appropriate and realistic?
- Do the boundary conditions and initial conditions match?
- Are the active physics models appropriate?

- Are physics model and component coefficients correct? Check units and values carefully.
- If all of the above considerations have been checked *carefully* and the problem is not resolved, refer to the *User Manual > Troubleshooting* chapter.

Continue adjusting and re-running the coarse case model until you are confident in the simulation and ready to commit to longer run times.

Select and calibrate numerical methods

In this step, you will test the model sensitivity to different numerical methods. Whenever a higher-order method gives significantly different results, you must make a choice as to whether or not to use it. This step is optional but recommended whenever the coarse test runs quickly and there is sufficient time in the budget. Test as many of the following as you can.

Check sensitivity to the VOF method

Models with free surfaces have several options for volume-of-fluid methods. The default option includes the most robust surface reconstruction method, but higher-order methods are available that may be more accurate. In particular, if you have a curved free surface that moves (*e.g.*, waves), a higher-order option may be desirable. A guideline is to try Split Lagrangian VOF and compare the results to the default One Fluid, Free-Surface VOF. If your output doesn't change significantly, then use the more robust and faster default option.

Check sensitivity to the momentum advection method

Higher-order momentum advection methods are generally more accurate, but may be less stable and/or generate more noise in the solution.

- If your model has a free surface, try 2nd-order, monotonicity-preserving momentum advection. If there is no free surface, try pure 2nd-order momentum advection.
- Compare the results against the default 1st-order method, and if the results are significantly different use the higher-order method.

Special considerations:

- 1st-order should be used with wave boundaries (along with Split Lagrangian VOF).
- There is some anecdotal evidence that 1st-order momentum advection may give more accurate diffusion for highly turbulent flow than 2nd-order, monotonicity-preserving momentum advection. This has not been demonstrated for all cases.

Check sensitivity to recommended implicit methods

Implicit methods are indicated whenever they will result in a significant speed-up of the run time. The *mentor*, for example, suggests an implicit method when the speed-up will be at least a factor of 10. If an implicit method will not result in a significant speed up, then explicit methods are preferred because they are more reliably accurate. Some physics models (*e.g.*, *Moving Objects*) may need implicit methods for stability.

- Compare results from indicated implicit methods to results from explicit methods.
 - If the results are sufficiently similar and the implicit method(s) are noticeably faster, consider using the implicit method for the rest of the project.
- Check sensitivity to boundary conditions**
- Move the boundary location (mesh block extents) away from the region of interest and re-run the coarse case until the results do not vary significantly.
 - Each expansion should be an even multiple of the cell size in that direction. Check for each expansion that the cell size remains the same.
 - Take the smallest model that is free from boundary effects as the working case.
 - Turbulence at inlet boundaries can affect some models. Try a large turbulence intensity (e.g., 5%) at the boundary to see if it has an effect. If it does, use values that correspond to physically-realistic turbulence intensities at the boundaries.
- Check results sensitivity to physics model coefficients**
- Physics models may include *numerical coefficients* (as opposed to physically-based *input values*) that affect the results. The default settings of these coefficients are often acceptable, but can be calibrated to better match experiment.
- For example, the two-equation turbulence models use a purely numerical *maximum turbulent mixing length* parameter (*TLEN*). Test the sensitivity of the results to constant vs. dynamic *TLEN*. Start with 7% or 10% of the controlling length scale (often max fluid depth) and increase the value (e.g., to 15%, 30%, 50%, 100%) until the results cease to vary significantly. Take the setting that gives results that are closest to the experiment.
 - Repeat the process for any other empirically-based physics models that have numerical knobs that can be shown to affect the model results.

Minimize & quantify the mesh dependency

Your model is working as expected and has been proven to run. Now you can begin running the case on more accurate meshes. Running the same case with different mesh resolutions demonstrates that the solution converges when the mesh and time step are refined. This step is sometimes called *solution verification* and is discussed at length in ASME (2009) and (ASCE) 2009. This step is optional, but highly recommended when time and resources allow it.

Add copies or restarts with refined meshes

The easiest way to do this is to refine the cell sizes of *all mesh blocks*, in *all directions*, by the *same factor*. A refinement factor of 2 is common (increasing the cell count by a factor of 8), but any ratio larger than 1.3 is acceptable. A *minimum* of three grids is recommended.

- Create a copy or a restart copy of the simulation. If it ran to steady state, then a *restart* copy with a shorter finish time is recommended for faster results.
- Refine the mesh uniformly in all directions by the selected refinement factor.

- Adjust the finish time (if desired). The adjustment (if any) depends on (1) how long it took the coarse case to reach steady state and (2) if the new simulation is a *restart*.
 - Repeat the process to create a third case with an even finer mesh. The third case may be set up as a *restart* from the second case even if the second case has not run yet.
- Run and post-process the finer-mesh cases**
- Add all the simulations to the run queue and run them in order.
 - Post-process the results. Use the same variables of interest as in the coarse test case.
- Calculate the mesh-sensitivity of the results**
- The results should be converging on a constant value. Demonstrating this means the solution is “verified”. See ASCE (2009) and ASME (2009) for details and calculations.
 - A simple guideline: the model is “mesh-converged” when refining the mesh causes the output to change less than 3% (for output with magnitude greater than 1000) and less than 1% (for output magnitude less than 1000). This is a general guideline, and not always appropriate. Use more rigorous methods if in doubt.
 - The design mesh is either the converged mesh or the finest mesh that is feasible.

Quantify the validation uncertainty

This step quantifies *input*, *numerical*, and *experimental uncertainty* for cases that involve experimental data. This step is optional. It is included here because it results in 95% and standard-deviation confidence intervals which are more defensible than simple error comparisons. The interested user is referred to ASCE (2009) and ASME (2005, 2009) for guidance and details.

Numerical uncertainty

The results from mesh-dependency study described above can be used to calculate *numerical uncertainty*, *order of convergence*, *Grid Convergence Index*, and *extrapolated results*.

Input uncertainty

Input parameters represent physical reality, which is to some extent unknowable. The way to deal with this is to quantify the uncertainty of the CFD results due to the uncertainty in the input parameters. This means estimating the distributions of likely input values, picking values from the distributions, and recording and interpreting the results when those values are used. This is time-consuming. It is preferable to find input uncertainty using the design mesh, but using a coarse mesh is much faster.

Experimental uncertainty

You must make some estimate of the experimental data uncertainty whenever you want to compare CFD results to experimental results. Standard methods exist for quantifying experimental uncertainty; most involve estimating *measurement uncertainties* and combining

them to find *result uncertainties*. Detailed experimental uncertainties are ideal, approximate evaluations (e.g., $\pm 2\%$) are common in practice.

Validation uncertainty

Validation uncertainty is the *estimated standard deviation of the parent population of the combination of errors* (ASME 2009), that is, a function of numerical, input, and experimental data uncertainty. The uncertainty u_{val} can be compared to the disagreement $|E|$ between the CFD model and the experiment:

- Wherever u_{val} is greater than or close to $|E|$, the disagreement is within the level of “noise” of the experiment, and the model can be considered “good”.
- When the disagreement $|E|$ is greater than u_{val} , it suggests that the model setup could be improved (e.g., by calibrating mesh-dependent parameters as described below) or that the CFD itself is of limited reliability (e.g., turbulence models).

Calibrate mesh-dependent input parameters

CFD models often include some variables that must be calibrated *after* the design mesh is selected. An example is component surface roughness: when the physical roughness is used the results become realistic as the mesh is refined, but usually require resolutions that are not time-feasible. The usual working method is to instead use either the converged mesh or the best possible mesh and then *calibrate* component roughness to whatever value minimizes the difference between CFD and experiment. This value may not be physically realistic!

Select parameters to calibrate or keep fixed

When possible, adjust only one parameter for each physics model (e.g., *surface roughness* for *turbulence*, *Critical Shields Parameter* for *sediment scour*, *rate coefficient* for *air entrainment*). Numerical options that were justified in previous steps should stay fixed.

Vary the calibration parameters and record the results

When possible, adjust only one parameter per run. Use restart simulations to minimize the run time, if possible. Record the results and compare them to the experiment.

Check for trends

This helps reduce the amount of guesswork that goes into selecting the next calibration value. For example, results for three roughness values can be fit logarithmically to predict the value that will minimize the model/experiment disagreement.

Keep the values that minimize disagreement

It's important to remember that these design values are valid only for the given mesh, geometry, flow velocity, and dimensionless numbers (e.g., Re , Fr , We , k^+). The parameters may need to be recalibrated for each new case, unless they are similar. If the calibrated values are far from physically-realistic values, be prepared to explain why (e.g., the physics model used is

mesh-dependent, and a physically-realistic input value requires a mesh resolution that is not feasible for the project time constraints).

Model the case of interest

Use everything you have learned from the test case.

- Start with the coarsest-possible mesh**
- Follow the simulation setup guide (on this site) and prove the case runs**
- Use the numerical options and input parameters from test case**
- Refine the mesh to minimize error**
 - If the test case and case of interest are of similar scales and dimensionless numbers, use the design mesh from the test case immediately.
 - If no test case was run first, or if the test case is significantly different from the design case, then perform a mesh dependency study as described above for the test case.
- Post-process the results and check that they appear reasonable**

Report results

- Describe the physical situation and why CFD modeling is needed**
- Describe the test case & report results**
 - Describe the physical situation and experimental data.
 - List any non-standard numerical methods (and note if only default values were used).
 - Describe the selection of boundary conditions and their locations.
 - Describe the mesh dependency study, including sizes used and convergence results.
- Describe the case-of-interest & report results**
 - Report model options, whether they were the same or different as in the test case.
 - Report results for the variables of interest from the design mesh resolution.
 - Justify the design mesh resolution (e.g., mesh dependency results or prior experience).
 - Report any conclusions drawn or lessons learned from the simulation.
- Report validation uncertainty (optional)**
 - Report estimated *numerical, input, and experimental uncertainty* and methodology.
 - Report any additional calculated verification and validation results (e.g., GCI).
 - If no uncertainties were estimated, justify the omission. Rigorous uncertainty estimation is an ideal that is not always necessary or feasible in practice (e.g., when CFD is used to qualitatively compare potential designs).