


Runtime Optimizations Workshop

1 Objective



The goal of this workshop is to take a base hydraulic and sediment model that runs slow and has some issues and improve it so that it runs faster and more stable.

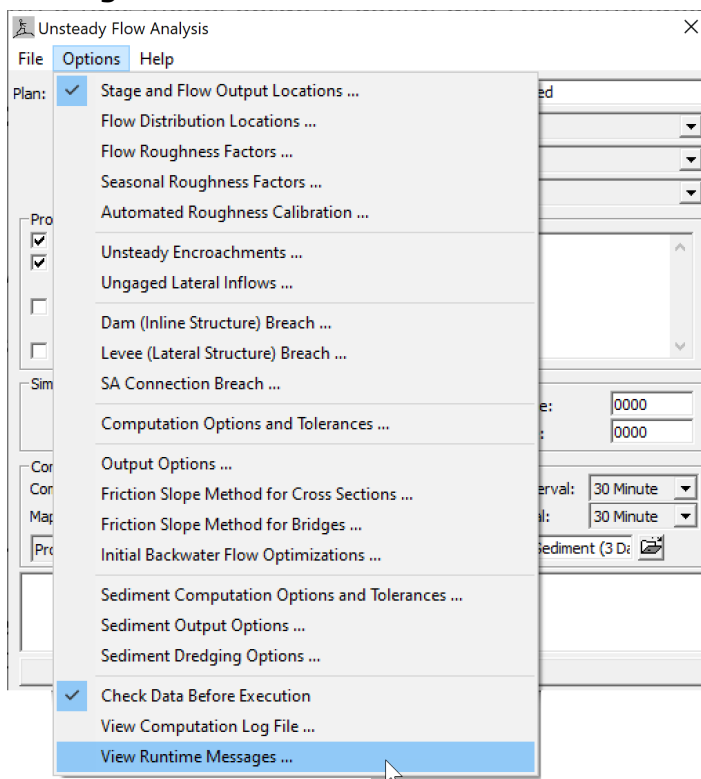
2 Introduction

The model is based on a small reach of the Arkansas River. For the purposes of the workshop, the model is setup is relatively simple and fast. The model extent is relatively small model, and the mesh has a coarse resolution. The computational mesh, boundary condition lines, and terrain are shown in the figure below.



The base model has a constant flow upstream of 70,000 cfs, and a constant stage downstream of 172.936 ft. It is important to note, that though the current model setup may not be good for final production runs of a project, the sensitivity analysis as applied here with a fast model setup is useful and recommended for projects.

1. **Start HEC-RAS** .
2. **Open** the HEC-RAS project names "ArkansasPool2.prj" in the "Advection-Diffusion" Workshop folder.
3. **Open** the **Unsteady Flow Analysis**  editor, and **Run** the plan called "Base Sed".
4. **Do NOT** run the plan. **Click** on the menu **Options | View Runtime Messages...** as shown below



5. The message window should look something like this

```

Geometric Preprocessor HEC-RAS 6.x Dev
Finished Processing Geometry
Writing Event Conditions ...
Completed Writing Event Condition Data

Performing Unsteady Sediment Flow Simulation HEC-RAS 6.x Dev

Unsteady Input Summary:
  2D Unsteady SWE-EM Equation Set (stricter momentum)
  2D number of Solver Cores: 6

Sediment Input Summary:
Transport Function:
Sorting Method: Active Layer
Fall Velocity Method: Soulsby
Global deposition/erosion methods:
  Soulsby

Overall Volume Accounting Error in Acre Feet: 0.000000
Overall Volume Accounting Error as percentage: 0.000000
Please review "Computational Log File" output for volume accounting details

Writing Results to DSS
Finished Unsteady Flow Simulation

Computations Summary

| Computation Task                         | Time(hh:mm:ss) |
|------------------------------------------|----------------|
| Completing Geometry                      | 1              |
| Preprocessing Geometry                   | <1             |
| Completing Event Conditions              | <1             |
| Unsteady Flow Computations with Sediment | 10:12          |
| Complete Process                         | 10:14          |



| Computation Speed                        | Simulation/Runtime |
|------------------------------------------|--------------------|
| Unsteady Flow Computations with Sediment | 1410x              |
| Complete Process                         | 1405x              |



```

- At the bottom of the messages, the run time for the unsteady flow and sediment as 10:12. This is how long it took to run a 10-day simulation. Dividing the simulation time by the runtime gives the computation speed of 1410x shown above. This is a good number to keep track of when optimizing the model performance because it can easily be used to compare the speed of different simulations of different lengths. It can also to compute how long a simulation will take to compute.

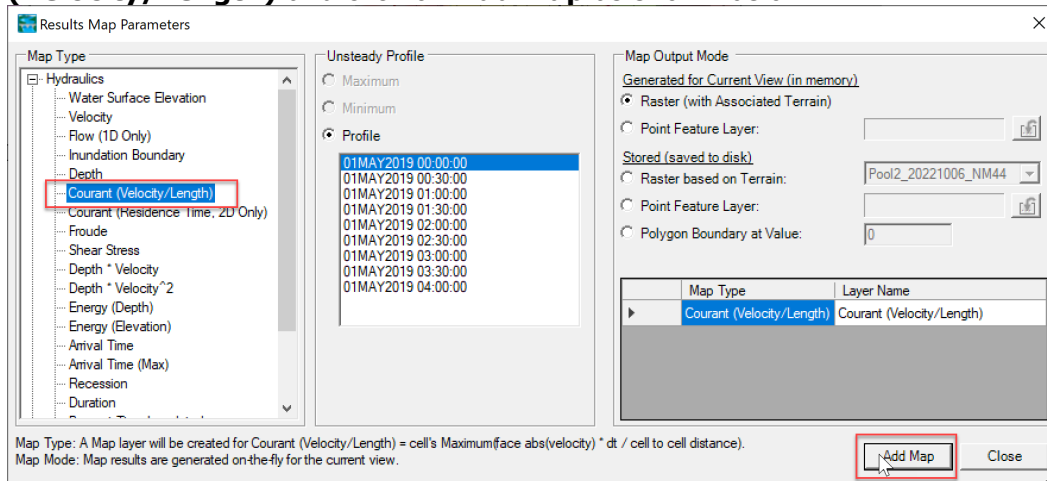
3 Flow Model Optimization

Before optimizing the sediment transport model, the flow will be optimized separately. It is often useful to optimize the model setup using a short simulation period. In this case a short steady state simulation is run to optimize the hydraulic simulation.


7. **Save As...** a new plan called "Flow Only".
8. **Change** the simulation time window to be 4-hours long. This should be sufficiently long for a steady simulation and long enough to compute a representative simulation computational speed.
9. **Run** the plan.
10. **Read** the simulation runtime and computational speed. This should be approximately **15-16 seconds** and **900x**, respectively.
11. **Add** the **Courant** number output layer by right-clicking on the Results layer **Flow Only** and selecting **Create a New Results Map Layer ...**



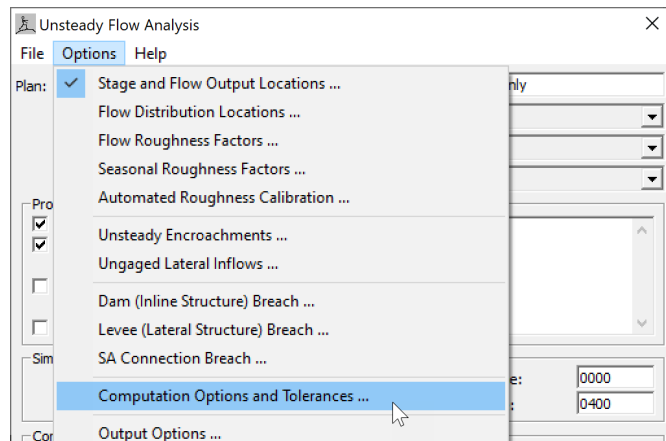
In the **Results Map Parameters** window, find and select **Courant (Velocity/Length)** and click on **Add Map** as shown below.



12. The Courant number layer should look something like the figure below



13. From the **Courant** results, it is clear that the time step is too small and can be increased by a factor of 3 to 4.
14. **Set** the time step to **30 seconds**, rerun the plan, and check how long the simulation takes to run. The simulation should take about **7 seconds** and have a computational speed of about **1900x**.
15. The next time is to optimize the computational options and parameters, In the **Unsteady Flow Analysis** editor, **click** on the menu **Options | Computation Options and Tolerances ...** as shown below.



16. The **Equation Set** is set to the **SWE-EM** solver, which is the slowest solver. For purposes of the workshop, the **SWE-ELM** should be fine. It is generally recommended to start with the SWE-ELM for river applications and only test the **SWE-EM** solver for comparison purposes. For most river applications, the two solvers will produce very similar results if the resolution and time step are appropriate. However, the **SWE-ELM** is generally faster and more stable. For these reasons, the **SWE-ELM** solver will be used here.

HEC-RAS Unsteady Computation Options and Tolerances

General | 2D Flow Options | 1D/2D Options | Advanced Time Step Control | 1D Mixed Flow Options

Use Coriolis Effects (not used with Diffusion Wave equation)

	Parameter	(Default)	Pool2
1	Theta (0.5-1.0)	1	1
2	Theta Warmup (0.5-1.0)	1	1
3	Water Surface Tolerance [max=0.2](ft)	0.01	0.01
4	Volume Tolerance (ft)	0.01	0.01
5	Maximum Iterations	20	20
6	Equation Set	Diffusion Wave	SWE-EM (stricter momentum)
7	Initial Conditions Time (hrs)		Diffusion Wave
8	Initial Conditions Ramp Up Fraction (0-1)	0.5	SWE-ELM (original/faster)
9	Number of Time Slices (Integer Value)	1	SWE-EM (stricter momentum)
10	Turbulence Model	Non-Conservative (original)	SWE-LIA (local inertia)

17. This dataset is relatively small, so other computational options don't make a big difference. For larger models
18. **Run** the new plan setup. The model should take about **1 second** to run. The model is now ready to add sediment transport

Question: What other kinds of model parameters and tolerances can be adjusted to speed up the flow model?

4 Sediment Model Optimization

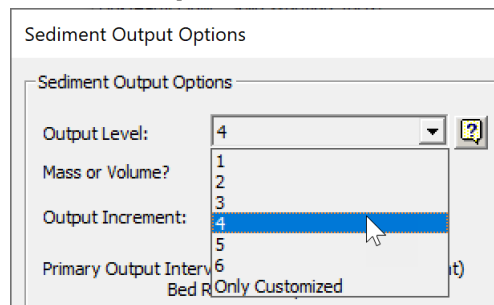
The next step is to turn sediment back on and optimize the sediment transport model. A large part of the sediment model optimization can be done with shorter simulations and some cases, smaller domains, and coarser grid resolutions.

19. **Save As** a new plan called "Sed 1day"
20. In the **Unsteady Flow Analysis** window.
21. **Check** the box for sediment transport.
22. **Change** the simulation window to be 1 day.
23. **Run** the plan as is and record the runtime and computational speed. It should take about **36 seconds** and have a computational speed of about **2347x**.

4.1 Sediment Model Convergence

Before continuing, it is useful to look at the model convergence, since a large part of the computational time can be spent on iterating within a time step. This is also important from a stability and model robustness point of view. The sediment model always outputs the number of iterations at the mapping output interval in the HDF5 output file (*.p##.hdf). However, to view this file, the user needs to either run a program to read the file or open it in a program such as HDFView. Another way of looking at the model convergence is through the **Computations Log File**.

24. **Open** the **Unsteady Flow Analysis** editor.
25. **Select** the menu **Options | Sediment Output Options ...**
26. **Set** the **Output Level** to **4**.



Note: For an Output Level of 4, only the status of the last iteration is logged. For Output Levels 5 or 6, the status of each iteration is written. For Output Level 3 (default) or lower, no iteration convergence information is written to the Computations Log File.

27. **Save** the plan
28. **Run** the plan.
29. **Open** the **Computations Log File** by opening the **Unsteady Flow Analysis** editor, and selecting the menu **Options | View Computation Log File ...**
The output for the very first timestep should look something like this:

```

STATUS: Sediment: Pool2, Time (hrs): 1.00000, Time Step (s): 30.0000
Transport Convergence:
Outer Convergence: Converged

```

Iter	State	MaxAbsCor(cell)	RMSCor	Units	Status
5	1	1.673E-03(484)	1.185E-04	mg/L	Continue
5	2	2.029E-05(484)	1.421E-06	mg/L	Converged
5	3	1.840E-05(484)	1.303E-06	mg/L	Converged
5	4	5.104E-07(484)	3.591E-08	mg/L	Converged
5	5	2.043E-08(484)	1.426E-09	mg/L	Converged

```

Active Layer Grain Class Fractions (%):

```

Iter	Grain	MaxAbsCor(cell)	Status
5	1	3.416E-04(482)	Converged
5	2	4.607E-05(482)	Converged
5	3	2.653E-04(482)	Converged
5	4	2.693E-05(482)	Converged
5	5	3.357E-06(482)	Converged

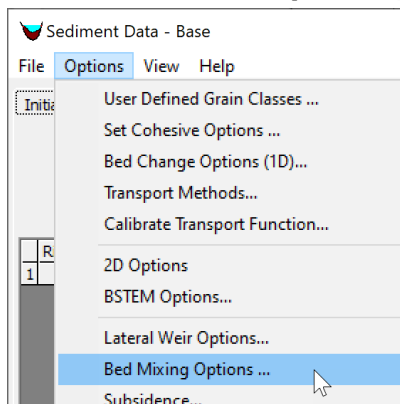
It is noted that the model went to maximum iterations and the first grain class (i.e. VFS) did not reach convergence. This could be avoided by increasing the maximum number of iterations. However, it is best to improve the model convergence and reduce the number of iterations. Inspecting the model convergence further, it is noted that the sediment model often went to 4 iterations.

For such a simple model as this one, the model is expected to have better convergence. This is an indication something about the model needs to be improved.

One of the common issues with model convergence is the failure to including hiding and exposure. The reason hiding and exposure improves model stability is because it decreases the erosion of the finer particles which represent a larger portion of the transport and therefore it tends to reduce the overall sediment transport and of finer particles especially (which is the one not converging in this case).

30. **Open the Sediment Data editor**

31. **Click on the menu Options | Bed Mixing Options.**



32. Select the **Wu et al.** hiding function and set the exponent to **0.6** as shown in the figure below and click **OK**.

33. **Save As** the sediment data with the name "Base HE" for hiding and exposure.
 34. **Save** a new plan with the Title and Short ID "1d HE"
 35. **Run** the plan.
 36. **Open** the **Computation Log File** by opening the **Unsteady Flow Analysis** editor, and selecting the menu **Options | View Computation Log File ...**
 The output for the very first two steps should look something like this:

```

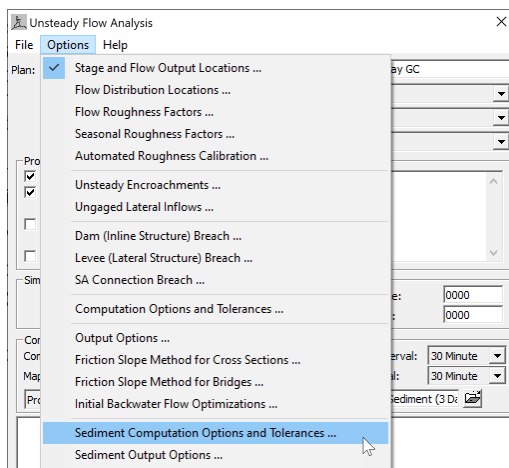
STATUS: Sediment: Pool2, Time (hrs): 1.00000, Time Step (s): 30.0000
Transport Convergence:
  Outer Convergence: Converged
  Iter  State  MaxAbsCor( cell)  RMSCor  Units  Status
  2     1     8.848E-05( 365)  1.608E-05  mg/L  Converged
  2     2     4.502E-05( 1023)  6.497E-06  mg/L  Converged
  2     3     1.427E-05( 485)  1.984E-06  mg/L  Converged
  2     4     1.356E-06( 484)  1.776E-07  mg/L  Converged
  2     5     1.404E-07( 484)  1.794E-08  mg/L  Converged
  Active Layer Grain Class Fractions (%):
  Iter  Grain  MaxAbsCor( cell)  Status
  2     1     1.531E-05( 363)  Converged
  2     2     9.069E-06( 1022)  Converged
  2     3     1.387E-05( 483)  Converged
  2     4     1.440E-06( 482)  Converged
  2     5     1.748E-07( 482)  Converged
  
```

Now all the grain classes are converging. Further inspections and comparisons between the convergence with and without the hiding and exposure shows that adding the corrections, significantly improved the model convergence and reduced the number of iterations by almost half.

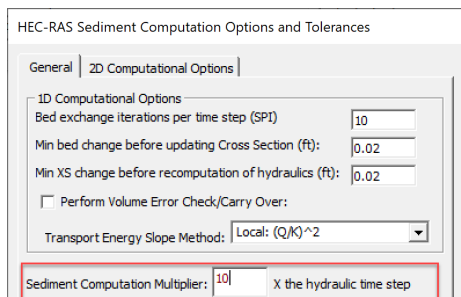
4.2 Sediment Computation Multiplier

The Sediment Computation Multiplier (CM) allows the model to use a larger time step for sediment transport compared to the hydraulic model. The sediment time step is computed as the hydraulic model time step times the CM. The default value for CM is 1, which means the hydraulic and sediment models have the same time step by default. The CM is a simple way to speed up the sediment model. However, the parameter can also make the sediment model unstable or cause the sediment model results to diverge from the results with a smaller time step. For this reason, the parameter should always be tested and verified against a base value of 1

37. **Save As** the plan with the Title and Short ID "1d HE CM".
38. **Open** the **Unsteady Flow Analysis** window.
39. **Click** on the menu **Options | Sediment Computation Options and Tolerances...**

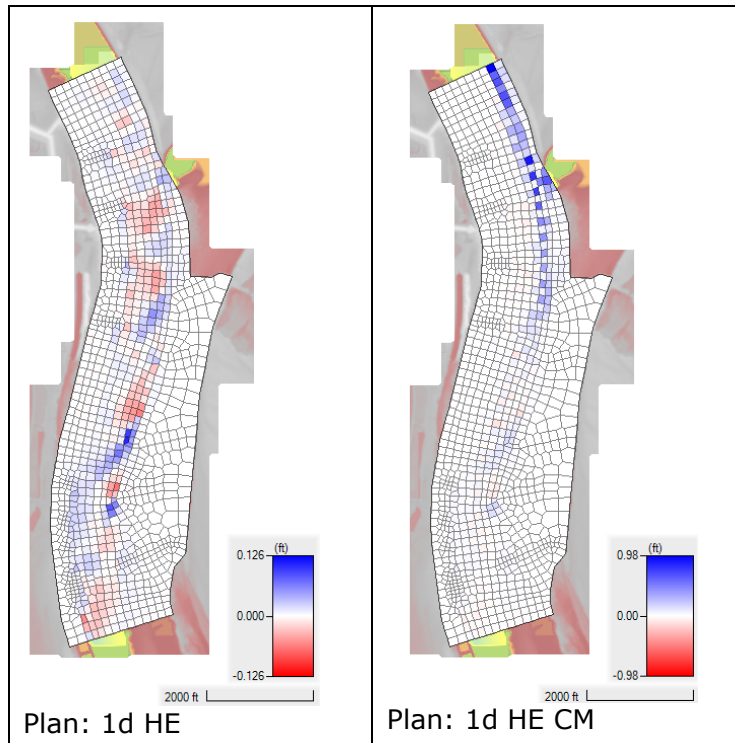


40. In the **General** tab, set the **Sediment Computation Multiplier** to 10.



41. **Save** the plan.
42. **Run** the model. The runtime should be about **14 seconds** with a computational speed of about **7443x**.

43. **Verify** that changing the CM, did not change the results significantly by comparing the **bed change** and other variables in **RAS Mapper**.



From the figure above it's clear that changing the computation multiplier caused the model to go unstable. This can be confirmed by looking at the **Computation Log File**. The figure below shows the convergence of the first two time steps.

```

External Equilibrium Boundary Condition:
Number of cells:      24

STATUS: Sediment: Pool2, Time (hrs): 0.50000, Time Step (s): 300.0000
Transport Convergence:
Outer Convergence: Converged
  Iter  State  MaxAbsCor( cell)  RMSCor  Units  Status
  5     1     4.340E-01( 355)  5.860E-02  mg/L  Continue
  5     2     1.709E-02( 355)  2.025E-03  mg/L  Continue
  5     3     2.783E-02( 355)  3.356E-03  mg/L  Continue
  5     4     1.671E-03( 355)  2.012E-04  mg/L  Continue
  5     5     1.677E-04( 355)  2.006E-05  mg/L  Converged
Active Layer Grain Class Fractions (%):
  Iter  Grain  MaxAbsCor( cell)  Status
  5     1     1.054E+00( 1080)  Continue
  5     2     1.557E-01( 1080)  Continue
  5     3     8.081E-01( 1080)  Continue
  5     4     8.016E-02( 1080)  Continue
  5     5     9.924E-03( 1080)  Continue

STATUS: Sediment: Pool2, Time (hrs): 1.00000, Time Step (s): 300.0000
Transport Convergence:
Outer Convergence: Converged
  Iter  State  MaxAbsCor( cell)  RMSCor  Units  Status
  5     1     5.893E-01( 1081)  7.573E-02  mg/L  Continue
  5     2     1.926E-02( 469)  2.413E-03  mg/L  Continue
  5     3     3.071E-02( 1081)  3.983E-03  mg/L  Continue
  5     4     1.816E-03( 1081)  2.379E-04  mg/L  Continue
  5     5     1.783E-04( 1081)  2.365E-05  mg/L  Converged
Active Layer Grain Class Fractions (%):
  Iter  Grain  MaxAbsCor( cell)  Status
  5     1     1.367E+00( 1080)  Continue
  5     2     2.027E-01( 1080)  Continue
  5     3     1.048E+00( 1080)  Continue
  5     4     1.037E-01( 1080)  Continue
  5     5     1.282E-02( 1080)  Continue

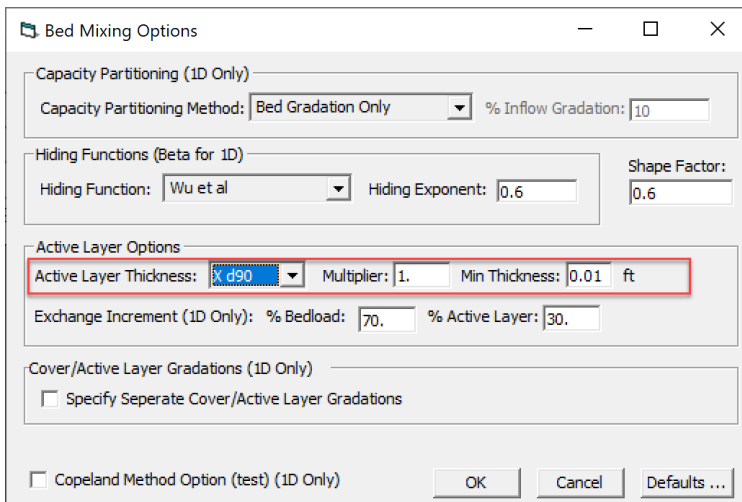
```

The sediment model reached maximum number of iterations and 4 of the 5 grain classes did not converge. Again, it's time to take a closer look at the model setup and identify any parameters which need to change or could be adjusted to improve model stability.

A common model parameter that affects the sediment model convergence is the active layer thickness. When the active layer is too thin, the whole layer can be eroded or exchanged within a single time step, and then can cause the convergence issues. The thicker the active layer is, the more stable the model will be because the active layer bed gradations will change more slowly.

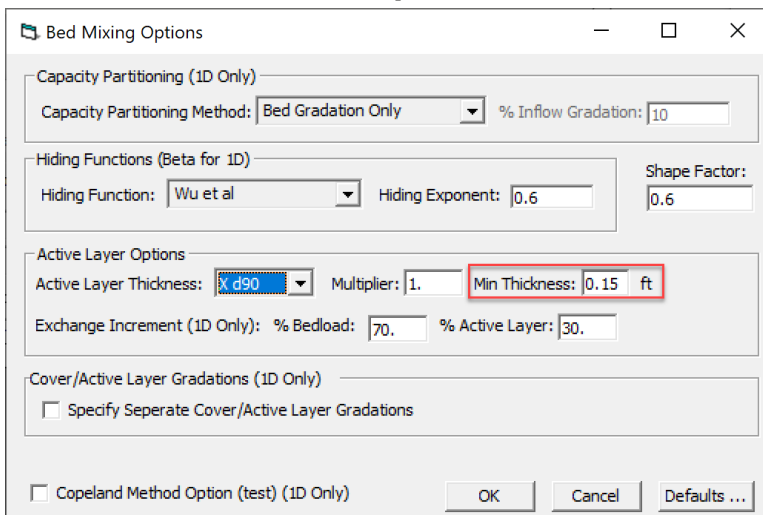
44. **Open** the **Sediment Data** editor.

45. **Select** the menu **Options | Bed Mixing Options...** the window should look like the figure below.



The active layer thickness is set to the d90 with a minimum value of 0.01 ft ~ 3 mm. This active layer is small and should be increased. The coarsest grain class has an upper limit of 2 mm so using the d90 multiplier is unlikely to make any significant difference unless it is to values in the range of 10 to 20. Instead, the minimum active layer thickness is increased.

46. **Set** the **Minimum Active layer Thickness** to **0.15 ft** ~ 5 cm.



47. **Save As** the **Sediment Data** with the name "Base HE AL".

48. **Save** a new plan with the Title and Short ID "1d HE CM AL".

Run the plan. The runtime should be about **8 seconds** with a computational speed of about **9735x**. Note that the computational speed increased by changing the active layer thickness because the sediment model is doing less iterations per time step. This illustrates the importance of good model convergence for computational speed.

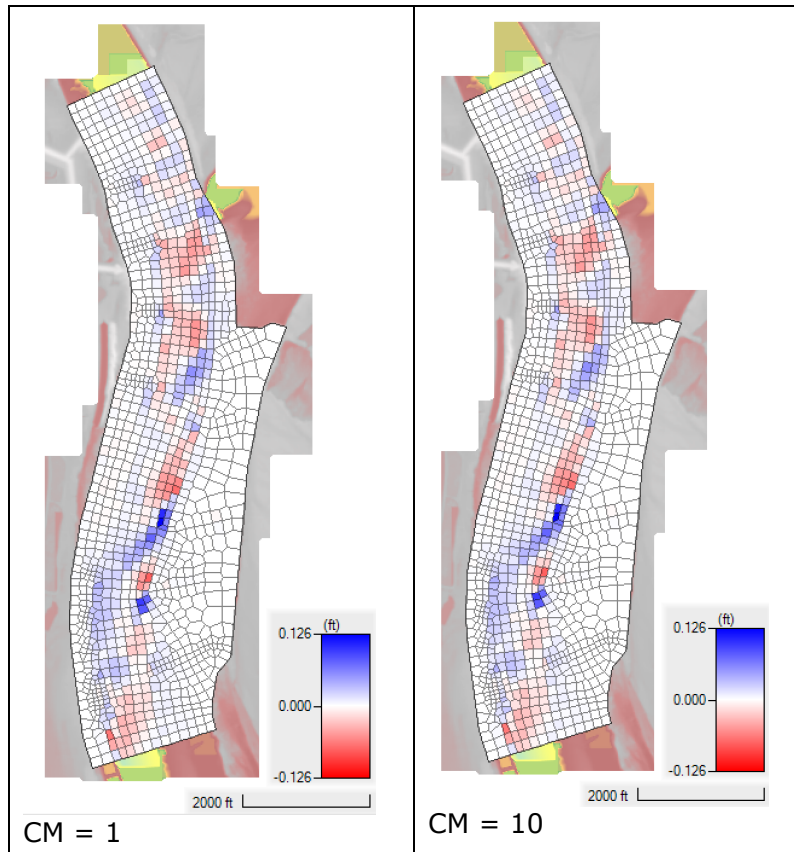
49. **Inspect** the model convergence in the **Computation Log File**. Below is an example for the first timestep.

```

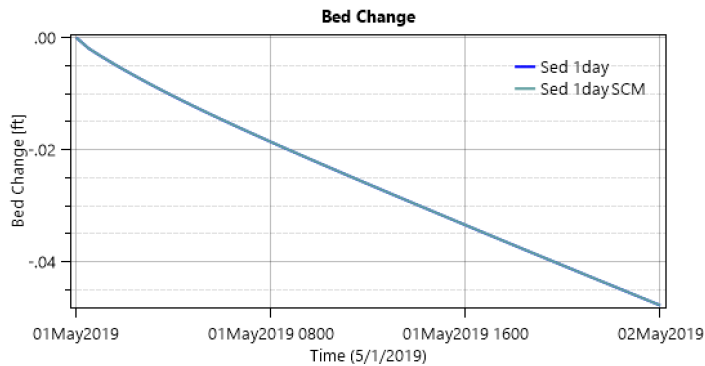
STATUS: Sediment: Pool2, Time (hrs): 0.50000, Time Step (s): 300.0000
Transport Convergence:
  Outer Convergence: Converged
  Iter  State  MaxAbsCor( cell)  RMSCor  Units  Status
  2     1     7.567E-04( 468)  9.875E-05  mg/L  Converged
  2     2     6.797E-05( 467)  8.690E-06  mg/L  Converged
  2     3     5.910E-05( 468)  7.685E-06  mg/L  Converged
  2     4     3.957E-06( 468)  5.104E-07  mg/L  Converged
  2     5     3.972E-07( 468)  5.045E-08  mg/L  Converged
Active Layer Grain Class Fractions (%):
  Iter  Grain  MaxAbsCor( cell)  Status
  2     1     5.148E-05( 1023)  Converged
  2     2     2.103E-06( 1024)  Converged
  2     3     4.552E-05( 1023)  Converged
  2     4     4.515E-06( 1023)  Converged
  2     5     5.568E-07( 1023)  Converged
  
```

Increasing the active layer thickness made a big difference in the model stability.

50. Below is a comparison of the bed change after 5 days, with a CM=1 and CM=10.



When verifying if a parameter has significantly changed the results, it's also useful to plot time series of variables at various locations. Below is an example of a bed change time series.



From these results it can be concluded that for the existing model setup, a CM of 10 did not significantly impact the accuracy of the results.


51. Now that the simulation is converging well, increase the simulation time window to 10 days.
52. **Run** the plan. The plan runtime should be about 1:00 with a computational speed of about **14000x**.

4.3 Grain Classes


Another thing that significantly impacts the runtime and model convergence is the number of grain classes. Each grain class requires solving another transport equation (i.e. unsteady advection-diffusion equation), bed sorting, and bed change equation. In addition, the grain classes are coupled through hiding and exposure and other processes. Having many grain classes can decrease the model convergence. Therefore, it is useful to reduce the number of grain classes to reduce the simulation runtime and the improve the model robustness. However, special care should be taken in combining grain classes or ignoring grain classes in areas of the domain which do not affect the area of interest.

The effect reducing the number of grain classes has on the runtime depends on the simulation and especially the sediment Computation Multiplier (CM). The larger the CM, the less effect reducing the number of grain classes has, because the sediment calculations become a smaller portion of the overall computation time. For this reason, a plan is first created without the CM, but with the bed mixing option changes. Then the number of grain classes is reduced from 5 to 3 to compare runtimes.


53. **Open** the **Unsteady Flow Analysis** editor.
54. **Open** the plan called "1d HE"
55. **Save As** the plan with the Title and Short ID "1d HE AL".
56. **Select** the Sediment file "Base HE AL". This sediment data has the hiding and exposure corrections and active layer thickness changes made previously. However, the CM is set to 1.0.



57. **Run** the plan. The runtime and computational speed should be approximately **30 seconds** and **2871x**, respectively.
58. **Open** the **Sediment Data** editor.
59. **Click** on the **Define/Edit Bed Gradation...** button.

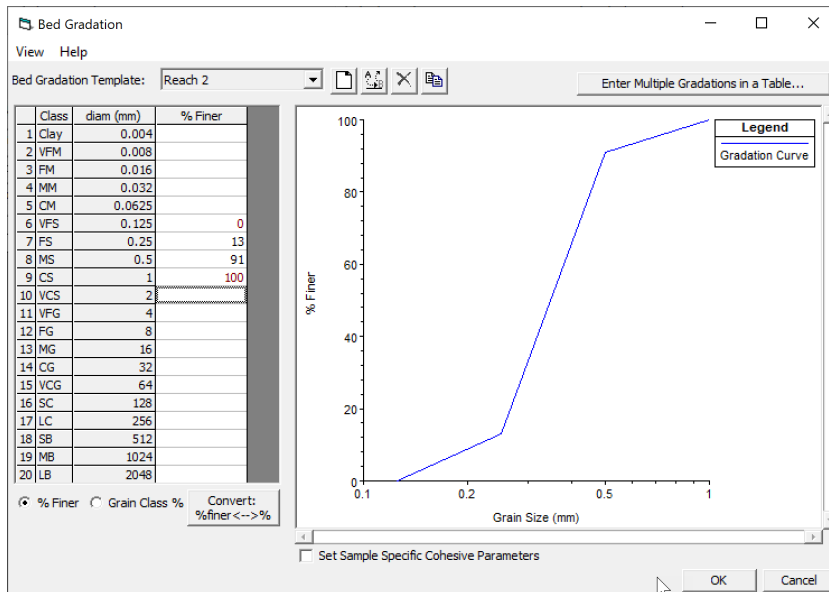


60. **Inspect** the **Reach 2** bed gradation which is used everywhere except for the non-erodible structures.



The sediment gradation has 5 grain classes. Two of these grain classes have 1% or less material (by weight). These grain classes can be ignored from the bed gradation.

- 61. **Save As** the sediment data as "Base HE AL GC". This sediment data file will contain the modified number of active grain classes. It's good practice to rename files in HEC-RAS before making changes to the data.
- 62. **Open** again the **Bed Gradation** editor.
- 63. **Remove** the VFS and VCS grain sizes.



- 64. **Save** the plan with the Title and Short ID "1d HE AL GC".
- 65. **Run** the model. The runtime should be about **24 seconds** with a computational speed of about **3489x**.
- 66. **Compare** the bed change and other variables for the model with 3 grain classes compared to 5. The differences in this case should be negligible.

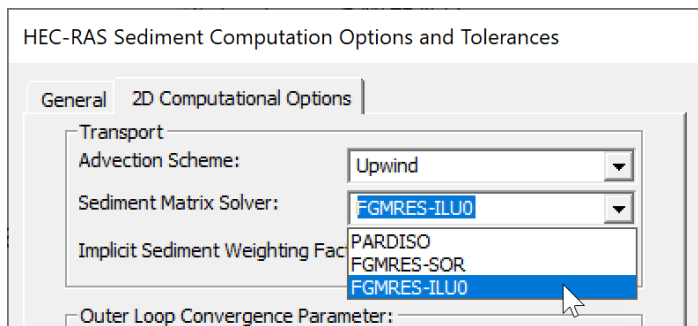
Note: For the purposes of the workshop, the comparison of using different number of grain classes is done with relatively short simulations. For a real-life project, this comparison should be done with a longer simulation comparable to the simulation length of the production runs.

Question: Since the number of grain classes is proportional to the simulation runtime, it seems reasonable to exclude grain classes which are part of the wash load since they do not contribute significantly to the bed change. However, this is not always the case. In what situations would it be useful to include wash load grain classes?

4.4 Iterative Solvers

A large portion of the computational time in the sediment model is spent solving a large sparse matrix related to implicit solution of the transport equation.

67. **Open** the **Unsteady Flow Analysis** editor.
68. **Save As ...** the plan with the Title and Short ID "1d HE AL GC IS".
69. **Click** on the menu **Options | Sediment Computation Options and Tolerances**.
70. In the **HEC-RAS Sediment Computation Options and Tolerances** editor, click on the **2D Computational Options** tab
71. In the **Transport** section, select the **FGMRES-ILUO** from the drop-down menu for the **Sediment Matrix Solver**, as shown in the figure below.



Run the plan. The runtime and computational speed should be approximately **18 seconds** and **4779x**, respectively.

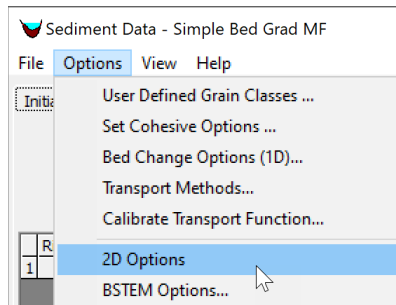
As the model performance improves, it is possible to run longer and longer simulations with reasonable runtimes. But before doing so, we'll reintroduce the sediment **Computation Multiplier**.

72. **Save As** a new plan with the name "10d CM HE AL GC IS".
73. **Change** the time window to **10 days**.
74. **Click** on the menu **Options | Sediment Computation Options and Tolerances**.
75. In the **General** tab, enter a **Sediment Computation Multiplier** of 10.
76. **Run** the plan. The runtime should be about **50 seconds** with a computational speed of about **17178x**.
77. **Check** the **Computation Log File**.
78. **Inspect** the results in **RAS Mapper** and compare with the previous plans.

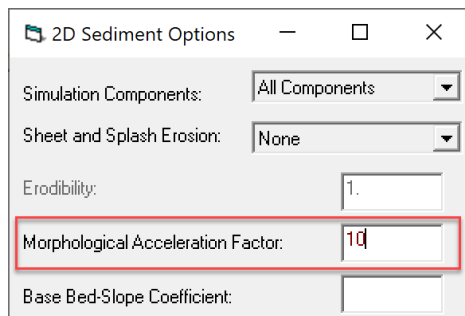
4.5 Morphologic Acceleration Factor

The Morphologic Acceleration Factor (MF) is a useful parameter to speedup sediment simulations. The MF can be used several ways, in this workshop, the goal is to simulate the bed change corresponding to design flow of 70,000 cfs for 10 days. By setting MF =10, 1 day of simulation time will represent 10 days of bed change. Since the upstream flow and downstream stage boundary conditions are constant, no changes are needed for these boundary conditions. Since the upstream sediment boundary condition is an equilibrium boundary condition, no changes are required for this boundary condition type either.

79. **Open** the **Sediment Data** editor.
80. **Save As ...** the sediment data as "Base HE AL GC MF".
81. **Click** on the menu **Options | 2D Options** as shown below.

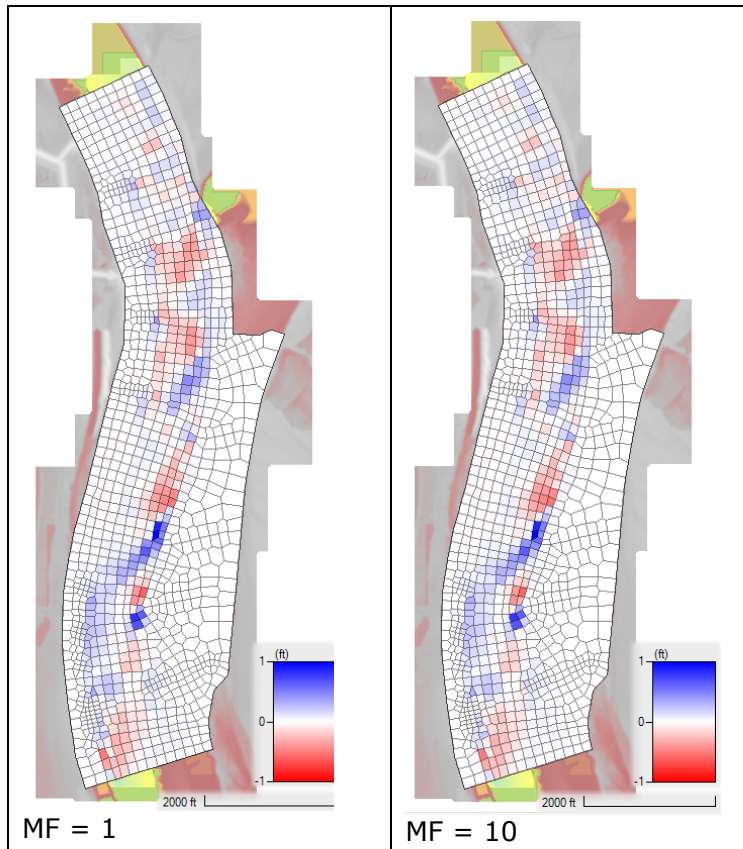


82. In the **2D Sediment Options** window, set the **Morphologic Acceleration Factor** to **10** as shown below.



83. **Save** the **Sediment Data**.
84. **Open** the **Unsteady Flow Analysis** window.
85. **Save** the plan with the Title and Short ID "10d CM HE AL GC IS MF".
86. **Set** the simulation window to be **1 day**. Note that since the morphological acceleration factor was set to 10, 1 day of model simulation time will represent 10 days of bed change.
87. **Set** the **Mapping Output** interval to **6 min**. This is so that the output is at an equivalent interval with the MF. 5
88. **Run** the plan. The simulation should take about **8 seconds** to run. When interpreting the computational speed it is necessary to multiply the number by the morphologic acceleration factor in this case. Therefore, the computational speed is approximately **103750x**.

89. **Verify** that the MF did not significantly affect the results by comparing the new results with plan "10d CM HE AL GC IS". Below is an example of the computed bed change. Below is an example of the bed change after 10 days with and without morphologic acceleration factor.



When comparing time-series of bed change of results which have the morphologic acceleration factor, it is necessary to scale time with the morphologic acceleration factor. This can be done by copy-pasting the data to excel and plotting it there. However, the final values of the bed change can also be used to quickly identify differences in the results due to the MF.

Compared to the base sediment run without any optimizations, this last run is approximately **44x** faster.