

Tasks Completed:

This week, a fair amount of progress has been made, despite running into a few roadblocks along the way. For Paraview, the Cmake precompiler was never able to find the location of the Qt executable, even if it was set to point to it manually. Too much time was spent on trying to figure out how to use Cmake, and in the end Paraview was just installed on the Windows 2000 virtual machine available on the computers. It would be ideal to have it installed on Linux, as the rest of the project is done in a Linux environment. However, the important thing is that it works. The data from the samples that were being run on the supercomputer was downloaded to my local machine to visualize. The output looked exactly like the sample output from the tests I ran locally, as well as the pictures that were in the Gale manual.

Another thing that was completed this week is that the issue that was causing tibet3D.xml to have a segmentation fault was isolated. For some reason, using the default solver provided by Gale does not work correctly for this example, as it requires a lot of memory. Instead, a direct solver must be used. To do this in parallel, the MUMPS package had to be installed to allow for parallel direct solving. This was a source for a lot of trouble, because apparently the installation of MPI on the supercomputer was not configured in such a way that the required packages for MUMPS, Scalapack and BLACS, were able to find the mpi include files from the arguments passed. At first, the installation of BLACS that PETSc downloaded was made to work by modifying the make file to add the include path for the MPI include files, but this installation was not complete for some reason, and PETSc/MUMPS was unable to install correctly. After running into this roadblock for quite a while, I eventually decided to just have PETSc download its own MPI so that it would know where all of the includes were located, and the installation was completed successfully.

Ben was able to provide both the 2d and 3d versions of his model to testing. The 2d model is currently running on the supercomputer for 1000 time steps, which should take quite a while. Hopefully it completes successfully and provides reasonable output. One thing that was noticed about PETSc is that it apparently runs operations that are by nature very memory intensive, and therefore each node should only be running one process at a time.

Tasks for Next Week:

Next week I hope to accomplish a few things. First of all, I need to get the direct solver to work. This could solve a few of the problems that I have been faced with. It should make the tibet3d, as well as any other 3d model. Also, I believe that it will lead to increased performance on multiple nodes. One of the issues that comes up with using multiple processors is that when the data set is split upon multiple nodes, it can make the GMRES solver used by Gale take much longer to converge. Using a direct solver, which is generally slower on larger sets of data, but faster among smaller sets of data, should allow for proper scaling when the simulation is run on 16, 32, and more processors. By the end of next week, provided the direct solver works, Gale should be operational for 3d and 2d simulations, as well as more scalable to run on more nodes.