nvprof and nvvp with mpi applications on Blue Waters (cuda or openacc)

Profiling cuda or OpenACC codes with nvprof requires some extra syntax on Blue Waters (and probably on other linux cluster instances).

The following was gleaned from : http://docs.nvidia.com/cuda/profiler-users-guide/index.html#mpi-nvprof , and https://bluewaters.ncsa.illinois.edu/openacc-and-cuda-profiling.

1. Construct a wrapper that will be executed by mpirun/mpiexec/aprun or your local MPI launch mechanism
    a. aprun -N 1 -n 2 ./wrap.sh # for the example below
2. The wrapper should invoke nvprof with appropriate options and then start your MPI executable
3. Use Nvidia's nvvp (https://developer.nvidia.com/nvidia-visual-profiler) to analyse the resultant profiles.

This example was built on Blue Waters based on the PGI tutorial at: http://www.pgroup.com/lit/articles/insider/v4n1a3.htm

PrgEnv-pgi and cudatoolkit modules were loaded for the build and at runtime.

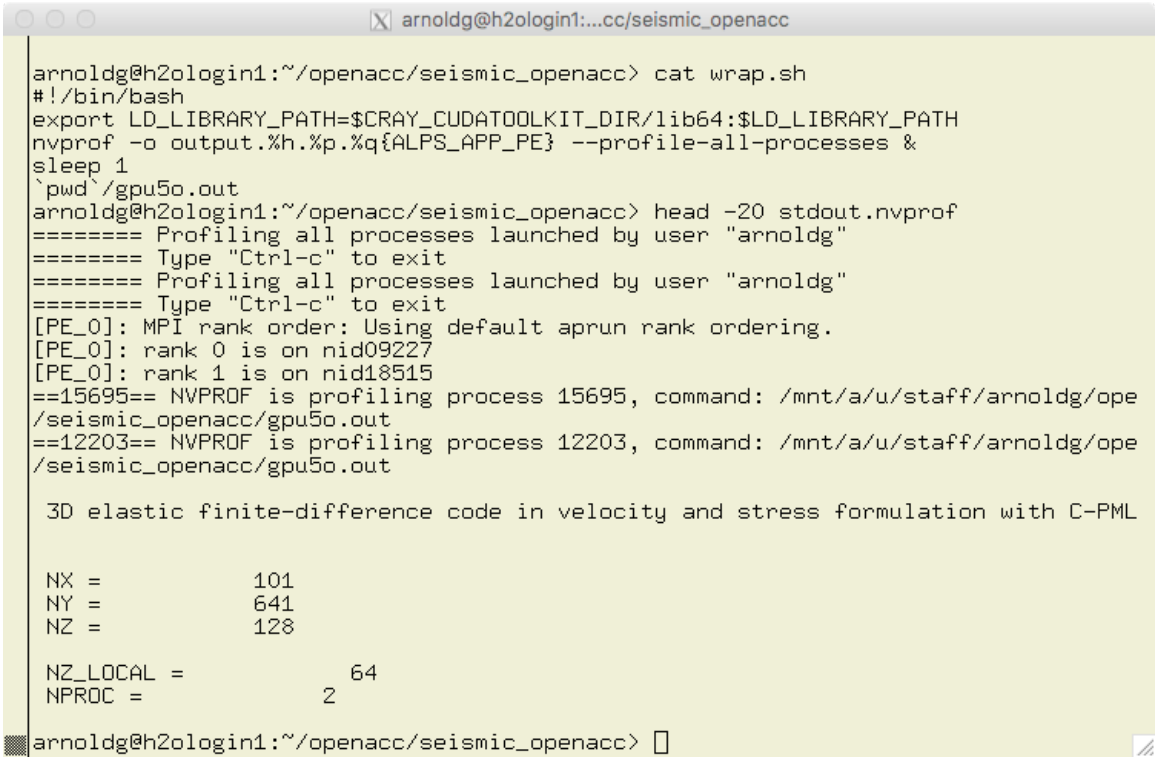

After successfully generating the output files , the nvvp profiler is run from a login node (or a local machine if you have cudatoolkit installed and have copied the files):

```
arnoldg@h2ologin1:~/openacc/seismic_openacc> nvvp output.nid09227.15695.0
```

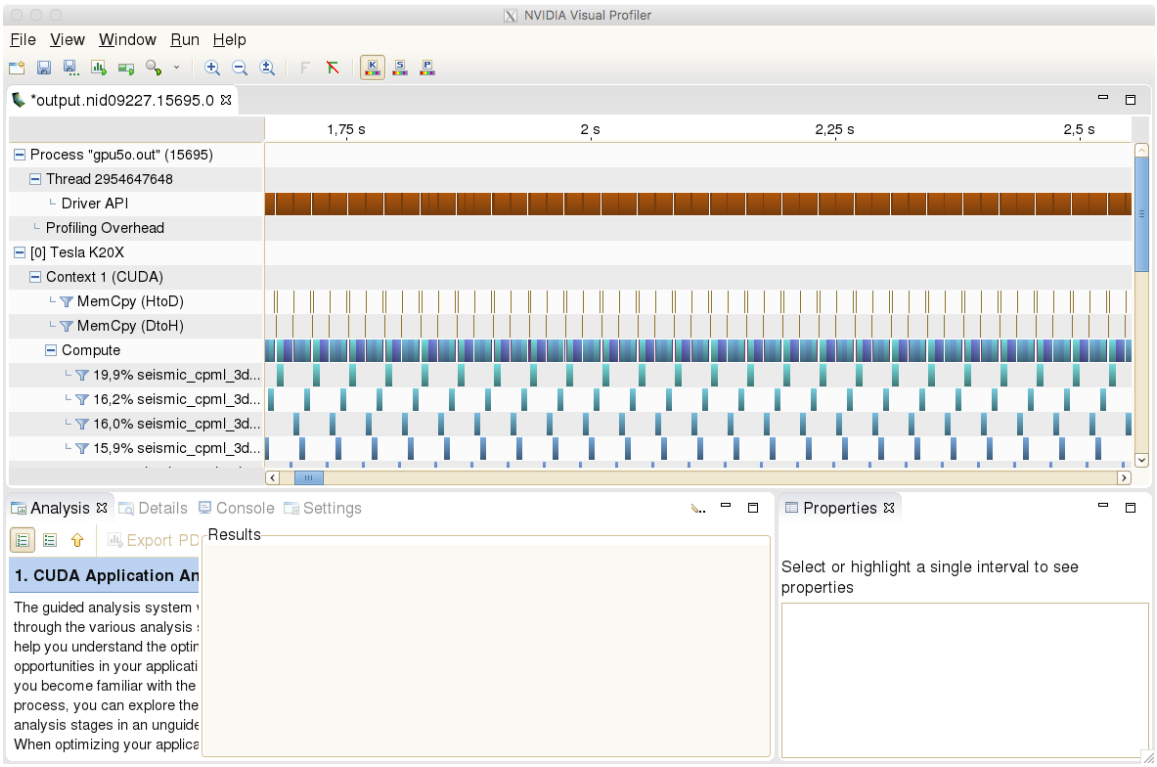
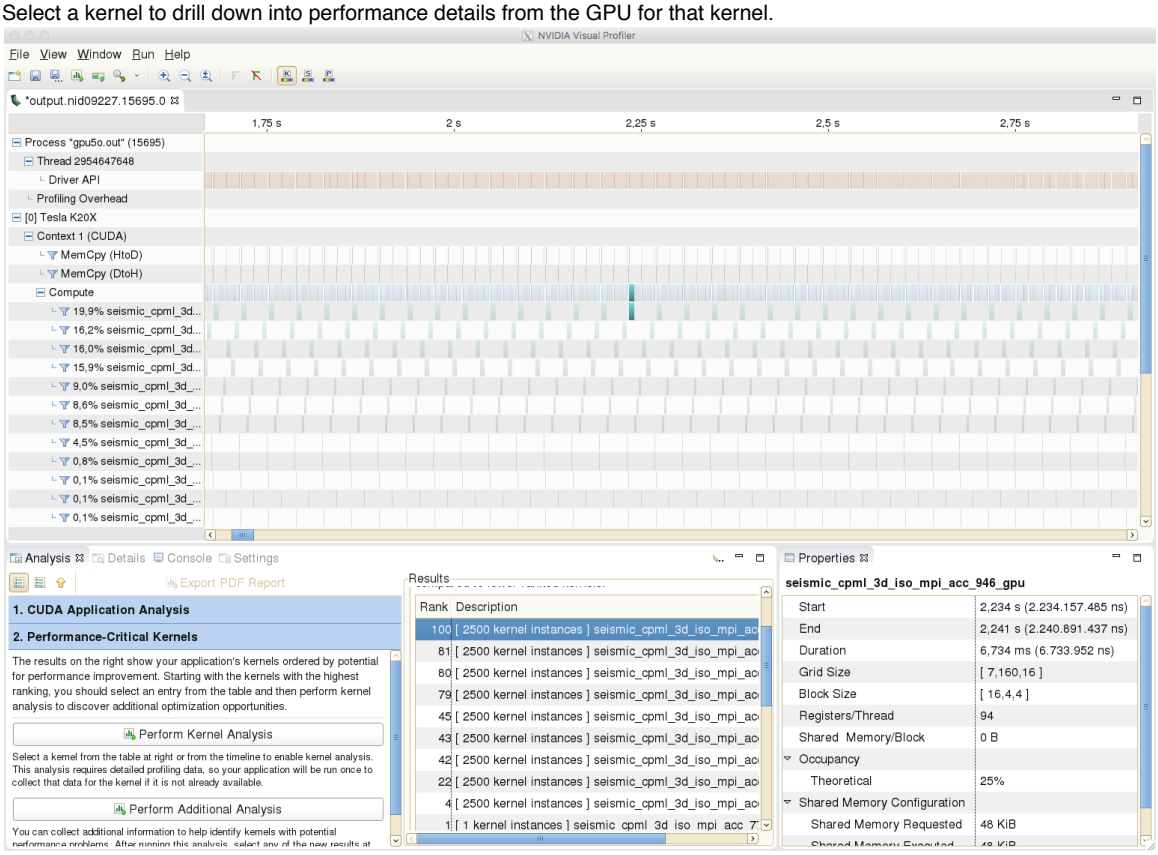
Select a kernel to drill down into performance details from the GPU for that kernel.