

OpenSees on Pan

Notes about the execution of several scripts on Pan.

Lamb's problem

Based on Seokho's Lamb's script, I created a Python script to generate an arbitrary sized mesh with quad elements (https://github.com/ucgmsim/OpenSees_script-generation). Using this, I generated a reasonably sized problem to execute on Pan. The model has 500K elements, which totals 1M degrees of freedom.

On Pan, OpenSeesSP has been compiled using SCOTCH for mumps as well as the Intel Maths library (IMKL). I wanted to see if this would make any difference.

Execution times for 100 iterations of the above mentioned problem:

| Cores | Time |
|-------|-----------------|
| 8 | 16m44.187s |
| 16 | 10m21.278s |
| 32 | 8m2.496s |

We see that it is not worth using 32 cores, as the reduction of execution time is marginal.

I have tried to investigate about the initialization time, so I have executed 1 iteration on each case. The numbers I obtained do not seem to agree with what I would have expected (namely: the more cores the more time initializing):

| Cores | Time | Node(s) |
|-------|-----------|----------------|
| 8 | 3m3.356s | compute-gpu-a1 |
| 16 | 2m48.556s | compute-phi-d1 |
| 32 | 2m44.745s | compute-gpu-c1 |

I need to probably re-run this test.

Some thoughts here:

- I was sharing the node with some other jobs in the first case
- Maybe different nodes, I should stick to a single one
- The job on 32 cores actually decided to go on two nodes because of the memory requirements.