

## Smoke and ventilation simulation using the CHPC

Wednesday, 4 December 2019 11:20 (20 minutes)

Kenneth Allen - Greenplan Consultants - kennethguyallen@greenplan.co.za

Greenplan Consultants undertook a combined smoke, ventilation and wind study for a basement/underground parking area of approximately 27 000 m<sup>2</sup>. This presentation gives an overview of the project and the experience of running it at the CHPC. In order to reduce the size of the transient model, the full domain (about 0.25 km<sup>3</sup>) was modelled under steady-state conditions using OpenFoam. The wind flow patterns around the basement were then imposed as boundaries on the transient model of the basement. Even with this approach, the transient model was too large to simulate on a small office network, so high performance computing (HPC) was essential.

Fire Dynamics Simulator (FDS) 6.7.0. was used for the transient simulations. It is purpose-written for simulating fire and smoke, and uses the computationally intensive Large Eddy Simulation (LES) method. The FDS model CHPC requirements were as follows - RAM:  $\approx$  100 GB; Nodes used: 10-15; Cores used: 240-360; Simulated time per case:  $\approx$  1-2 min; Wall time per case:  $\approx$  20-60 hours; Total cells: max. 70-80 million; Cell size: 100 mm; Data output per simulation: 200-250 GB.

FDS can make use of OpenMP (Multi-Processing) and MPI (Message Passing Interface). Tests on the CHPC with OpenMP enabled showed little-to-no improvement, so OpenMP was set to 1 (disabled). An experiment was made whereby all cores per node were booked on PBS but only half were used to run MPI process. This did not give better performance – possibly because of ghost processes running on one or two of the cores. Subsequently, all models were run with 1 MPI process per core, with 24 MPI processes per node. This seemed to be the best option for cost-effective performance.

FDS makes use of manually-specified rectilinear meshes. Unfortunately, where parallel simulation is desired, each MPI process requires at least one mesh, which means the mesh domain must be manually split up into sub-meshes. Any mismatch in mesh size leads to uneven loading on the cores, which means that the less heavily loaded cores have to wait. Due to the rectilinear grid, there are also “wasted” cells in walls/floors/roofs which have no function but contribute to the computational load. Thus, although the initial CSIR scaling tests on the CHPC (6 million cells) were promising for cell counts as low as 30 000 cells per core, scaling tended to be less efficient than expected.

On a number of occasions the simulations progressed at different speeds despite having the same configuration, flow speeds, and boundary conditions. It is possible that this might have been caused by ghost processes on individual cores and the CHPC architecture effect – in particular, the blocking ratio between racks. As an experiment on our final model geometry, we reduced the number of nodes in use from 15 (360 cores) to 10 (240 cores). The cell count per core was a factor of 2.5 higher, the area modelled was larger, and there were more jet fans and extraction fans than in the previous model. Despite this, the CHPC wall time required per unit simulated time increased by a factor of only 2. Rigorous testing is necessary before any conclusions are drawn, as a rough test like this does not provide sufficient data and there might be factors not taken into account. While the above meant that simulations did not run as fast as desired, they ran far faster than they would have on a small office network (it would have taken years, if they ran at all). Greenplan had a good experience with the CHPC and are keen to use it for future projects of this nature.

### Supported Student

**Primary author:** Dr ALLEN, Kenneth (Greenplan Consultants)

**Presenter:** Dr ALLEN, Kenneth (Greenplan Consultants)

**Session Classification:** HPC Applications

**Track Classification:** Computational Mechanics