


FALL3D WP2

Heribert Pascual (GEO3BCN - CSIC)
Leonardo Mingari (GEO3BCN - CSIC)
Arnau Folch (GEO3BCN - CSIC)

hpascual@geo3bcn.csic.es
lmingari@geo3bcn.csic.es
afolch@geo3bcn.csic.es



ChEESE WP-2 Next steps

- 1 - The execution issues observed at very large node counts have been identified and will be further investigated, as they currently limit scalability beyond ~400 nodes
- 2 - Portability work on LUMI will continue, since recent changes in the software stack have led to reduced GPU performance and incorrect results
- 3 - Potential improvements suggested during code audits will be assessed and selectively implemented, taking into account their impact on portability across the different target platforms
- 4 - Output file compression strategies will be investigated and prototyped to facilitate post-processing workflows at large scales, while assessing their impact on the overall computational cost

FALL3D High-Scalability Use Case



A medium-sized problem was created that could efficiently utilise GPUs, while avoiding excessively long execution times, given the limited wall-time available for large node runs

Physical scenario:

- Etna volcano eruption (2018)
- Grid: 512 x 216 x 32
- Simulation period: 1 week
- Runtime: about 15 minutes on GPUs

Additional reduced use case:


- Shortened simulation covering the first 10 hours.
- Used for CPU vs GPU performance comparisons
- Test analysis for TALP and other tools



FALL3D High-Scalability Issue

Reference use case tested from 1 node up to 800 nodes:

- 1, 100, 250 and 400 nodes working OK
- 450 nodes: UCX problem, solved with flags
 - << Transport retry count exceeded on mlx5_1 >>
- 500 nodes: Solved reducing simulation time 24h
- 800 nodes: Always fails due to MPI retry




Current understanding: The issue appears to be linked to increasing load imbalance between ensembles at large GPU scales, where some ensembles finish significantly earlier than others, triggering MPI timeouts. Reproducibility is currently limited due to the lack of access to comparable GPU node counts.

FALL3D MN5 TALP ensemble weak scaling


From 1 node to 500 nodes

TALP Execution Time — Ensemble Weak Scaling



Wall time per ensemble size (reduced use case)

TALP Efficiencies — Ensemble Weak Scaling



Obtained TALP metrics (reduced use case)

END

Heribert Pascual (GEO3BCN - CSIC)
Leonardo Mingari (GEO3BCN - CSIC)
Arnau Folch (GEO3BCN - CSIC)

hpascual@geo3bcn.csic.es
lmingari@geo3bcn.csic.es
afolch@geo3bcn.csic.es