Walberla

April 21, 2023

Jayesh Badwaik, Chelsea Maria John (JSC)
Radim Vavrik, Kristian Kadlubiak (IT4I)

Single Node Performance

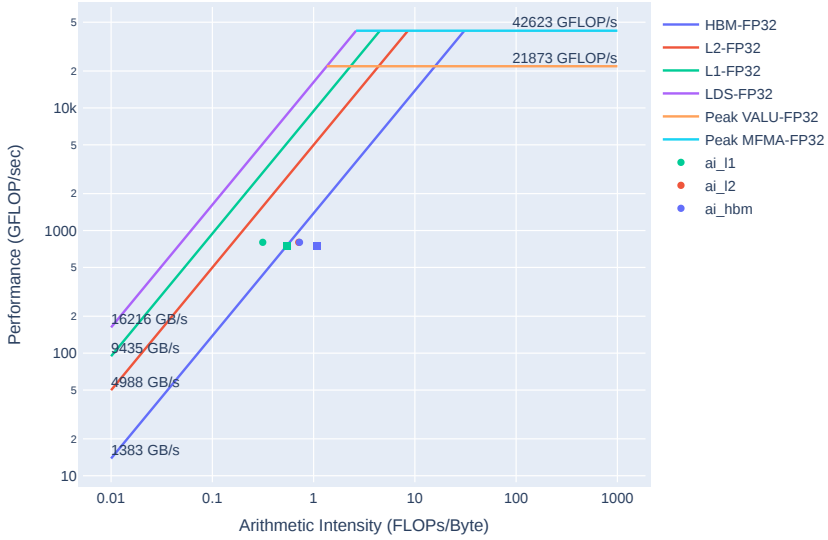

Figure: Peak Performance $= 2500$ MLUPS

Multi-Node Performance

- CPU-GPU Bindings Important
  - 300 MLUPS to 2200 MLUPS
- `rocprof` and ExtraE profiling
  - Compute kernels not the bottleneck (as should normally be)
  - Increase the number of cells on in a block by 2
  - 2200-2300 MLUPS (this is more than acceptable)
  - Scale to thousand nodes

Observations

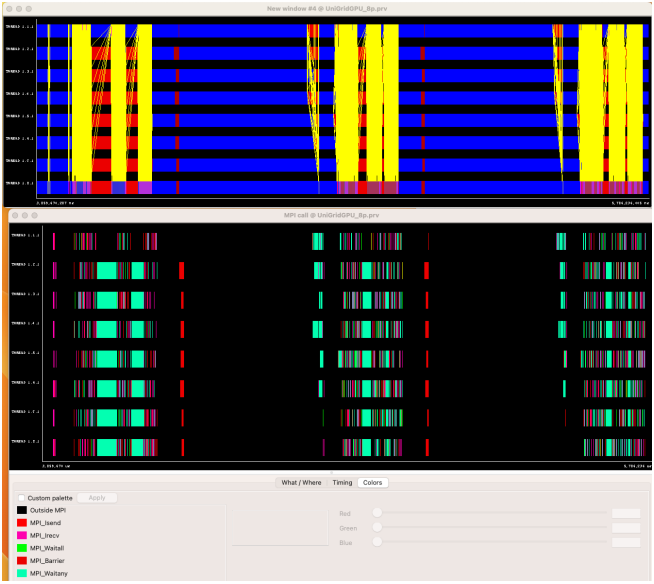

Observations

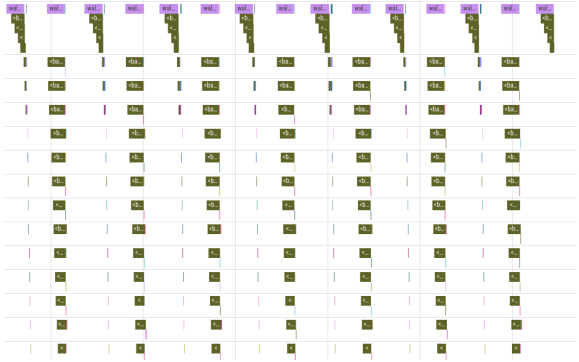

Observations

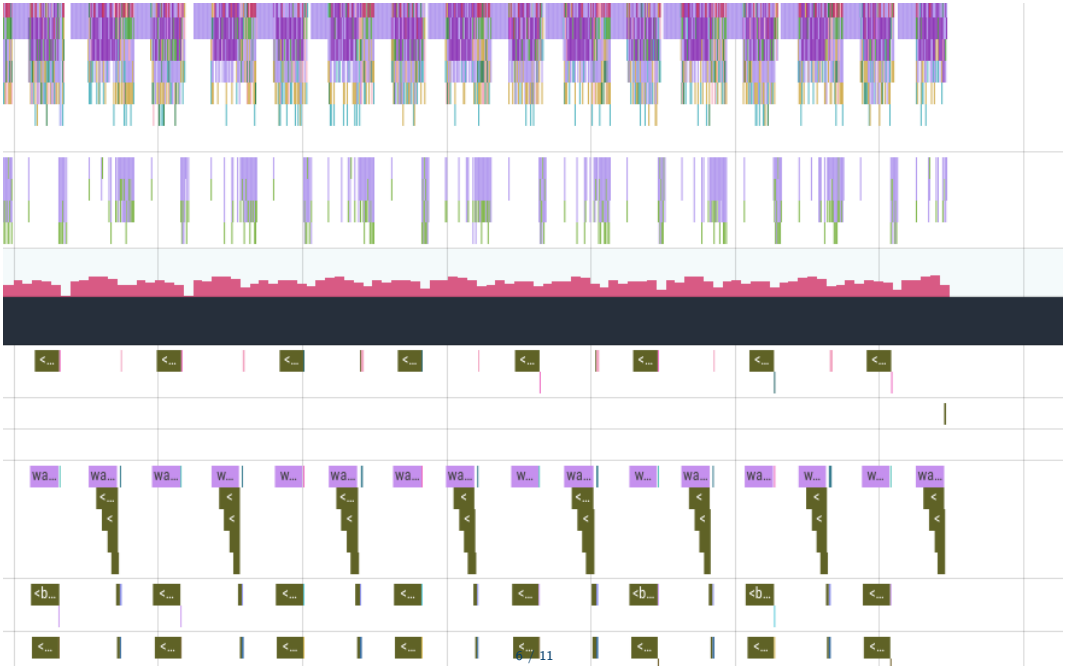

Observations

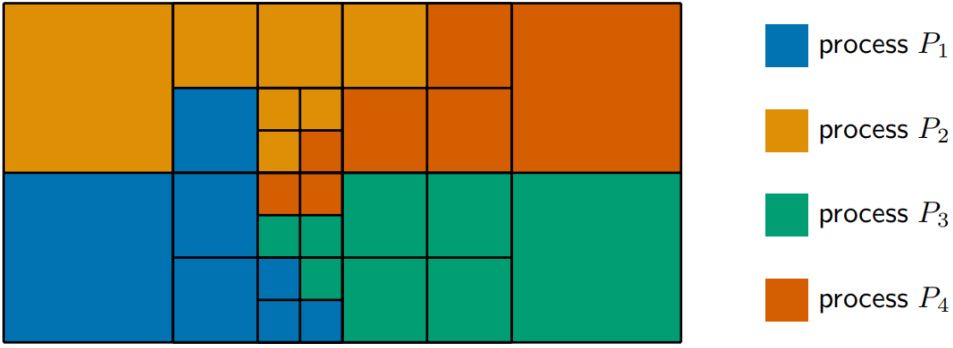

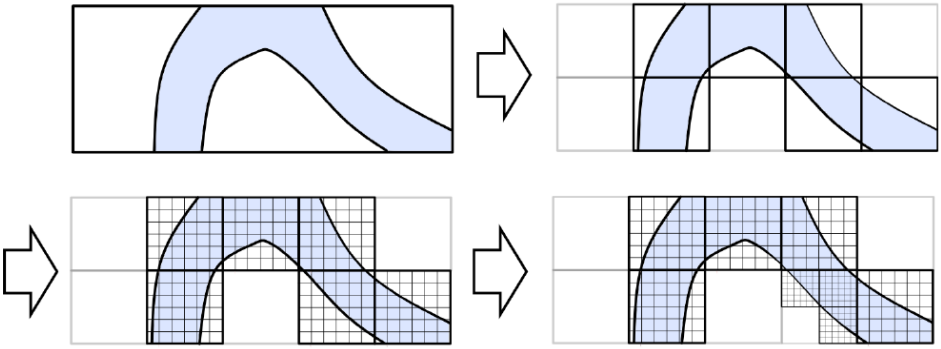

Figure 4: Block-structured domain setup. From left to right: defining domain using a surface mesh, decomposition into coarse blocks, allocation of cells in blocks, block refinement

Next Steps

- Profile

Achievements in the Hackathon

- Get help in profiling tools to work
- Verify that the single-node performance is near-optimal
- Identify the correct configuration to run the multi-node code
- Use profiling of multi-node to identify the performance bottleneck

Jayesh Badwaik, Chelsea Maria John (JSC)
Radim Vavrik, Kristian Kadlubiak (IT4I)

www.scalable-hpc.eu

@scalable_hpc

The SCALABLE project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No. 956000.