ISC High Performance

TRANSFORMING THE FUTURE

MAY 29 – JUNE 2, 2022 | HAMBURG, GERMANY

Eulerian-Lagrangian cloud model

cloud length scales

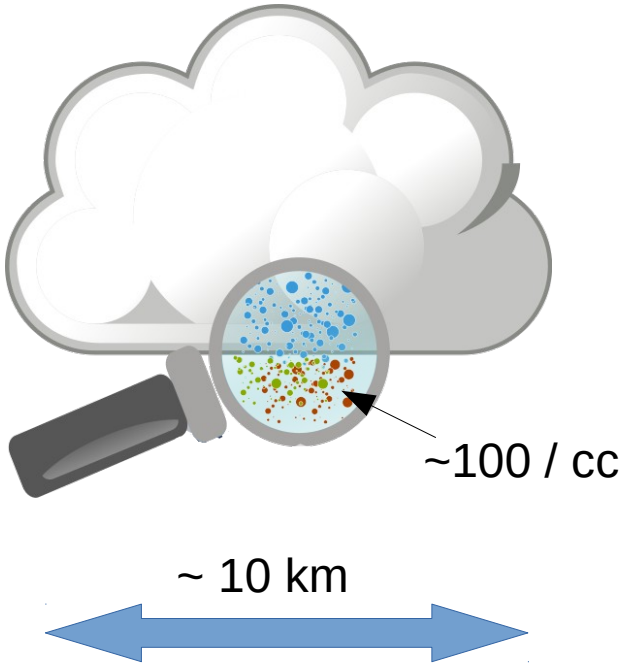

Lagrangian particles (aerosols, droplets)
in an Eulerian grid (thermodynamic fields)

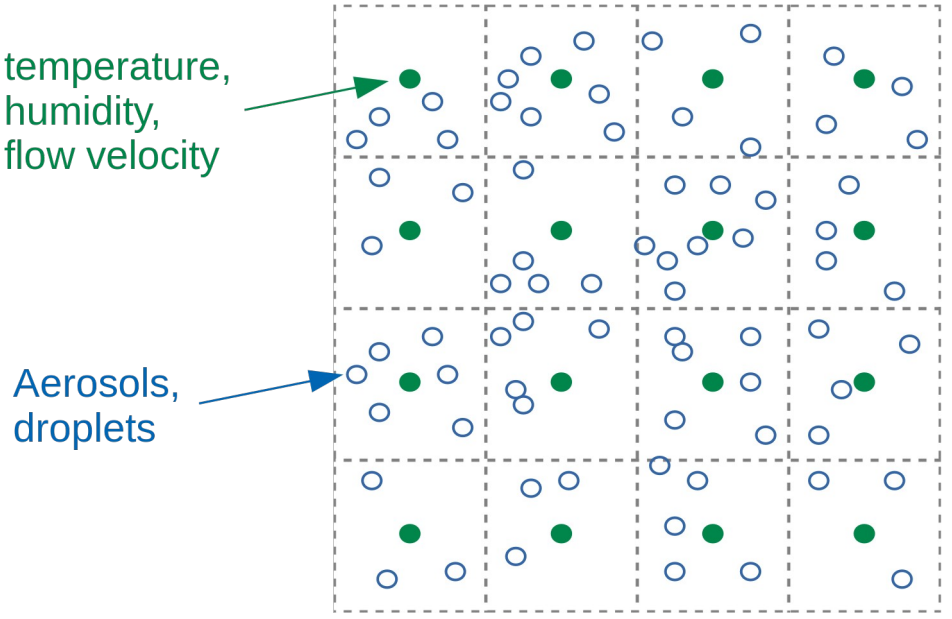

Adaptation for heterogeneous computing clusters

- Simultaneous computation of Lagrangian (GPU) and Eulerian (CPU) components
- Domain decomposition
- Each MPI process controls multiple threads (OpenMP) and GPUs (CUDA memory copy)

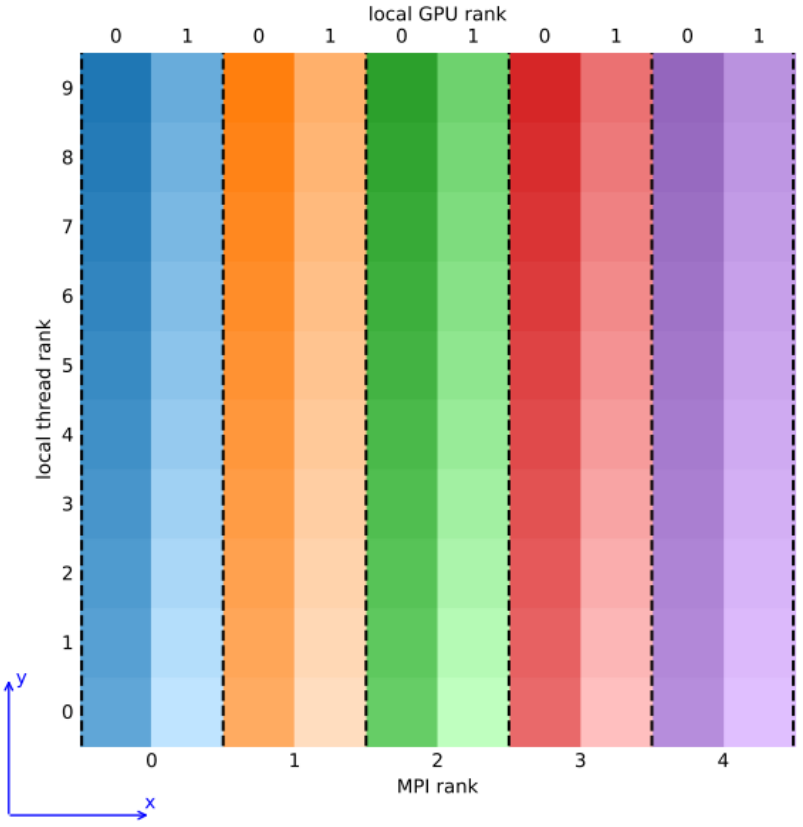

Top-down view of modeled domain; squares are Eulerian grid cells; coloring shows domain decomposition.

Weak scaling test

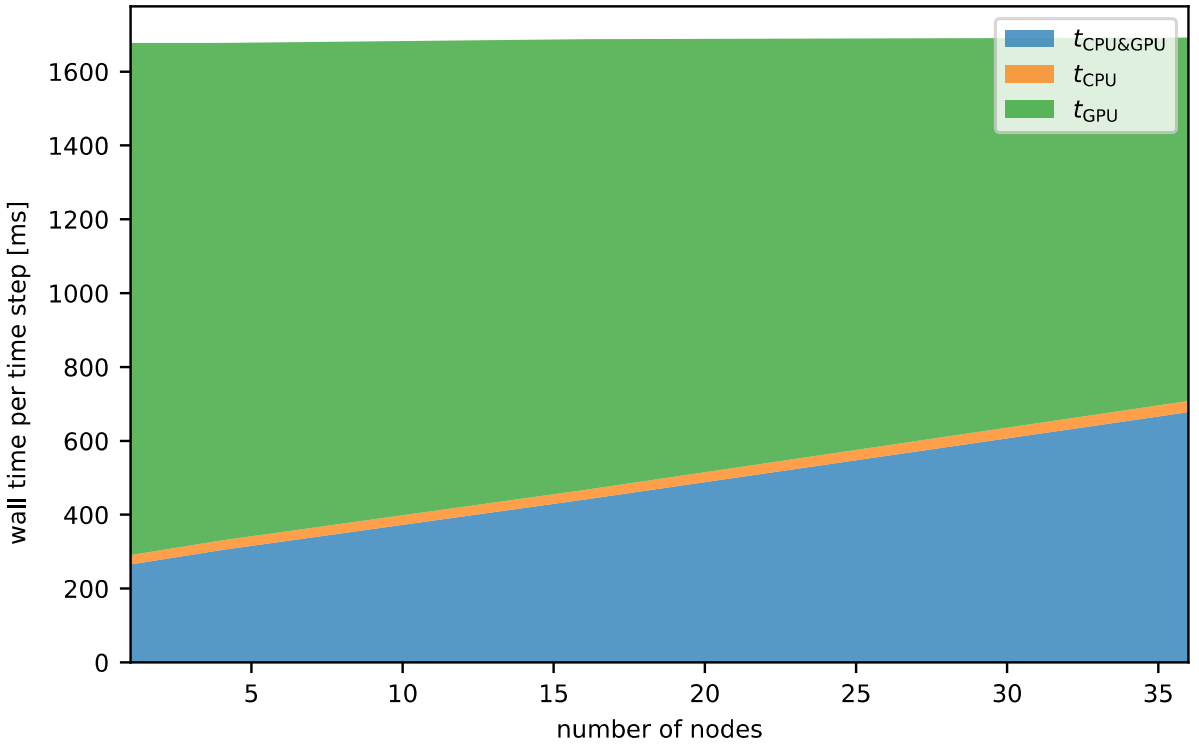

Wall time per time step vs number of nodes. Timings of simultaneous CPU and GPU computations (blue), CPU-only computations (orange) and GPU-only computations (green) are stacked.

- GPU time scales better than CPU time
- Simultaneous CPU and GPU usage should be maximized for an optimal number of nodes (larger than shown)
- Up to the optimal number of nodes, scaling efficiency of the total wall time is ca. 100%

Strong scaling on CPU, weak on GPU

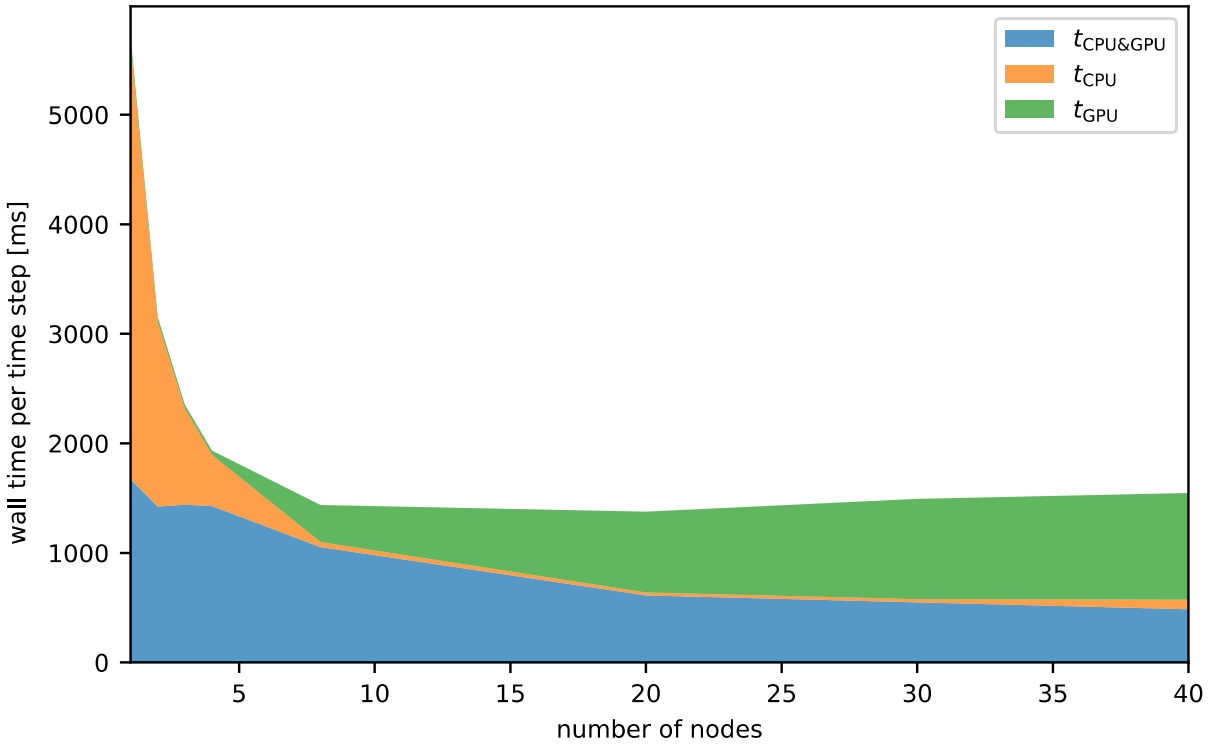

Wall time per time step vs number of nodes. Timings of simultaneous CPU and GPU computations (blue), CPU-only computations (orange) and GPU-only computations (green) are stacked.

- Good balance of CPU and GPU computations (ca. 80%) for an optimal number of nodes (5-10 in this case)