University of Warsaw Lagrangian Cloud Model (UWLCM) - Eulerian-Lagrangian Cloud Model for Heterogeneous Computing Clusters

P. Dziekan, P. Zmijewski, H. Pawlowska

Faculty of Physics, University of Warsaw, Warsaw, Poland

Numerical cloud modeling

- Why?
  - Clouds are important for weather and climate
  - Cloud observations are challenging
  - Laboratory experiments do not cover all length scales

Numerical cloud modeling

- Why?
  - Clouds are important for weather and climate
  - Cloud observations are challenging
  - Laboratory experiments do not cover all length scales
- How?
  - Modeling air flow (CFD) and cloud droplets (microphysics)

Numerical cloud modeling

- Why?
  - Clouds are important for weather and climate
  - Cloud observations are challenging
  - Laboratory experiments do not cover all length scales
- How?
  - Modeling air flow (CFD) and cloud droplets (microphysics)
- Challenges
  - Large range of important spatial and temporal scales

Cloud modeling across scales

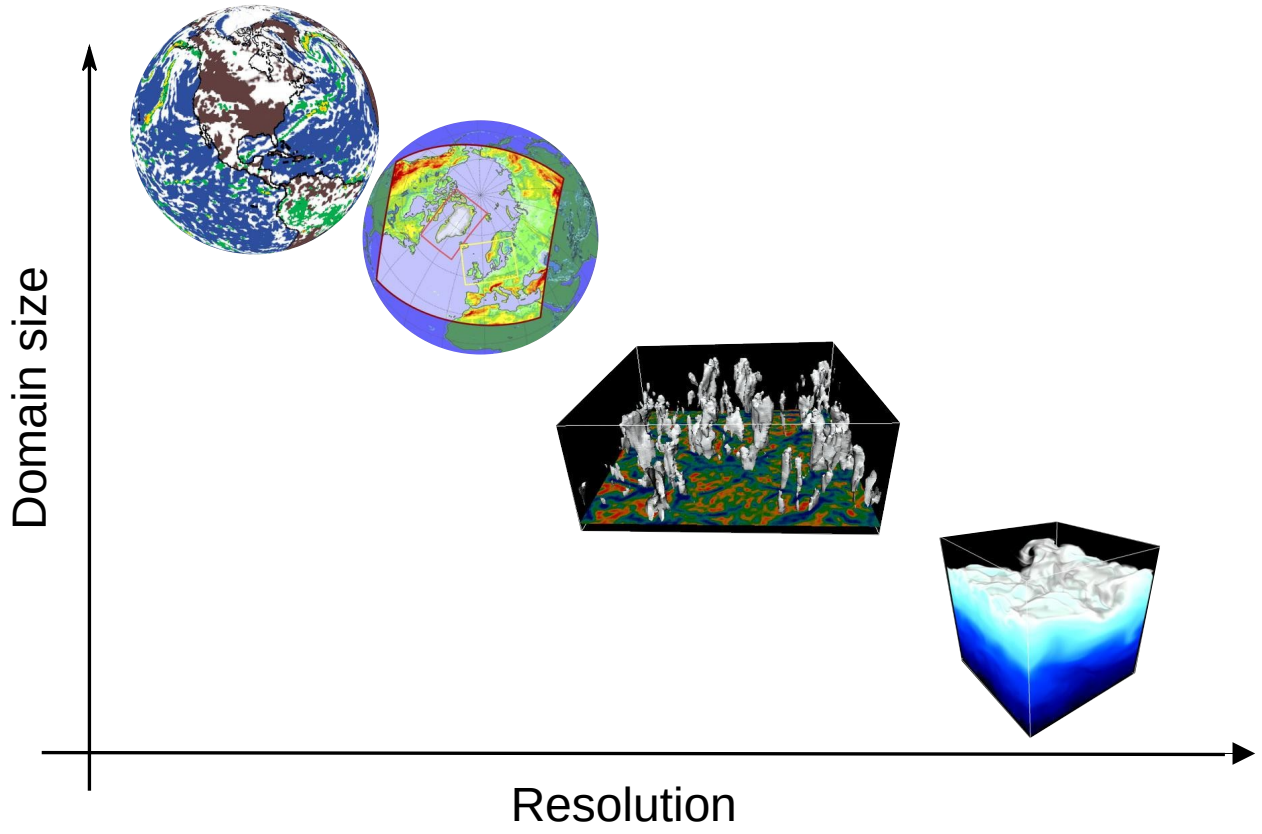

Cloud modeling across scales

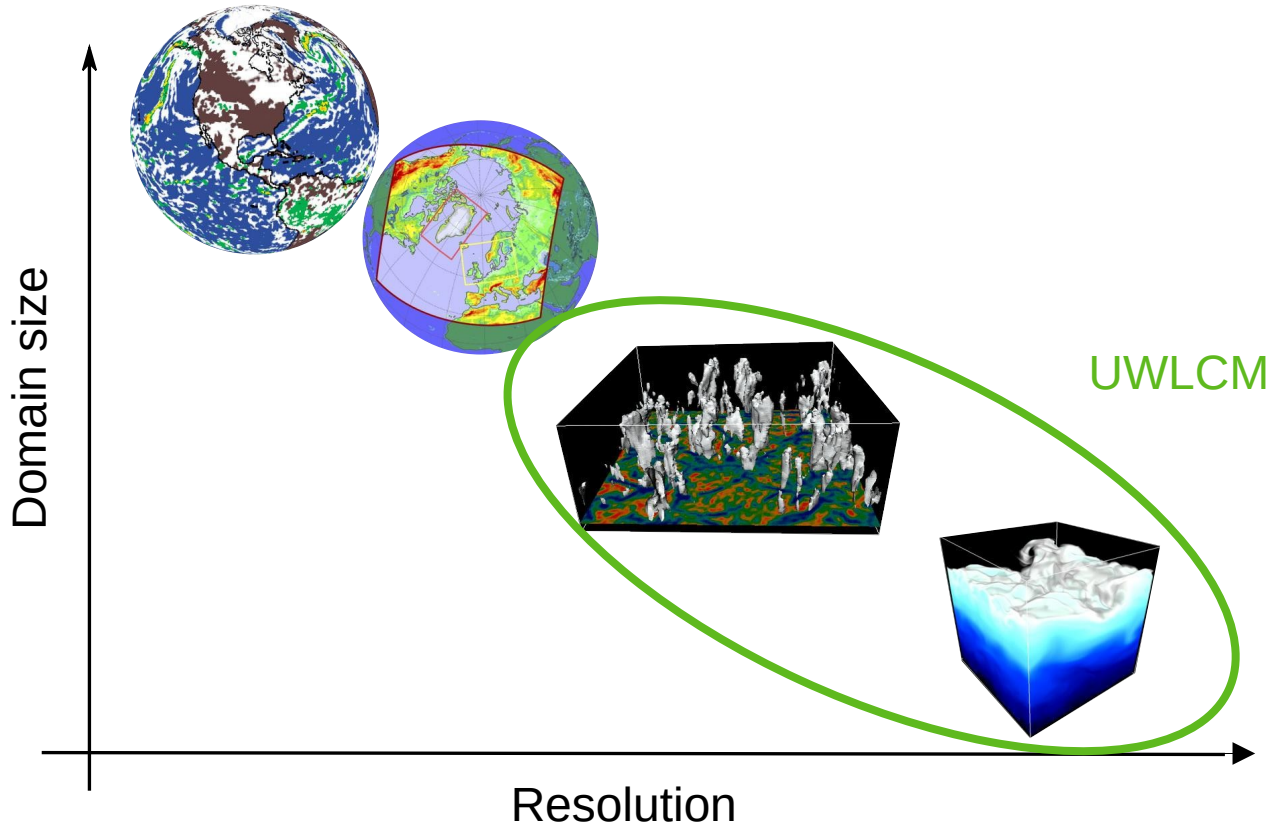

Eulerian-Lagrangian simulation

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

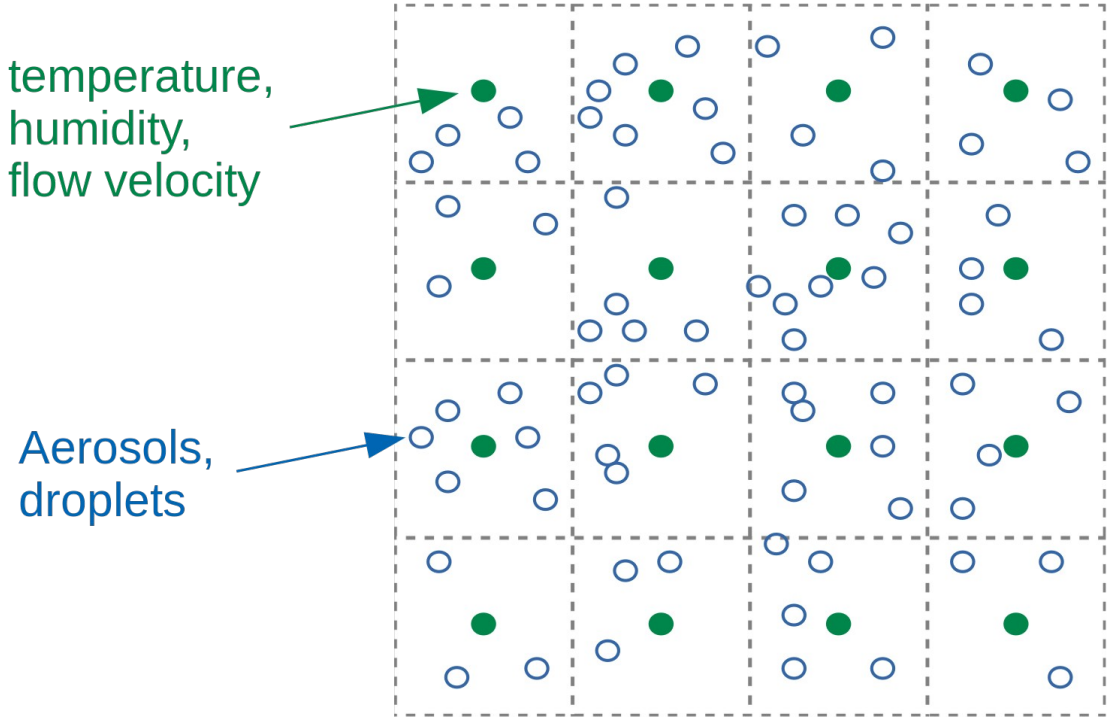

Use of heterogeneous (CPU+GPU) clusters

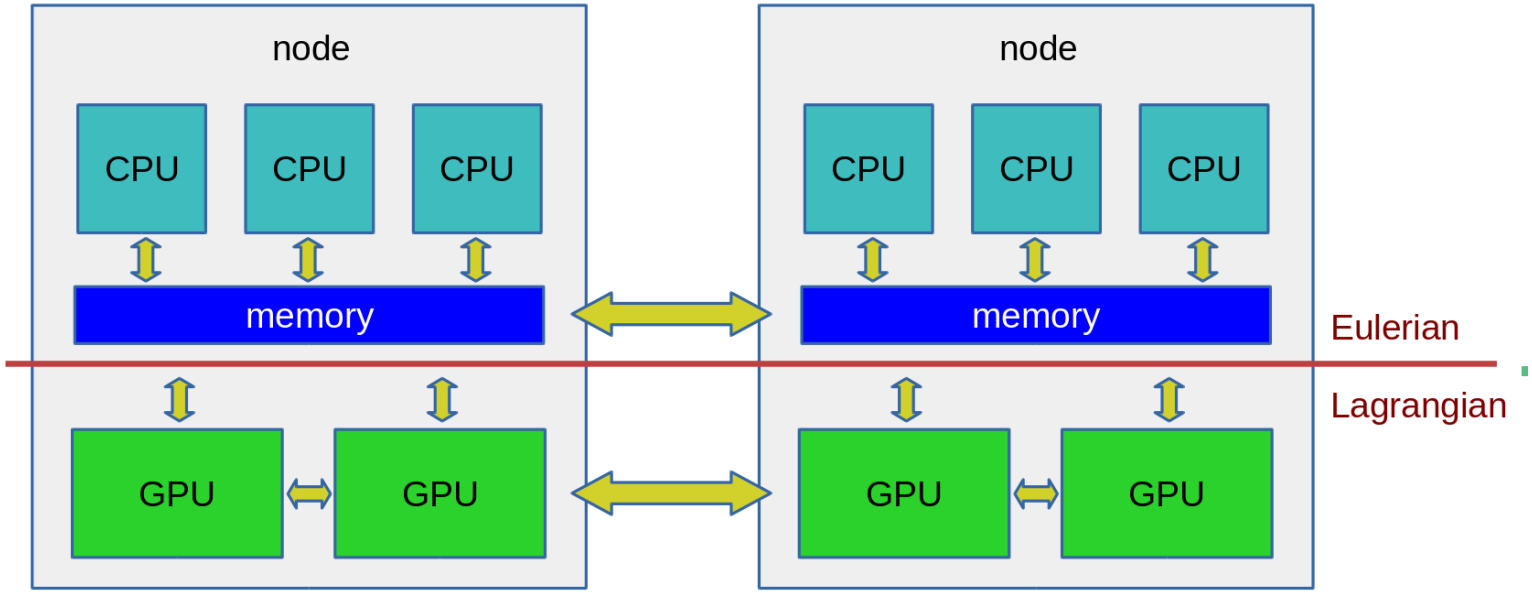

- Eulerian component: resides in RAM, computed by CPUs
- Lagrangian component: resides in GPU RAM, computed by GPUs

Domain decomposition

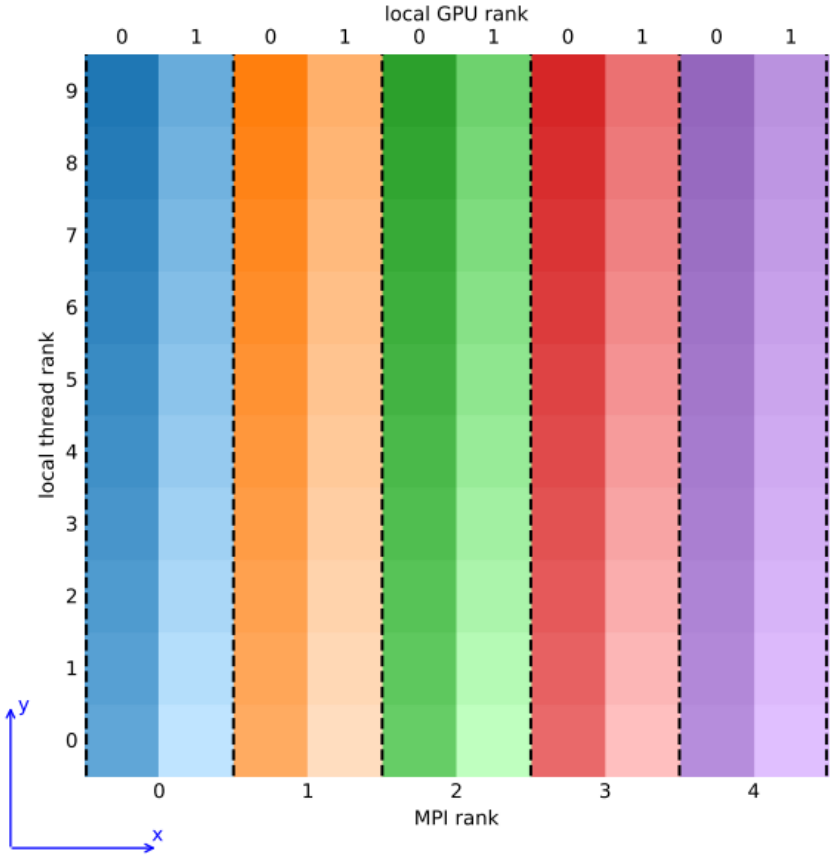

Top-down view of modeled domain; squares are Eulerian grid cells; coloring shows MPI, thread and GPU ranks.

Scaling tests on Prometheus

P. Dziekan and P. Zmijewski, *Geoscinetific Model Development (*2022)

Weak scaling

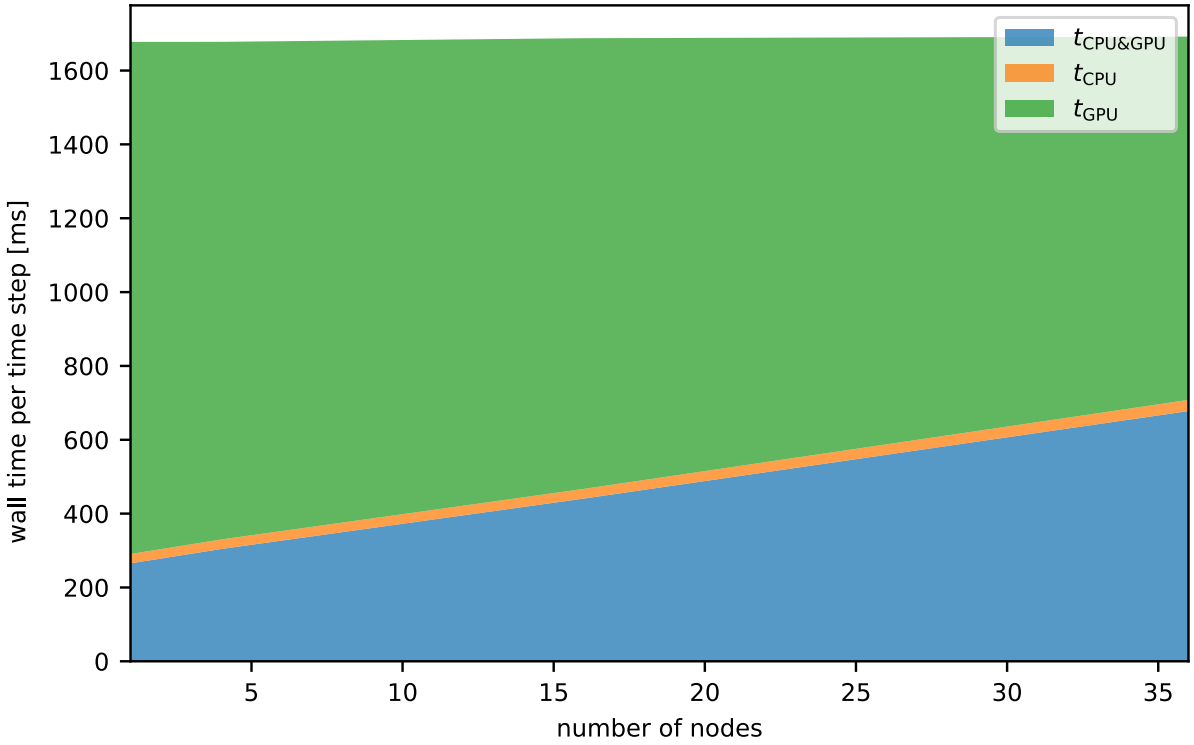

Wall time per time step vs number of nodes. Timings of simultaneous CPU and GPU computations (blue), CPU-only computations (orange) and GPU-only computations (green) are stacked.

- GPU time scales better than CPU time.
- For an optimal number of nodes, simultaneous CPU and GPU usage is maximized.
- Up to the optimal number of nodes, scaling efficiency of the total wall time is ca. 100%

Strong scaling on CPU, weak on GPU

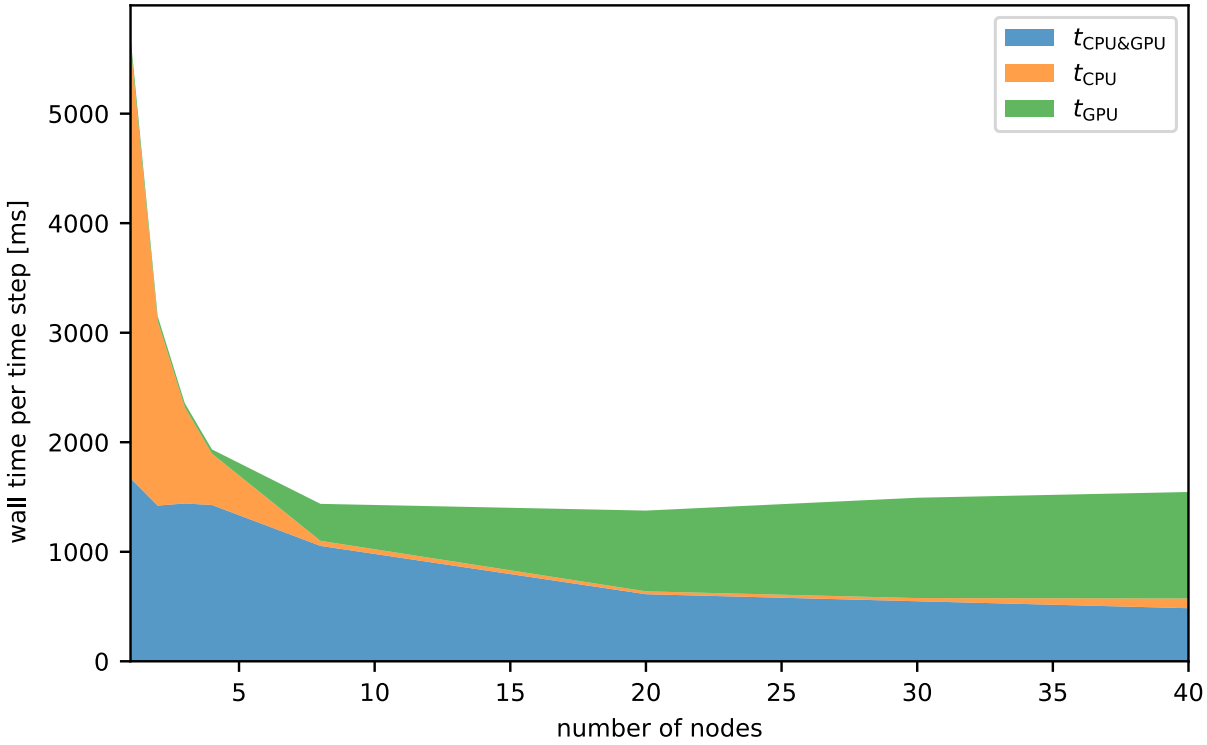

Wall time per time step vs number of nodes. Timings of simultaneous CPU and GPU computations (blue), CPU-only computations (orange) and GPU-only computations (green) are stacked.

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

Get involved!

- github.com/igfuw

Conclusions

- CPUs and GPUs can work simultaneously for up to 80% of time in a meaningful scientific simulation, but this requires tuning simulation parameters and the number of nodes.
- GPU benefits: Same wall time could be achieved by replacing each GPU by between 10 to 60 CPUs, depending on the architecture.