Barcelona Supercomputing Center
Centro Nacional de Supercomputación

EXCELENCIA SEVERO OCHOA

Optimization of the DT for Weather and Climate Predictions: An exciting battleground

Stella V. Paronuzzi Ticco

10/12/23

GOBIERNO DE ESPAÑA
MINISTERIO DE CIENCIA E INNOVACIÓN

Funded by the European Union
Destination Earth
implemented by ECMWF esa EUMETSAT

The on-demand extremes digital twin

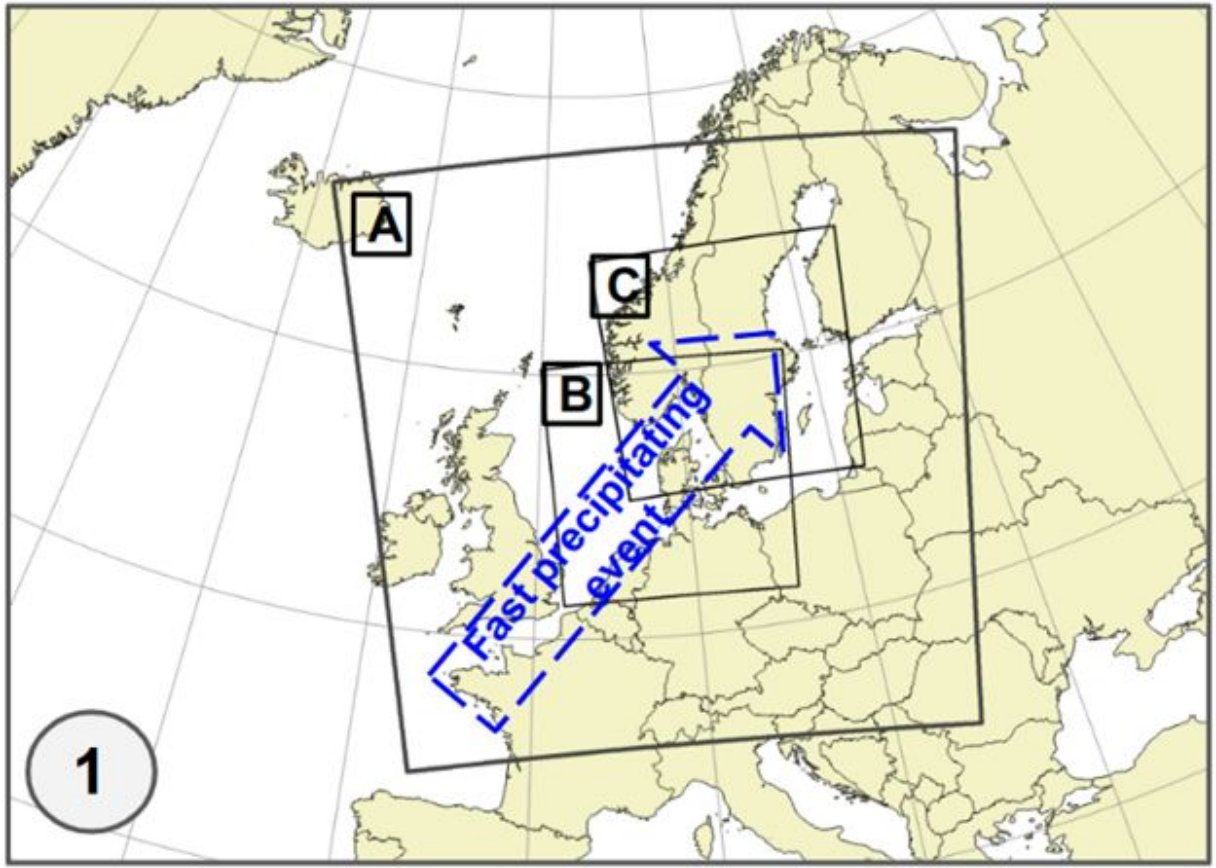

From CPU to GPU

Credits: Daan Degrauwe, Denis Haumont

- LAM and global model have different computation
- Still we could built on previous effort

- OpenACC + Highly optimized vendor libraries
  (cuFFT/cuBLAS in case of NVIDIA GPUs; rocFFT/rocBLAS in case of AMD GPUs).

- FFTs taken in batches

Initial performance

Credits: Daan Degrauwe, Denis Haumont

- 1 GPU is ~20% slower than 1 CPU
- Scaling across nodes seems okay

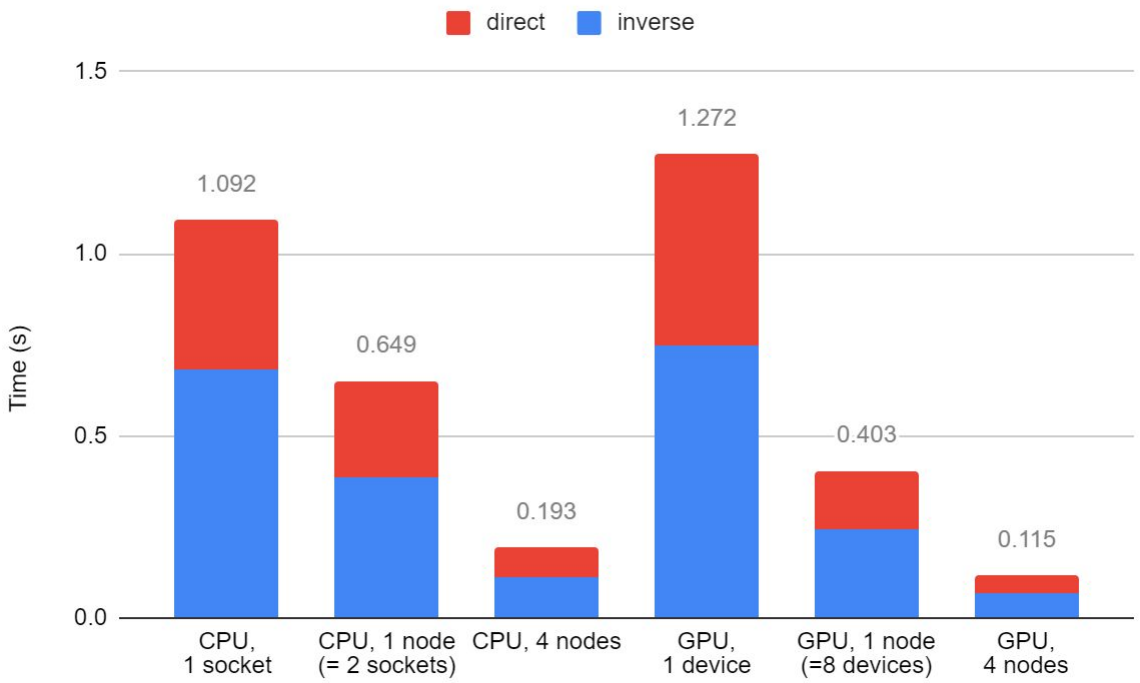

Initial performance (single AMD GPU)

Credits: Daan Degrauwe, Denis Haumont

```
--nproma 32 --nlat 1024 --nlon 1536 --vordiv --nlev 100
     --nfld 3 --scders --uvders --niter 4
```

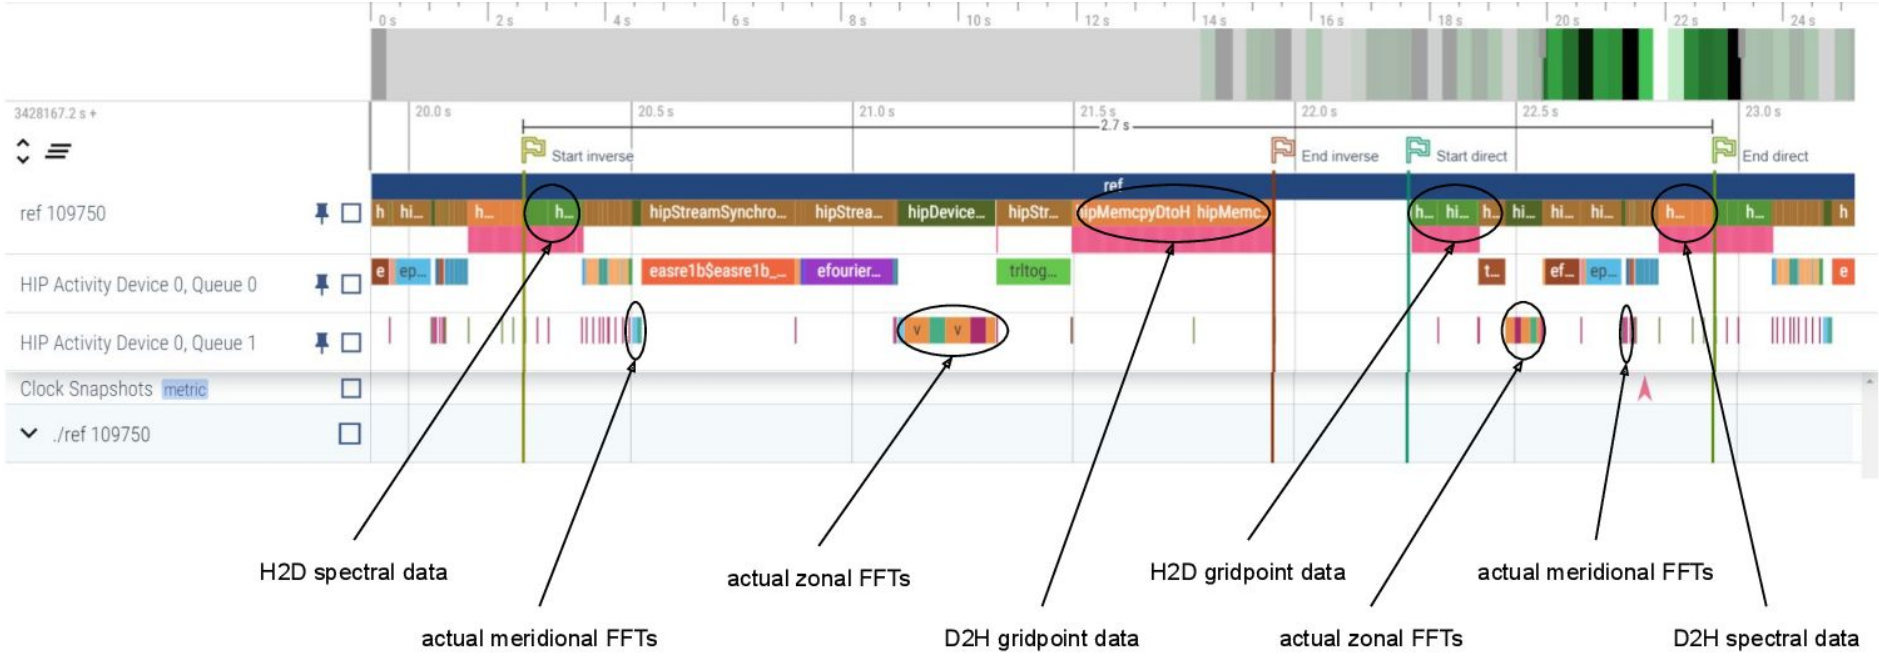

Overlapping data transfers and computations

Credits: Daan Degrauwe, Denis Haumont

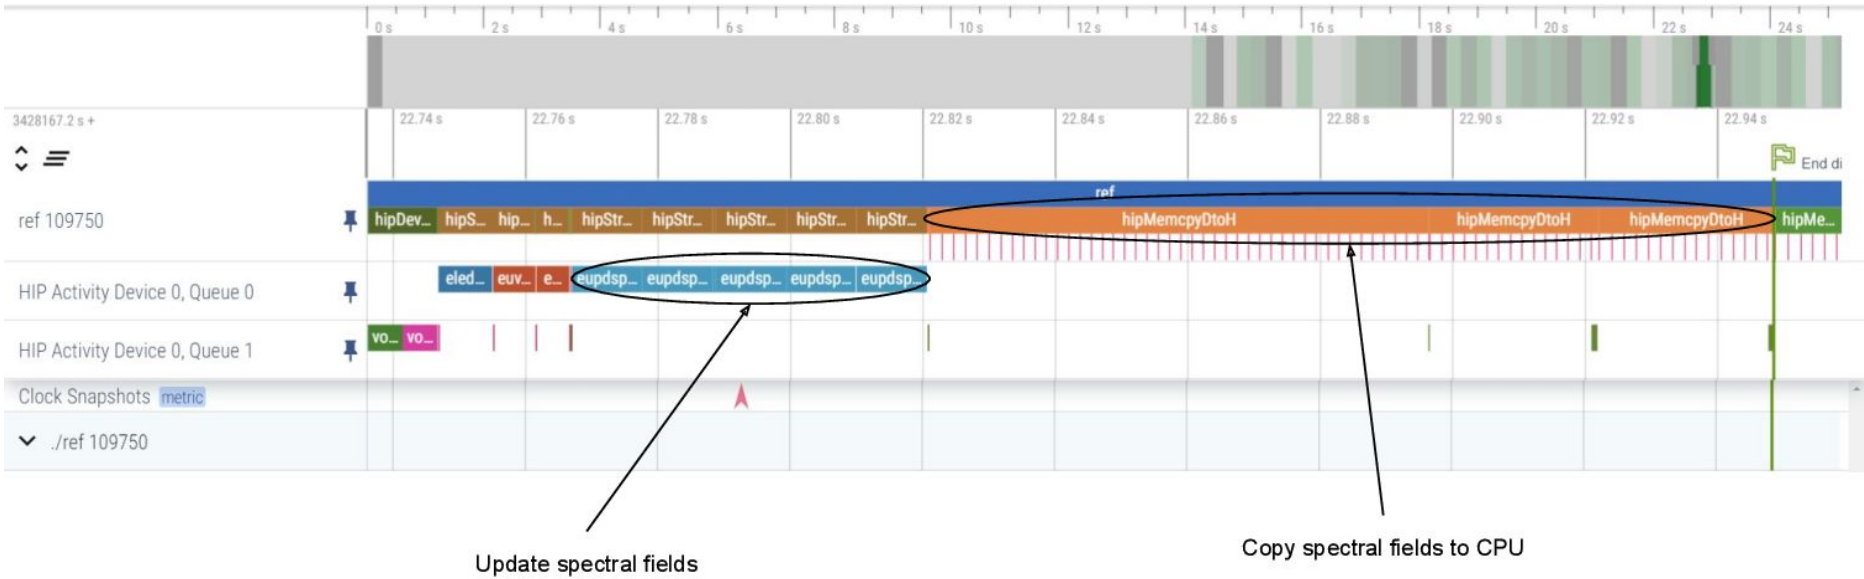

Overlapping data transfers and computations

Credits: Daan Degrauwe, Denis Haumont

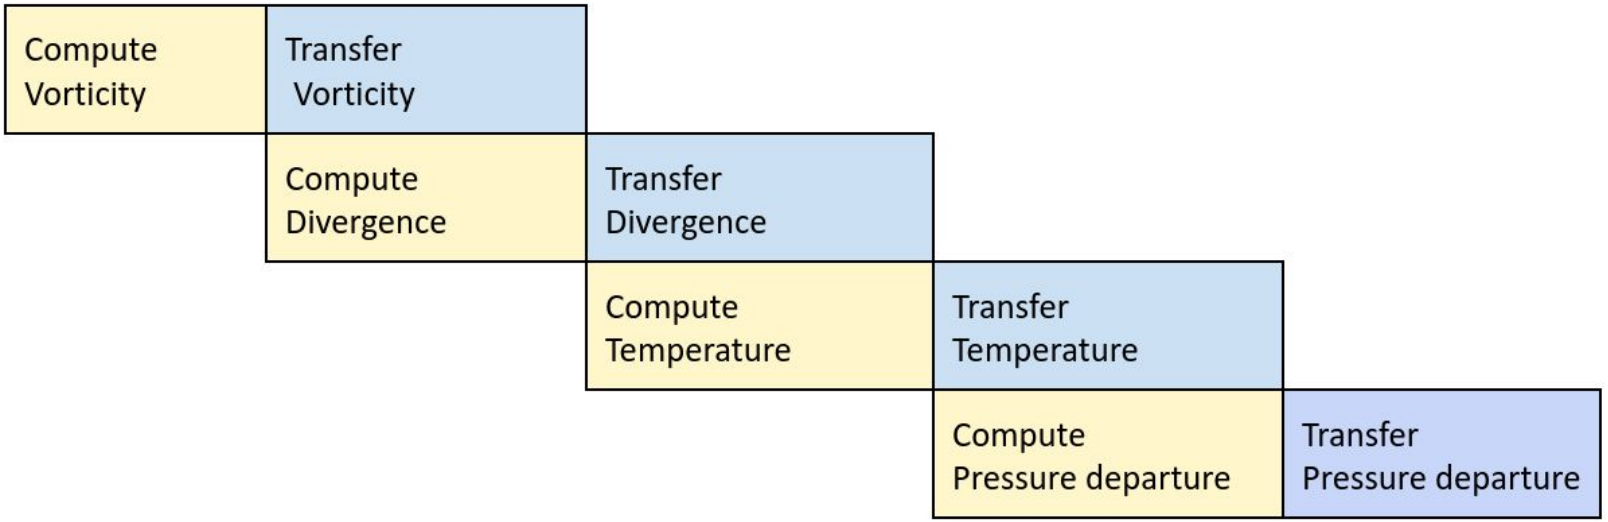

Overlapping data transfers and computations

Credits: Daan Degrauwe, Denis Haumont

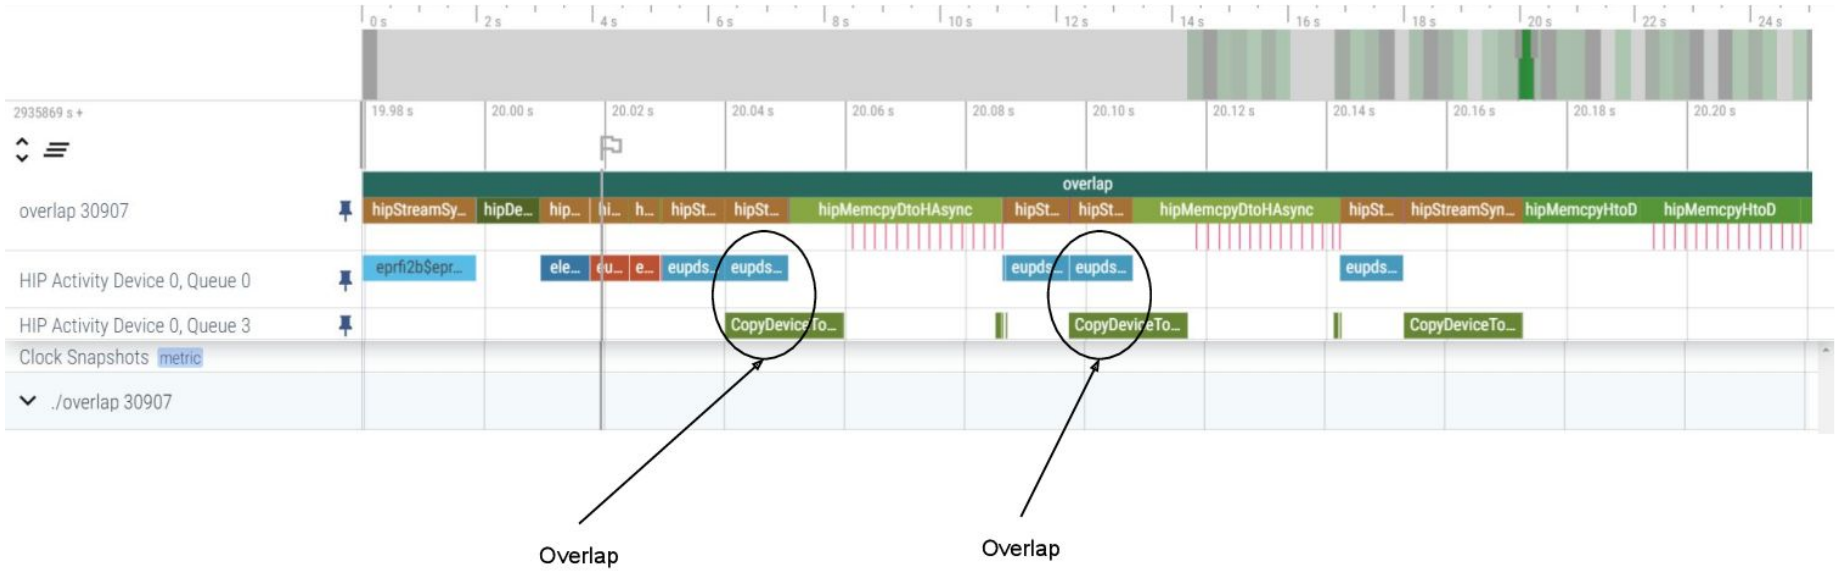

Avoiding Transpositions

Credits: Daan Degrauwe, Denis Haumont

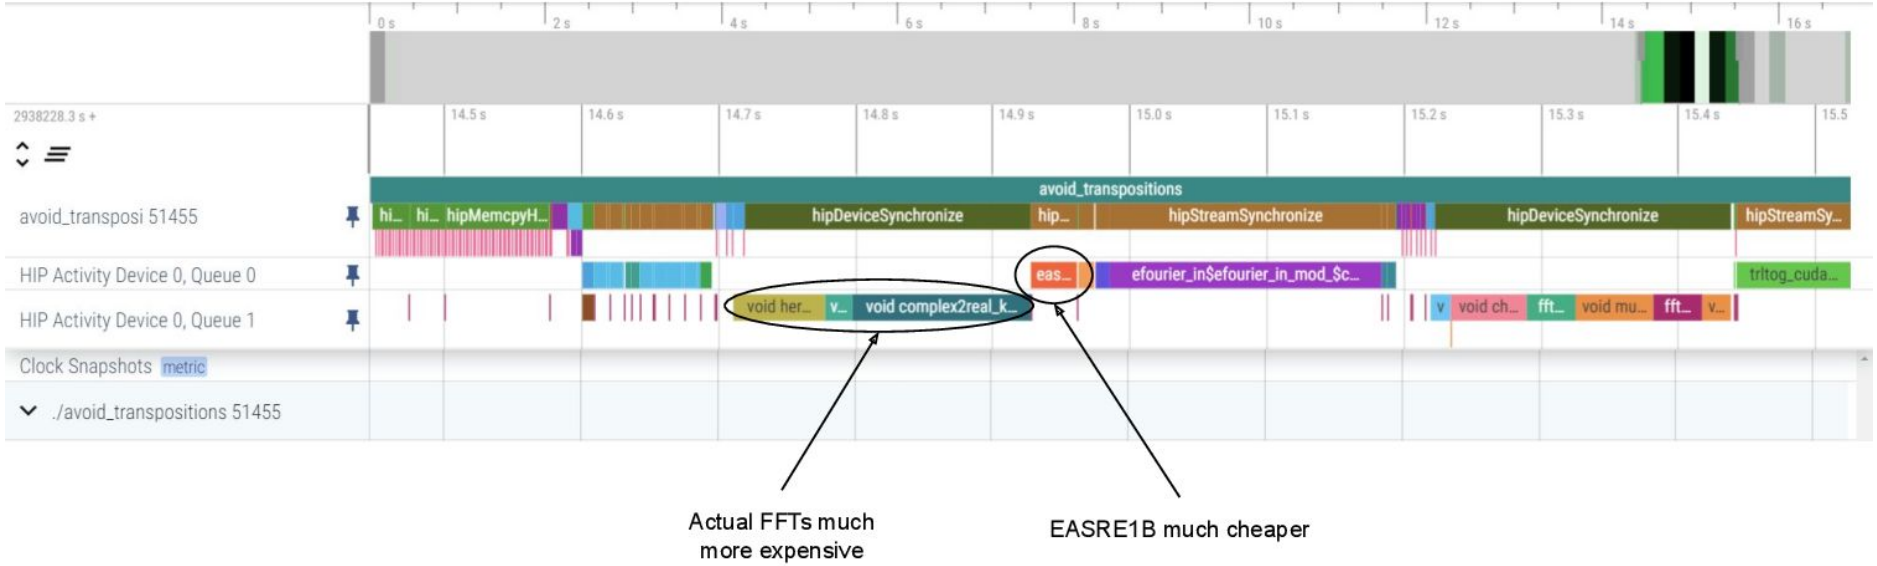

Conclusion

- An overall improvement in performance has been achieved (~20%)
- Data layout changing routines remain the most expensive ones.

Future work

- Replace the OpenACC loops in the data transposition routines by HIP kernels.

The Climate Adaptation Digital Twin

Development **C**onfiguration (~ 10 km)
End of February 2023

Production **C**onfiguration (~5 km)
August 2023

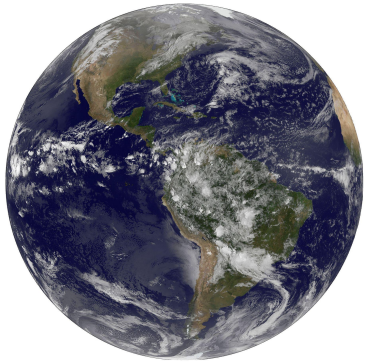

φ2 **C**onfiguration (~2.5 km)
Summer of 2024

Development Configuration (OpenMP scaling)

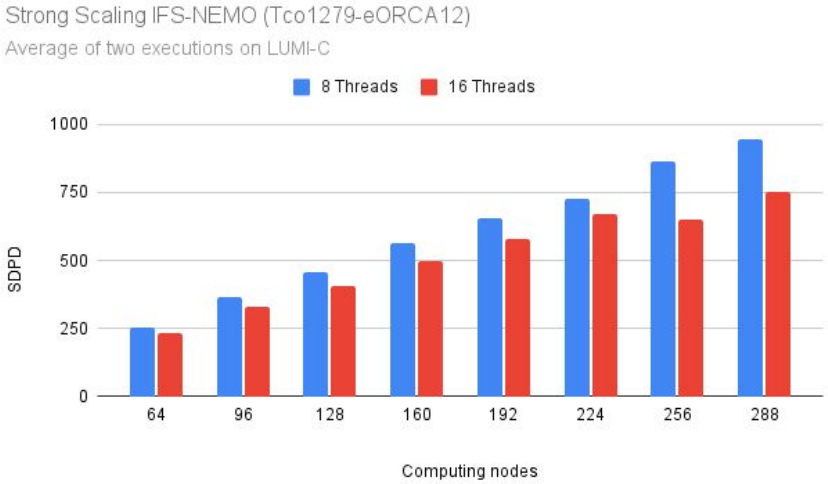

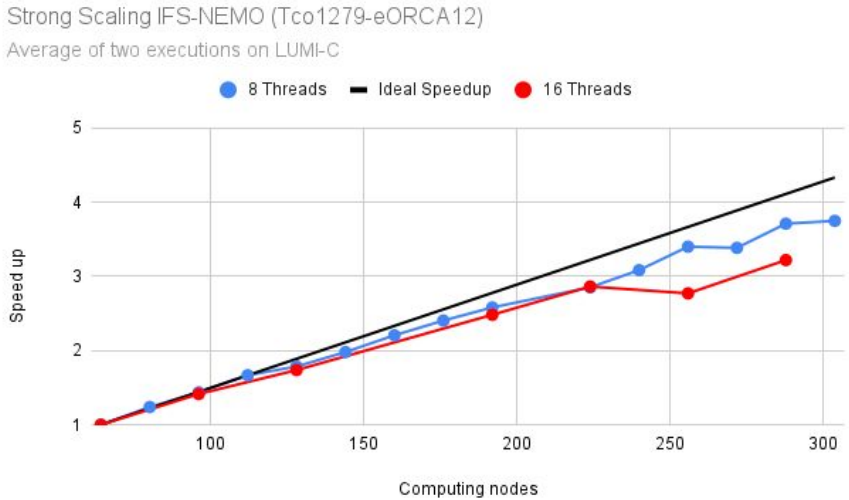

Development Configuration

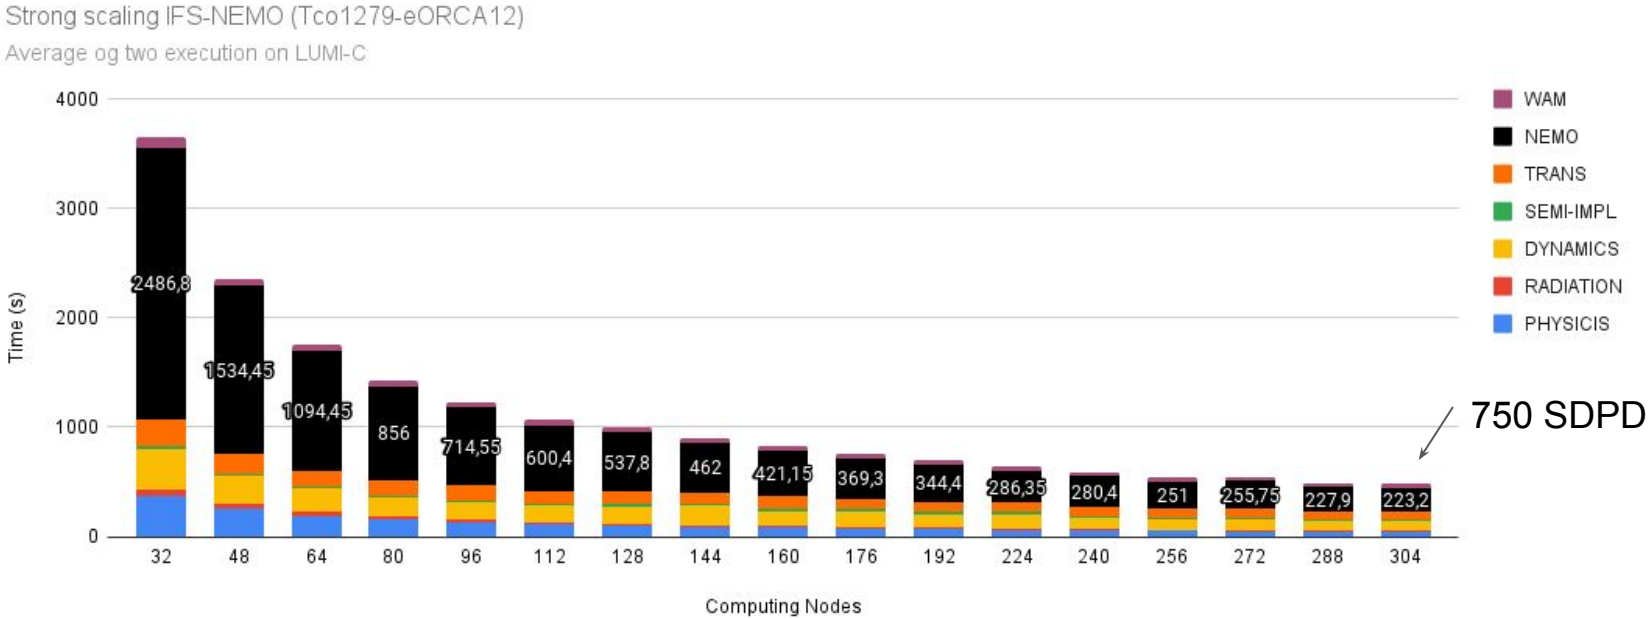

Development Configuration (I/O servers)

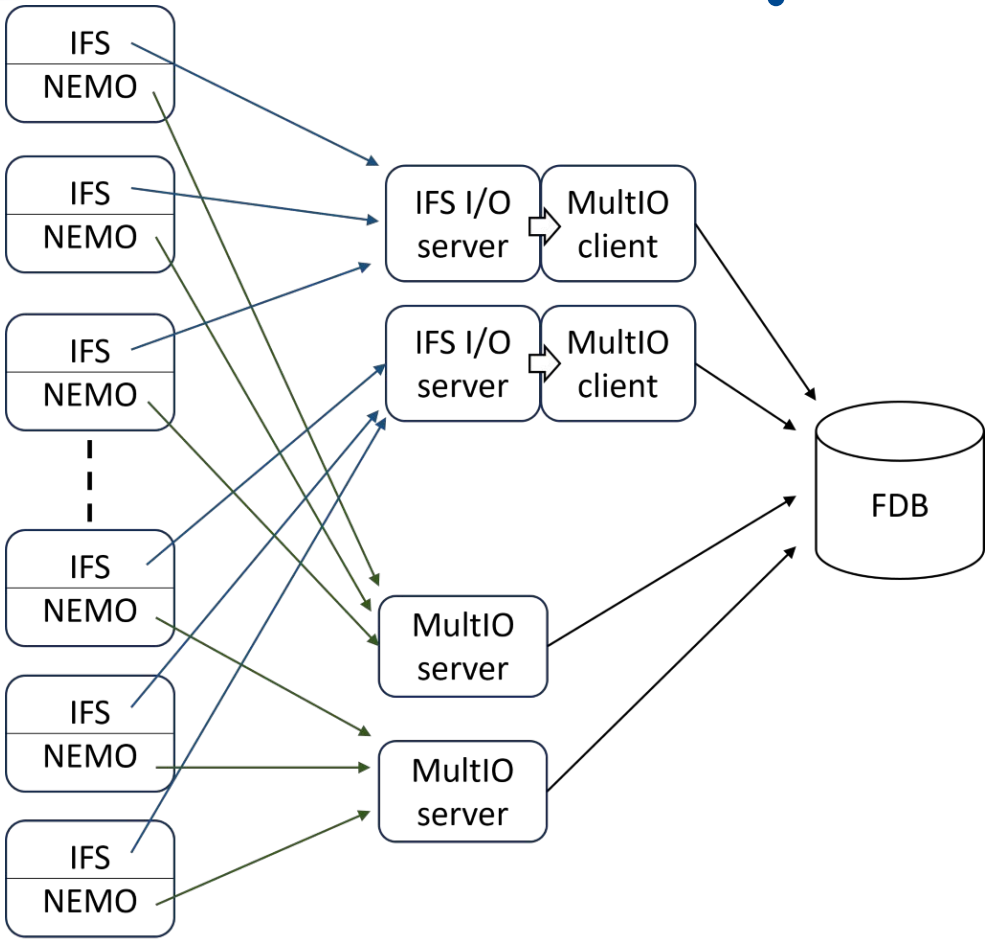

Development Configuration (I/O servers)

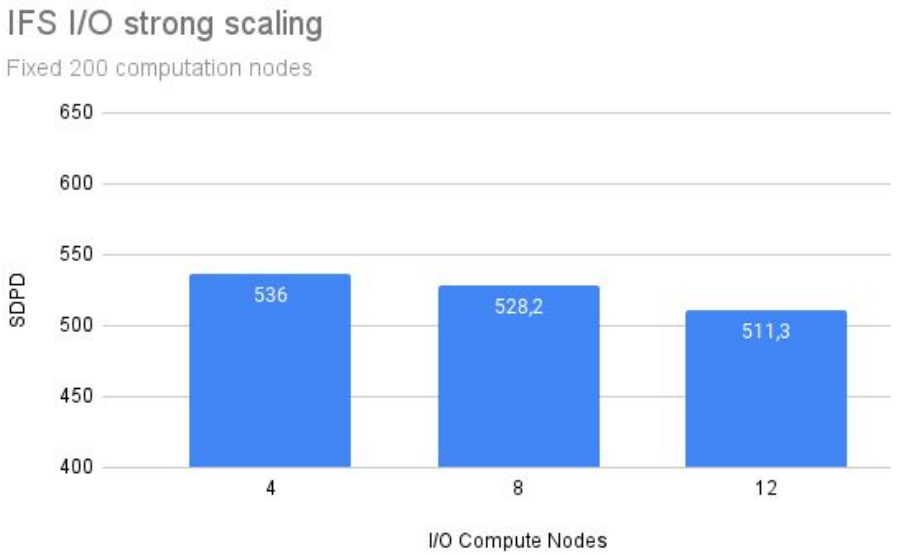

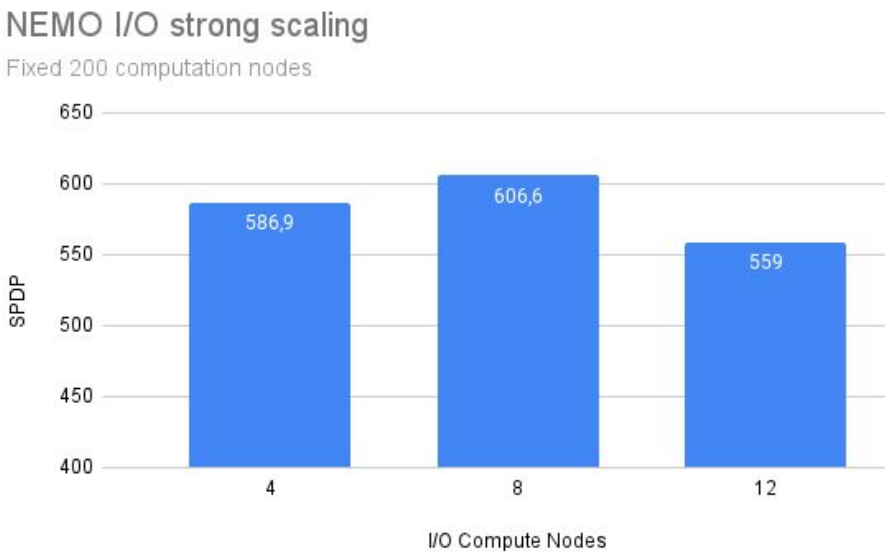

Development Configuration (I/O servers)

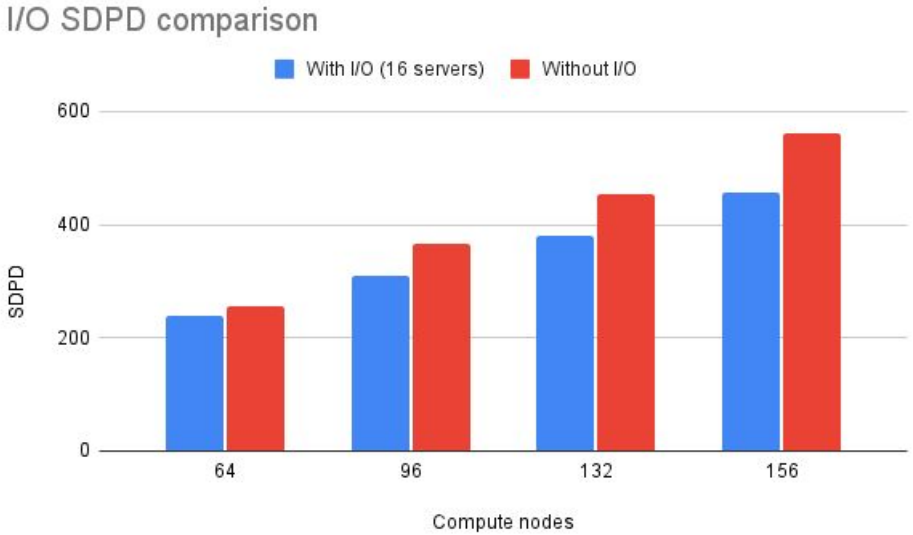

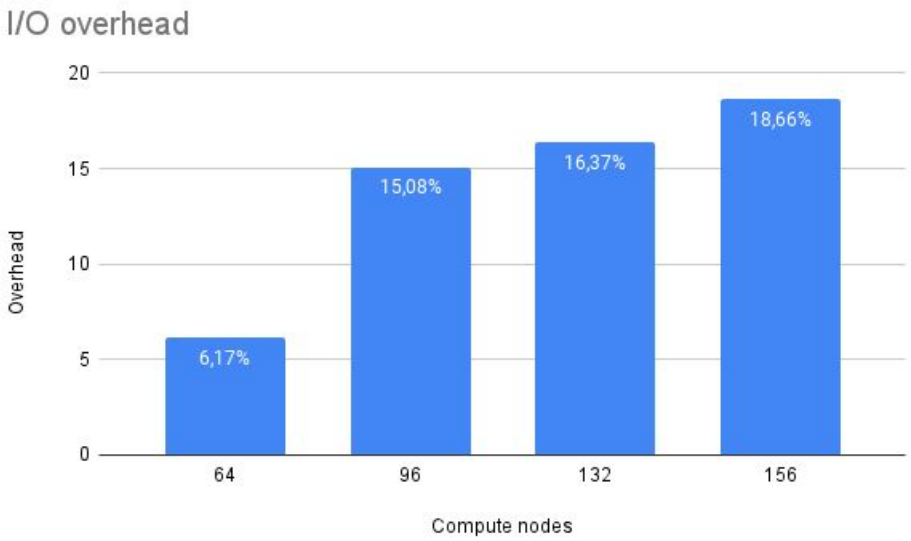

Optimizations

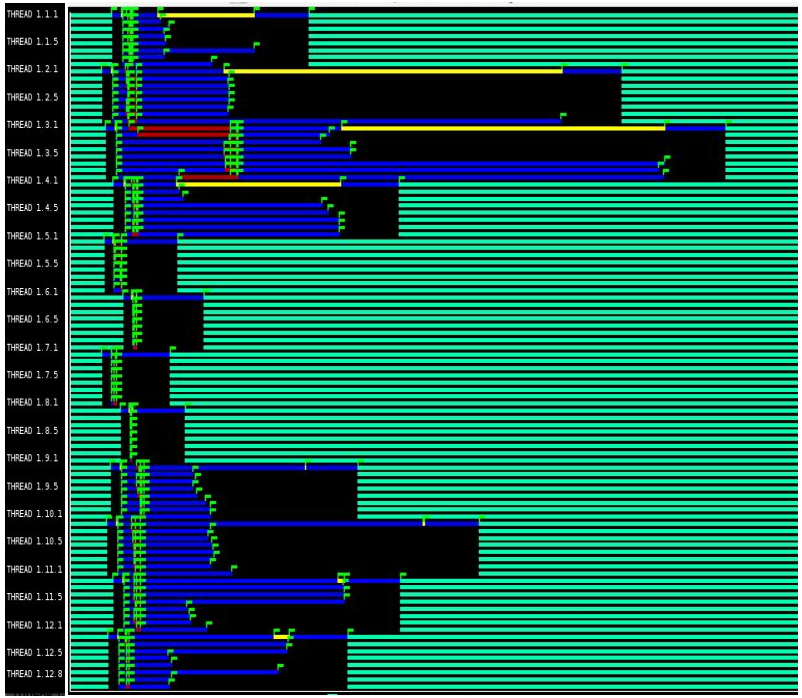

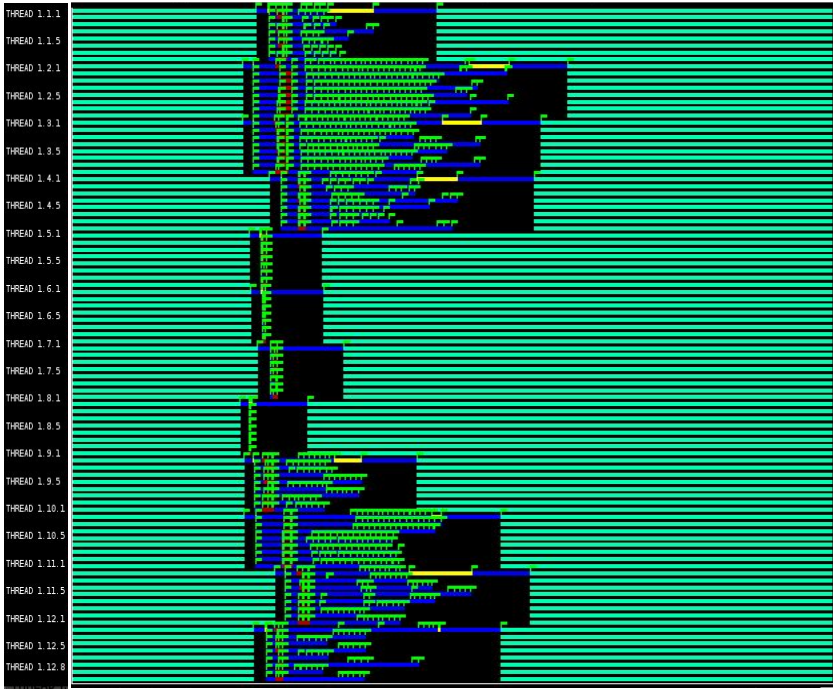

Optimizations

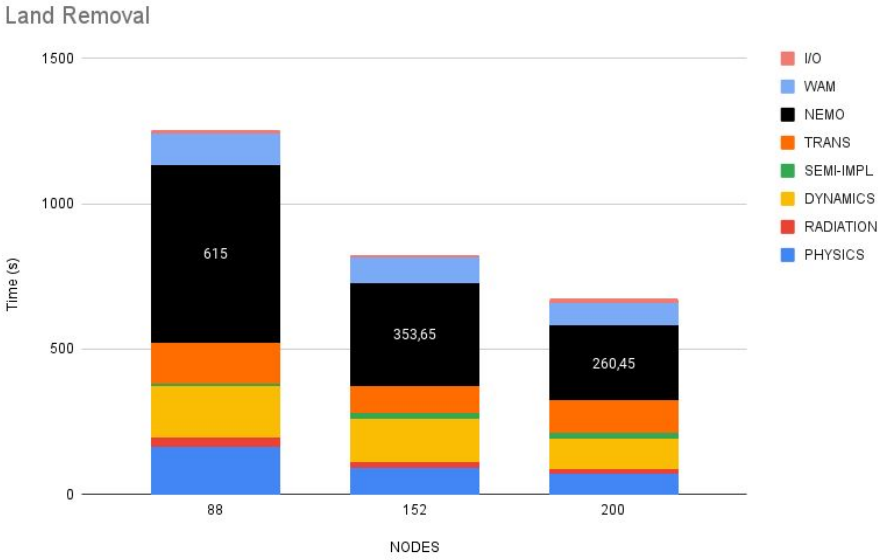

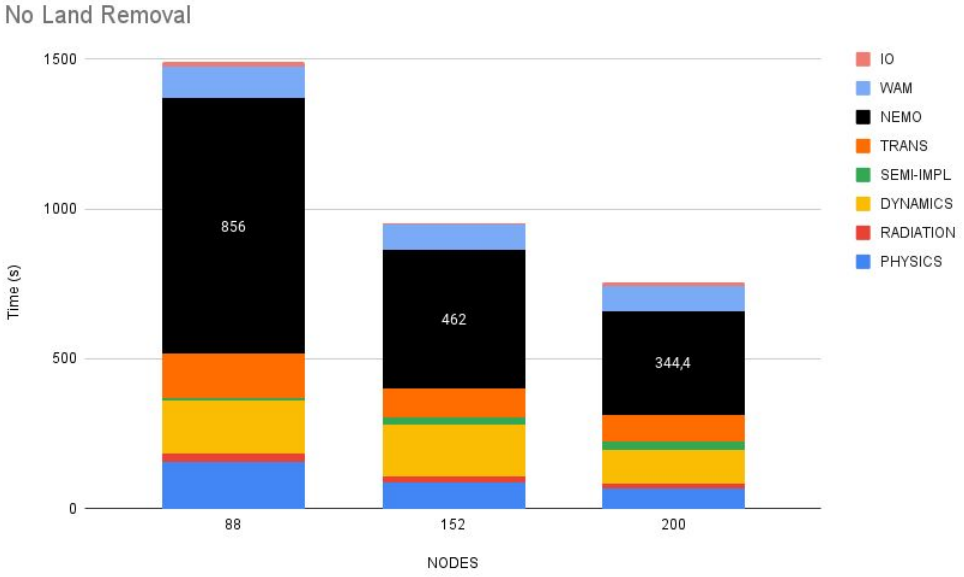

Optimizations

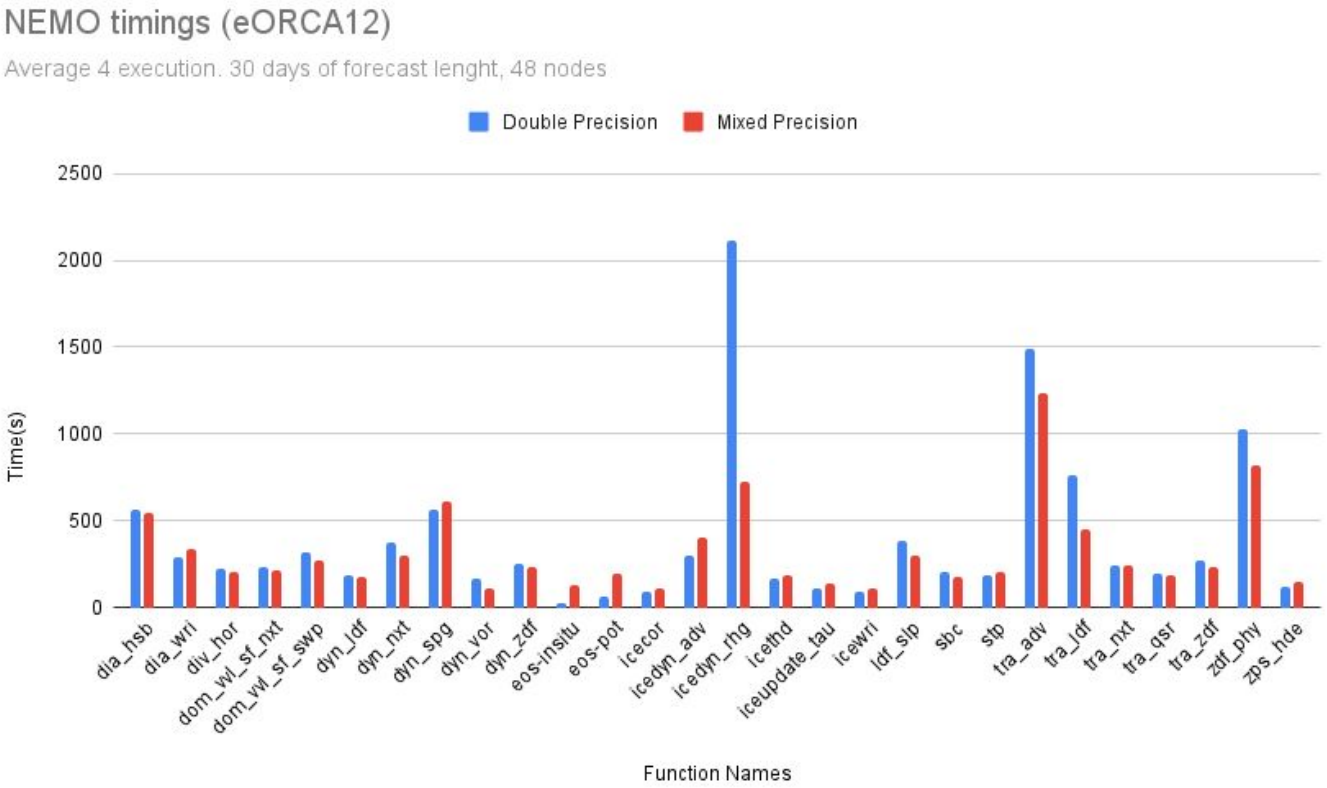

Credits

- Victor Correal & Vijendra Singh (IFS-NEMO Scaling + Profiling)
- Carlos Peña & Xavier Yepez (I/O profiling + Land Removal)
- Gladys Utrera & Joan Vinyals (OpenMP optimizations)
- Stella Paronuzzi & Oscar Michel (Mixed precision)
- Alexey Medvedev & Rommel Quintanilla (GPU porting)

Production Configuration

Using 222 computing nodes, we observed 220 simulated days per day (equivalent to 0.6 SYPD)

Running on LUMI-G

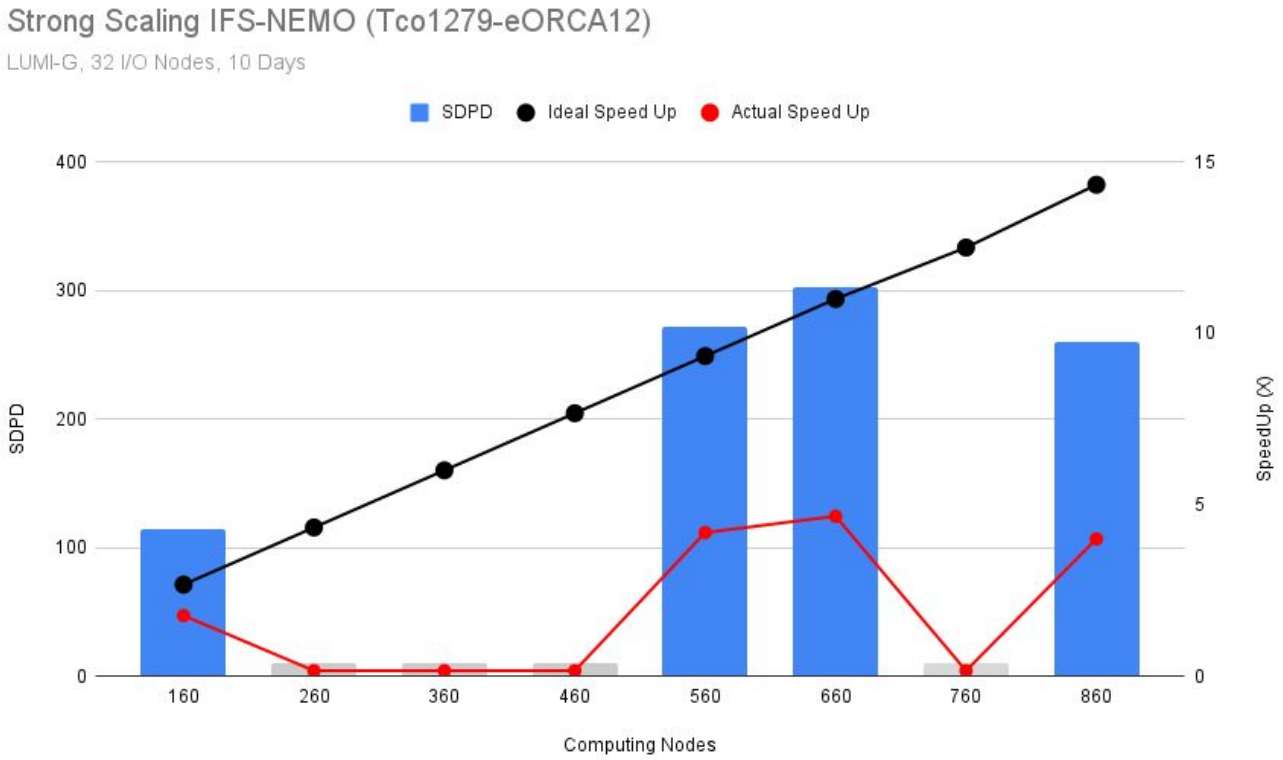

Conclusion

- An overall improvement in performance has been achieved (~20%)
- Data layout changing routines r remain the most expensive ones.

Future work

- Replace the OpenACC loops
- in the data transposition routines by HIP kernels.

- IFS-NEMO Scaling + Profiling (Victor Correal & Vijendra Singh)
- I/O profiling + Land Removal (Carlos Peña & Xavier Yepez)
- OpenMP optimizations (Gladys Utrera & Joan Vinyals)
- Mixed precision (Stella Paronuzzi & Oscar Michel)
- GPU porting (Alexey Medvedev & Rommel Quintanilla)

Mixed precision validation