



ICON-Modelling Climate and Weather

LUMI hackathon April, 2023 re- sults

Luis Kornblueh, Lukas Kluft, Sergey Kosukhin
Rene Redler and Claudia Frauen (DKRZ)
April 21, 2023

- Hadley test (dynamical core and advection), runs and got faster in the course of the week, access to profiles by AMD tools work
- Aquaplanet (reduced radiation, no land), radiation speed up by more than 90% (by workarounds due to compiler bugs), in case with full radiation it runs when switching off async mode.
- running test with land opens a new compiler bug for which we could not find a workaround yet.
- fixed bug in logical expression in radiation code


Profiling

Model resolution R2B9 (5 km)

Nodes	Lumi initial	Lumi fixed	Levante
2	130 s	98 s	-
4	67 s	51 s	-
8	35 s	27 s	28 s
16	-	14 s	15 s
32	-	7.2 s	10 s
48	-	5.7 s	7.7 s
64	-	4.8 s	-
128	-	3.8 s	-

Model resolution R2B1 (1.25 km)

Nodes	Lumi fixed
32	113 s
64	63 s
128	28 s
256	15 s
512	29 s



Additional steps forward

- speeding up compilation of the land source code (bad for coffee farmers)
- access to hardware counters via papi (and within contained AMD profiling libraries for GPU counter access) in a model specific timing library. Added an additional feature unit-test for this
- SLURM hetjobs on lumi still broken
- the weak CI based on buildbot is now enrolled and working on lumi (Thanks to Ralf Müller,DKRZ and Ulf Tigerste, CSC)
- Knowledge gain: we know now much more on the problem details.

Still open questions

- copies generated in subroutine calls requiring contiguous arrays in the callee are not created on the device but on the host. Compiler bug!
- warp size used in ACC directives exposed as TILE(...) requires differentiation between NVidia and Cray but solution is not clear yet as changing the Cray TILE(...) does show hardly an effect - solution open: runtime difference A100/MI250 300%.

Acknowledgments

Thanks a lot to John Levesque, and of course Harvey, (HPE/Cray) as well as Samuel Antao (AMD) including all their support in the background. Jussi Enkovaara, Tuomas Lunttila for all code and GPU related support, and Jenni Kontkanen for her actions in the background (CSC).

CSC support for the workshop is excellent!