

Porting AROME on AMD platform

Ryad El Khatib
DR/CNRM/GMAP/ALGO

31.03.2020

Plan

- 1. Configuration of tests**
- 2. Raw comparison of AMD vs Intel on AROME**
- 3. Sensitivity to certain parameters like NPROMA**
- 4. Global impact of recent optimisations made on Intel**
- 5. Best of both cycles (46T1 et 47T1) for computation parts**
- 6. Best of both worlds on cycle 47T1 for computation parts**
- 7. Conclusion**

Configurations of tests

- Compiler Intel 18 + IntelMPI 18 + MKL ~19
- 25 nodes AMD for 80 nodes Intel,
Giving the same total number of cores
- Quasi-operational namelist but I/O disabled
=> identical configurations
in number of MPI tasks
in number of Open-MP threads Open-mp per task (4)
- 24 hours forecast, but longer time-to-solution
=> more weight on compute parts than communications
- Cycles : 46T1 (next candidate for operations) ; 47T1 (optimized)


AROME cycle 47T1 on AMD

Following ATOS recommendations :

- Light depopulation (31 task cores instead of 32)
- Use of FFTW instead of FFT992
- NPROMA reduced to 16

AROME cycle 47T1 on AMD Rome - 25 nodes (4 threads per MPI task)

Speedup against reference NPROMA=-50



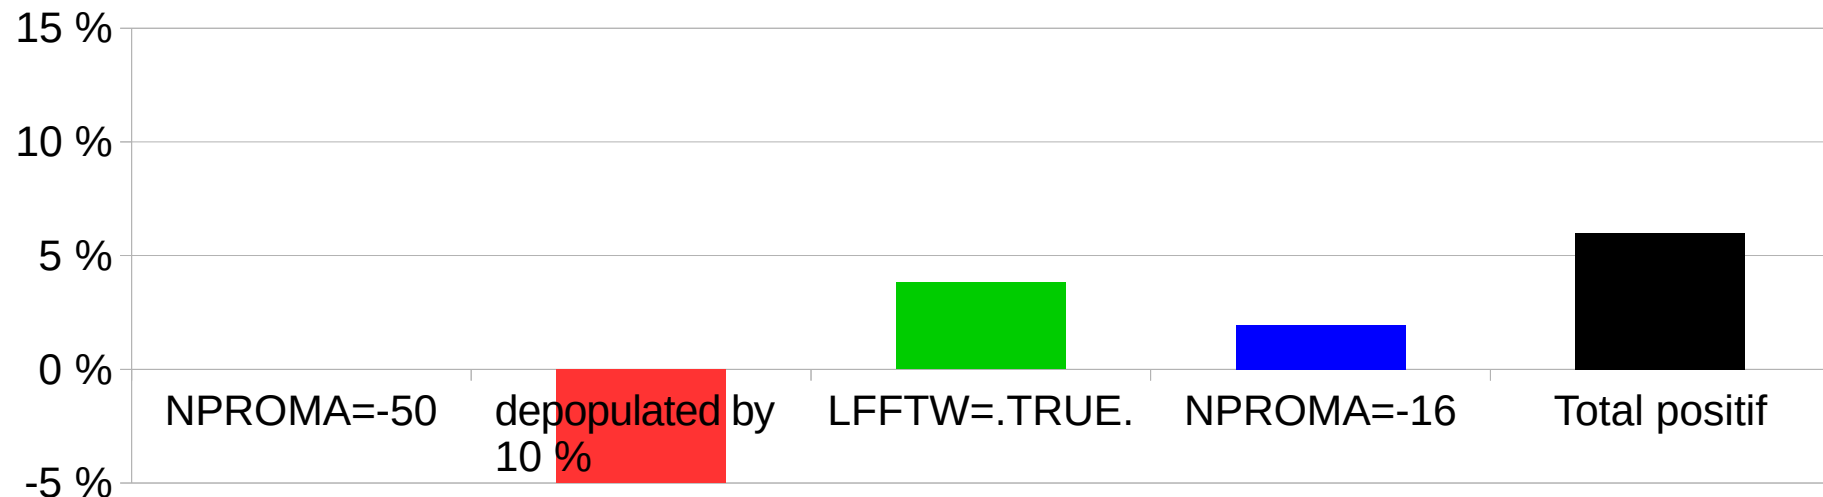
AROME cycle 47T1 on Intel

Application of ATOS recommendations for AMD on Intel :

- **Depopulation : detrimental**
- **FFTW : (unexpectedly) positive**
- **NPROMA reduction : slightly positive**

AROME cycle 47T1 on Intel Broadwell - 80 nodes (4 threads per MPI task)


Speedup against reference NPROMA=-50



AROME cycle 47T1 : AMD vs Intel

With ATOS recommendations the speedup of a core is 8 to 12 %

AROME cycle 47T1 - 800 MPI tasks x 4 OMP threads - I/O disabled




Warning : the impact of MPI communications is ignored !


Fine tuning of NPROMA on AROME (cycle 47T1)

**Prefer multiples of 4
(cf AVX2 and real in double précision)**

relative speedup against NPROMA=16 (Depopulation : 6,25%)




NPROMA dependency to cycle and cpu



Sensitivity to depopulation

Impact of depopulation on speedup




- Depopulate, but not too much !

Benefit for AMD of recent optimisations made for Intel

Do the optimisations made on **Intel** for cycle 47T1 benefit to **AMD** ?

Cycle 47T1+ against Cycle 46T1

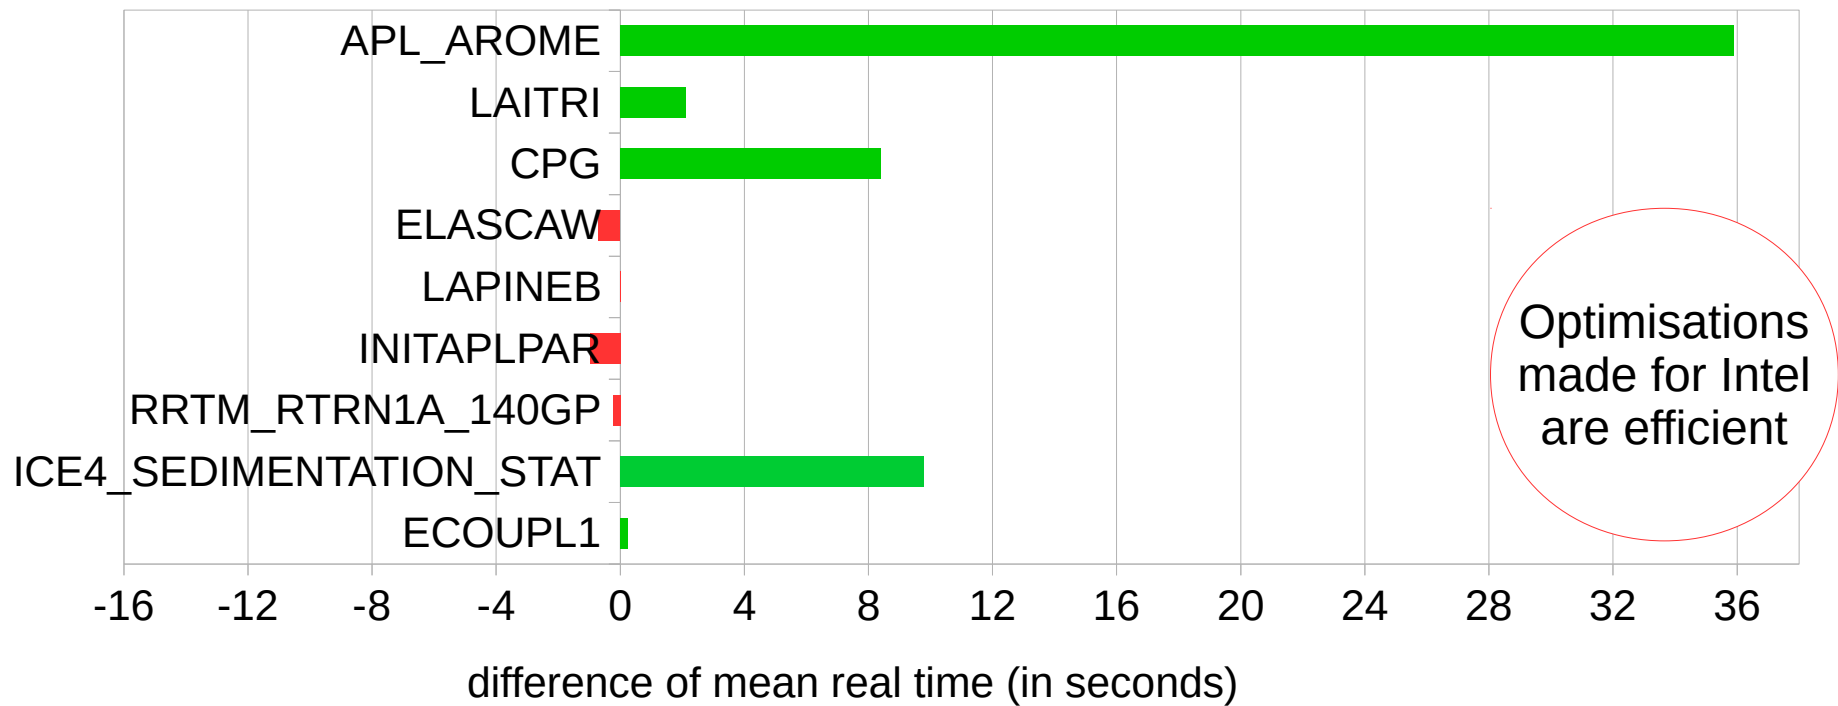


The answer is : yes ... but marginally

Best of both cycles on AMD with fixed NPROMA

Best of both cycles 46T1 and 47T1+ on AMD (NPROMA=-16)

reduction of mean real time (s) for most expensive calculation subroutines (from top to bottom)

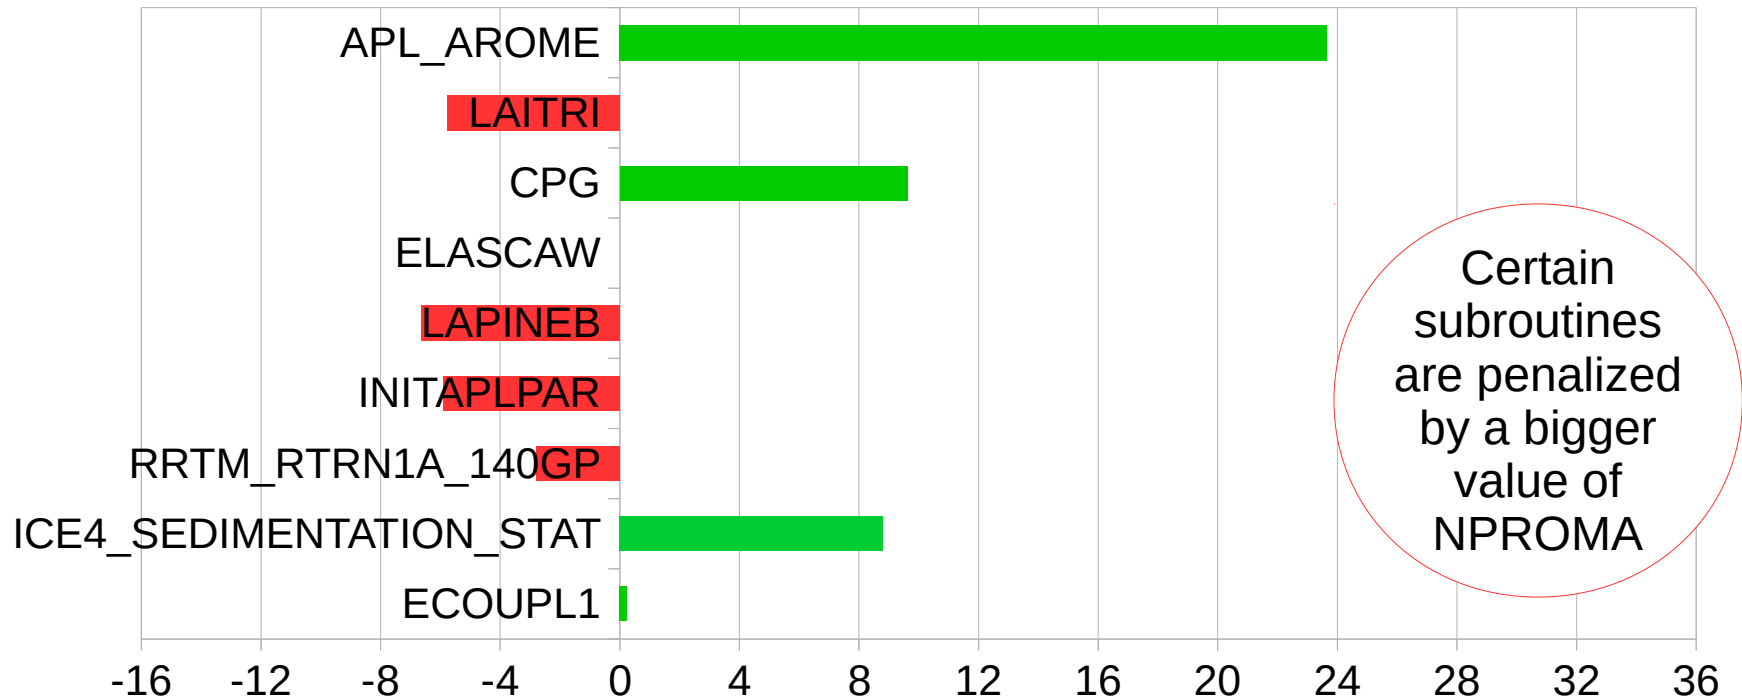


Penalization of the code

Optimization of the code

Best of both cycles on AMD with best NPROMA

NPROMA=-16 for cycle 46T1, NPROMA=-24 for cycle 47T1




Negative contribution by raising NPROMA

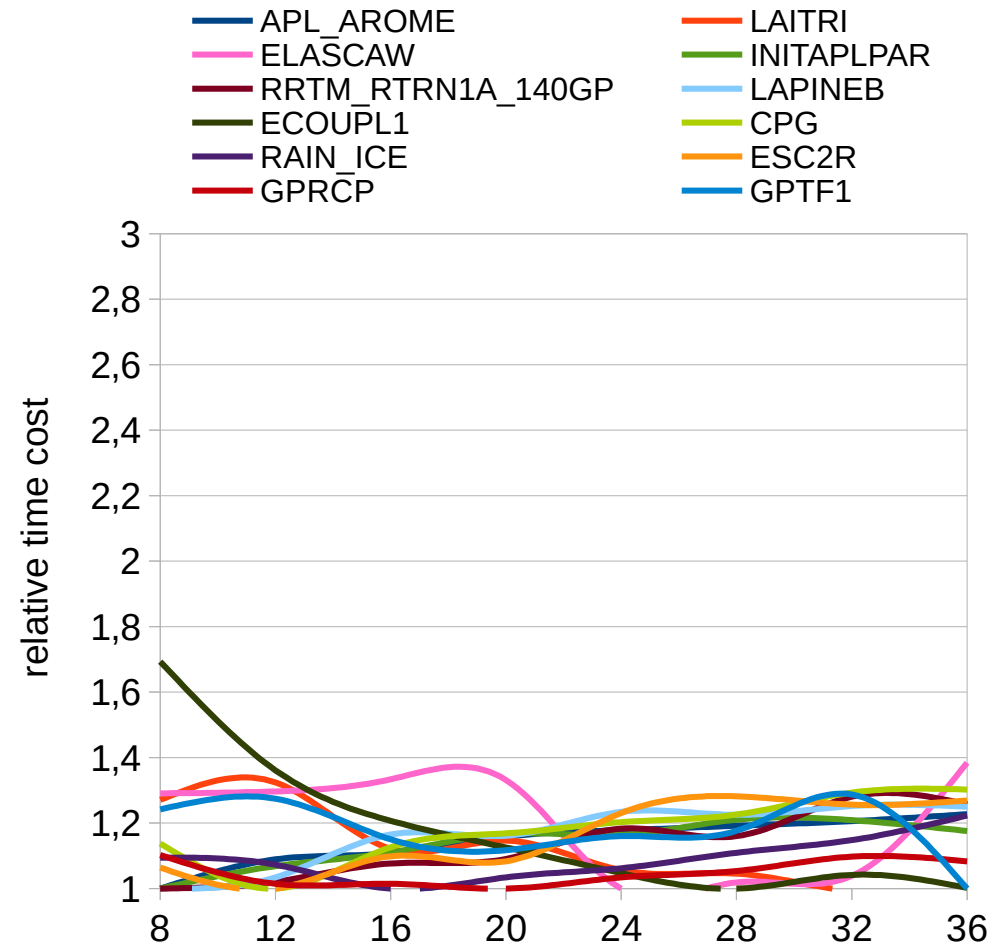
Optimisation of the code

Sensitivity to NPROMA (AROME CY47T1) on Intel

Absolute sensitivity

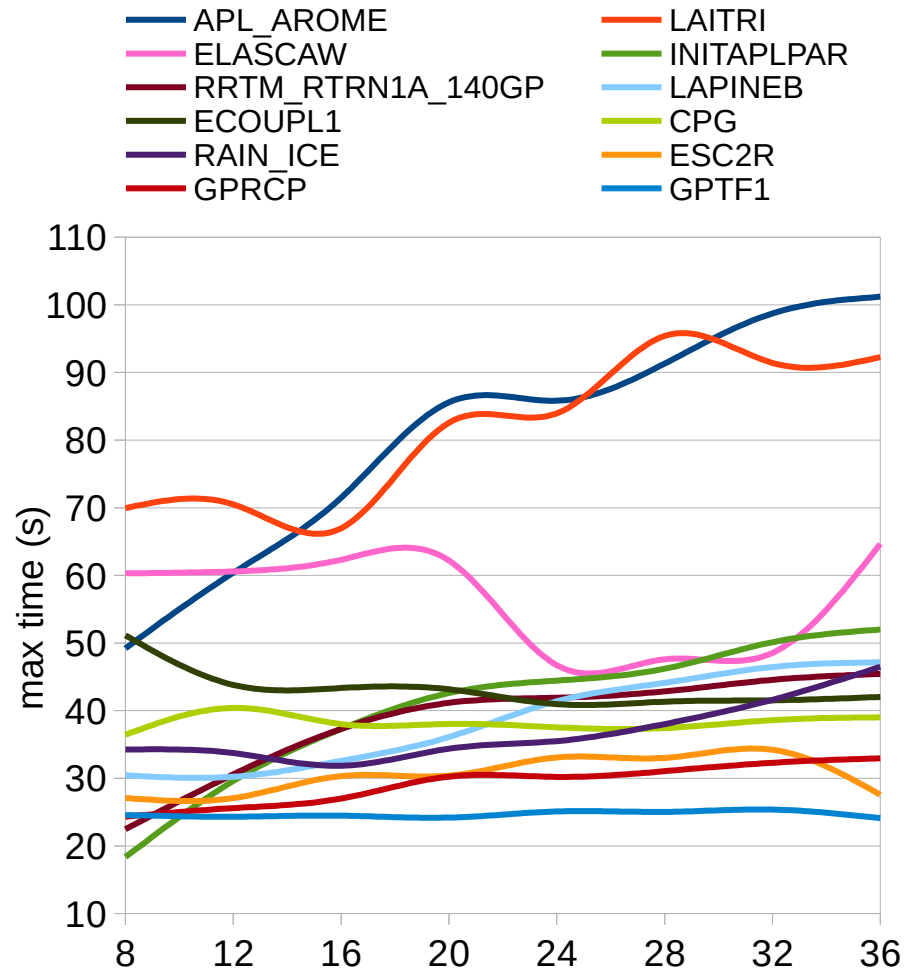


Relative sensitivity

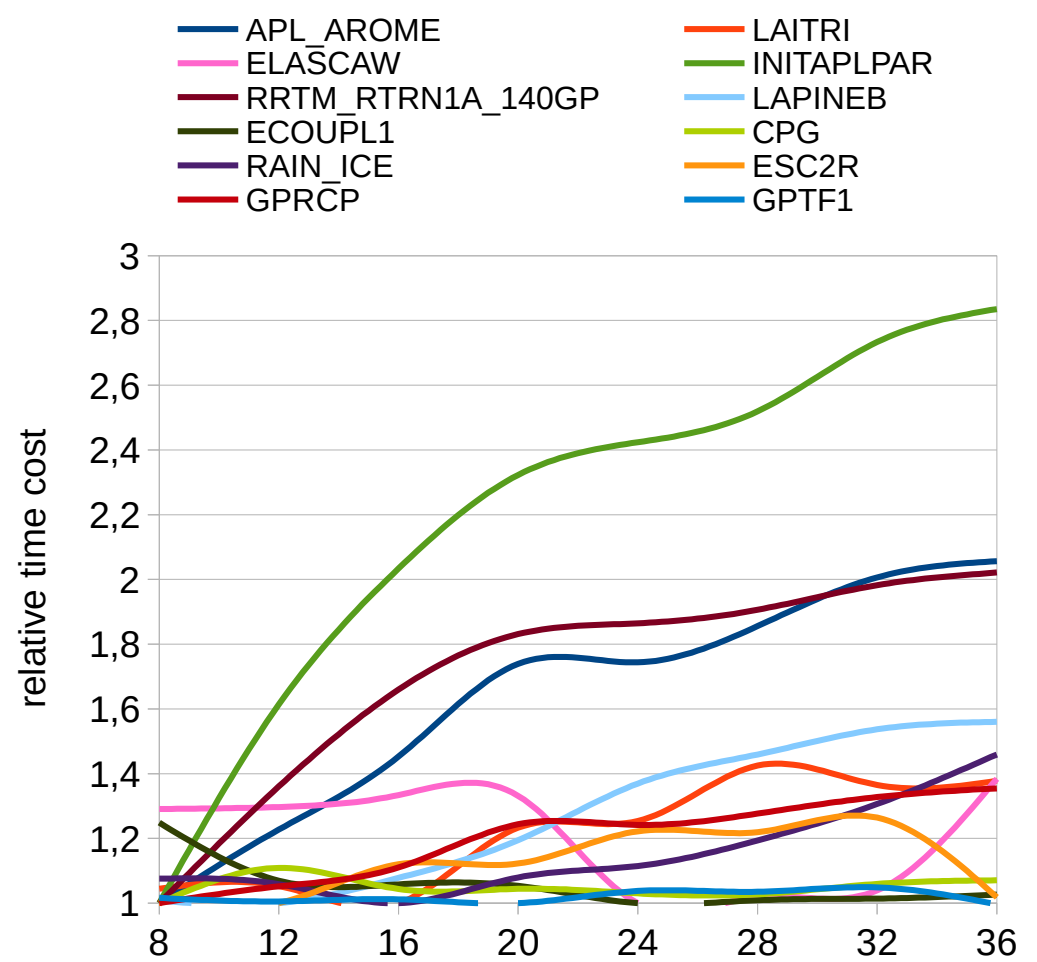


Sensitivity to NPROMA (AROME CY47T1) on AMD

Absolute sensitivity




Relative sensitivity



Best of both worlds Intel et AMD

Best of both worlds Intel and AMD on AROME cycle 47T1+
Reduction of real time (s) for most expensive calculation subroutines (from top to bottom)



Conclusions

- **FFTW is back (but beware of licensing or reproducibility)**
- **A light depopulation of nodes is beneficial to AROME**
- **Optimizations made for Intel fit AMD**
- **Optimal values for NPROMA on AMD : -16 or -24
On Intel : -16, -24 or -32**
- **Certain subroutines may be excessively sensitive to the value of NPROMA**



Météo-France

ryad.elkhatib@meteo.fr

www.meteofrance | [🐦@meteofrance](https://twitter.com/meteofrance)