



Development of RAMSES in the context of the EuroHPC project SPACE

Tine Colman, Benoît Commerçon (CRAL, CNRS)

ASNUM 2025



Co-funded by
the European Union

Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441.



EuroHPC
Joint Undertaking

SPACE

- Scalable Parallel Astrophysical Codes for Exascale
- 4 year European Center of Excellence project
- Goal: prepare state-of-the-art astro codes to efficiently use exascale computation resources
- 7 codes: OpenGadget, ChaNGa, PLUTO, iPIC3D, RAMSES, FIL, BHAC
- Collaboration: research institutes, computing centers, vendors



RAMSES in SPACE

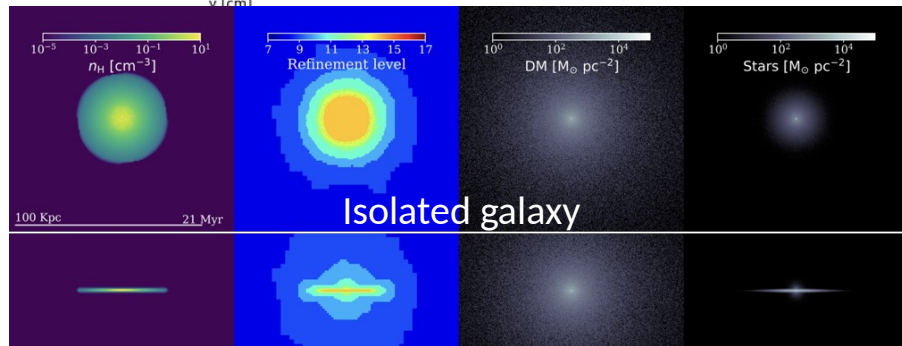




Goal: improve time-to-solution and scaling

Progress:

- Benchmarking on all EuroHPC machines
- Low level optimisations
 - neighbor searching, godunov solver
- OpenMP
 - sedov & cosmo done, amr wip
- Also tests, doc, code cleaning

Selected representative use cases:



1 Open ✓ 48 Closed Author Label Projects Milestones

- Overhaul of compiler flags in Makefile ✓ SPACE
#185 by tinecolman was merged 3 days ago • Approved
- Combine hydro, mhd and rhb versions of hydro_parameters.f90 ✓ code quality SPACE
#184 by tinecolman was merged 5 days ago • Approved
- Combine hydro and mhd versions of output_hydro.f90 ✓ code quality SPACE
#183 by tinecolman was merged 5 days ago • Approved
- further document the order of hydro variables ✓ documentation SPACE
#178 by tinecolman was merged last week • Approved
- combine the hydro and mhd versions of init_flow_fine.f90 into one ✓ code quality SPACE
#174 by tinecolman was merged last week • Approved
- fix after merge ✗ openmp SPACE
#171 by tinecolman was merged on Jul 15
- change slope_type in imhd-tube and imhd-tube-nener tests ✓ CI SPACE
#169 by tinecolman was merged 5 days ago • Approved
- Add 1D advection test ✓ SPACE test case
#166 by tinecolman was merged on Jul 9 • Approved
- Improvements to parameter sweep of sod-tube test ✓ CI SPACE
#165 by tinecolman was merged on Jul 9 • Approved
- regenerate reference solution for stellar-HII with 2 cpus ✓ CI SPACE
#160 by tinecolman was merged on Jul 1 • Approved
- optimizations for hydro solver on uniform grid ✓ performance SPACE
#159 by tinecolman was merged 4 days ago • Approved
- Optimize ctotrim ✓ performance SPACE
#158 by tinecolman was merged 5 days ago • Approved

Commits from

46 PRs merged on dev

+ 10 PRs merged on openmp

Various stuff in development on

<https://github.com/tinecolman/ramses>


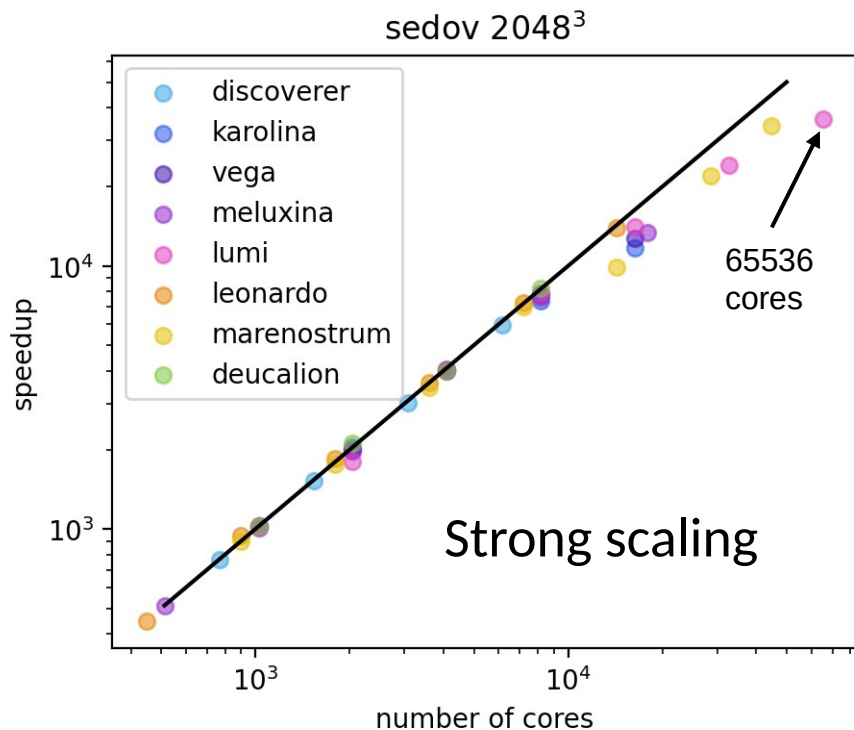
Benchmarking

EuroHPC systems (CPU partitions)




system (country)	#nodes	RAM/node [GB]	processor (cores)	chosen compiler
Discoverer (Bulgaria)	1128	256, 1024	2 x AMD EPYC 7H12 (2x64)	GNU + OpenMPI
Karolina (Czech rep.)	720	256	2 x AMD EPYC 7H12 (2x64)	GNU + OpenMPI
Vega (Slovenia)	960	256, 1024	2 x AMD EPYC 7H12 (2x64)	GNU + OpenMPI
MeluXina (Luxembourg)	573	512, 4096	2 x AMD EPYC 7H12 (2x64)	GNU + OpenMPI
Lumi (Finland)	2048	256, 512, 1024	2 x AMD EPYC 7763 (2x64)	GNU + MPICH (Cray)
Leonardo (Italy)	1536	512	2 x Intel Xeon Platinum 8480+ (2x56)	Intel + Intel MPI
Mare Nostrum (Spain)	6192	256, 1024	2 x Intel Xeon Platinum 8480+ (2x56)	Intel + Intel MPI
Deucalion (Portugal)	-	-	-	-

Scaling on EuroHPC systems



Performance on EuroHPC systems



Smaller AMD systems similar

Intel systems typically slightly faster

Leonardo best timings (but often node problems)

Performance benchmarking



GOAL: create script to make benchmarking ramses easy...

- for users, that need scaling plots for proposals
- for developers, that need to verify performance
- for maintainers, to verify code quality on PR

Benchmark package for RAMSES



<https://github.com/tinecolman/ramses-benchmarks-prototype2>

```
✓ RAMSES-BENCHMARKS-PROTOTYPE2
> analysis
> doc
> HPCclusters
> results
> setups
◆ .gitignore
$ launch_benchmark_suite.sh
① README.md
```

- **Main script:** compile code, create and launch jobs
- Presets for different clusters
- Collection of predefined setups
- Scripts to analyze results
- Manual
- Results and figures
- Front page of repo

Using the benchmark script





```
$ ./launch_benchmark_suite.sh -c meluxina -l "1" -h refactor_3cube_nbor_utils -t 1,2 -a p200525 -i 3
```

Parameters:

- -c: name of the preset cluster
- -l or -n: specify a list or maximum number of nodes to use
- -t: specify which benchmark setups to run
- -i: how many times to run the same benchmark (1 by default)
- -a: allocation number (can also be determined automatically)
- -w: also do weak scaling
- -m: openMP threads

```
Switched to branch 'openmp_hydro_unigrid'
Your branch is up to date with 'origin/openmp_hydro_unigrid'.
#####
#   Launching RAMSES performance benchmarks   #
#####
Multiple projects found. Please select one:
1) nocredit
2) p200525
#? 2
You selected allocation ID: p200525
Will launch the following benchmarks:
[ 2] setups/sedov
-----
Test 1/1: sedov
Compile job submitted with Job ID: 2577666
Waiting for compile job to finish...
Compile job completed successfully.
Launched benchmark sedov on 1 nodes [JOB ID 2577669]
Launched benchmark sedov on 1 nodes [JOB ID 2577670]
Launched benchmark sedov on 1 nodes [JOB ID 2577671]
Launched benchmark sedov on 2 nodes [JOB ID 2577672]
Launched benchmark sedov on 2 nodes [JOB ID 2577673]
Launched benchmark sedov on 2 nodes [JOB ID 2577674]
Launched benchmark sedov on 4 nodes [JOB ID 2577675]
Launched benchmark sedov on 4 nodes [JOB ID 2577676]
Launched benchmark sedov on 4 nodes [JOB ID 2577677]
Submitted batch job 2577678
```

Benchmarking for users

 tinicolman / ramses-benchmarks-prototype2

Benchmark setups

- [SEDOV](#): classic Sedov explosion in 3D
- COSMO: dark matter-only cosmological box on a uniform grid

Benchmark results per cluster

Cluster	Benchmark results
MeluXina	SEDOV
MareNostrum	


How to do your own benchmarking

- [How to use the benchmark script](#)
- [How to process the results](#)
- [How to add a setup](#)
- [How to add a cluster](#)

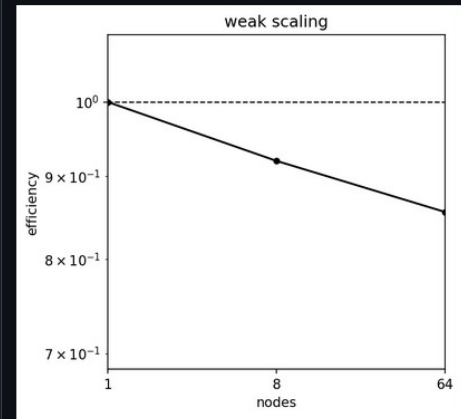
Benchmark results meluxina

SEDOV

Strong scaling



Weak scaling



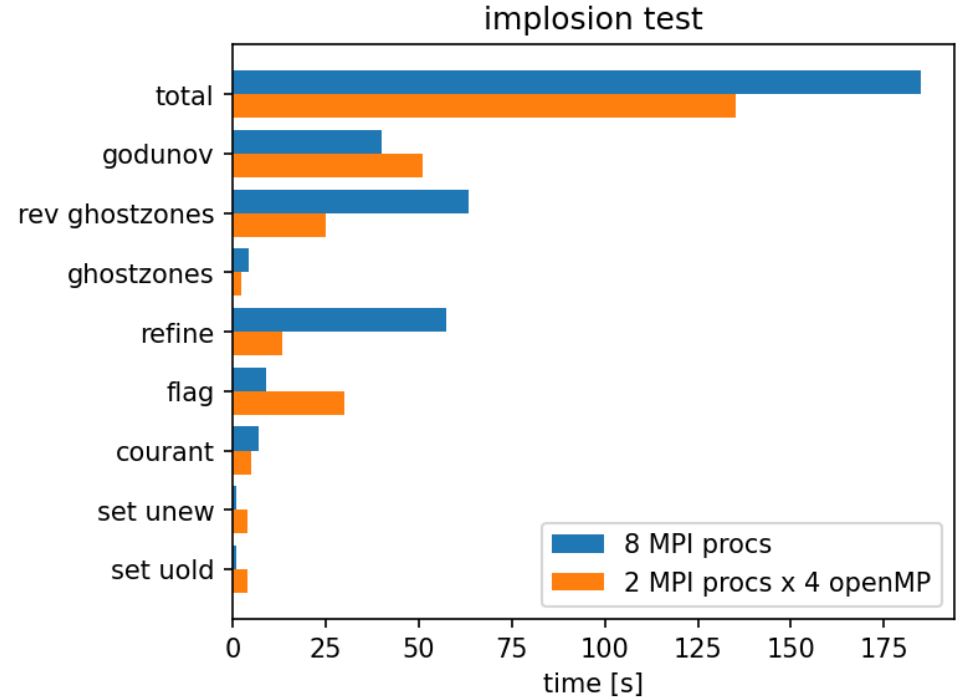

Strong scaling

nodes	efficiency
1	100%
2	97%
4	97%
8	95%
16	91%
32	92%
64	79%


Weak scaling

nodes	resolution	efficiency
1	256	100%
8	512	92%
64	1024	86%

Benchmarking for developers



Benchmarking for maintainers




Optimization

Optimization of neighbor searching



Elementary routine used in many parts of the code
Identified as bottleneck in COSMO use case
=> refactor for efficiency

cosmo profiling		Inclusive Time w.r.t. Wall Time(s)	
Name		old	new
gauss_seidel_mg_fine		15.31	15.17
interpolate_and_correct_fine		12.08	12.15
move1		9.89	9.75
sync		9.32	9.60
mca_btl_vader_poll_handle_frag		9.32	7.83
get3cubefather		17.86	0.28
cic_amr		6.92	6.64
cmp_residual_mg_fine		6.23	6.15
gauss_seidel_mg_coarse		4.16	4.03
check_tree		4.19	3.97
get3cubepos		3.85	3.79



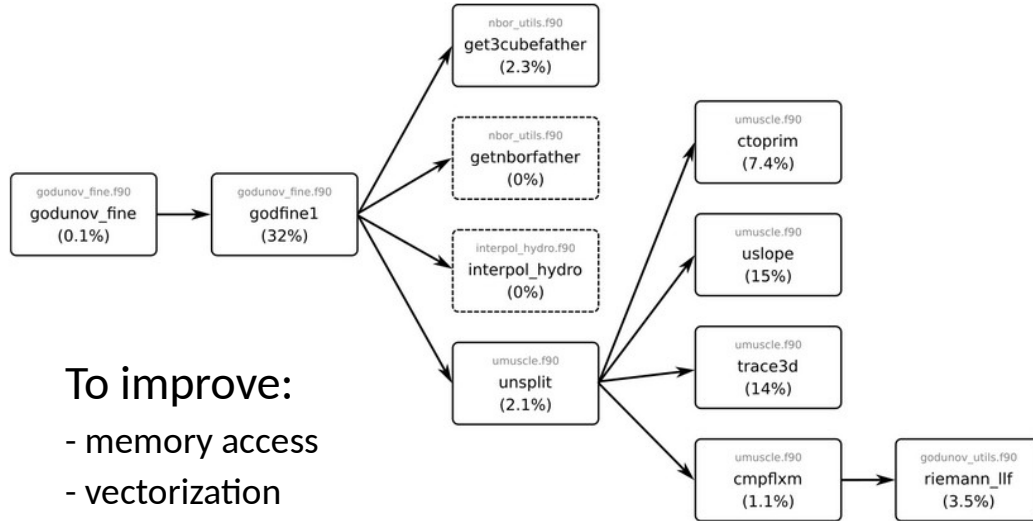
[SPACE] Optimize and document nbor utils ✓

#108 by tinecolman was merged 3 weeks ago • Approved


Optimization of Godunov solver



Initial profiling of sedov use-case



Progress



23% speedup for sedov test

MPI + OpenMP implementation

Shared memory parallelism with
OpenMP inside nodes



=> reduce number of MPI domain

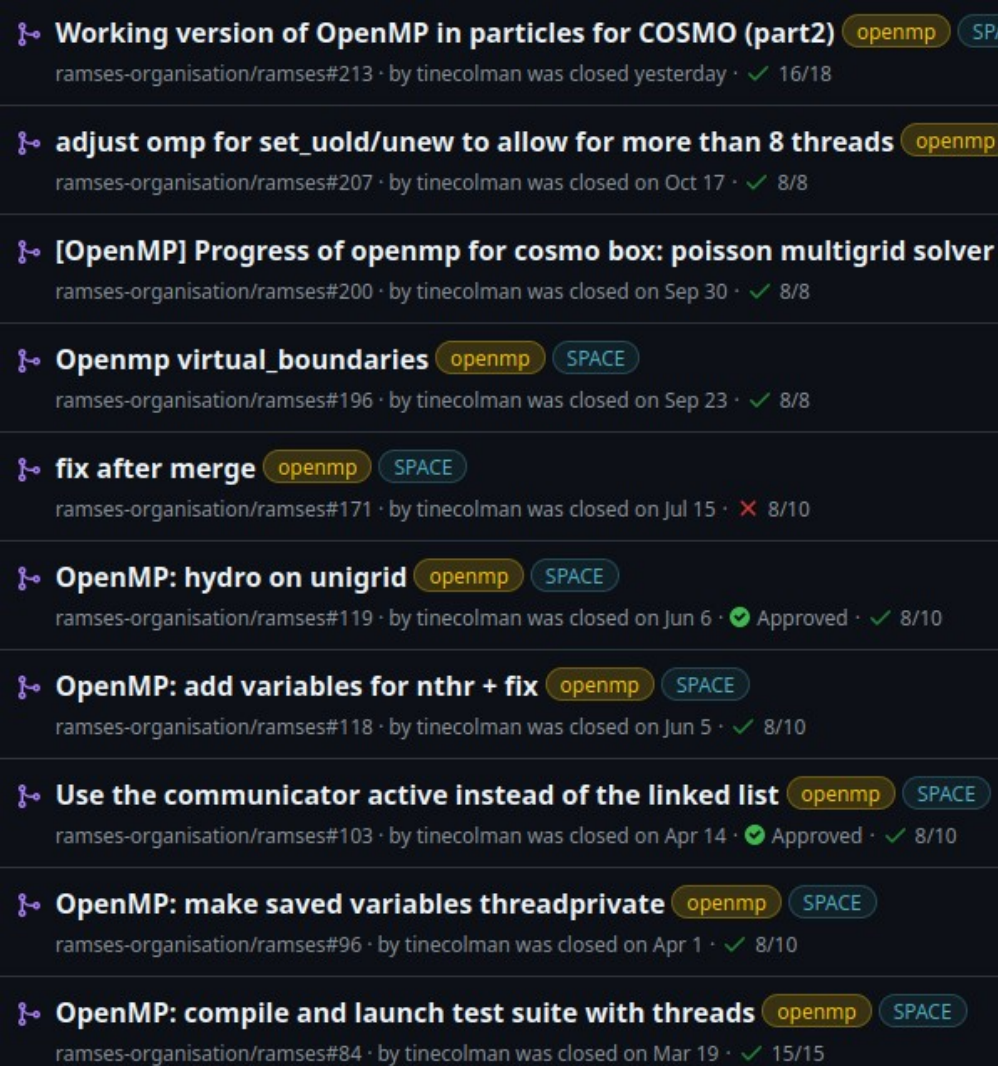
=> decrease communication &
reduced memory imprint ghost zones

=> less time spent communicating

=> **improved scalability**

Starting from RAMSES-yOMP





openmp branch



Dedicated branch on RAMSES GitHub

Done:

- Hydro on uniform grid (SEDOV)
- DM-only on uniform grid (COSMO)


WIP:

- Refinement
- MHD
- RT & cooling (STROMGREN)

OpenMP for SEDOV




Full MPI vs OpenMP speedups



AMD Milan architecture, 128 cores (256 threads) / node, AOCC compiler


sedov 1024³ on MeluXina



OpenMP for COSMO



Cosmo on MeluXina



Modules:

- Poisson multigrid solver
- Force calculation
- Particle-mesh
- Moving of particles

Result on 64 nodes:

x2 faster

scaling 37% → 74%

CI/CD improvements



- SNO working group: test suite
- 41 PR related to tests/CI merged
- 16 new tests added, 5 tests updated
- Various infrastructure improvements
- Test suite coverage
 - Initially: 47.63 % (16828/35334 lines)
 - Current: **63.02 %** (21428/34000 lines)

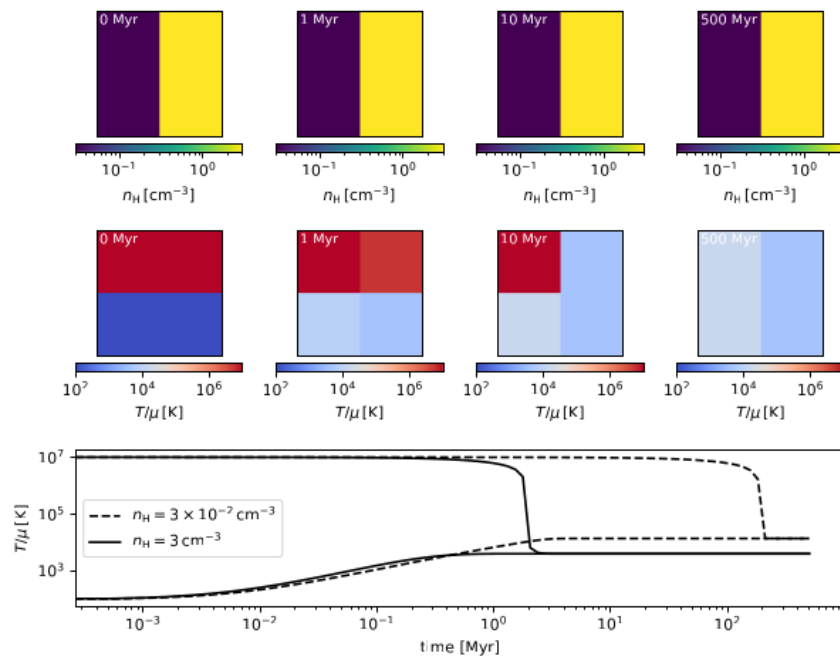


Figure 3: hydro/cooling-eq test

Summary



Benchmarking RAMSES


- on all EuroHPC systems
- Development of scripts

Optimizations: low level + OpenMP

- Hydro uniform grid:
 - 20% base speedup
 - 5-10% scaling improvement
- Cosmo uniform grid:
 - 10% base speedup
 - 30-40% scaling improvements

Test suite improvements

COSMO



Acknowledgement & Disclaimer



Funded by the European Union. This work has received funding from the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain under grant agreement No 101093441.

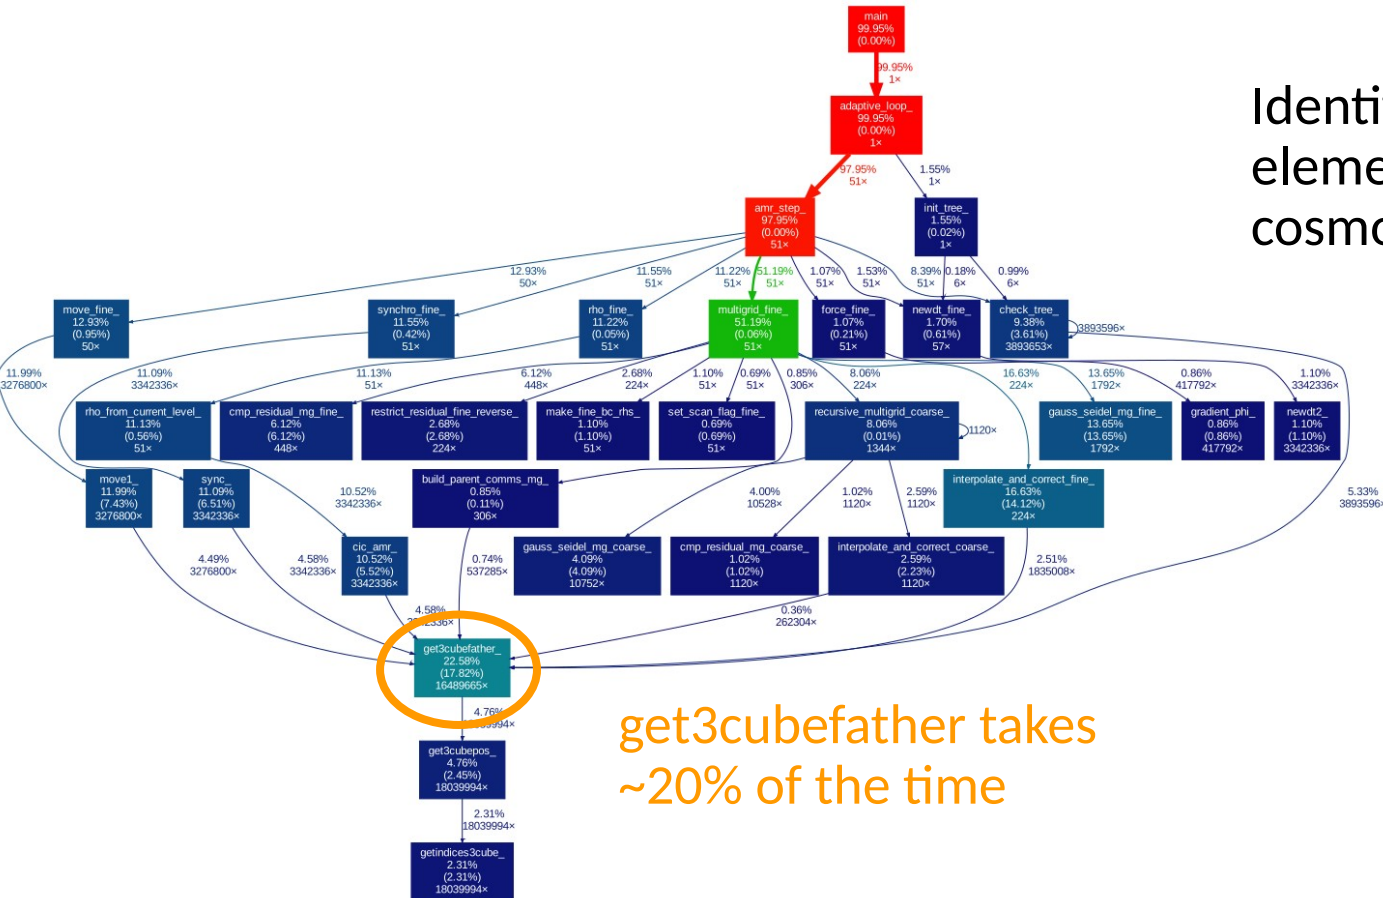
Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European High Performance Computing Joint Undertaking (JU) and Belgium, Czech Republic, France, Germany, Greece, Italy, Norway, and Spain. Neither the European Union nor the granting authority can be held responsible for them



Optimisation of neighbor searching



Identified bottleneck in an elementary routine for cosmo benchmark



get3cubefather takes
~20% of the time



```

5  subroutine get3cubefather(ind_cell_father,nbors_father_cells,&
142 else !,else, more complicated...
147 end do
148 ! Get father grid
149 do i=1,ncell
150   ind_grid_father(i)=ind_cell_father(i)-ncourse-(pos(i)-1)*ngridmax
151 end do
152
153 ! Loop over position
154 do ind=1,twotondim
155
156   ! Select father cells that sit at position ind
157   iok=0
158   do i=1,ncell
159     if(pos(i)==ind)then
160       iok=iok+1
161       ind_grid_ok(iok)=ind_grid_father(i)
162     end if
163   end do
164
165   if(iok>0)&
166   & call get3cubepos(ind_grid_ok,ind,nbors_father_ok,nbors_grids_ok,iok)
167
168   ! Store neighboring father cells for selected cells
169   do j=1,threetondim
170     iok=0
171     do i=1,ncell
172       if(pos(i)==ind)then
173         iok=iok+1
174         nbors_father_cells(i,j)=nbors_father_ok(iok,j)
175       end if
176     end do
177   end do
178
179   ! Store neighboring father grids for selected cells
180   do j=1,twotondim
181     iok=0
182     do i=1,ncell
183       if(pos(i)==ind)then
184         iok=iok+1
185         nbors_father_grids(i,j)=nbors_grids_ok(iok,j)
186       end if
187     end do
188   end do
189
190 end do
191
192 end if

```

Same filtering with if inside loop done trice

Both cells and grids are outputted but only one is used

```

48 subroutine get3cubefather(ind_cell,nbor_cells,ncell,ilevel)
70   if(ilevel==1)then ! Easy...
126   else ! else, more complicated...
127
128     ! Get the cell's position in its grid, that is the
129     ! index ind, between 1 and twotondim.
130     do i=1,ncell
131       pos(i)=(ind_cell(i)-ncoarse-1)/ngridmax+1 !integer deviation
132     end do
133
134     ! Convert the cell's index to the index of the grid to which the
135     do i=1,ncell
136       iskip=ncoarse+(pos(i)-1)*ngridmax
137       ind_grid_father(i)=ind_cell(i)-iskip
138     end do
139
140     ! Using the grid index and cell position, get the neighbor cells
141     call get3cubepos(ind_grid_father,pos,nbor_cells,ncell)
142
143   end if

```


```

170 > subroutine get3cubefather_grids(ind_cell,nbors_grids,ncell,ilevel)
170 >   if(ilevel==1)then ! Easy...
223   else ! else, more complicated...
224
225     ! Get the cell's position in its grid, that is the
226     ! index ind, between 1 and twotondim.
227     do i=1,ncell
228       pos(i)=(ind_cell(i)-ncoarse-1)/ngridmax+1 !integer deviation
229     end do
230
231     ! Convert the cell's index to the index of the grid to which the cell
232     do i=1,ncell
233       iskip=ncoarse+(pos(i)-1)*ngridmax
234       ind_grid_father(i)=ind_cell(i)-iskip
235     end do
236
237     ! Using the grid index and cell position to get neighbor grids
238     call get_grids_of_nbor_cells(ind_grid_father,pos,nbors_grids,ncell)
239
240   end if

```

Solutions:

- Split into two routines
- Remove filtering and artificial loop over ind. Pass array pos instead of integer ind to get3cubepos.



Implementation

Place !\$omp on loops over grids

Beware of updates to neighbors or other levels

```
! Loop over active grids by vector sweeps
ncache=active(ilevel)%ngrid
!$omp parallel do private(ngrid,i)
do igrid=1,ncache,nvector
  ngrid=MIN(nvector,ncache-igrig+1)
  do i=1,ngrid
    ind_grid(i)=active(ilevel)%igrig(igrig+i-1)
  end do
  call godfine1(ind_grid,ngrid,ilevel)
end do
```

```
! Update particles position and velocity
!$omp parallel private(ig,ip,jgrid,igrig,npart1,ipart,local_c
  ig=0
  ip=0
  ! Loop over particles that are not tracers
!$omp do
  do jgrid=1,active(ilevel)%ngrid
    igrig=active(ilevel)%igrig(jgrid)
    npart1=numbp(igrig) ! Number of particles in the grid
    if(npart1>0)then
      ...
      ! Loop over particles in current grid
      do jpart=1,npart1
        ! Save next particle <---- Very important !!!
        next_part=nextp(ipart)
        ...
        if(ip==nvector)then
          ! send gathered nvector particles to calculation
          call move1(ind_grid,ind_part,ind_grid_part,ig,ip,il
        ...
        end if
        ipart=next_part ! Go to next particle
      end do
      ! End loop over particles
      ...
    end if
  end do
!$omp end do nowait
! End loop over grids
if(ip>0)call move1(ind_grid,ind_part,ind_grid_part,ig,ip,il
!$omp end parallel
```