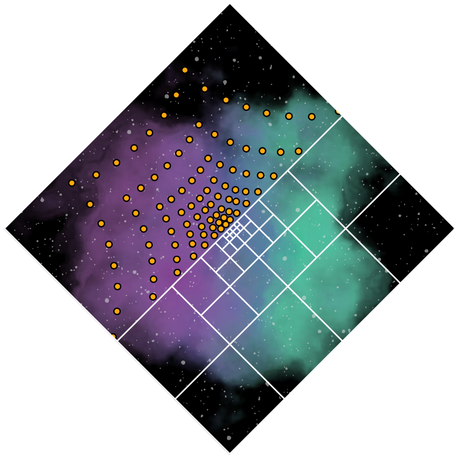

FRANCE 2030

PROGRAMME DE RECHERCHE

ORIGINES

UNIVERSITÉ DE LYON

ENS DE LYON

CRAL

erc

PHAST
PHYSIQUE
ET ASTROPHYSIQUE
UNIVERSITÉ DE LYON

ÉCOLE DOCTORALE
52

Why ??????????????????

Why a new code ?

Context & Aims

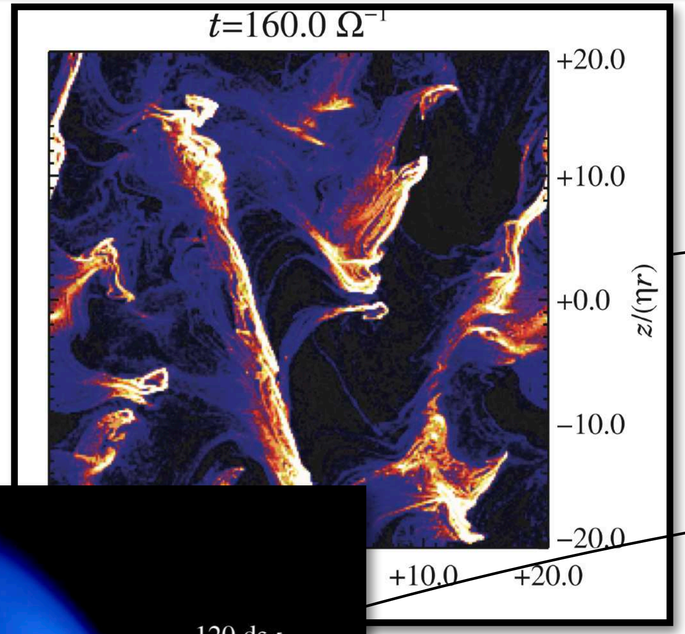

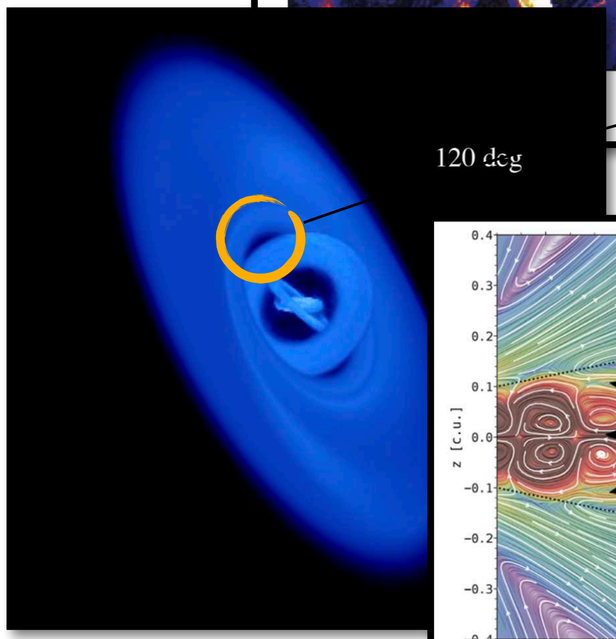

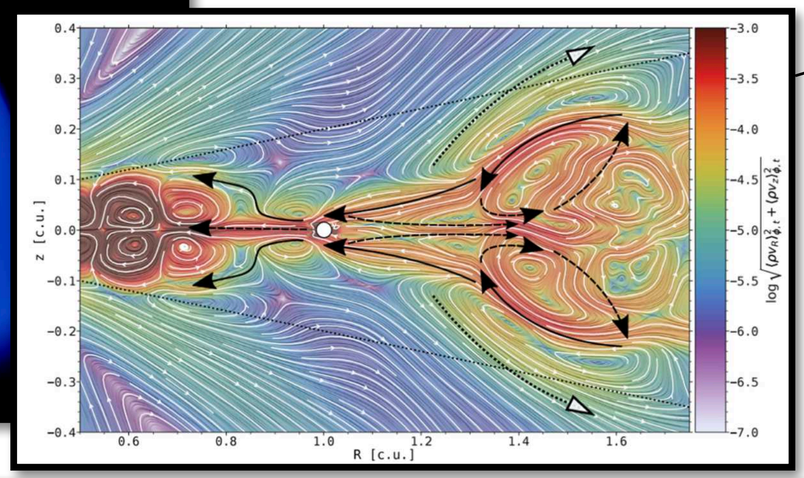

S.I. In global discs ?

What happens here ?

And close to a planet ?

Resolution is mandatory !!!

Context & Aims

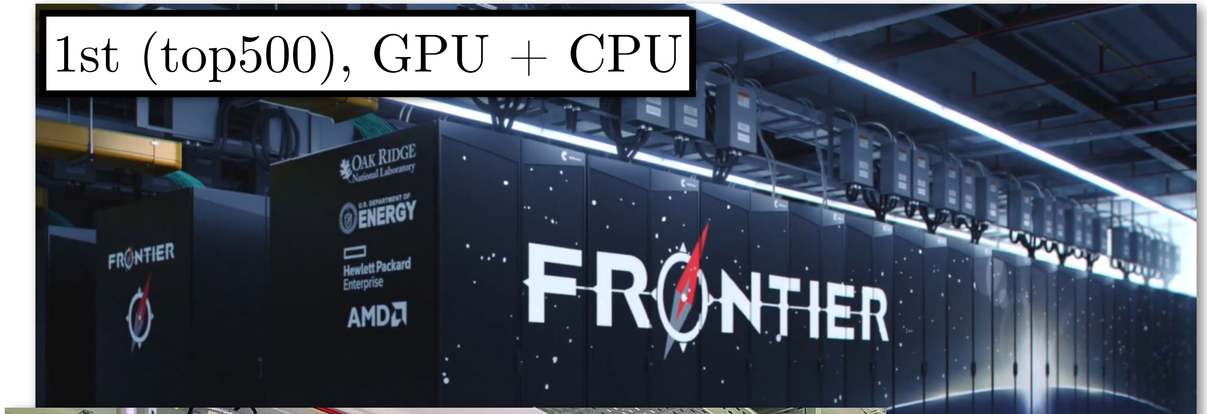

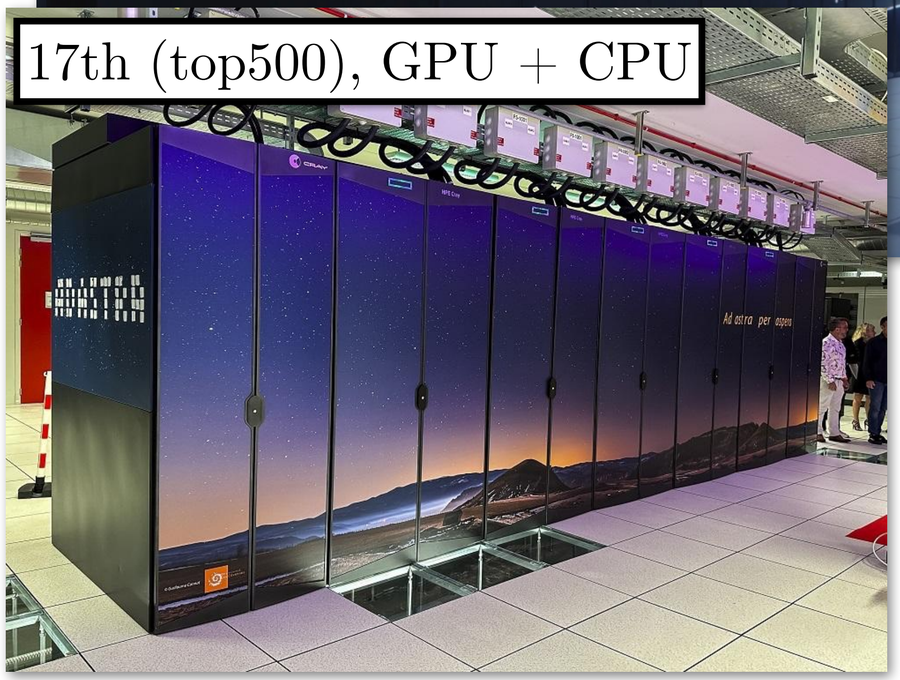

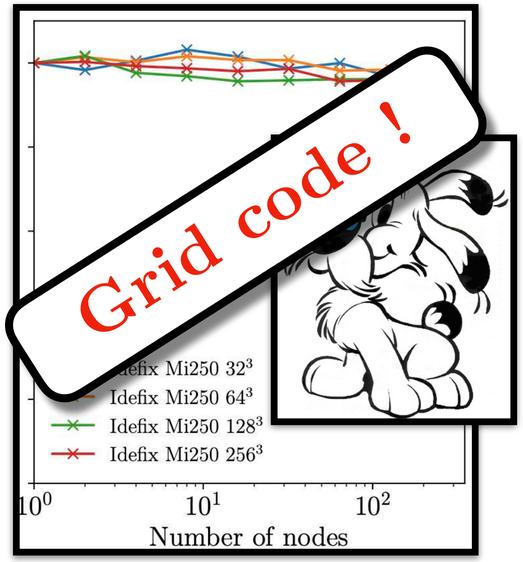

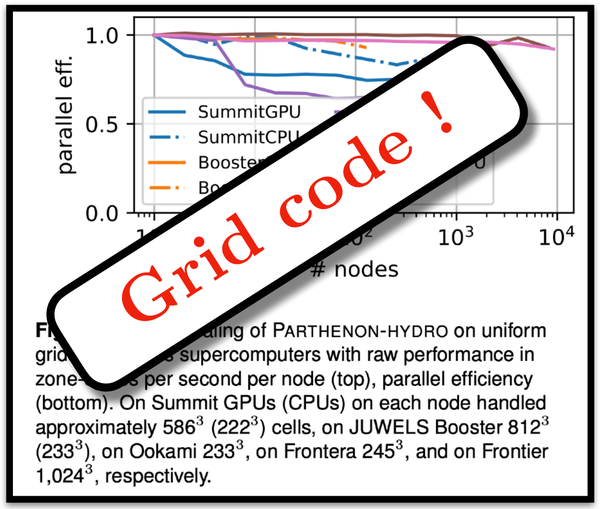

ling of PARTHENON-HYDRO on uniform
supercomputers with raw performance in
s per second per node (top), parallel efficiency
(bottom). On Summit GPUs (CPUs) on each node handled
approximately 586^3 (222^3) cells, on JUWELS Booster 812^3
(233^3), on Ookami 233^3, on Frontera 245^3, and on Frontier
$1{,}024^3$, respectively.

Where is SPH ???

So we need something ....

Well you know where this is going ....

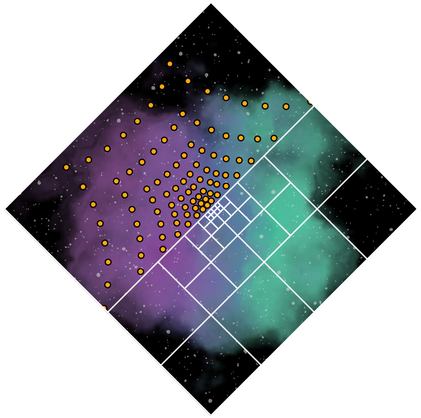

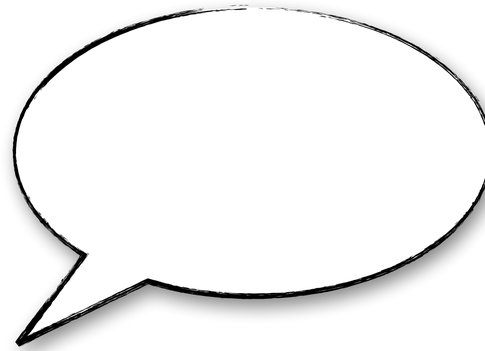

Technical Locks

- `>1G part` in a simulation

  How do you store a lot of particles on a cluster ?

- Fast on GPU

  Can you optimise SPH on GPU

- Neighbours finding

  How do you work with neighbours on GPU

- Portability

  But you need also to run well on CPU

The plan !

Distributed computing

Software engineering

neighbours & GPUs

Portability

& Surprise

Source : passeportsante.net

Portability

Our choice : SYCL™

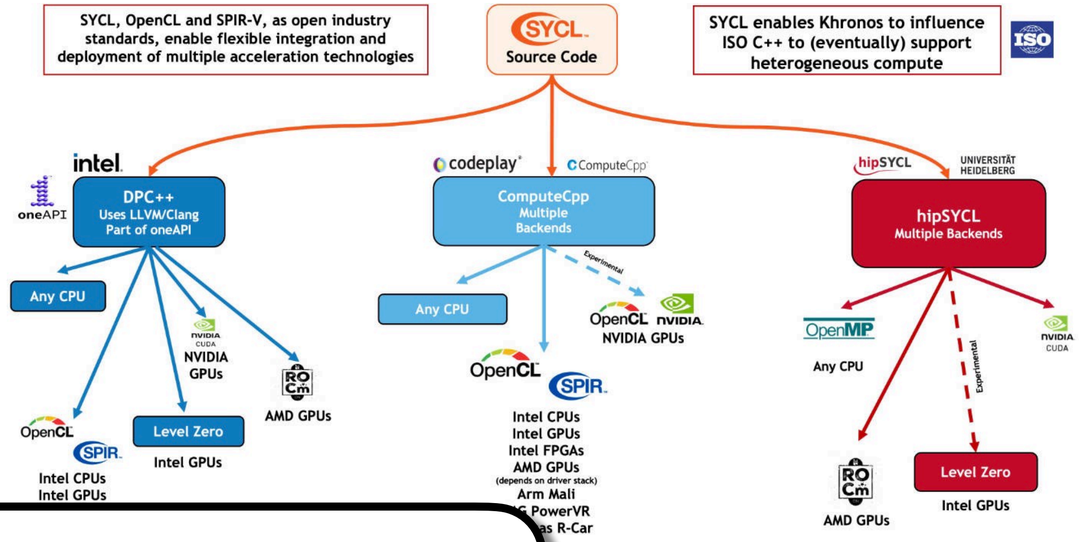

- C++17
- Can be compiled to native CUDA, Hip, OpenMP, ...
- Support **all** CPU & GPUs

- Portability

  But you need also to run well on CPU

Domain decomposition

The global simulation is divided
into multiple patches
(**abstract decomposition**)

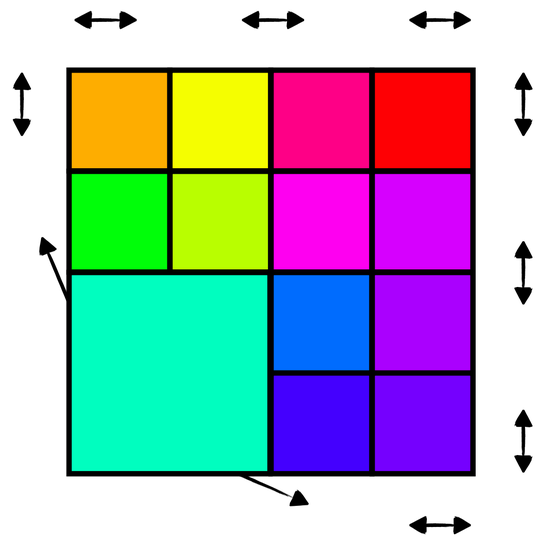

We only have to communicate
the interface to itera

- `>1G part` in a simulation

Patches are scattered across the nodes

If patches are too large they split
If they are too small they merge

What is that ?

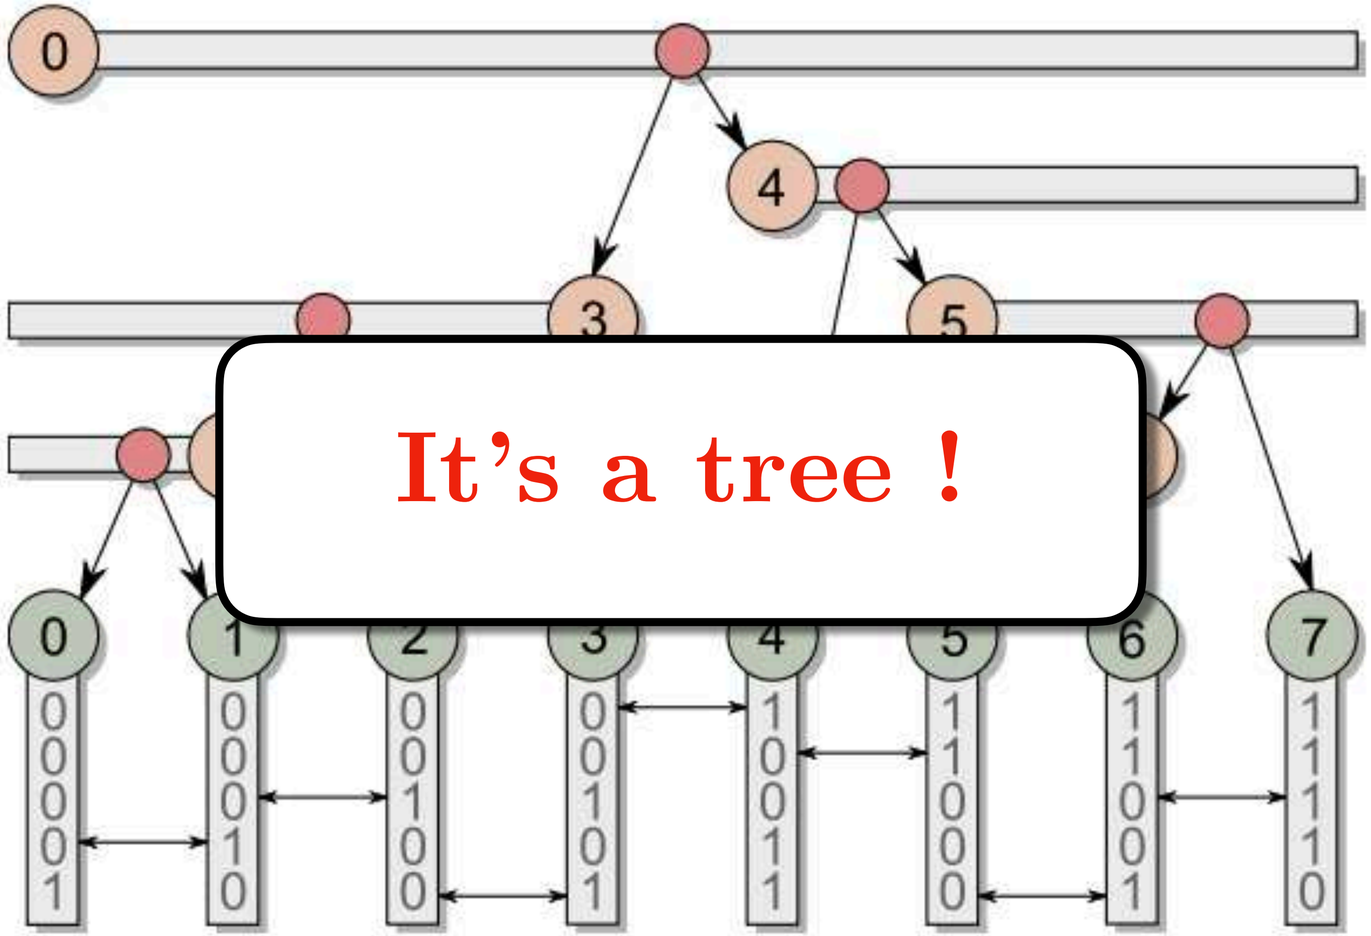
0
4
3
5
It’s a tree !
0
1
2
3
4
5
6
7
00001
00010
00100
00101
10011
11000
11001
11110

Neighbour finding

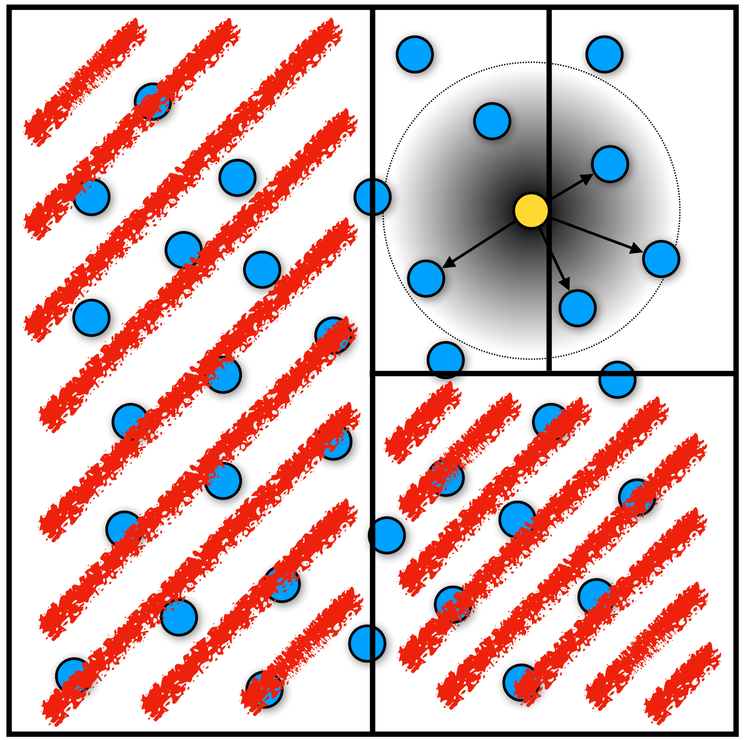

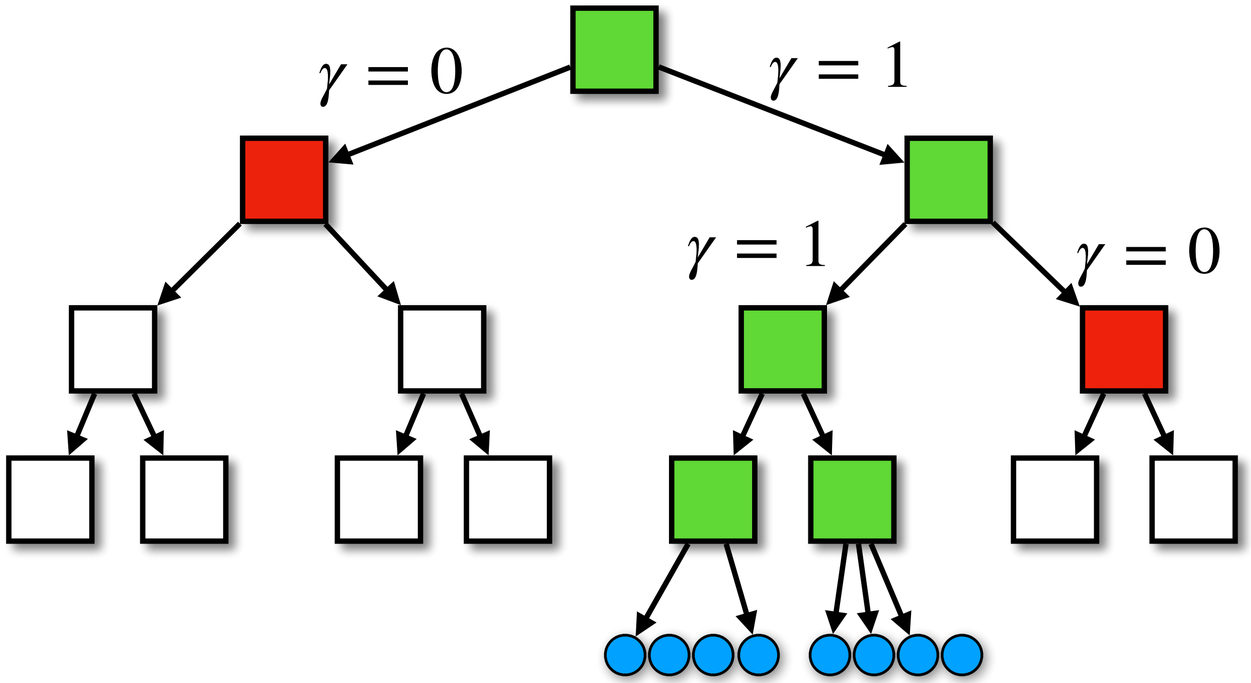

Tree based version :

- Exclude half of the space every steps
- The remaining particles are likely to interact

Complexity : $\mathcal{O}(n \log(n))$

It took a year ....

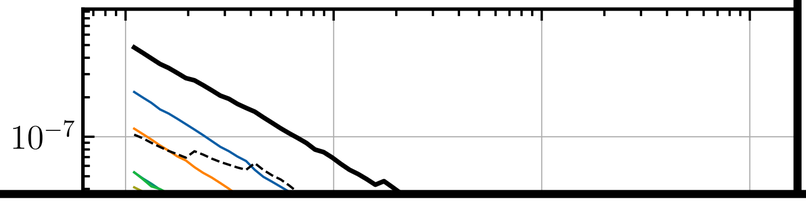

Tree algorithm :

- Using binary arithmetics (fast)
- Morton code to convert position to binary
- Internal layout using binary tree
- just rebuild it !

- Fast on GPU

  Can you optimise SPH on GPU

- Neighbours finding

  How do you work with neighbours on GPU

Shamrock paper 1 (in prep)

A100 SXM4 : >400.000.000 objects/ 1 sec

y faster than a timestep !

But .....

A lot of code $\rightarrow$ Software engineering ....

Github test pipeline + regression ✅

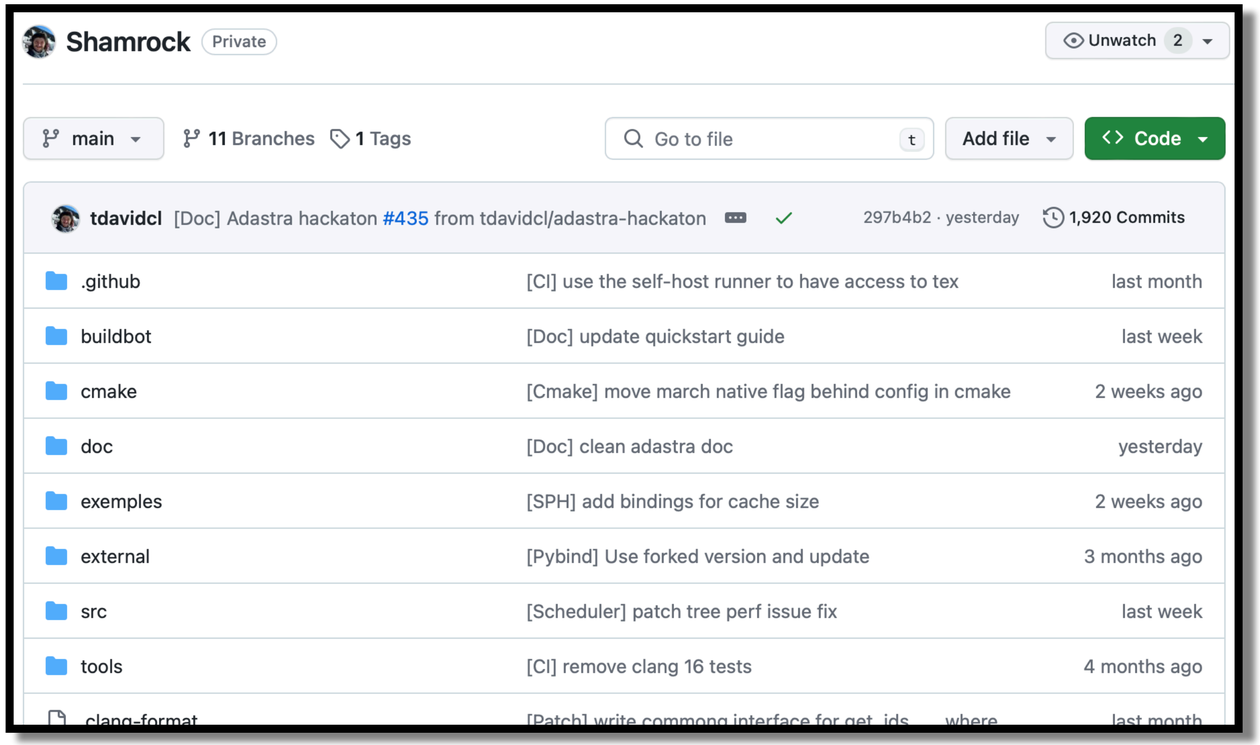

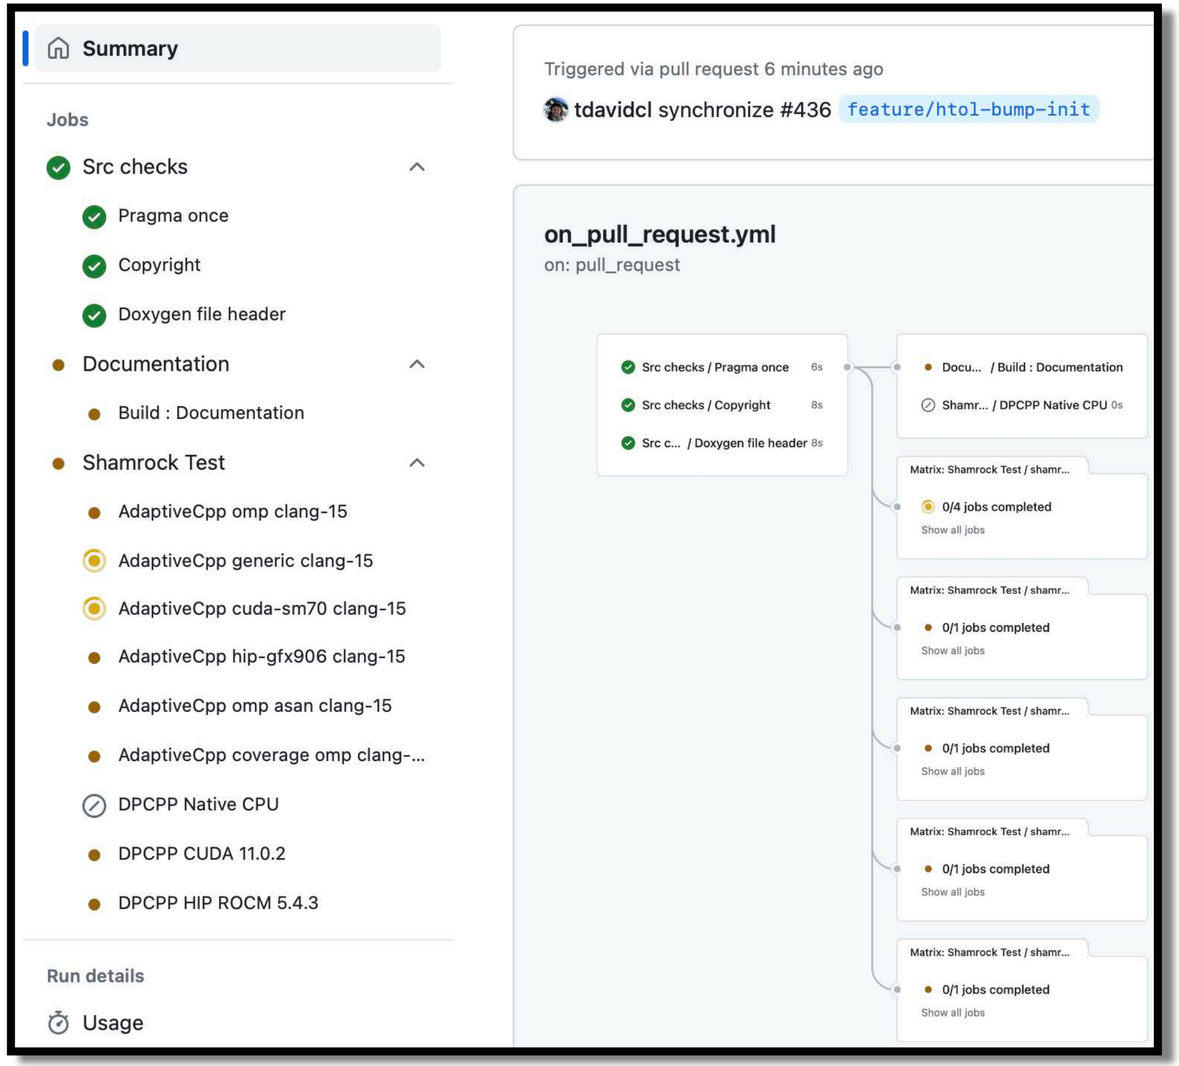

Testing library

Testing :

- Integrate
- Can outp
- Can run
- Can run

Test → TestResult → TexReport

```
TestStart(Unittest, "coo

    using namespace sham

    auto & dataset = tes

    dataset.add_data("X", {0,1,2,3,4});
    dataset.add_data("Y", {0,1,2,3,4});

    asserts().assert_bool("why not", true);
}
```

SHAMROCK test suite report

Timothée David–Cléris*
CRAL ENS de Lyon
October 15, 2023

Global status : FAIL

Contents

1 Shamrock config

1.1 Git info

```
commit : 4ef3885d437cf951dc767440dac76ca8377dbc2e
HEAD   : refs/heads/281-test-improve-testing, refs/remotes/orig
modified files (since last commit):
   src/main_test.cpp
   src/shammodels/generic/tests/fmmTests.cpp
   src/shamtest/details/TestResult.cpp
   src/shamtest/details/TestResult.hpp
   src/shamtest/shamtest.cpp
   src/shamtest/shamtest.hpp
   src/tests/experimental/syclEventUSMTests.cpp
   src/tests/shammath/sphkernelsTests.cpp
   src/tests/shamsys/syclLoopsTests.cpp
   src/tests/test_pythonscript_handle.cpp
```

1.2 c++ flags

```
--hipsycl-cpu-cxx=g++
--hipsycl-targets=omp
--hipsycl-config-file=/home/tim/Documents/Travail/ShamrockWorkspace/
-Werror=return-type
-pedantic-errors
-fcolor-diagnostics

-DSYCL_COMP_ACPP
-isystem
/home/tim/Documents/Travail/ShamrockWorkspace/sycl_compilers/opensyc
-isystem
/home/tim/Documents/Travail/ShamrockWorkspace/sycl_compilers/opensyc
-DSHAMROCK_LOOP_DEFAULT_PARRALEL_FOR_ROUND
-DSHAMROCK_LOOP_GSIZE=256
-DSHAMROCK_USE_PROFILING
```

*timothee.david-cleris@ens-lyon.fr

1

Test suite report SHAMROCK

2 Unittests

Test name	Status	Asserts
core/patch/base/patchdata_field:constructor	OK	5/5
core/patch/base/patchdata_field:patch_data_field_check_matc		
core/tree/kernels/karras_alg (32)		
core/tree/kernels/karras_alg (64)		
core/tree/kernels/key_pair_sort		
made to fail		
models/generic/fmm/multipole_moment_offset		
patchdata.cpp/patch_data_check_match		
shamalgs/algorithm/details/bitonicSort_fallback		
shamalgs/algorithm/details/bitonicSort_legacy		
shamalgs/algorithm/details/bitonicSort_updated		
shamalgs/algorithm/index_remap		
shamalgs/algorithm/sort_by_key		
shamalgs/atomic/DynamicIdGenerator		
shamalgs/collective/distributedDataComm		
shamalgs/collective/reduction/allreduce_sum		
shamalgs/collective/sparseXchg		
shamalgs/container/ResizableBuffer		
shamalgs/container/ResizableUSMBuffer		
shamalgs/container/ResizableUSMBuffer:synchronisation		
shamalgs/memory/SerializeHelper		
shamalgs/numeric/details/exclusive_sum_atomic_decoupled_v5		
shamalgs/numeric/details/exclusive_sum_fallback		
shamalgs/numeric/details/exclusive_sum_gpugems39		
shamalgs/numeric/details/exclusive_sum_in_place_fallback		
shamalgs/numeric/details/exclusive_sum_sycl_jointalg		
shamalgs/numeric/details/stream_compact_fallback		
shamalgs/numeric/stream_compact		
shamalgs/reduction/max		
shamalgs/reduction/min		
shamalgs/reduction/sum		
shambase/Constants		
shambase/stacktrace/print_stack		
shambase::has_inf		
shambase::has_nan		
shambase::has_nan_or_inf		
shambase::parralel_for		
shammath/sphkernels/C2		
shammath/sphkernels/C4		
shammath/sphkernels/C6		
shammath/sphkernels/M4		
shammath/sphkernels/M5		
shammath/sphkernels/M6		
shammodels/amr/AMRBlock		
shamrock/patch/Patch.cpp:SplitCoord		
shamrock/patch/Patch.cpp:SplitMerge		
shamrock/patch/PatchData::serialize_buf		
shamrock/patch/PatchDataField::serialize_buf		
shamrock/patch/PatchDataField::serialize_full		
shamrock/scheduler/SchedulerPatchData::apply_change_list		
shamrock/sfc/morton/min-max		

3

Test suite report SHAMROCK

Test name	Status	Asserts
shamrock/tree/RadixTree::serialize	OK	1/1
shamrock/tree/RadixTree:bounding_volume_inclusion	OK	14/14
shamrock/tree/TreeCellRanges::serialize	OK	1/1
shamrock/tree/TreeMortonCodes::serialize	OK	1/1
shamrock/tree/TreeReducedMortonCodes::serialize	OK	1/1
shamrock/tree/TreeStructure::serialize	OK	1/1
shamsys/Log	OK	0/0
shamsys/comm/CommunicationBuffer/constructor	OK	4/4
shamtest/PyScriptHandle(plot)	OK	0/0
shamtest/PyScriptHandle(run)	OK	2/2
shamtest/PyScriptHandle(shamrock)	OK	1/1
syclhandler/test_sycl_mpi_types_sizes	OK	66/66
testamrtmp	OK	0/0
vtk-write-cells	OK	0/0
vtk-write-parts	OK	0/0

Test : made to fail
- Assert [1/3]
_Assert(false) : FAIL
→ failed assert logs :

```
failed assert location : /home/tim/Documents/Travail/ShamrockWorkspace/Shamrock/src/tests/test_pythons
```

- Assert [2/3]
_Assert(true) : OK
- Assert [3/3]
_Assert(false) : FAIL
→ failed assert logs :

```
failed assert location : /home/tim/Documents/Travail/ShamrockWorkspace/Shamrock/src/tests/test_pythons
```

3 Tex report

3.1 shammath/sphkernels_plotall

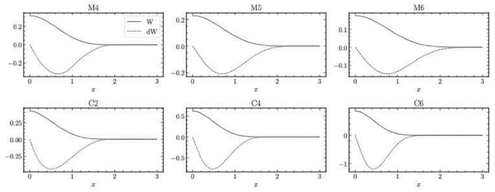

Figure 1: SPH kernels implemented in shamrock

4

```
) {

)"
}
```

Python interop

Runscripts :

- Flexible for simulation inputs
- No C++ for the user
- Compatible with restarts
- One job multiple simulation
- Can be run interactively (IPytho

Quick exemple ...

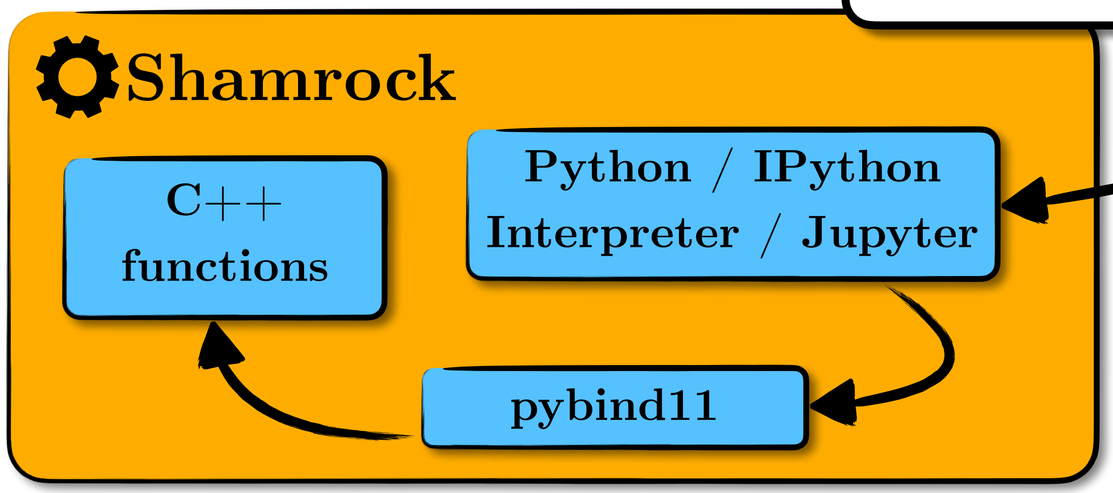


```python
import shamrock

ctx = shamrock.Context()

ctx.pdata_layout_new()
ctx.pdata_layout_add_field("xyz",1,"f64_3")
ctx.pdata_layout_add_field("vxyz",1,"f64_3")
ctx.pdata_layout_add_field("u",1,"f64")

ctx.init_sched(int(1e6),1)
ctx.set_box_size((0,0,0),(1,1,1))

rho, V, dr = 1 , 1, 0.02

setup = shamrock.SetupSPH(kernel = "M4")
setup.init(ctx)
setup.set_boundaries("periodic")
setup.add_cube_fcc(ctx,dr, (0,0,0),(1,1,1))
setup.set_total_mass(1)
setup.clear()

model = shamrock.BasicSPHGas(kernel = "M4")
model.init()
model.set_cfl_cour(0.3)
model.simulate_until(ctx, ti = 0, te = 1 ,"jisco_")
model.clear()
```

Organisation

CMake

Simplified :

(12 libs in reality)

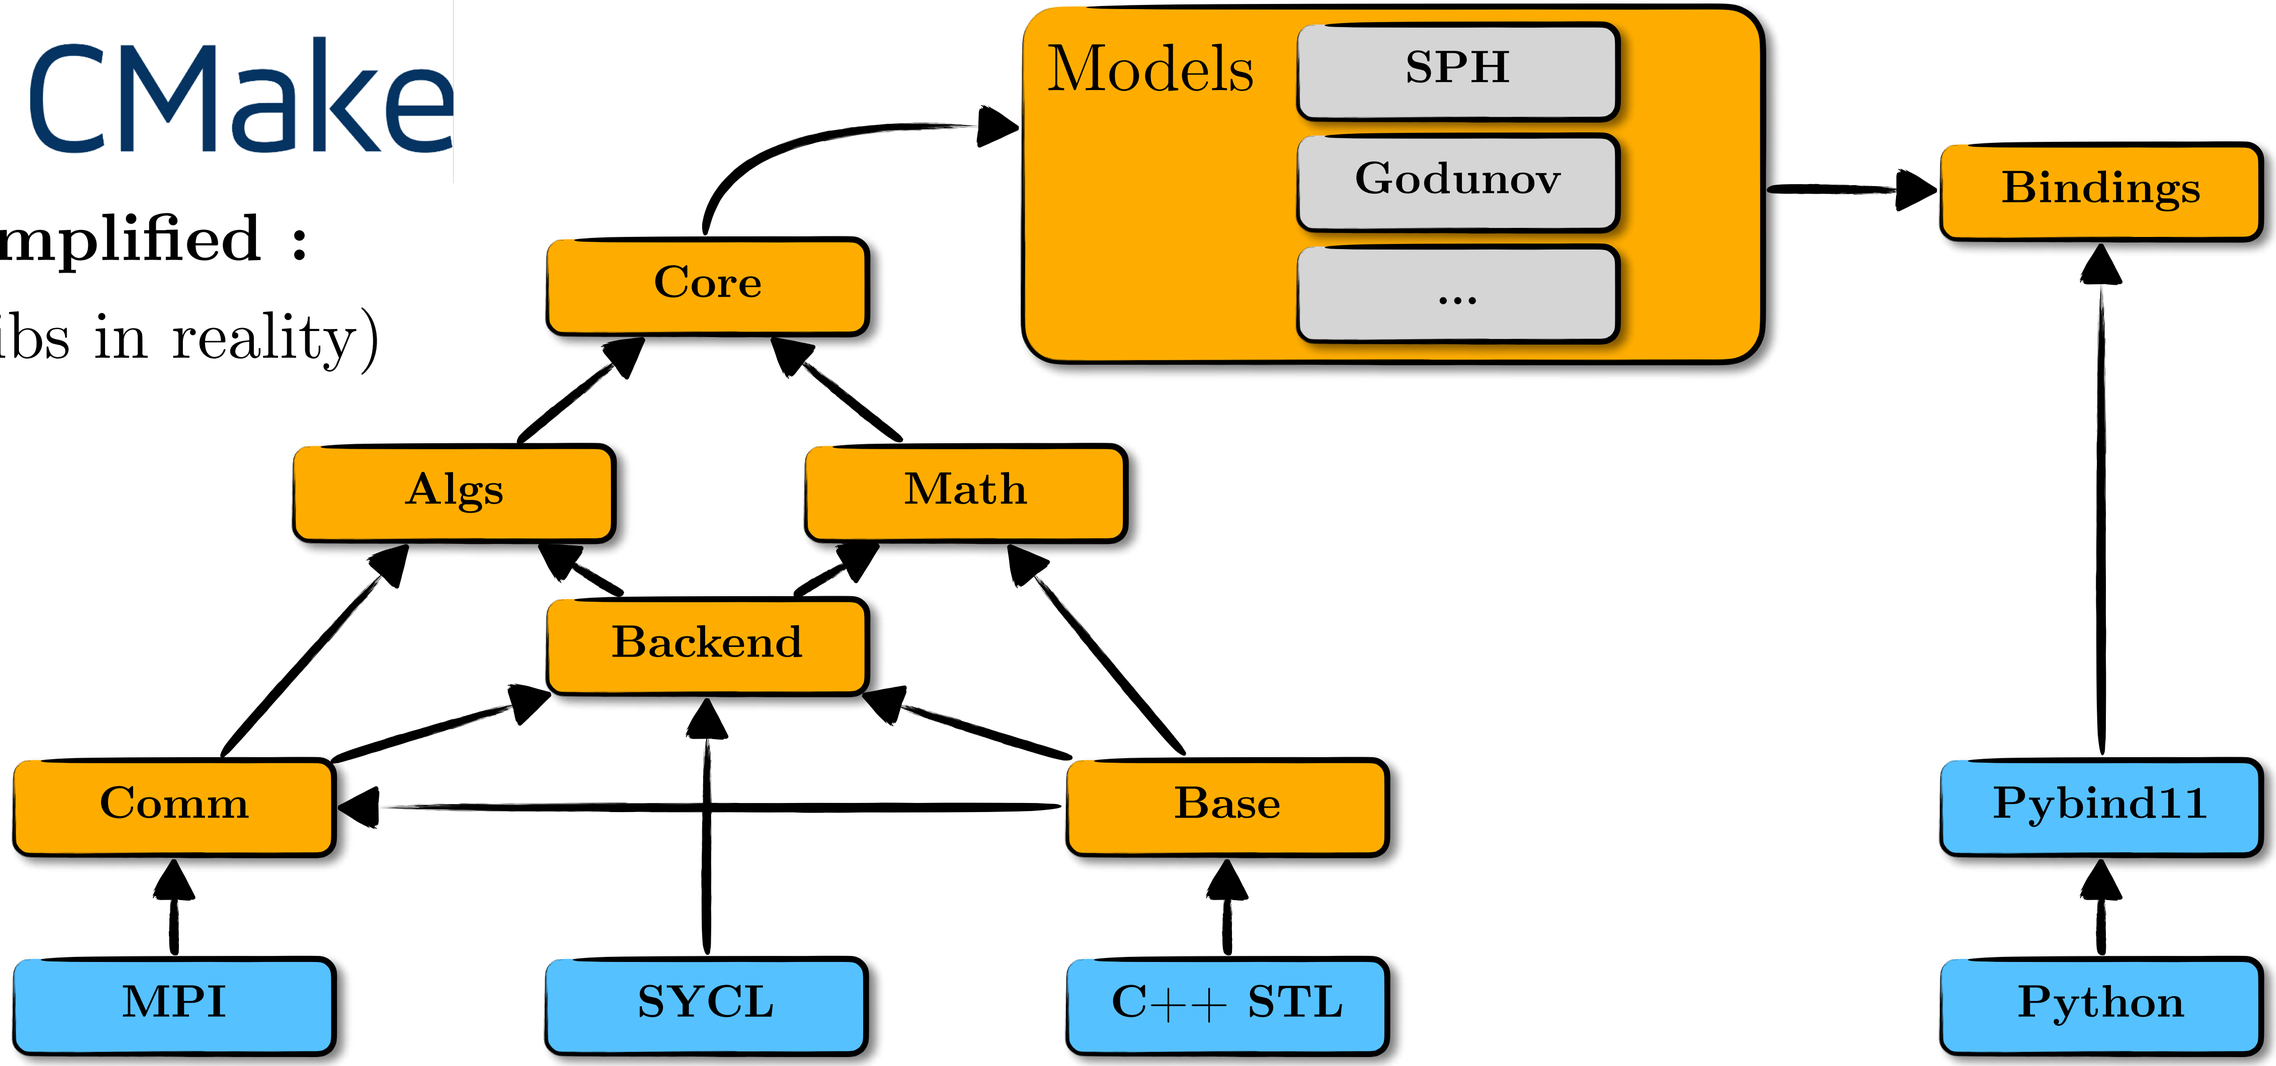

Units library (Public)

https://github.com/Shamrock-code/Shamrock-units

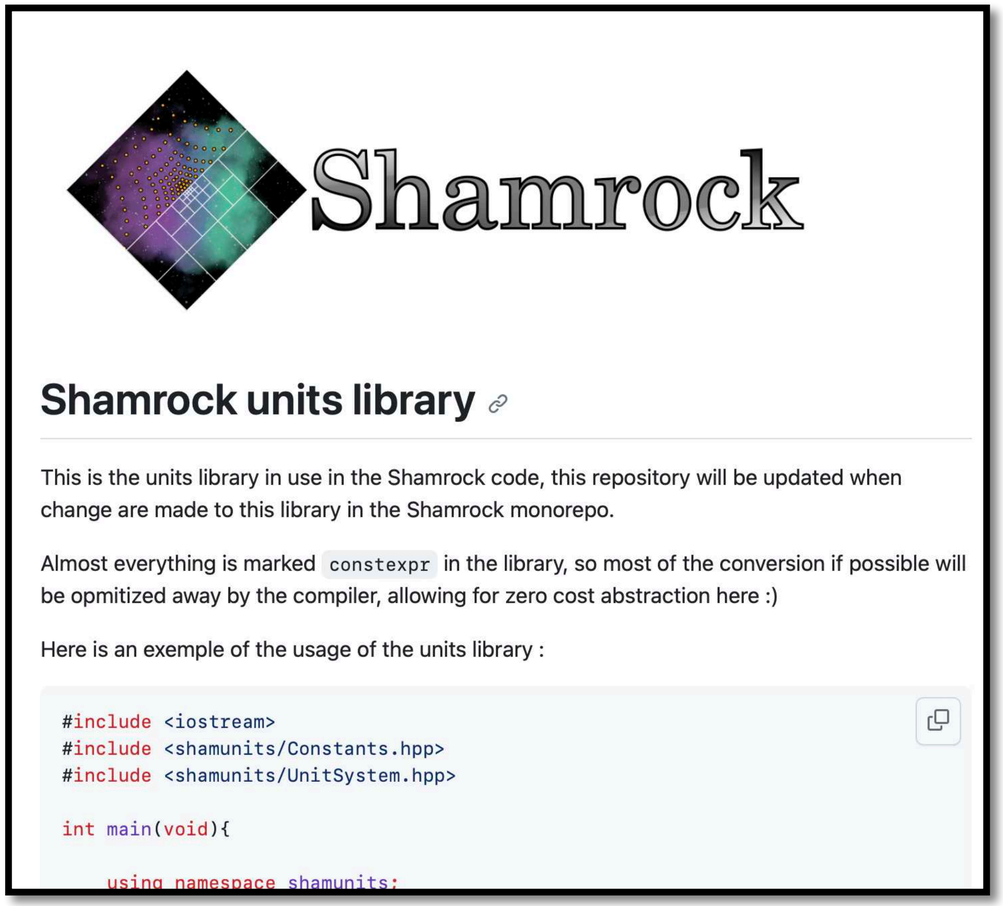

- Constexpr
- Simple
- Header only
- Can be used in GPU kernels
- Python bindings

Try it : https://godbolt.org/z/5zjGMea57

Results

Kelvin Helmotz

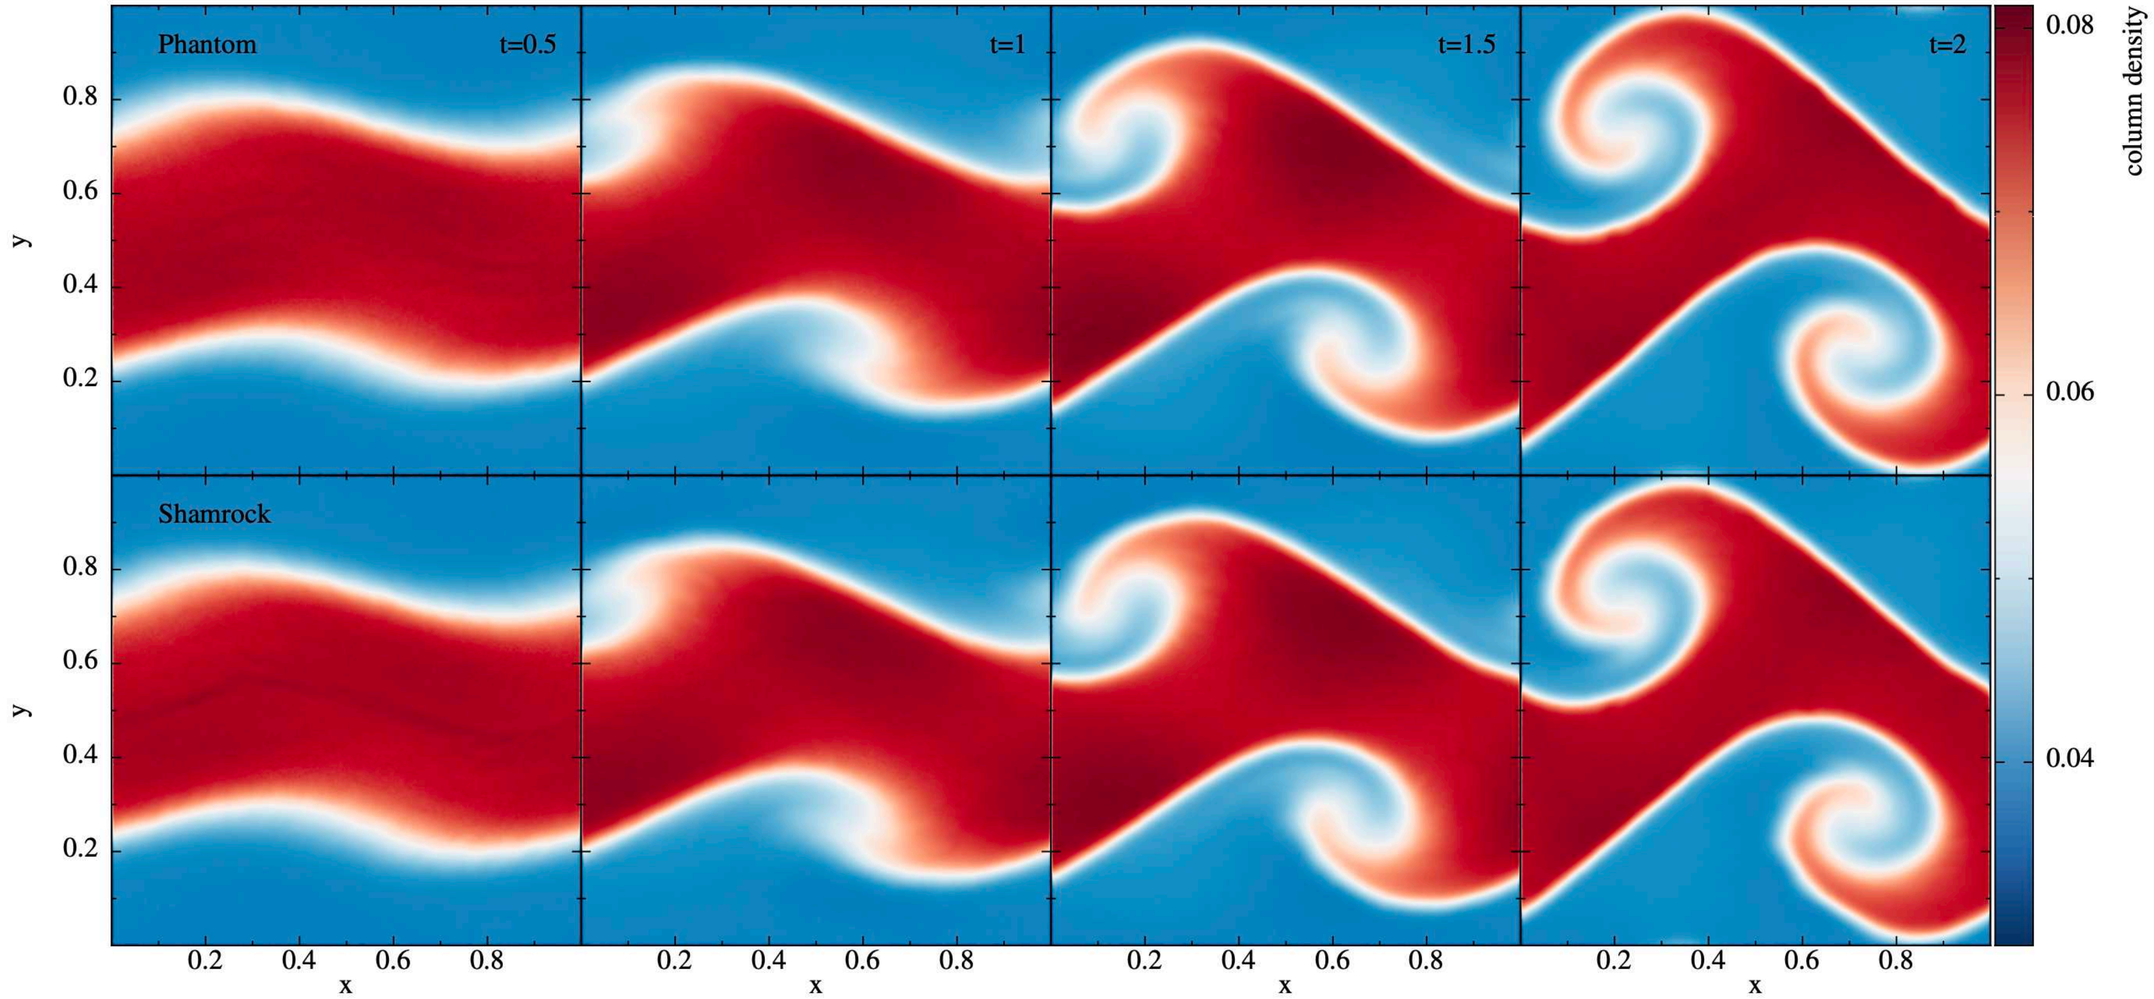

Looks similar : is the viscosity the same ?

Kelvin Helmotz

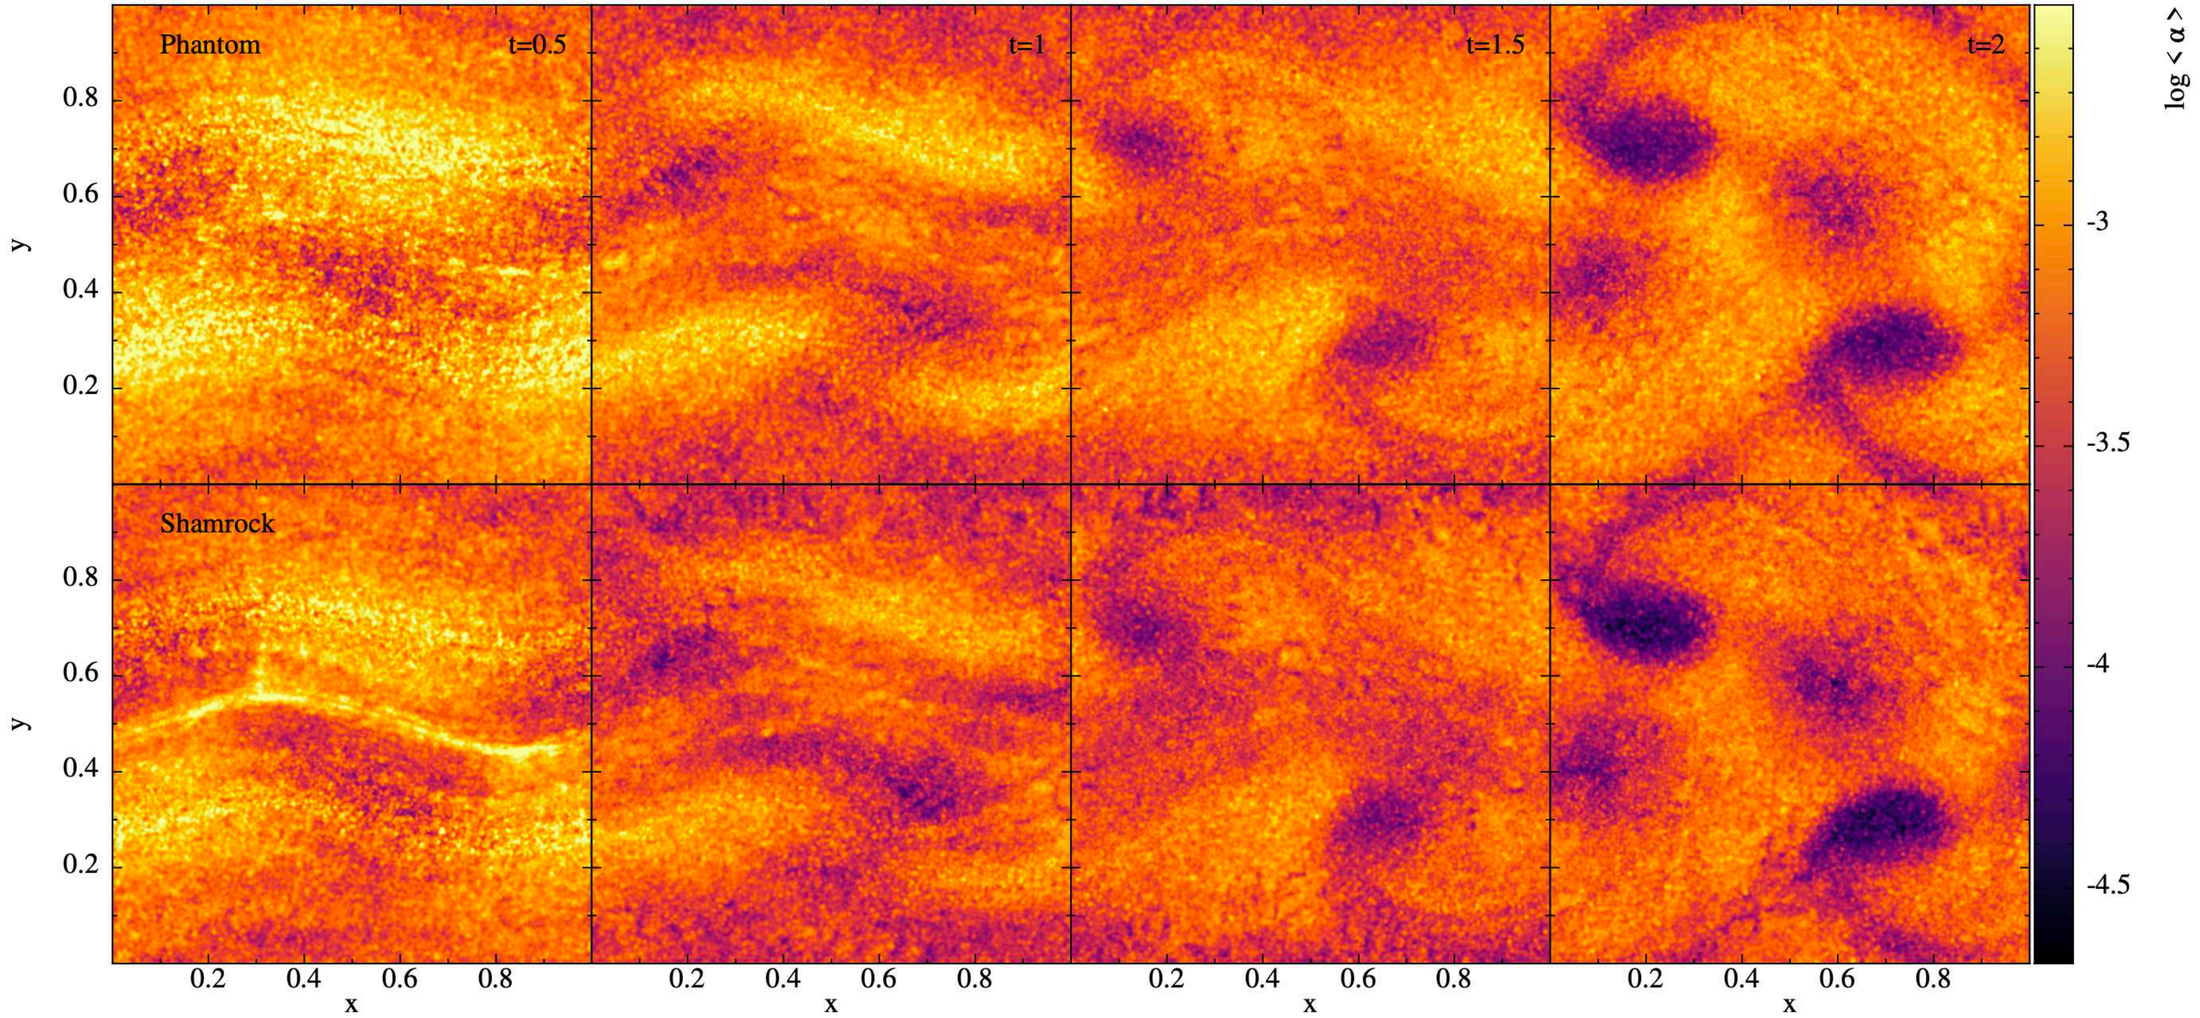

Yes ! (Differences due to difference in optimisation) **Is the solver the same ?**

Sedov-Taylor Blast

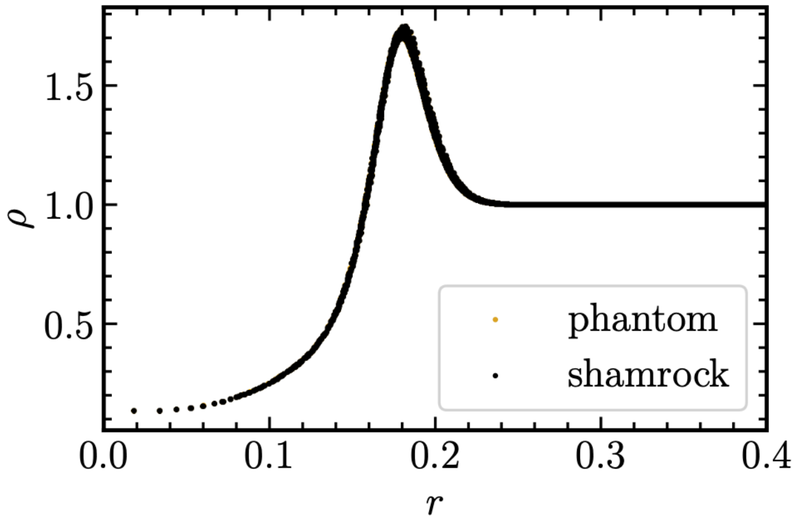

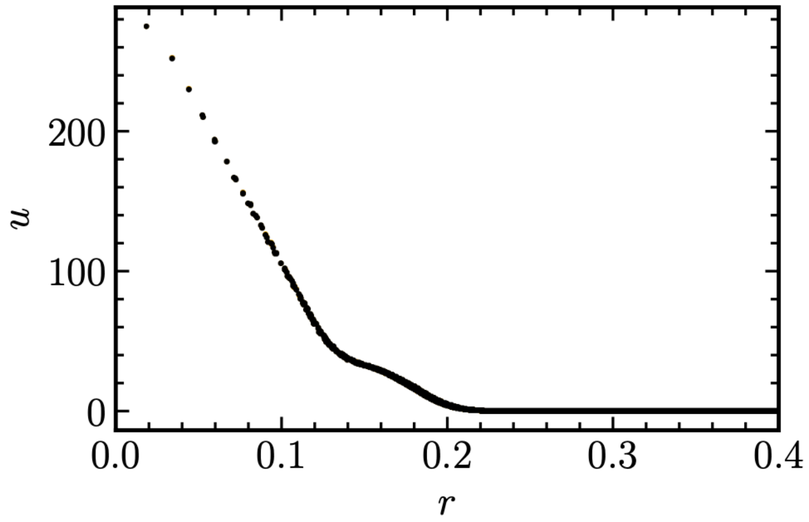

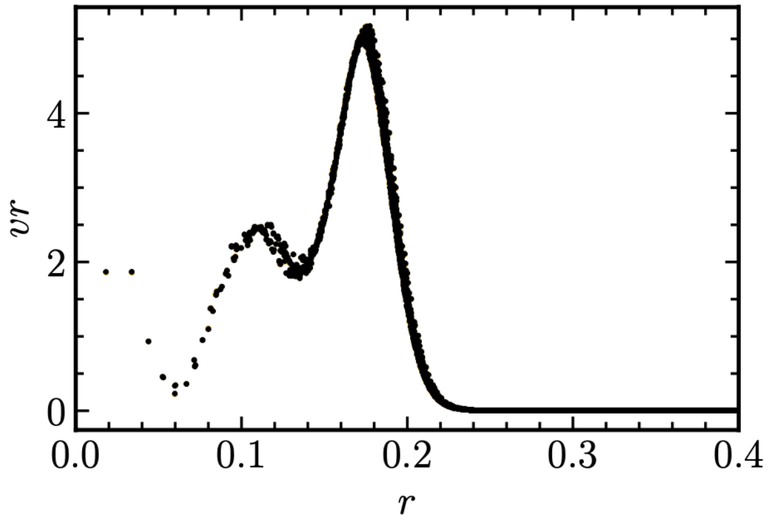

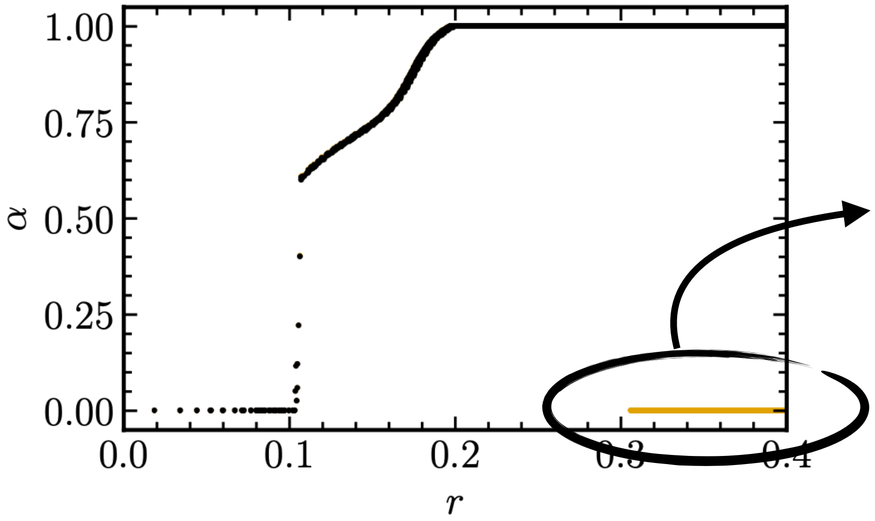

Marginal differences :

- L2 distance r : 1.323e-05
- L2 distance h : 7.687e-06
- L2 distance vr : 0.0133
- L2 distance u : 0.03803

Rounding off when $c_s = 0$

Yes !

Performance

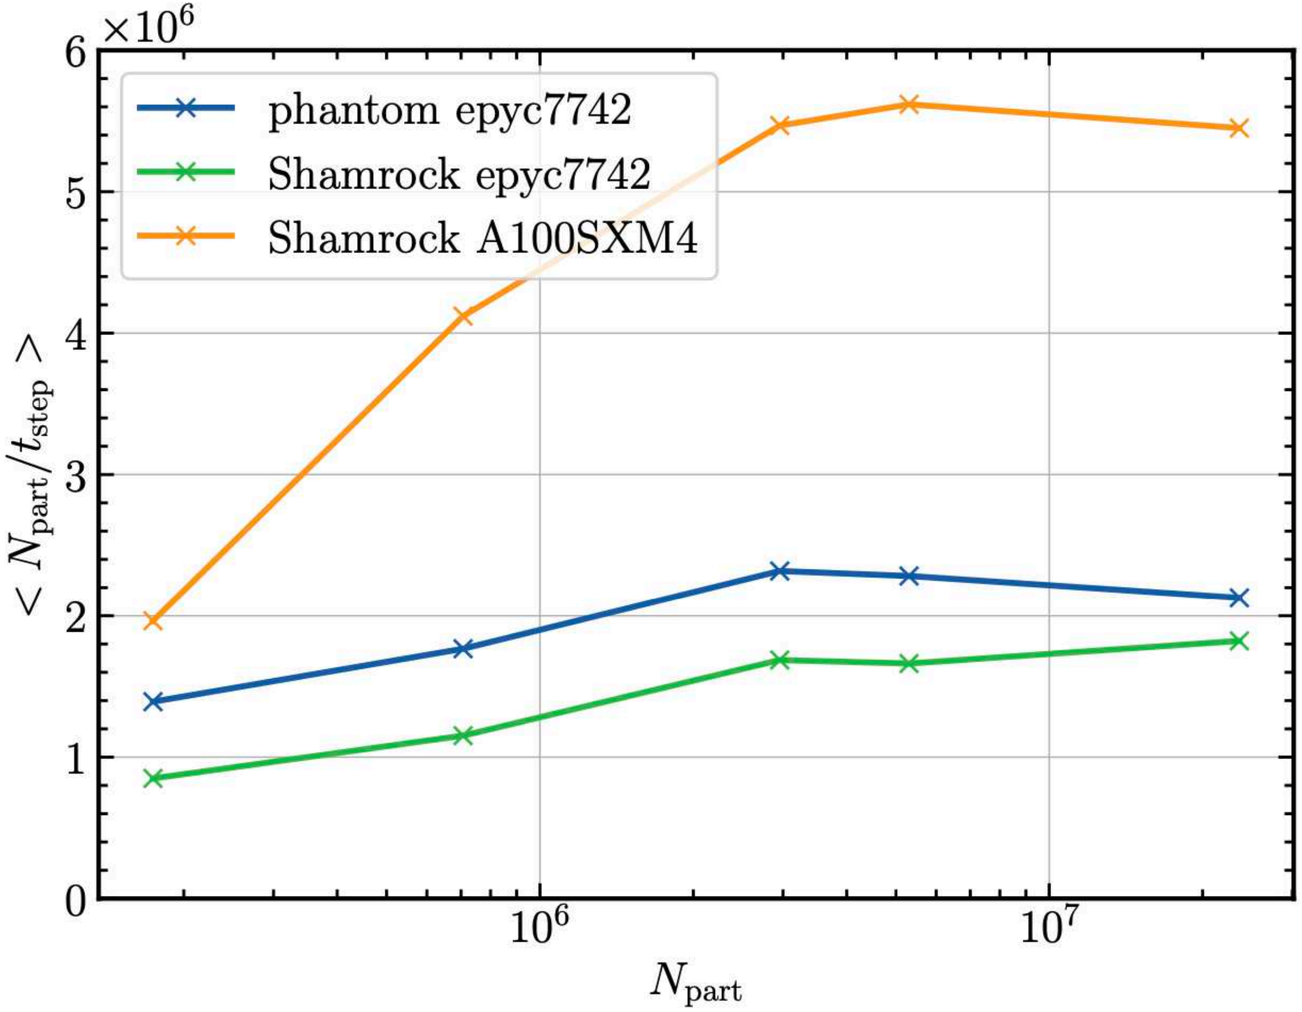

Sedov taylor blast :

- Sightly slower on CPU
- Way faster with GPU

Discs

Same performance in discs !

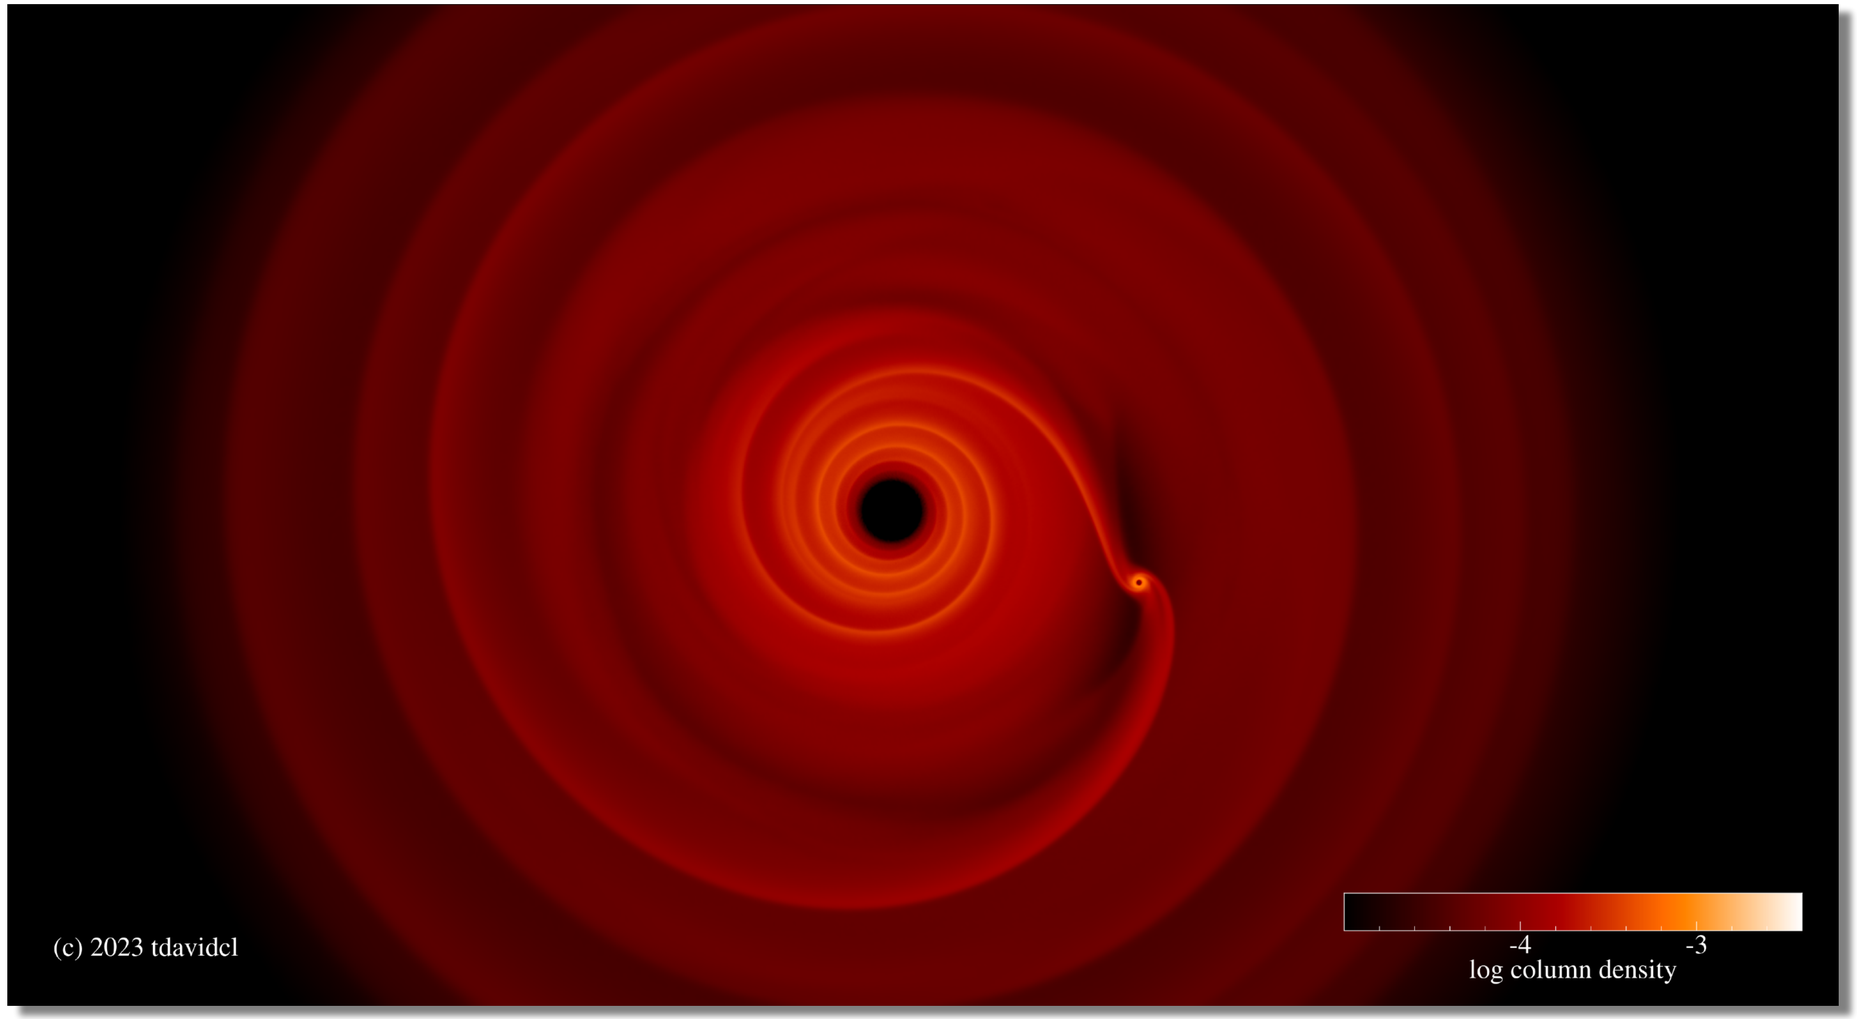

Discs (also black holes ones)

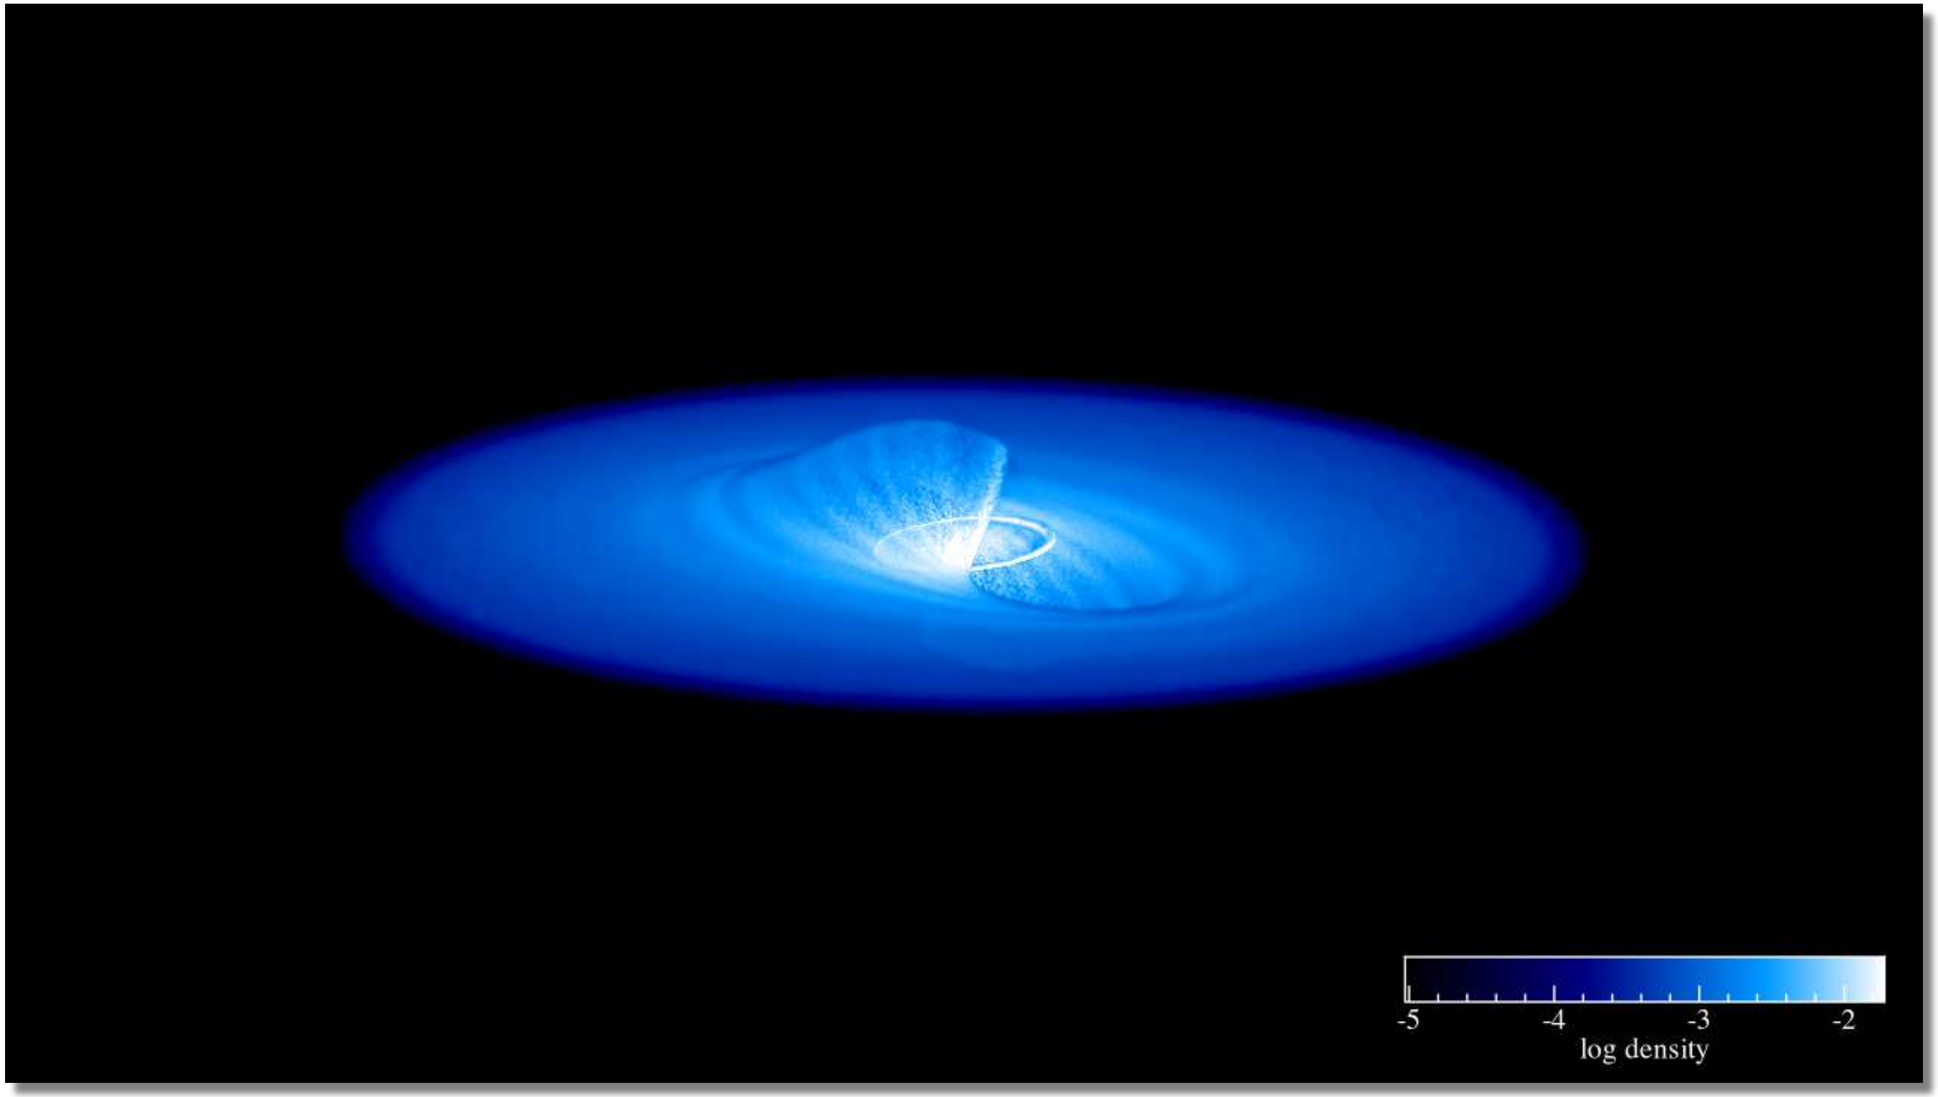

Now if you take a big cluster

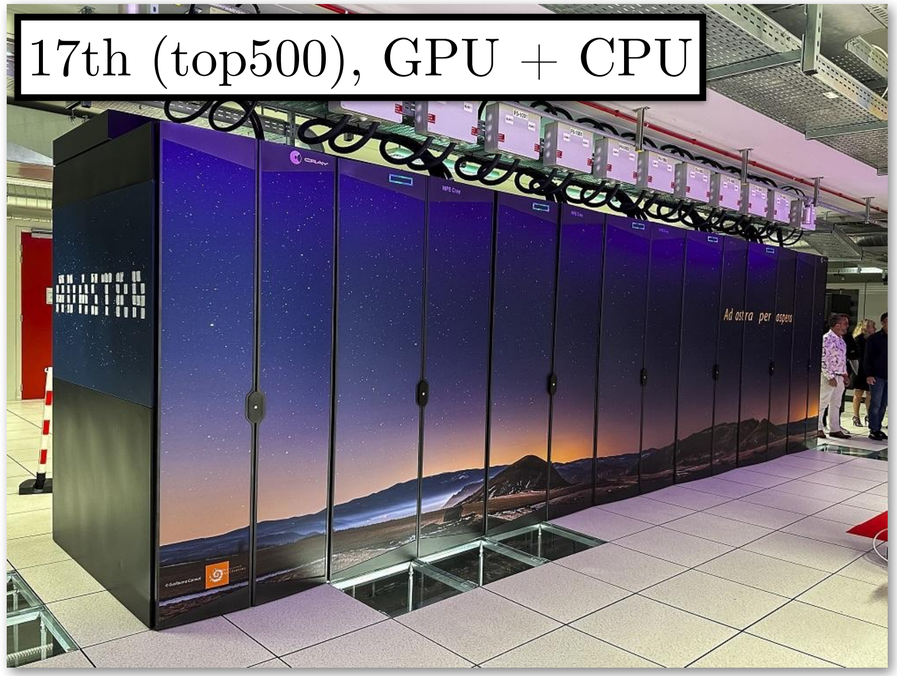

RANKING

List	Rank	System					
11/2023	17	HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11	HPE	319,072	46.10	61.61	921.48

HPC wire

Since 1987 - Covering the Fastest Computers in the World and the People Who Run Them

Home

Topics

GENCI's Adastra Marks a New Step Towards a More Sustainable HPC

November 25, 2022

Nov. 25, 2022 — The energy efficiency of GENCI (Grand Equipement National de Calcul Intensif)'s new supercomputer Adastra has been improved to reach now 58.2 GF/W, ranking Adastra at #3 position on the new Green500 list announced during SC22. Adastra is ranking #11 in the November 2022 Top500 list with 46.10 PFlops measured performance.

338 nodes, 4 MI250x / Nodes

921kW (max)

61.6 PFlop/s

It scales !!! (Weak calling)

Sedov Blast :

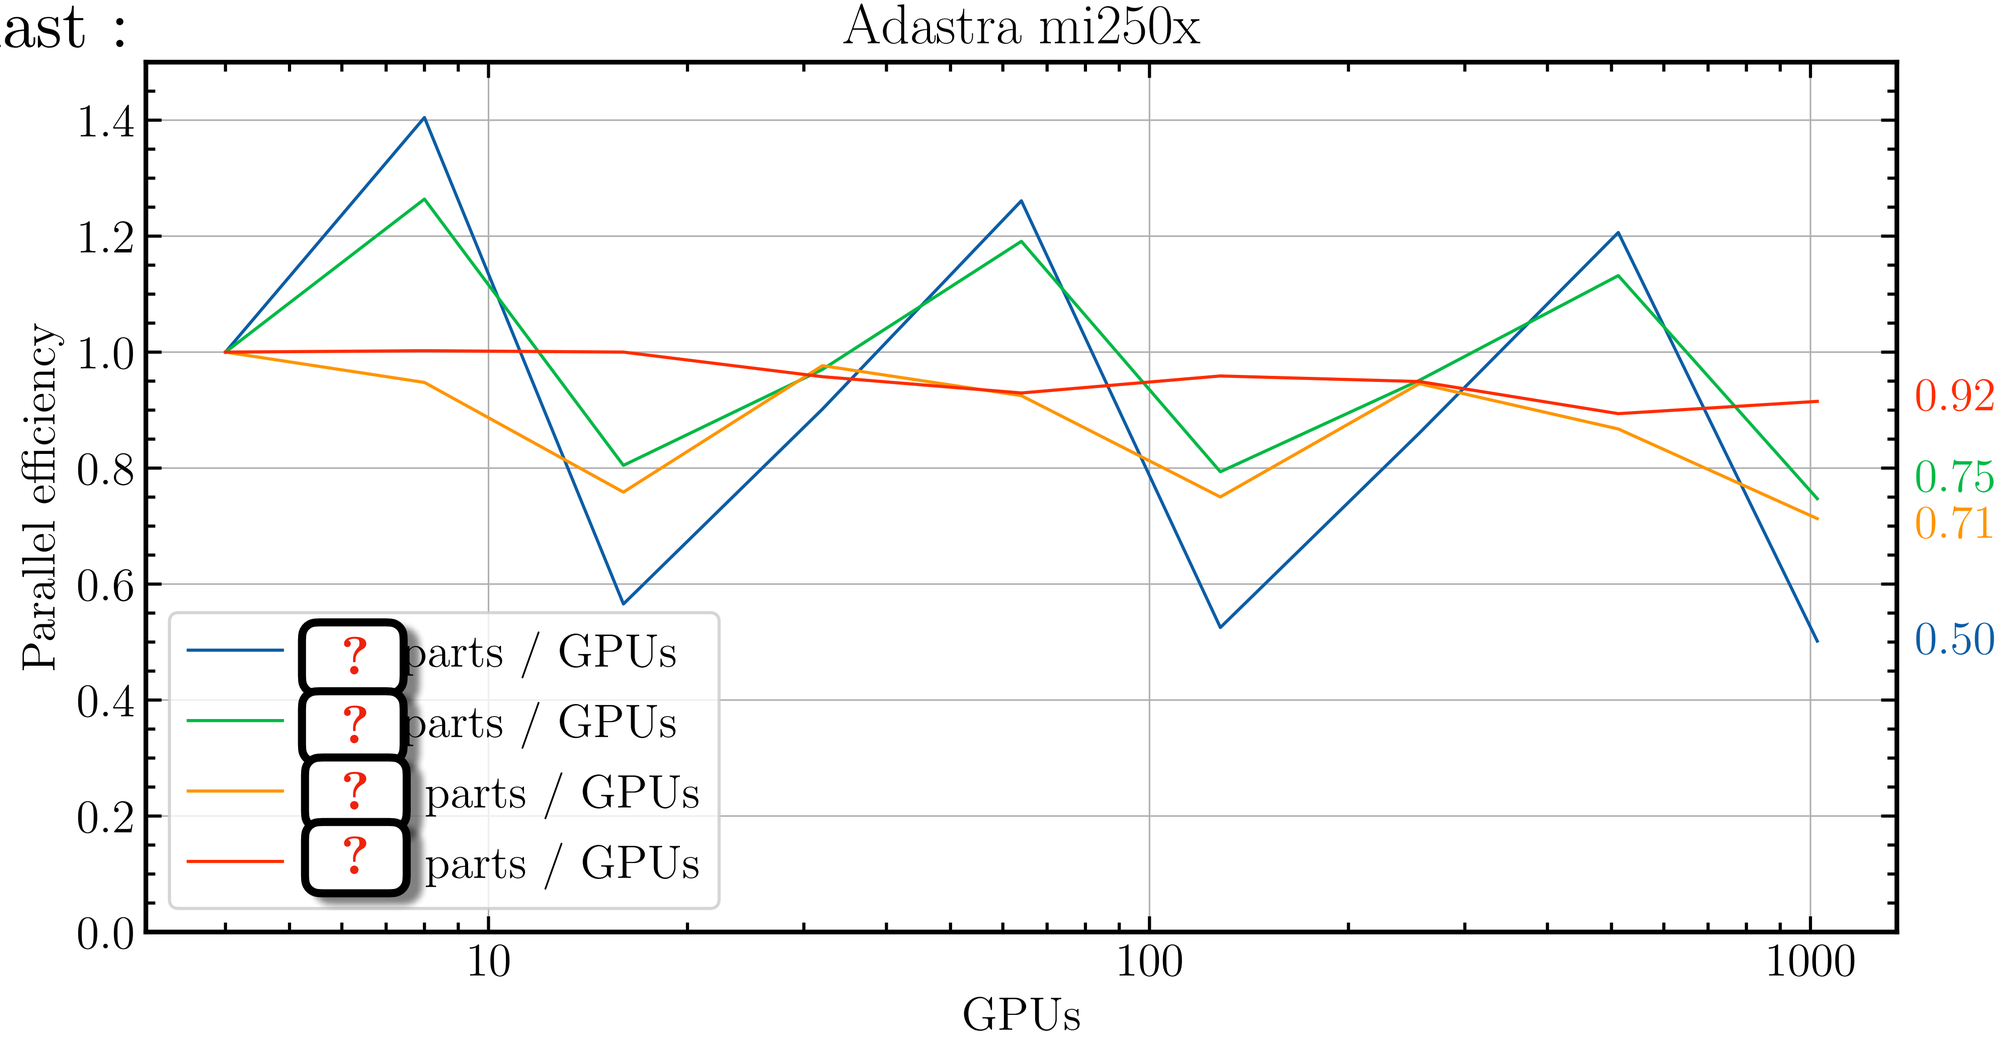

0.92 efficiency !!!

Now if you take a big cluster

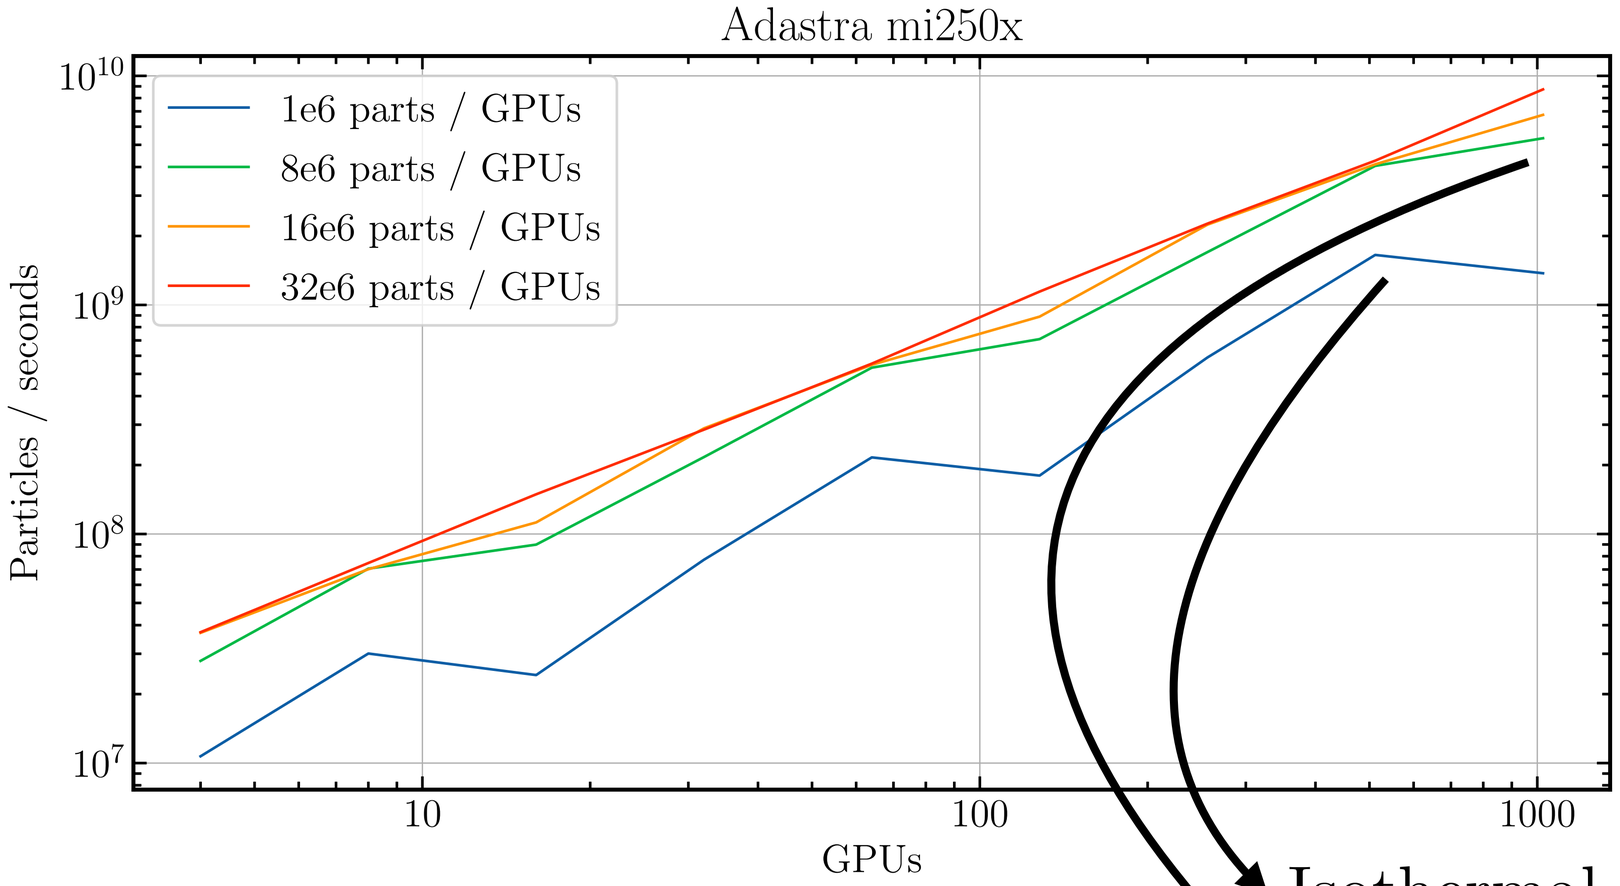

- 9G part / seconds
- 65.000.000.000 particles
  (4000^3)
- 7 sec / iterations

Global time steps :

(Estimate 10k steps at 1M part)

Isothermal disc : 1G part = 1 orbit/16h

Isothermal disc : 16G part = 1 orbit/4day

But it is also efficient

Adastra MI250 :

- 350W / GPU
- 90W CPU
- 80W Memory
- 1900W / Node (70% max power)

256 nodes tests : 486kW

(2/3 Adastra mi250)

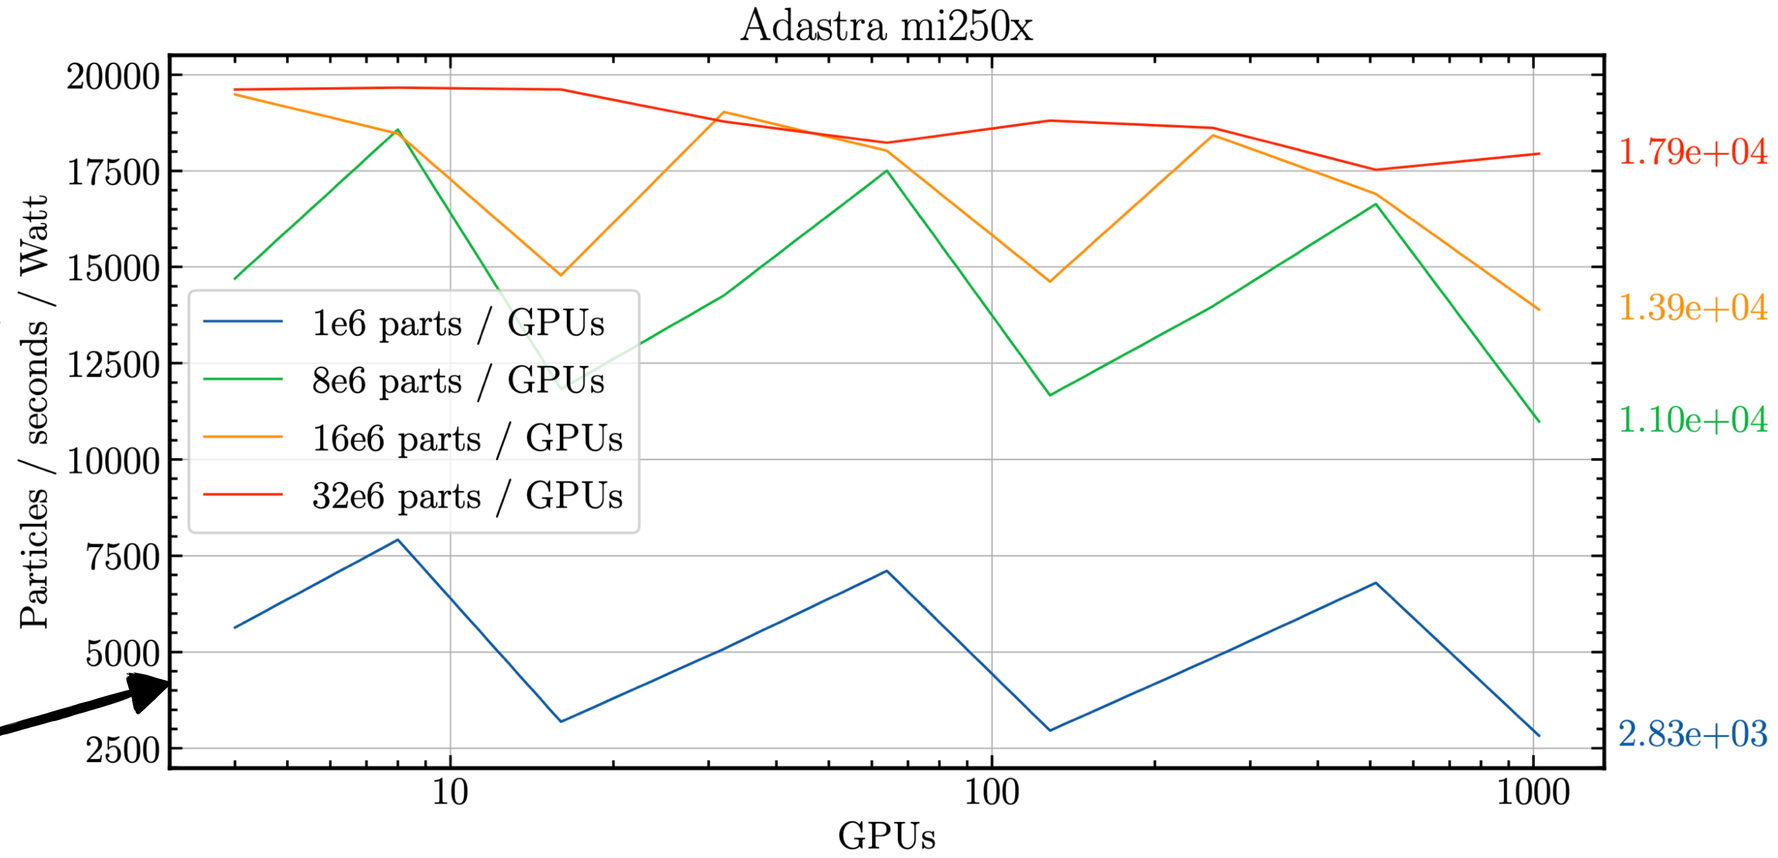

Epyc 7742 (Rome) Rack

4700 part/s/W

Large simulations required !

Context & Aims

1st (top500), GPU + CPU

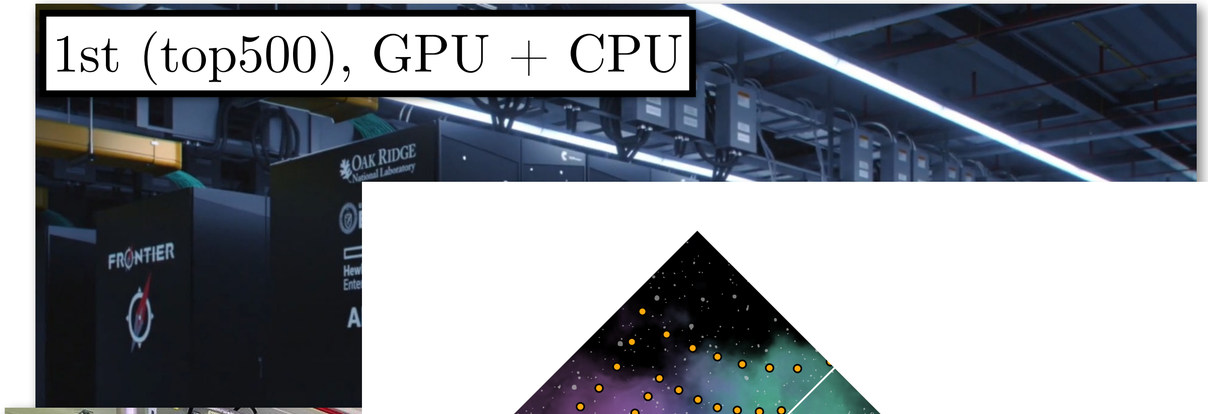

17th (top500),

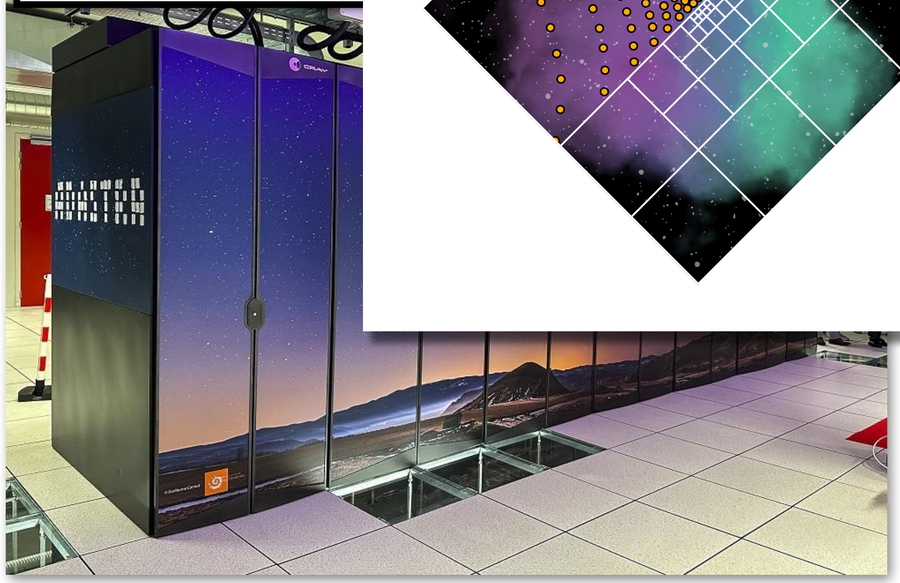

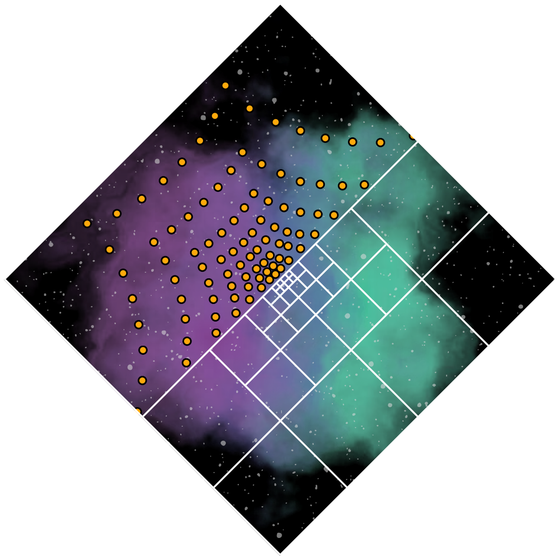

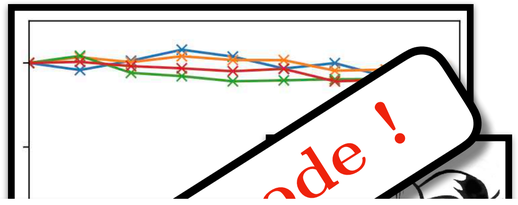

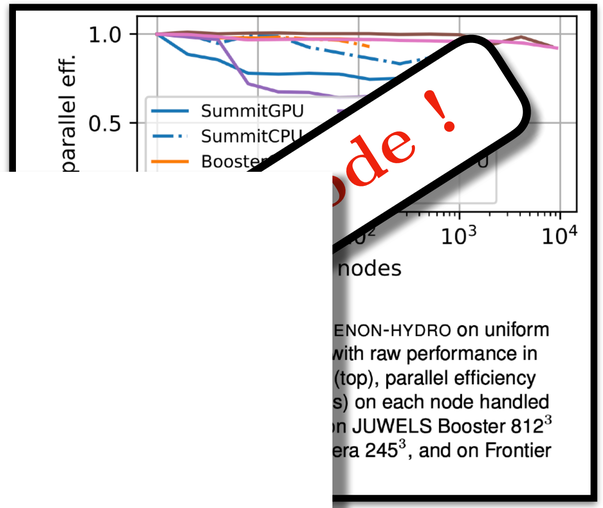

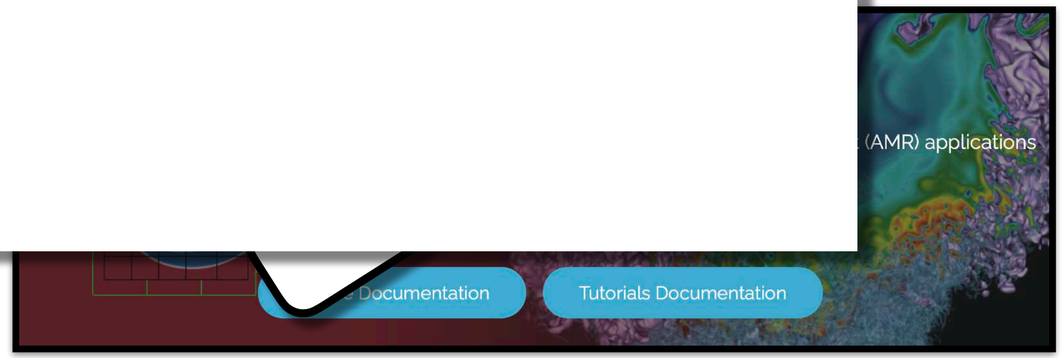

Where is SPH ???


HPE Cray
HPE Cray
Ad astra per a


Thanks !