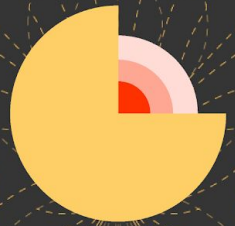

ChEESE

www.cheese-coe-eu

Center of Excellence for Exascale in Solid Earth

Successful interaction of ChEESE and POP: CoEs cross-collaboration

Mauricio Hanzich and Claudia Rosas

Barcelona Supercomputing Center

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823844

Our Story Today

- CoEs
- What is ChEESE?
- Flagship Codes
- Fall3D: from level 3 and beyond

1

What is a CoE?

Center of Excellence

A premier organization providing an exceptional product or service in an assigned sphere of expertise and within a specific field of technology, [...], consistent with the unique requirements and capabilities of the CoE organization." --- The Software Engineering Institute of the Carnegie Mellon University

Centers of Excellence on HPC

POP + ChEESE

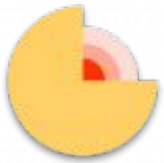

ChEESE

Center of Excellence for Exascale in Solid Earth

1

What is ChEESE?

Center of Excellence in Solid Earth

ChEESE is a new initiative to integrate HPC and data across the Solid Earths' related disciplines in Europe under the same umbrella."


Consortium

€ 7 683 241,25
EU Contribution

12 partners
7 countries

42 deliverables
15 Public

5 peer-reviewed publications
Up to date… and more to come

18 international events
Workshops, Meetings, Conferences, Hackatons

12 Pilot demonstrators
10 Flagship codes

36 months
31 October 2021

ChEESE Structure

WP1 - Management

Administrative and financial management, Operational decision making, Internal communication.

WP2 - HPC Software Engineering

Building blocks of the project, solve performance bottlenecks, prototype testing.

WP3 - Modelling workflows and tools

Address challenges: data management, complex workflows, post-processing and visualization.

WP4-Exascale Pilot Demonstrators

Geophysical simulations, reduce gap between HPC algorithms and end-users, end-to-end solution.

WP5 - Service Validation and Enabling

Make Pilot demonstrators available as services to a broader user community; implementation and validation

WP6 - Dissemination, Training and Training

Communication and dissemination, identify exploitable results, training activities.

ChEESE Structure

WP1 - Management

Administrative and financial management, Operational decision making, Internal communication.

WP2 - HPC Software Engineering

Building blocks of the project, solve performance bottlenecks, prototype testing.

WP3 - Modelling workflows and tools

Address challenges: data management, complex workflows, post-processing and visualization.

WP4-Exascale Pilot Demonstrators

Geophysical simulations, reduce gap between HPC algorithms and end-users, end-to-end solution.

WP5 - Service Validation and Enabling

Make Pilot demonstrators available as services to a broader user community; implementation and validation

WP6 - Dissemination, Training and Training

Communication and dissemination, identify exploitable results, training activities.

ChEESE Structure

WP1 - Management

Administrative and financial management, Operational decision making, Internal communication.

WP2 - HPC Software Engineering

Building blocks of the project, solve performance bottlenecks, prototype testing.

3 - Modelling workflows tools

ress challenges: data agement, complex kflows, post-processing visualization.

WP4-Exascale Pilot Demonstrators

Geophysical simulations, reduce gap between HPC algorithms and end-users, end-to-end solution.

WP5 - Service Validation and Enabling

Make Pilot demonstrators available as services to a broader user community, implementation and validation

WP6 - Dissemination, Training and Training

Communication and dissemination, identify exploitable results, training activities.

Flagship Codes

Flagship Codes

Level 1

Codes already at the Petascale, with current scalability proven up to about 10^4 -10^5 cores or above.

Level 2

Codes currently at a pre-Petascale, with scalability proven up to about 10^3 -10^4 cores.

Level 3

Codes with scalability bottlenecks at around or below 10^3 cores.

Codes Vs. Pilots

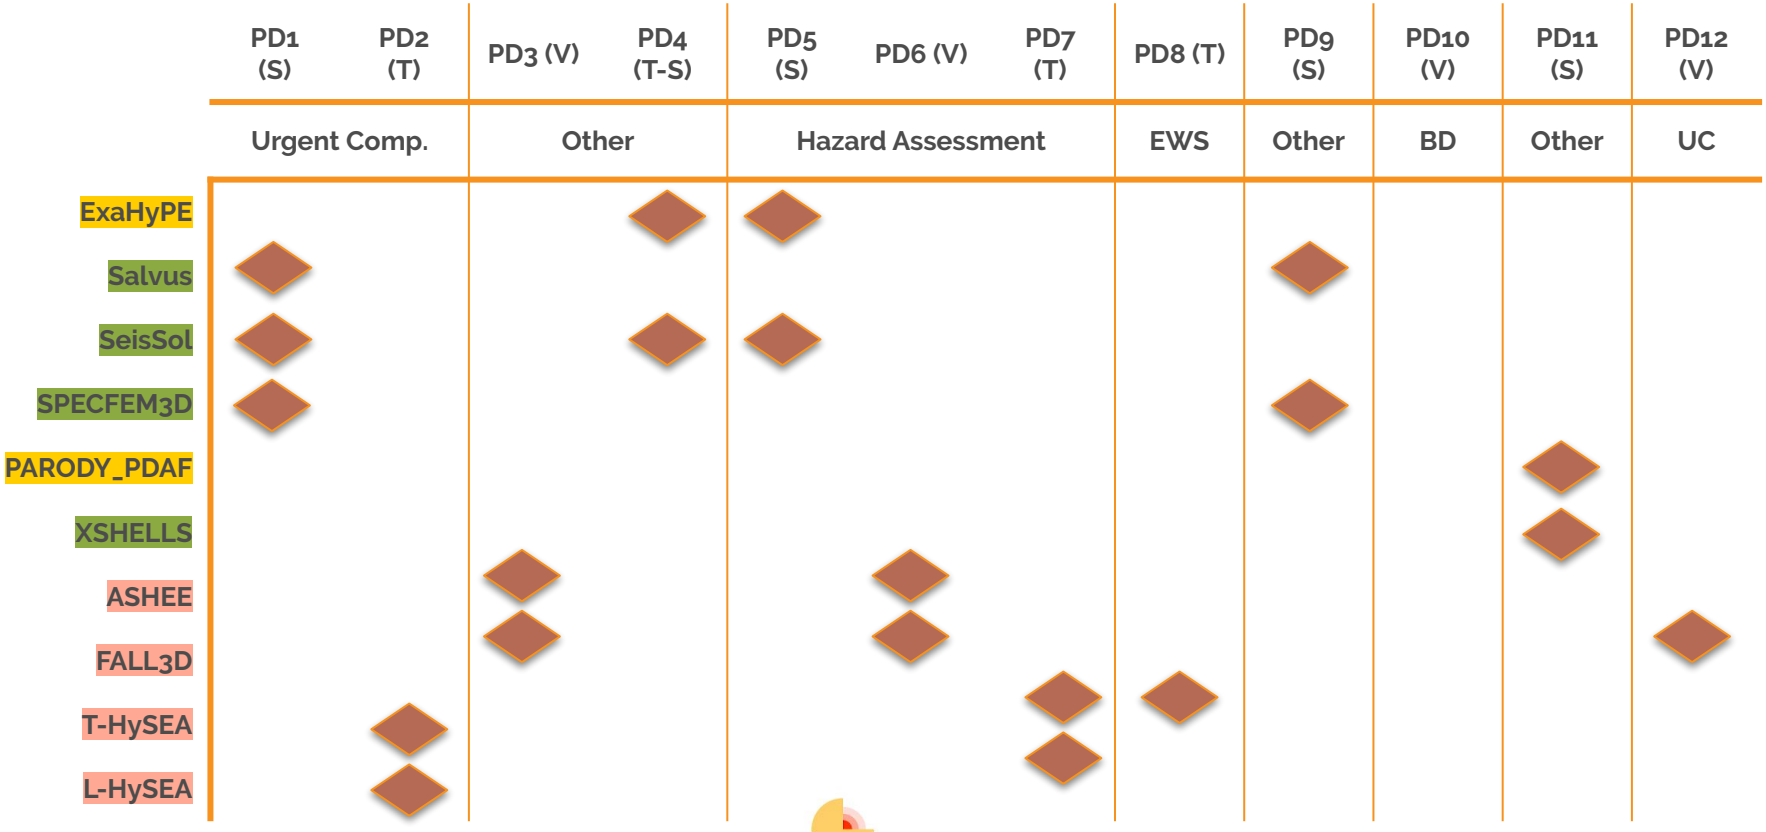

Codes Vs. POP

	Who?	Where?	What?
ExaHyPE	TUM (Developers)	SuperMUC	Almost perfect weak and strong scaling
Salvus	JSC (POP)	Piz Daint	I/O inefficiencies with scale.
SeisSol	JSC (POP)	MareNostrum IV	Load imbalances but good comp. and comm. effic.
SPECFEM3D	JSC (POP)	Irene	Good performance in all aspects
PARODY_PDAF	NAG (POP)	Joliot-Curie SKL	Hybrid performance
XSHELLS	NAG (POP)	KNL	Hybrid parallel efficiency pending to determine
ASHEE	INGV (Developers)	MareNostrum IV	Low weak scaling efficiency
FALL3D	BSC (POP)	MareNostrum IV	Fine granularity exposes low computation/comm ratio
T-HySEA	BSC (POP)	CTE Power	Scalability problems from MPI waiting time and overhead
L-HySEA	Reuse of T-HySEA		

Our target today... Fall3D

	Who?	Where?	What?
ExaHyPE	TUM (Devs)	SuperMUC	Almost perfect weak and strong scaling
Salvus	JSC (POP)	Piz Daint	I/O inefficiencies with scale.
SeisSol	JSC (POP)	MareNostrum IV	Load imbalances but good comp. and comm. effic.
SPECFEM3D	JSC (POP)	Irene	Good performance in all aspects
PARODY_PDAF	NAG (POP)	Joliot-Curie SKL	Hybrid performance
XSHELLS	NAG (POP)	KNL	Hybrid parallel efficiency pending to determine
ASHEE	INGV (Devs)	MareNostrum IV	Low weak scaling efficiency
FALL3D	BSC (POP)	MareNostrum IV	Fine granularity exposes low computation/comm ratio
T-HySEA	BSC (POP)	CTE Power	Scalability problems from MPI waiting time and overhead
L-HySEA	Reuse of T-HySEA		

Fall3D

From level 3 and beyond


"FALL3D is a Eulerian model for the atmospheric transport and ground deposition of volcanic tephra (ash)"

V8 Implemented in ChEESE

V7.3.4

- Atmospheric dispersal and deposition of tephra particles
- Lax-Wendorff (LW) central-difference scheme
- Parallelization on particle bins (expensive in comms)
- Master responsible for I/O

V8.0

- Extended model for additional types of particles, aerosols and radionuclides
- High-resolution central-upwind scheme; optionally a 4th-order Runge-Kutta explicit scheme
- Parallelisation strategy based on a full 3D domain decomposition
- Parallel model I/O using netCDF-Par and parallel model pre-process.

Interaction with POP

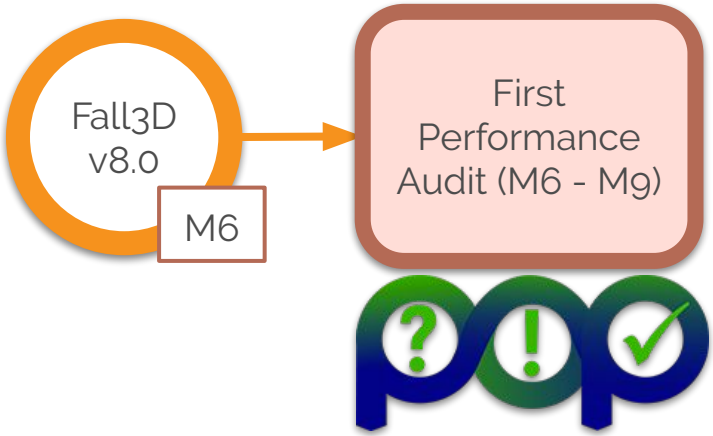

Interaction with POP

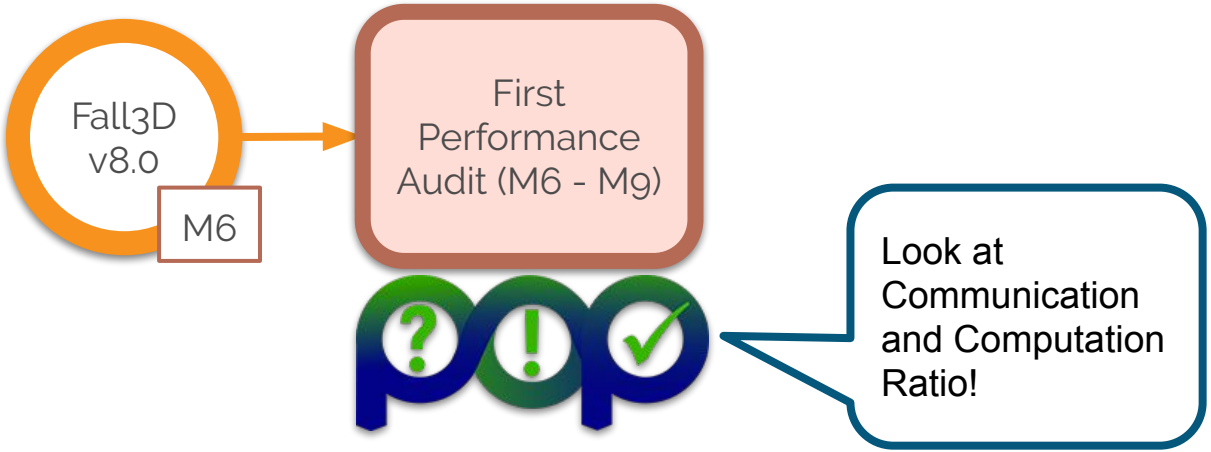

Interaction with POP

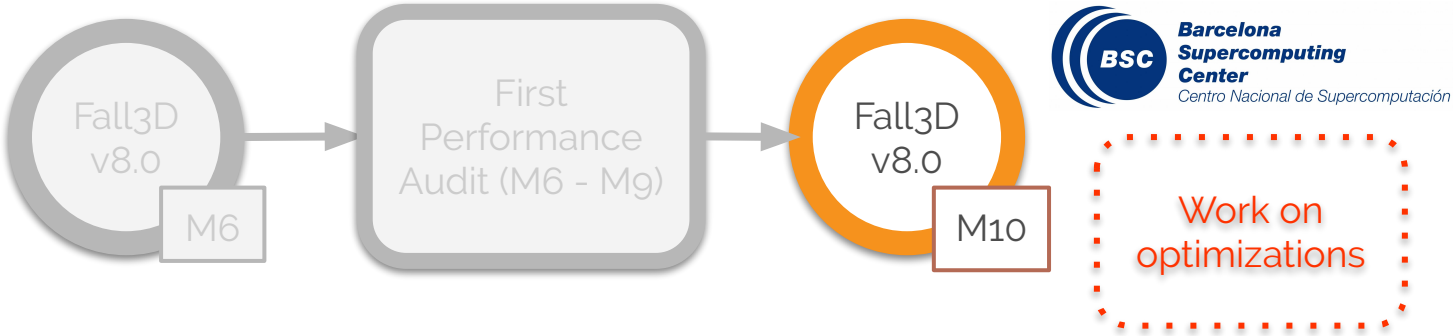

Interaction with POP

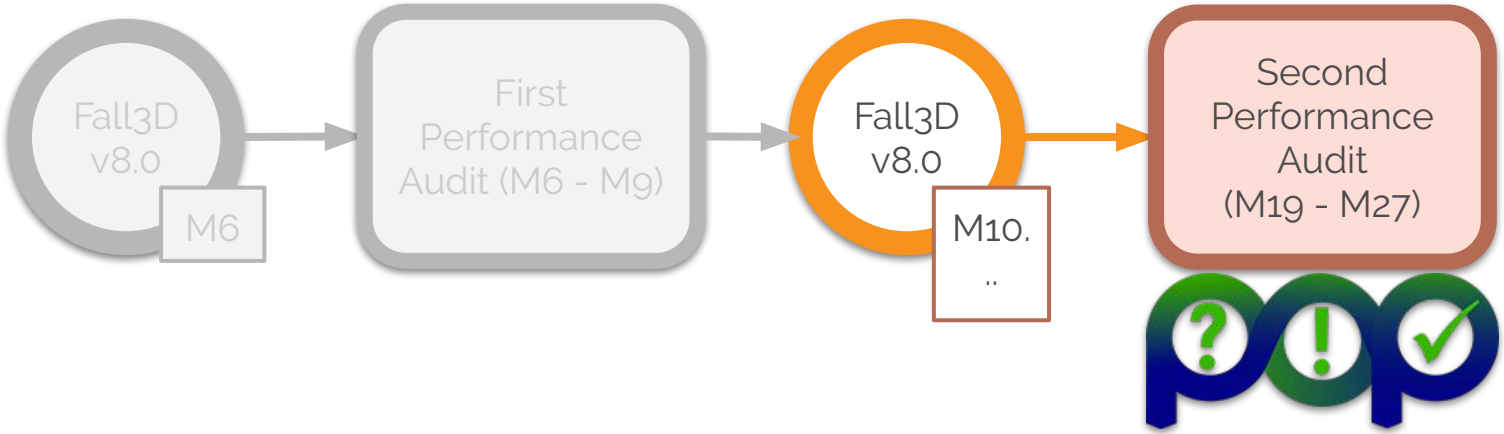

Interaction with POP

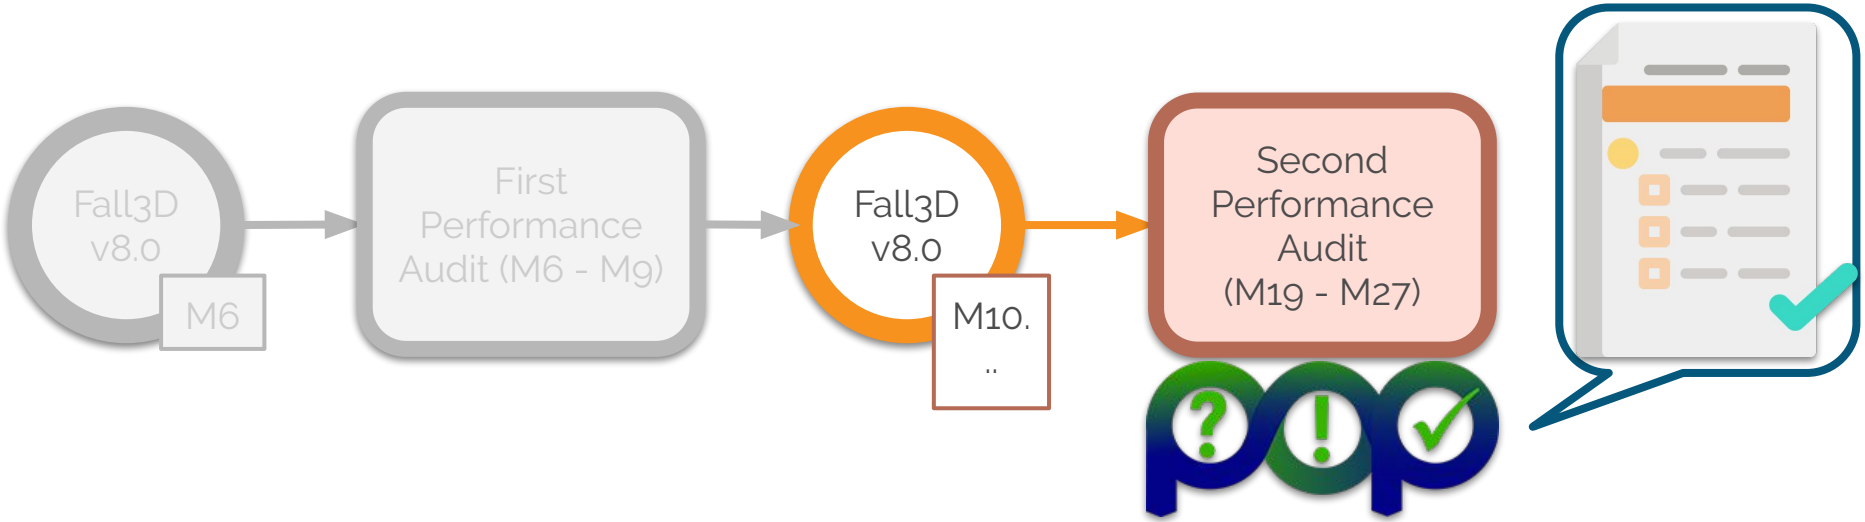

Interaction with POP

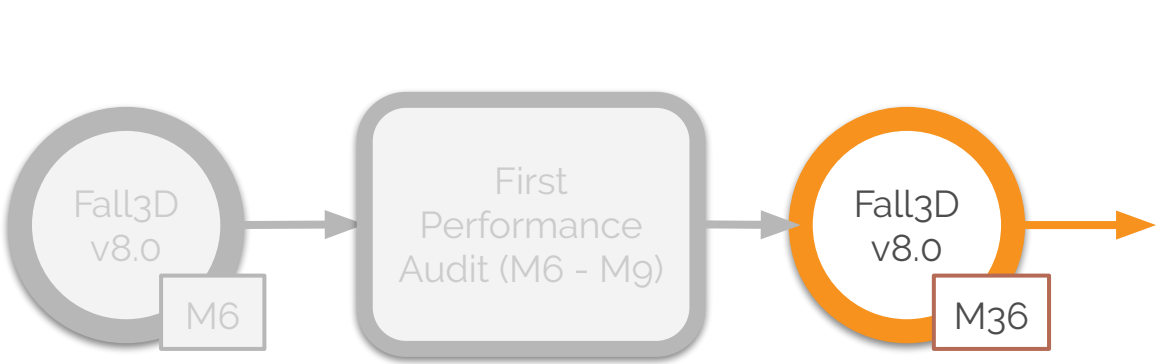

Optimization

Improve computational and communication ratio

Which are the communication problems?

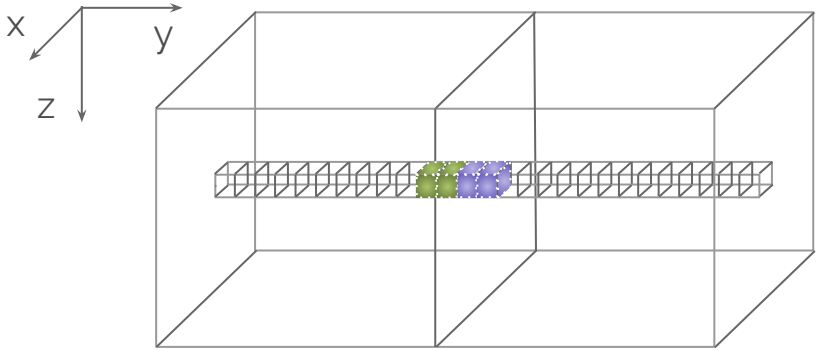


```
for x:
    for z:
        solve1D(y) →  step1
                      comm1
                      step2
                      comm2
                      step3
                      comm3
                      step4
```

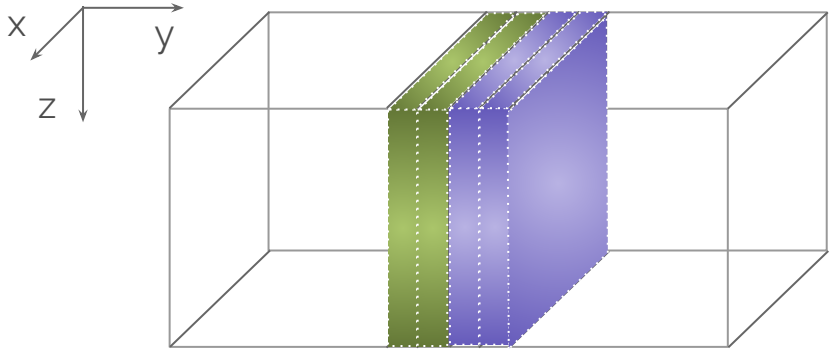


```
for n in (1, 4):
    for x:
        for z:
            solve1D_step_n(y)

    comm_n
```

Optimization issues

- Higher computing/communication ratio
- Data rearrangement, aligned with computing dimension
- Reduced decision statements in iterative code region
- Removed memory allocations from iterative code region
- Chance of using 1st-degree Euler solver

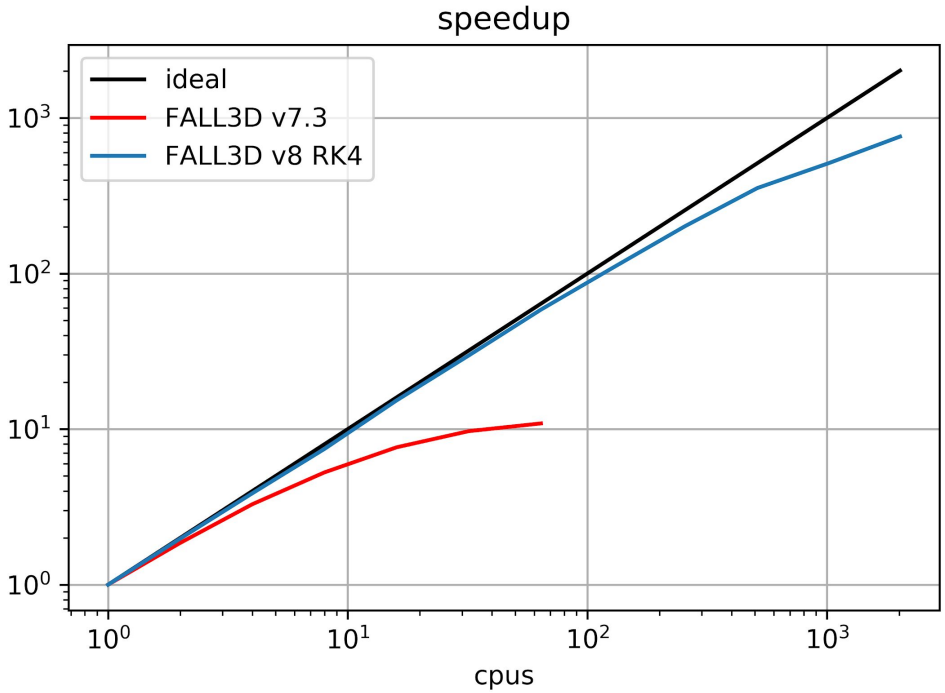

V7.3 serie vs V8(RK4) serie

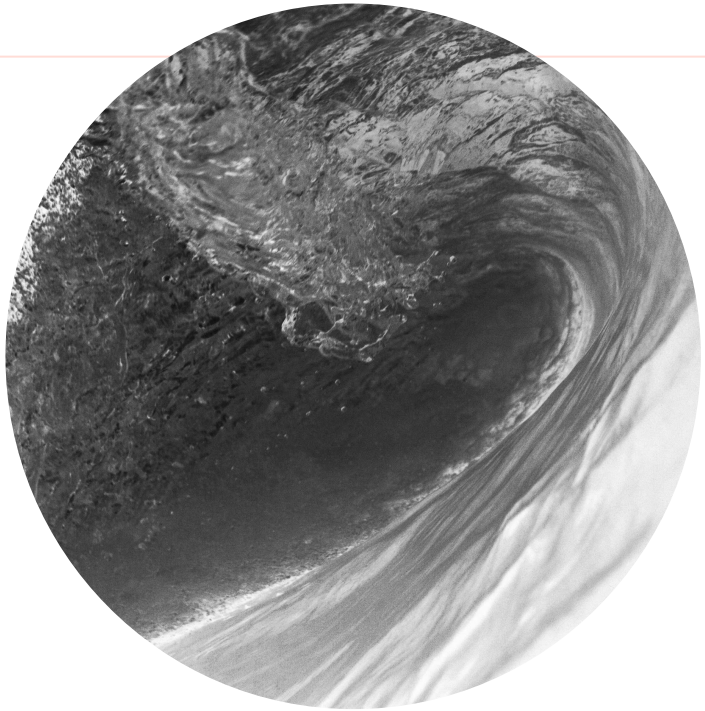

Euler

Uses same amount of resources yet sacrifices precision

Simplified computation from order 4 to order 1

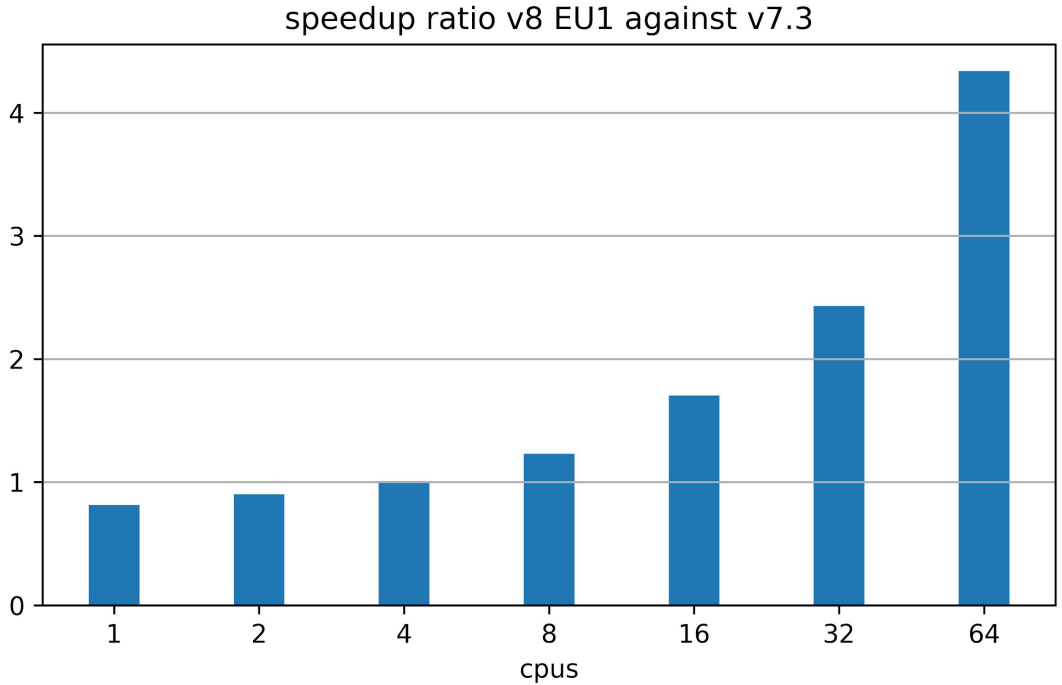

V7.3 serie vs V8(EU1) serie

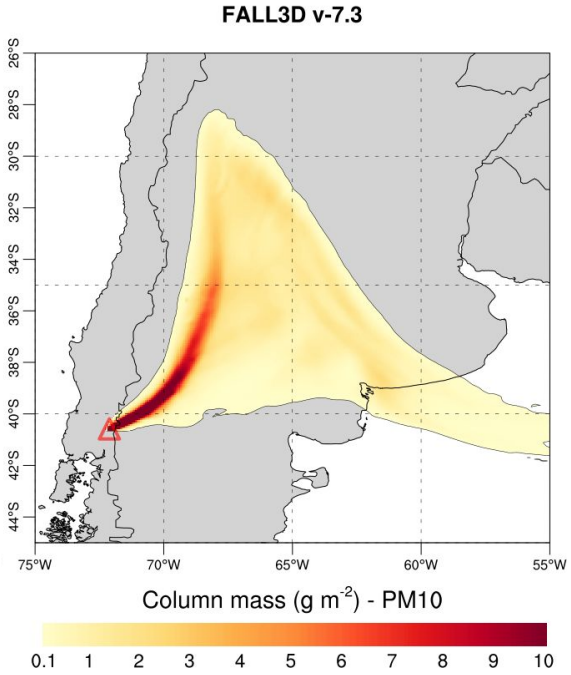

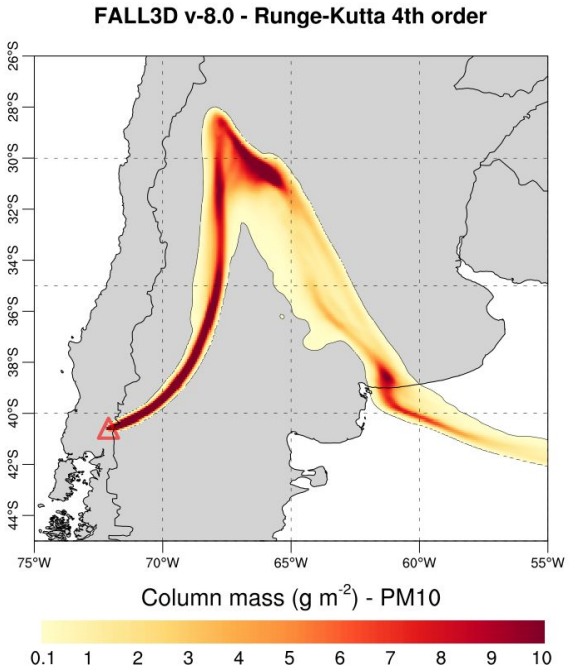

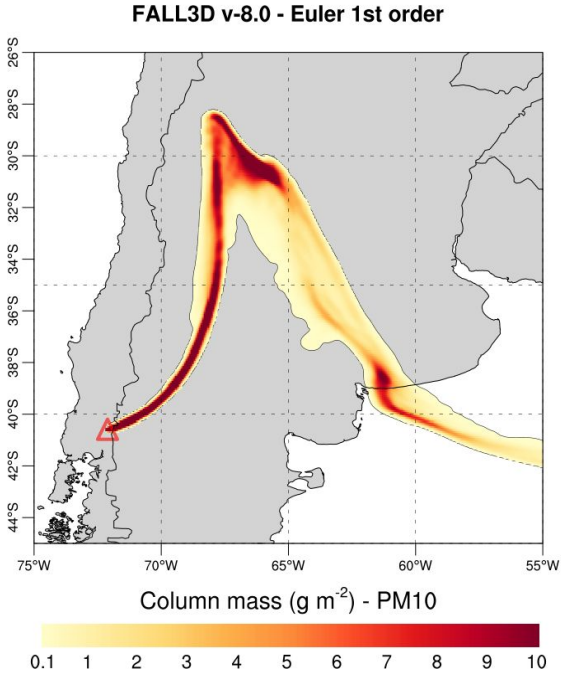

Dispersion results from FALL3D depending on the version

Next steps in Fall3D

- In last year of ChEESE, there will be another performance audit for Fall3D
- Continue improving scalability and efficiency of current new version of Fall3D

To publish the current performance of the code (already under review).

Next steps in ChEESE

- D2.1b Code Audit (M27) same methodology for the ten flagship codes
- Performance improvements in scalability in most of them
- Promote the successful interaction with POP as transversal CoE.
  - Already used in other H2020 projects (e.g. ENERXICO)

*Performance optimization of HPC codes following POP's methodology have proven to be **efficient** in time and use of resources and **significantly useful** for scientific software developers."*

*"The interaction has been **successful** not only within CoEs but also easily **exported** to international project with strong participation of Latin America."*

Thanks!

Any questions ?

You can find us at

- mhanzich@bsc.es
- crosas@bsc.es

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823844