ChEESE

CASTIEL-2 code of the month vol.8

(FALL3D and OpenPDAC)

29th May 2024

GEO3BCN

ISTITUTO NAZIONALE DI GEOFISICA E VULCANOLOGIA

Project funded by EuroHPC under the grant agreement No 101093038.

1

Introduction

Arnau Folch
Geociencias Barcelona (CSIC), Spain

2

FALL3D

Arnau Folch
Geociencias Barcelona (CSIC), Spain

3

OpenPDAC

Mattia De' Michieli Vitturi
Istituto Nazionale Geofisica e Vulcanologia (INGV-Pisa), Italy

Project funded by EuroHPC under the grant agreement No 101093038.

ChEESE flagship codes: overview

- ChEESE is preparing **11 open-source flagship codes** in different domains of Solid Earth in terms of performance, scalability, CI/CD in EuroHPC systems, and portability across current and emerging hardware architectures.
- In addition, **7 mini-apps** for co-design on OpenSequana (EuPEX) and RISC-V (EuPilot) exascale hardware prototypes.

Solid Earth Domain	No	Code	Accelerated	Mini-app
Computational Seismology	1	SeisSol	CUDA, SYCL	yes
	2	SPECFEM3D	CUDA, HIP	yes
	3	ExaHyPE	on-going	no
	4	Tandem	on-going	yes
Magnetohydrodynamics	5	xSHELLS	CUDA	yes
Tsunami modelling	6	HySEA	CUDA	yes
Volcanology	7	FALL3D	OpenACC	yes
	8	OpenPDAC	on-going	no
Geodynamics	9	LaMEM	on-going	no
	10	pTatin3D	CUDA	yes
Glacier modelling	11	Elmer/ICE	on-going	no

ChEESE flagship codes in volcanology

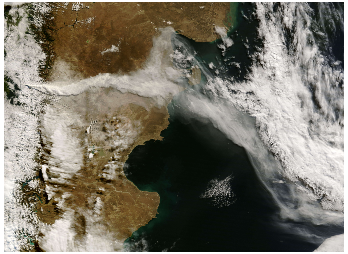

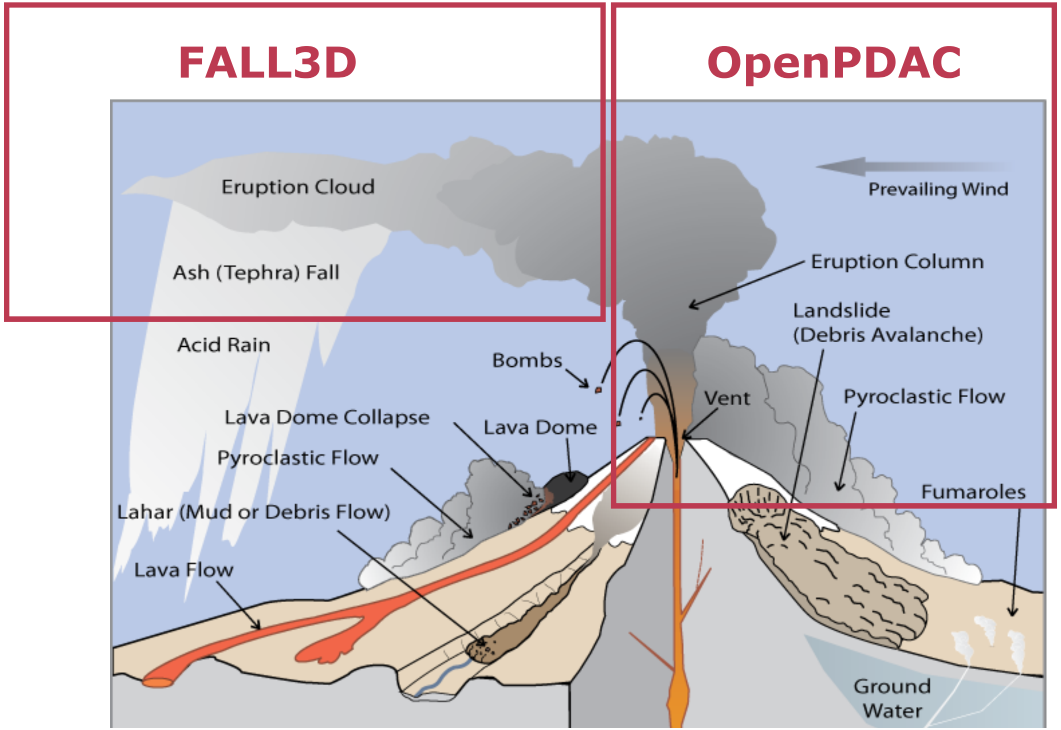

1

Introduction

Arnau Folch
Geociencias Barcelona (CSIC), Spain

2

FALL3D

Arnau Folch
Geociencias Barcelona (CSIC), Spain

3

OpenPDAC

Mattia De’ Michieli Vitturi
Istituto Nazionale Geofisica e Vulcanologia (INGV-Pisa), Italy

Project funded by EuroHPC under the grant agreement No 101093038.

FALL3D: code overview (1/2)

- FALL3D is an **Eulerian model** for the atmospheric transport and ground deposition of "passive particles".
- Code originally developed for volcanic particles, later on extended to other types of particles like mineral dust, aerosols (SO2 , H2O), or radionuclides (radioactive decay).
- FALL3D is a **multi-purpose model**, it can be used to compute:
  - Airborne concentration (e.g. at flight levels).
  - Fallout deposit (ground accumulation).
- FALL3D is a **multi-scale model**, can run from local-scales (few kms) to continental scales (1000s kms).
- FALL3D solves one Advection-Diffusion-Sedimentation (ADS) equation for each "particle class".

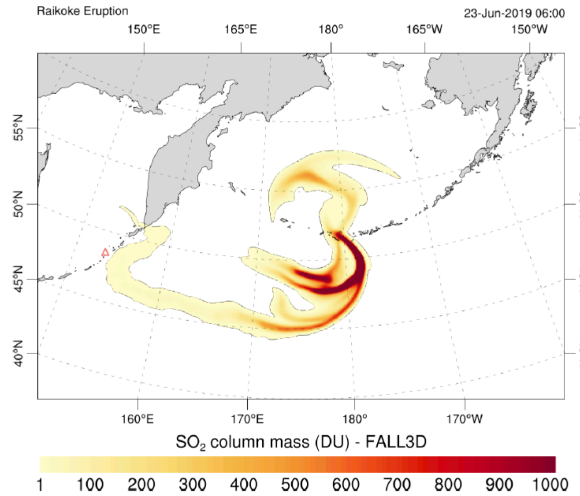

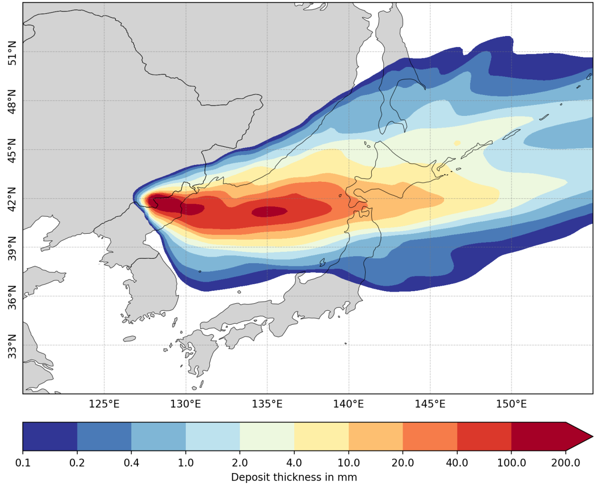

FALL3D: code overview (2/2)

- Code written in modern ("object oriented") FORTRAN.
- **Finite volumes** numerical scheme on a structured 3D grid (mapping) and explicit in time (R-K up to 4th order in time).
- Two levels of code parallelization:
    - Domain decomposition (1 MPI rank per subdomain)
    - Ensemble modeling (embarrassing parallelism) with an embedded workflow
- FALL3D supports data assimilation:
    - Data insertion from satellite retrievals (initial condition).
    - Data assimilation cycles using local/global Ensemble Transform Kalman Filters (EnTKF)

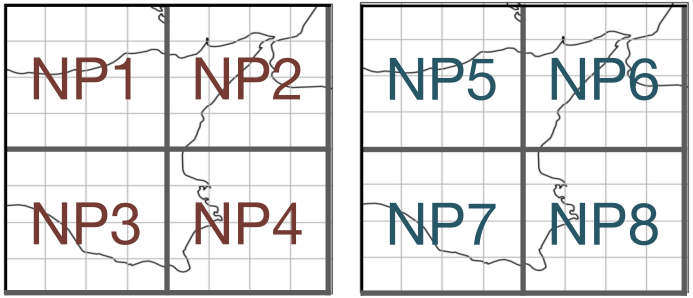

Example with 2 ensemble members and 8 processors (Ne=2, NPx=2, NPy=2)

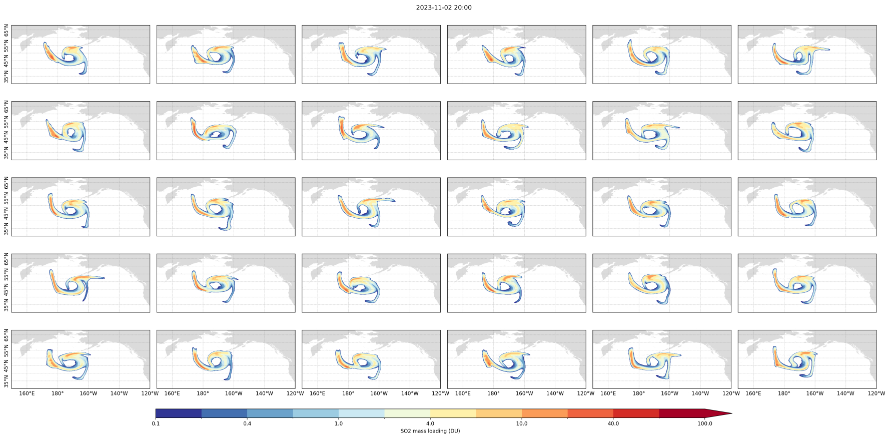

FALL3D: relevant code users

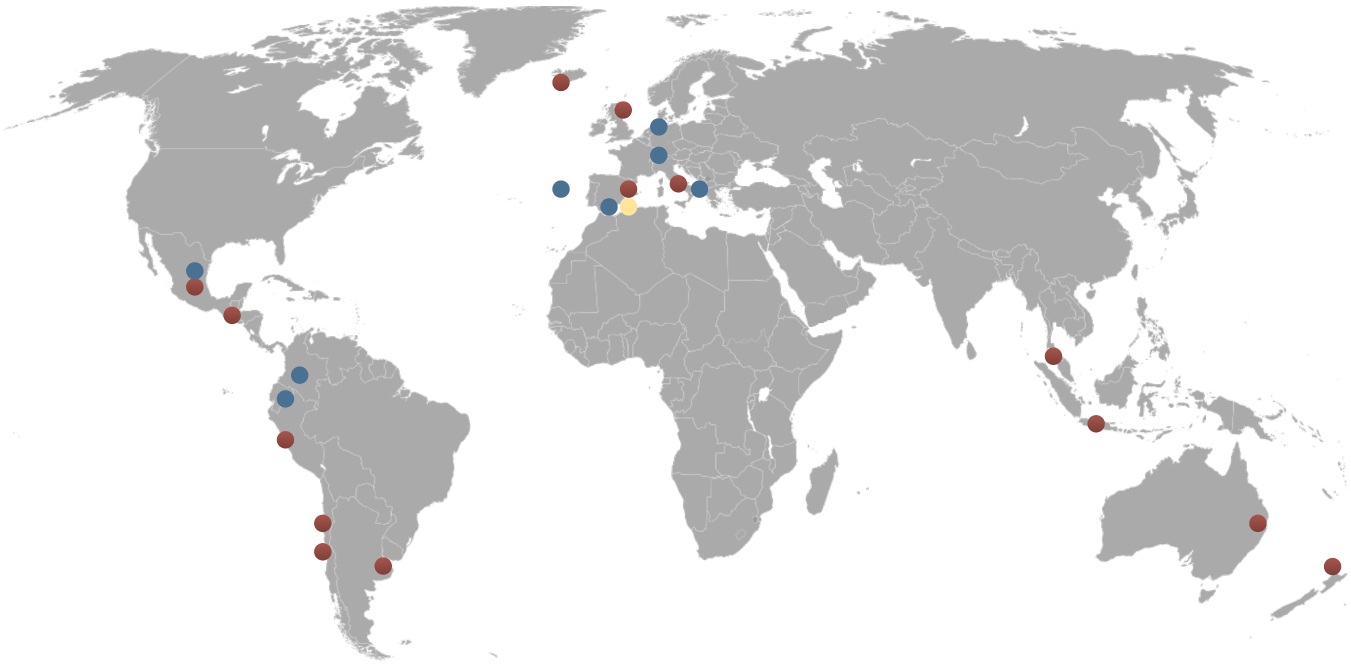

Public institutes (operational):
- INGV (Italy)
- CSIC (Spain)
- IMO (Iceland)
- BGS (UK)
- SMN (Argentina)
- IGP (Perú)
- SERNAGEOMIN (Chile)
- Dirección Meteorológica de Chile (Chile)
- BMKG (Indonesia)
- VAAC Buenos Aires (Argentina)
- BSC (Spain)
- GNS (New Zealand)
- Bureau of Meteorology (Australia)
- EOS (Singapore)
- Cenapred (México)
- INSIVUMEH (Guatemala)

Academic:
- University of Geneva (Switzerland)
- University of Bari (Italy)
- University of Bremen (Germany)
- University of Granada (Spain)
- UNAM (México)
- Univ. San Francisco de Quito (Ecuador)
- Universidad de Nariño (Colombia)
- University Azores (Portugal)

Private:
- Mitiga Solutions (Spain)

FALL3D: academic impact (proxy from peer-reviewed publications)

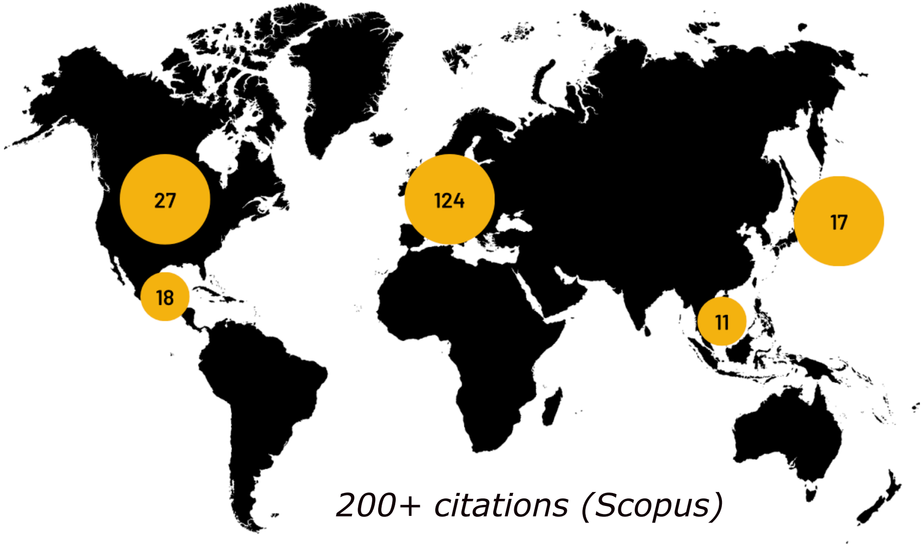

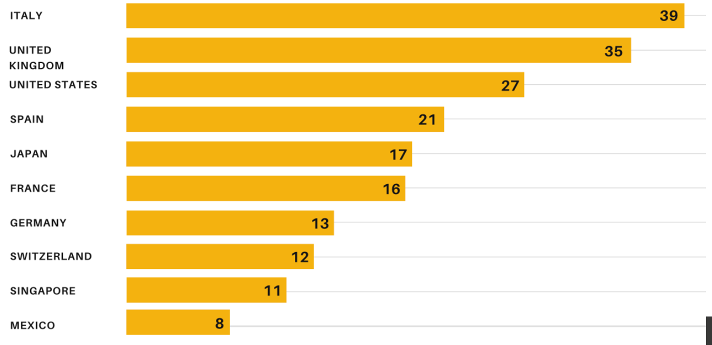

FALL3D CITATIONS

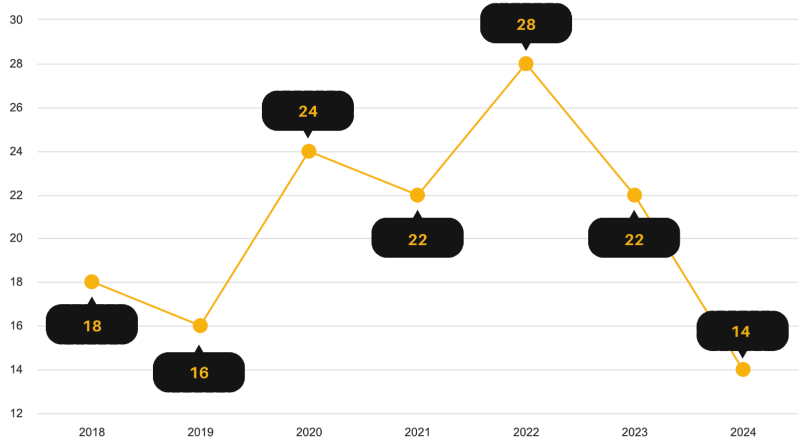

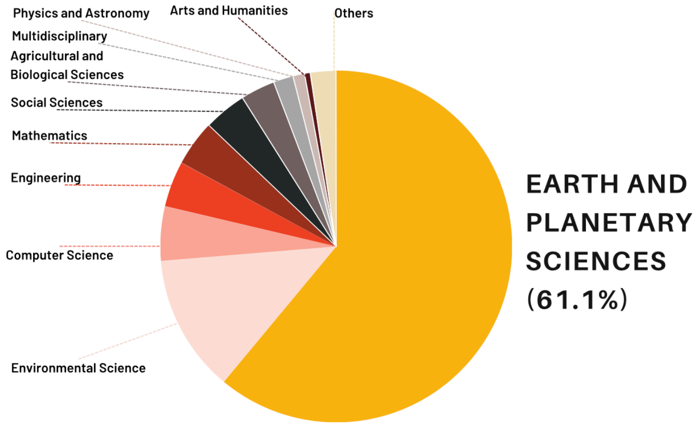

FALL3D: ChEESE-1P legacy

- During ChEESE-1P (2018-2022), FALL3D was heavily refactored.
- **CPU version** (MPI):
  - Overall speed-up was increased by 4.3x (e.g. above 50% parallel efficiency on 4096 cores at irene-rome@CEA).
  - Vectorization: speed-up up to 1.2x
  - Parallel I/O: performance increased by 2x
- **New GPU version** (MPI + OpenACC):
  - ChEESE-1P considered NVIDIA GPUs only.
  - Drastic decrease of time-to-solution
  - Poor scalability (50% on 16 GPUs at M-100@CINECA)

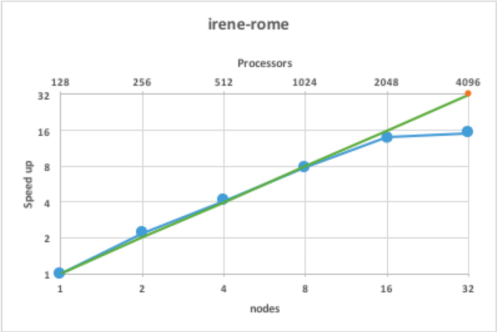

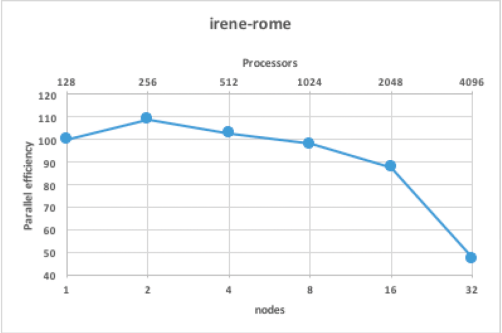

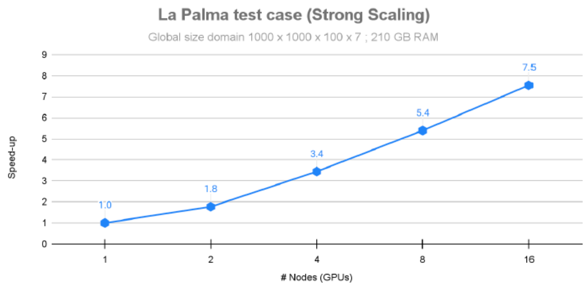

FALL3D: repository (CPU/GPU) https://gitlab.com/fall3d-suite/

F FALL3D suite

Subgroups and projects | Shared projects | Inactive

- T Tests
  - L Leonardo
  - L LUMI
- C course
- F FALL3D Owner
- F FALL3D GPU
- F fall3d-suite.gitlab.io
  FALL3D Suite - Main Website
- M Mini-FALL3D
- M Mini-FALL3D GPU

F FALL3D

Unstar 3 | Fork 6

master | fall3d / | + | History | Find file | Edit | Code

Increase size figure in User guide
lmingari authored 2 weeks ago — 22c8f995

Name	Last commit	Last update
Example	New source type: NONE	2 months ago
Manual	Added User Guide	4 months ago
Other	Remove prompt in user guide	2 weeks ago
Sources	Generate meteo profile only when is re...	2 months ago
autoconf	Added configure files for public distribu...	1 year ago
docs	Increase size figure in User guide	2 weeks ago
m4	Corrected configure.ac (compatible wit...	1 year ago
.gitignore	original gitignore	1 year ago
.gitlab-ci.yml	Use rust latest	4 weeks ago
AUTHORS.md	gold badge	1 year ago
CHANGELOG.md	json is optional now	4 months ago
CITATION.cff	Update CITATION.cff	4 months ago
CITING.md	gold badge	1 year ago
CODE_OF_CONDUCT.md	New CONTRIBUTING.md file	1 year ago
CONTRIBUTING.md	New CONTRIBUTING.md file	1 year ago
LICENSE	Modified mod_nc_IO for serial writing ...	4 years ago
Makefile.am	Corrected Makefile.am	11 months ago
Makefile.in	Merge branch 'master' into hotfix/gmac...	10 months ago
README.md	Update README.md	4 months ago
aclocal.m4	automake regenerated	10 months ago
build	New upgrade of the build script	4 years ago
configure	json is optional now	4 months ago
configure.ac	Update configure.ac	4 months ago

Project information

- 633 Commits
- 1 Branch
- 7 Tags
- 22 MiB Project Storage
- 2 Releases

- README
- GNU GPLv3
- CHANGELOG
- CONTRIBUTING
- CI/CD configuration
- GitLab Pages
- + Add Kubernetes cluster
- + Add Wiki
- + Configure Integrations

Created on
February 20, 2023

EOSC-SYNERGY * SOFTWARE * COMPLIANT *

FALL3D: repository (tests)

https://gitlab.com/fall3d-suite/

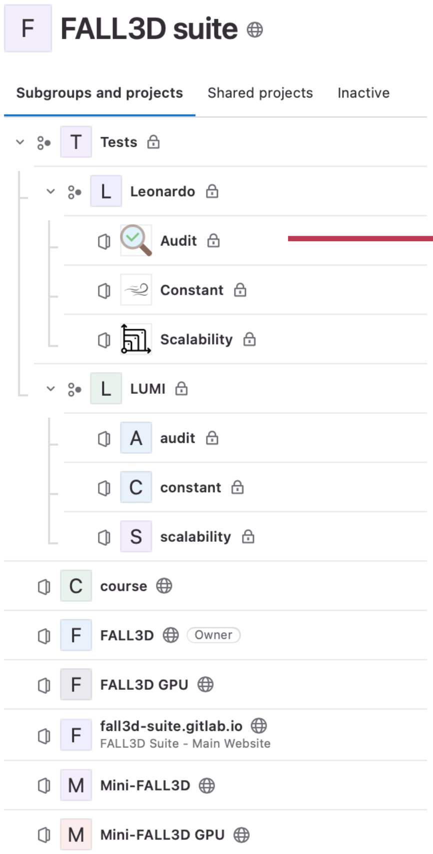

test_audit_L00.json 1.65 KiB


```json
{
    "test": {
        "testType": "test_audit",
        "testOwner": "Arnau Folch",
        "note": "Copy and modify this template in each new test_audit_ID folder",
        "comment": "Baseline version without Extrae instrumentation. At M09 of ChEESE-2P project"
    },
    "code": {
        "origin": "decube83/fall-3-d-gpu-cicd.git",
        "version": "8.2.1-gpu",
        "branch": "/cicd/hpascual",
        "commitID": "30349dce61c8ba3bd3f20ea79041c543a1471f76",
        "commitAuthor": "decube83 <hpascual@geo3bcn.csic.es>",
        "commitDate": "Mon Nov 06 09:15:29 2023 +0100"
    },
    "system": {
        "host": "Leonardo@CINECA",
        "partition": "booster",
        "nodeCPU": "32 cores Intel Ice Lake ",
        "nodeGPU": "4 NVIDIA Ampere 64GB",
        "nodeRAM": "512 GB"
    },
    "modules": {
        "fortran": "nvhpc/23.1",
        "cuda": "cuda/11.8",
        "mpi": "openmpi/4.1.4--nvhpc--23.1-cuda-11.8",
        "hdf5": "hdf5/1.12.2--openmpi--4.1.4--nvhpc--23.1",
        "netcdf-c": "netcdf-c/4.9.0--openmpi--4.1.4--nvhpc--23.1",
        "netcdf-fortran": "netcdf-fortran/4.6.0--openmpi--4.1.4--nvhpc--23.1",
        "netcdf-parallel": "parallel-netcdf/1.12.3--openmpi--4.1.4--nvhpc--23.1",
        "zlib": "zlib/1.2.13--gcc--11.3.0"
    },
    "compilation": {
        "FCFLAGS": "FCFLAGS=-g -gpu=pinned -cuda -fast -Minline -acc",
        "FFLAGS": "FFLAGS=-g",
        "LDFLAGS": "LDFLAGS= -gpu=pinned -cuda -acc",
        "enableGPU": "--with-gpu=pinned",
        "enableACC": "--with-acc=yes",
        "precision": "--with-r4=no",
        "binary": "/leonardo_work/EUHPC_D02_008/TESTS/test_audit/shared/binaries/Fall3d.gpu.231106.r8.init_time.x"
    }
}
```


Automatic reporting
(from the json file)

M11 scalability tests

Binary id: test

testType	test_scalability
testOwner	Heribert Pascual
note	Copy and modify this template in each new test_scalability_ID folder
comment	Test to write in the log the initialisation time

Binary id: code

origin	decube83/fall-3-d-gpu-cicd.git
version	8.2.1-gpu
branch	/cicd/hpascual
commitID	30349dce61c8ba3bd3f20ea79041c543a1471f76
commitAuthor	decube83 <hpascual@geo3bcn.csic.es>
commitDate	Mon Nov 06 09:15:29 2023 +0100

Binary id: system

host	Leonardo@CINECA
partition	booster
nodeCPU	32 cores Intel Ice Lake
nodeGPU	4 NVIDIA Ampere 64GB
nodeRAM	512 GB

Binary id: modules

fortran	nvhpc/23.1
cuda	cuda/11.8
mpi	openmpi/4.1.4--nvhpc--23.1-cuda-11.8
hdf5	hdf5/1.12.2--openmpi--4.1.4--nvhpc--23.1
netcdf-c	netcdf-c/4.9.0--openmpi--4.1.4--nvhpc--23.1
netcdf-fortran	netcdf-fortran/4.6.0--openmpi--4.1.4--nvhpc--23.1
netcdf-parallel	parallel-netcdf/1.12.3--openmpi--4.1.4--nvhpc--23.1
zlib	zlib/1.2.13--gcc--11.3.0

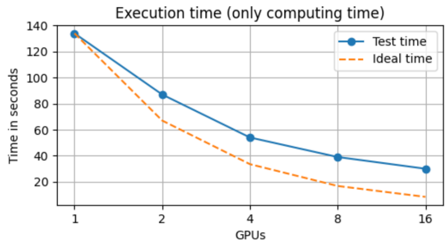

FALL3D: documentation

https://gitlab.com/fall3d-suite/

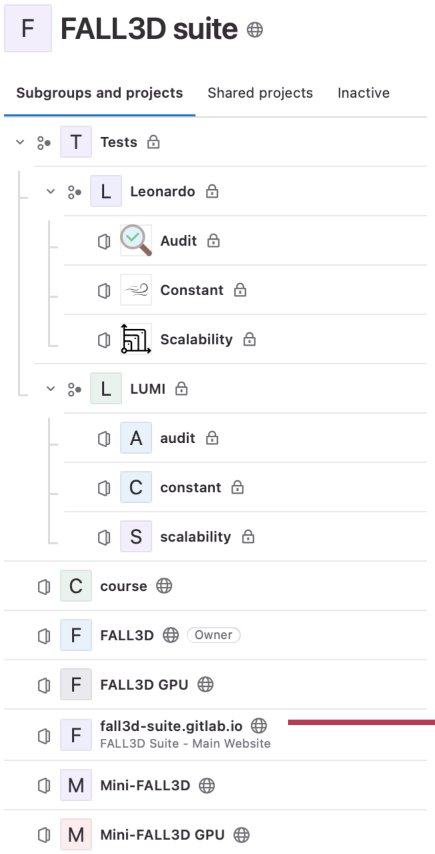

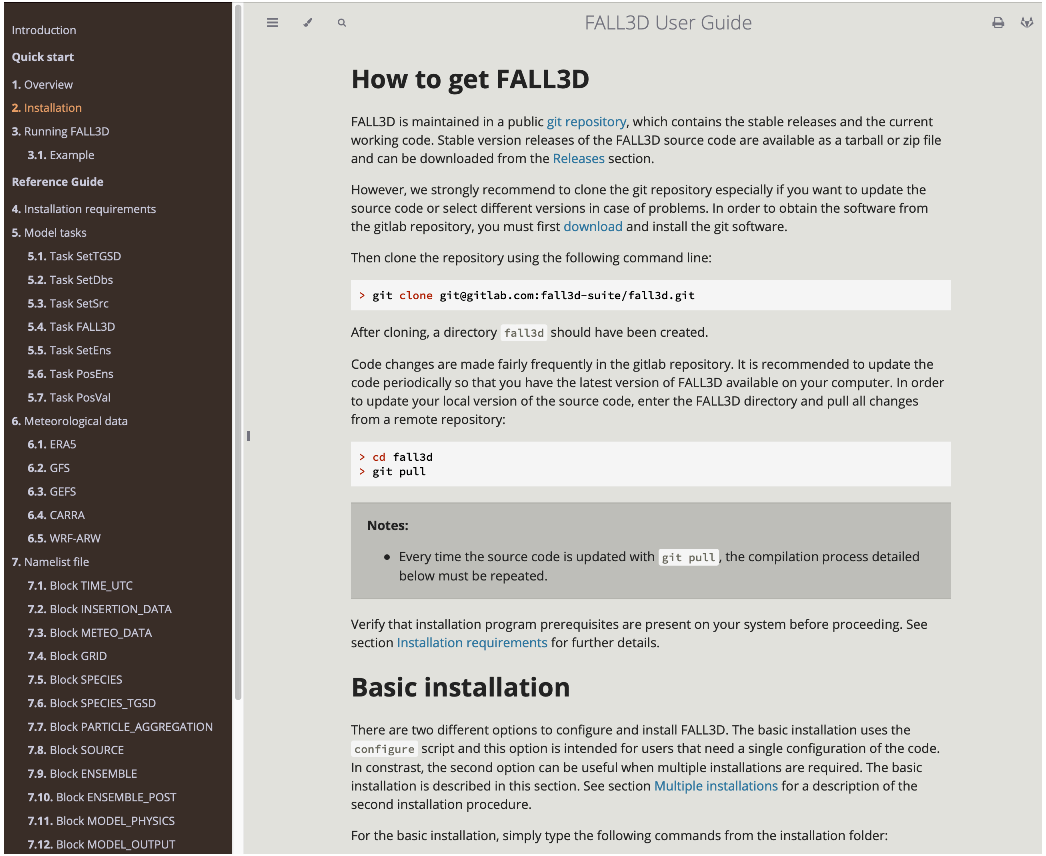

FALL3D User Guide

How to get FALL3D

FALL3D is maintained in a public git repository, which contains the stable releases and the current working code. Stable version releases of the FALL3D source code are available as a tarball or zip file and can be downloaded from the Releases section.

However, we strongly recommend to clone the git repository especially if you want to update the source code or select different versions in case of problems. In order to obtain the software from the gitlab repository, you must first download and install the git software.

Then clone the repository using the following command line:

```
> git clone git@gitlab.com:fall3d-suite/fall3d.git
```

After cloning, a directory `fall3d` should have been created.

Code changes are made fairly frequently in the gitlab repository. It is recommended to update the code periodically so that you have the latest version of FALL3D available on your computer. In order to update your local version of the source code, enter the FALL3D directory and pull all changes from a remote repository:

```
> cd fall3d
> git pull
```

Notes:

- Every time the source code is updated with `git pull`, the compilation process detailed below must be repeated.

Verify that installation program prerequisites are present on your system before proceeding. See section Installation requirements for further details.

Basic installation

There are two different options to configure and install FALL3D. The basic installation uses the `configure` script and this option is intended for users that need a single configuration of the code. In constrast, the second option can be useful when multiple installations are required. The basic installation is described in this section. See section Multiple installations for a description of the second installation procedure.

For the basic installation, simply type the following commands from the installation folder:

FALL3D: real use case (La Palma eruption)

- The eruption lasted for nearly 3 months (from 19 September to 13 December 2021).
- About 3,000 buildings were destroyed by lava flows and 8,000 people were evacuated (red zone).
- Occurrence of punctuated airport disruptions by ash fallout (imply runway cleaning, re-routing, etc).
- The crisis was successfully managed by the emergency committee (PEVOLCA); no fatalities.

FALL3D: real use case (La Palma eruption)

- A **daily operational forecast ChEESE service** delivered to the scientific committee of the PEVOLCA.
- Informed decision-makers about the next 48h in terms of civil aviation impacts and likelihood of low air quality scenarios (confinement of population even beyond the red zone).
- FALL3D simulations ran @MN4 on 2 different computational domains: archipelago (at 1 km grid resolution) and regional (at 5 km grid resolution).
- Showed the benefits of UC, informing authorities about expected scenarios and anticipating decision-making.

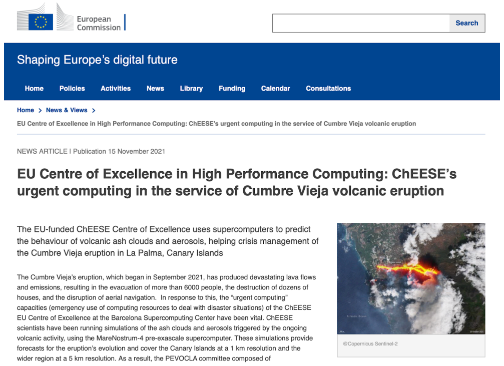

European Commission

Search

Shaping Europe's digital future

Home | Policies | Activities | News | Library | Funding | Calendar | Consultations

Home > News & Views >

EU Centre of Excellence in High Performance Computing: ChEESE's urgent computing in the service of Cumbre Vieja volcanic eruption

NEWS ARTICLE | Publication 15 November 2021

EU Centre of Excellence in High Performance Computing: ChEESE's urgent computing in the service of Cumbre Vieja volcanic eruption

The EU-funded ChEESE Centre of Excellence uses supercomputers to predict the behaviour of volcanic ash clouds and aerosols, helping crisis management of the Cumbre Vieja eruption in La Palma, Canary Islands

The Cumbre Vieja's eruption, which began in September 2021, has produced devastating lava flows and emissions, resulting in the evacuation of more than 6000 people, the destruction of dozens of houses, and the disruption of aerial navigation. In response to this, the "urgent computing" capacities (emergency use of computing resources to deal with disaster situations) of the ChEESE EU Centre of Excellence at the Barcelona Supercomputing Center have been vital. ChEESE scientists have been running simulations of the ash clouds and aerosols triggered by the ongoing volcanic activity, using the MareNostrum-4 pre-exascale supercomputer. These simulations provide forecasts for the eruption's evolution and cover the Canary Islands at a 1 km resolution and the wider region at a 5 km resolution. As a result, the PEVOCLA committee composed of

@Copernicus Sentinel-2

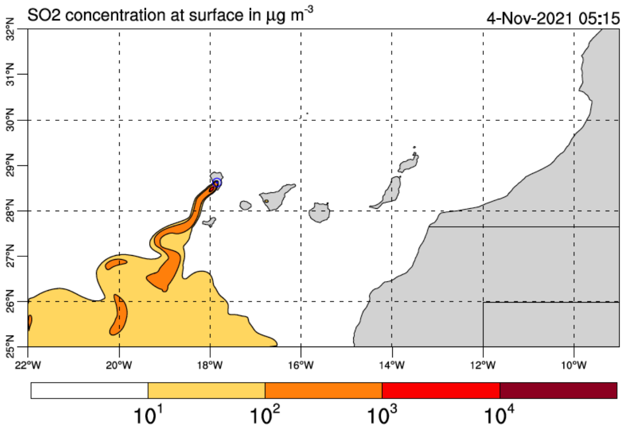

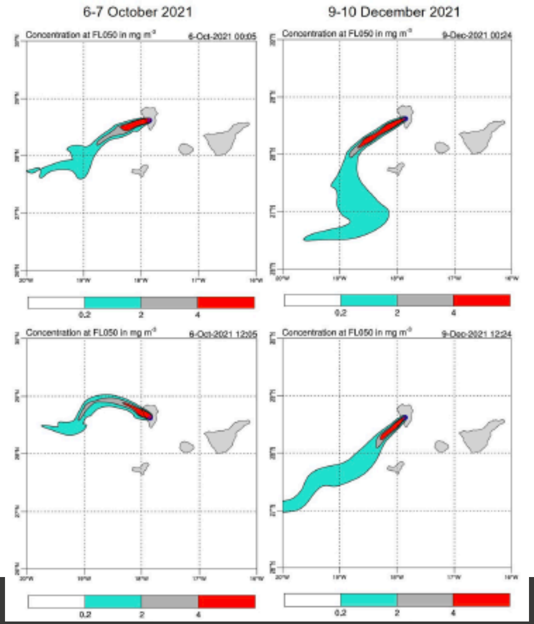

VOLCICE 2021

- On 12 March 2021, a VOLCICE exercise was scheduled to practice the response to an explosive eruption at Beerenberg volcano (Jan Mayen, Norway).
- The exercise is part of the VOLCICE series played by the Icelandic Meteorological Office (IMO) in collaboration with London VAAC and ISAVIA (the air navigation service provider in Iceland). The Jan Mayen exercise engaged also Met Norway.
- In order to exploit the resources developed within the ChEESE, IMO asked to test an eventual service in their operational response.

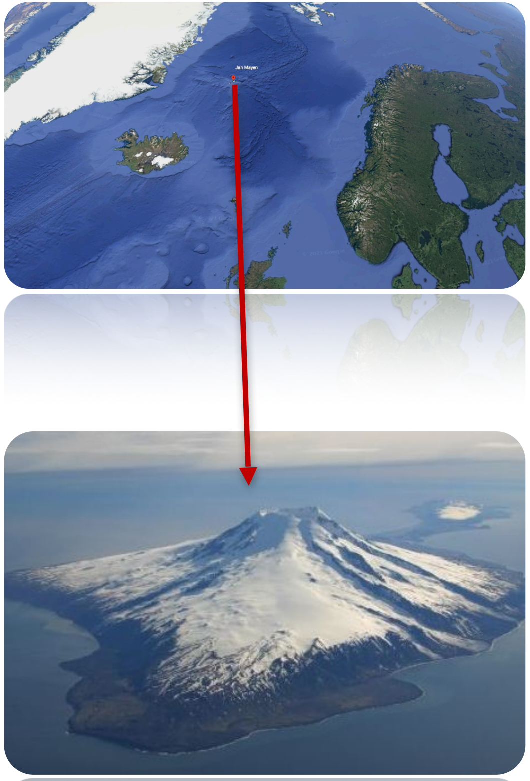

Comparison of results: VAAC (reprojected) vs PD12

Deterministic

Probabilistic

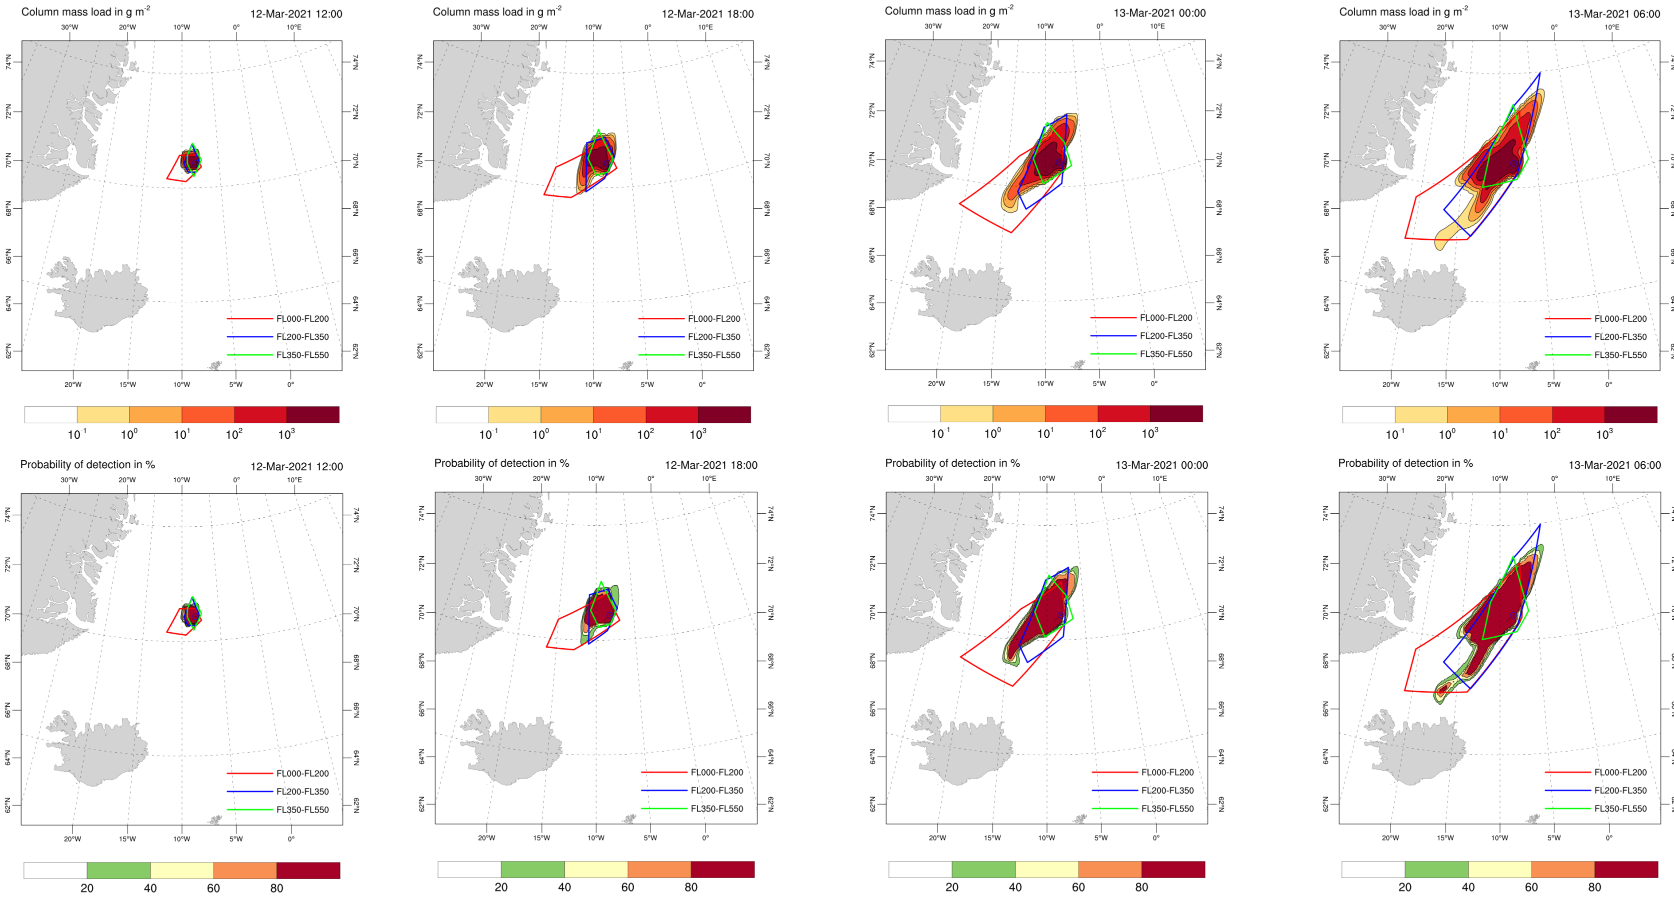

Towards the European tephra hazard model

ChEESE-2P will run the first homogeneous multi-volcano Probabilistic Volcano Hazard Assessment (PVHA) at European scale:

- A PVHA for airborne fine ash at the European scale and grid resolution of 10 km.
- A PVHA for tephra fallout at regional/national scales (nests) and resolution of 2 km.

The final datasets will be accessible through the Volcano Observations Thematic Core Service (TCS) in EPOS for their usage and distribution.

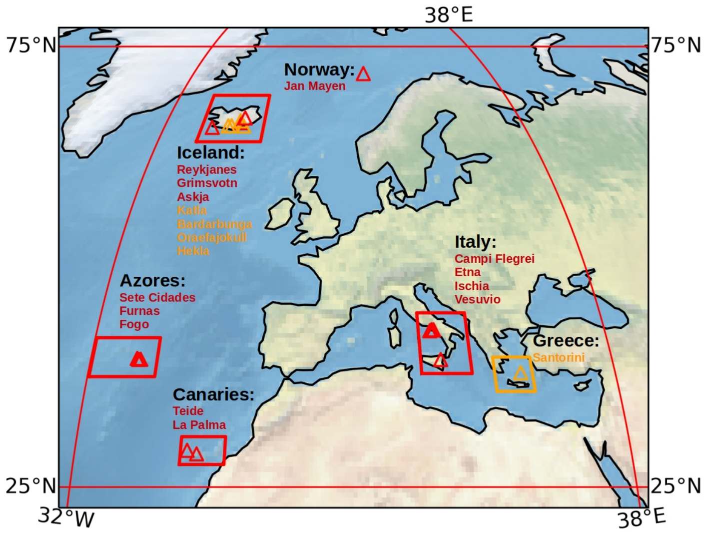

FALL3D: on-going ChEESE work and next steps

Performance portability on EuroHPC systems

- Adapt to AMD GPUs (LUMI).
- EuroHPC proposals to explore portability options (OpenACC, OpenMP off-loading)
  - EHPC-DEV-2023D07-008 (Leonardo, 3500 node/h).
  - EHPC-DEV-2023D10-028 (LUMI-G, 10000 node/h).
- Performance portability campaign

Audit-driven optimisations (CINECA)

- Reduce high data movement overhead relative to their computational workload.
- Reduce or aggregate communications to optimise GPU utilisation.
- Inspect the list of kernels flagged as suspicious provided by POP team.

1

Introduction

Arnau Folch
Geociencias Barcelona (CSIC), Spain

2

FALL3D

Arnau Folch
Geociencias Barcelona (CSIC), Spain

3

OpenPDAC

Mattia De' Michieli Vitturi
Istituto Nazionale Geofisica e Vulcanologia (INGV-Pisa), Italy

Project funded by EuroHPC under the grant agreement No 101093038.

OpenPDAC: code overview (1)

- OpenPDAC is a **Eulerian multiphase code** for simulating volcanic **gas-particle flows**, available via a public Github/Gitlab repositories.
- OpenPDAC is developed as an **OpenFOAM module** and is based on multiphaseEuler, distributed with OpenFOAM.
- Compared to the original module, in OpenPDAC the **kinetic theory equations for granular flows** are modified to model more than one dispersed solid phase, and by introducing models for particle-particle drag.
- Furthermore, a Lagrangian library (one-way coupling with the gas-solid mixture) is used to model large solid particles. OpenPDAC also implements the initialization of the hydrostatic pressure profile, necessary for simulations on large domains.
- OpenPDAC is well suited to simulate **phreatic explosions**, as those recently occurred at Ontake (Japan) and White Island (new Zealand).

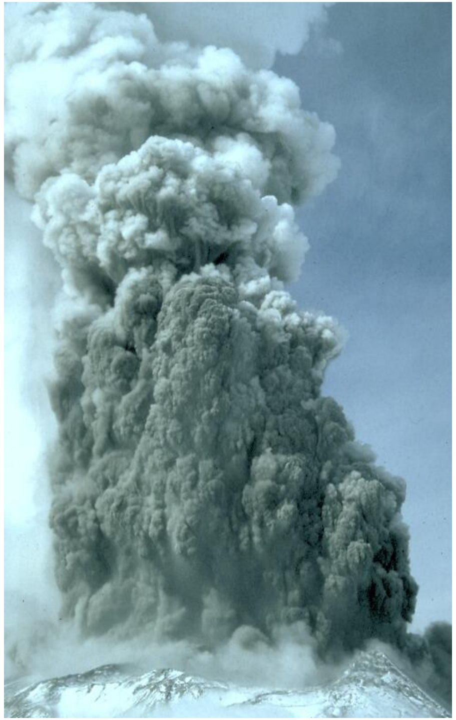

Phreatic eruption at the summit of Mount St. Helens, Washington. Credit: D.A. Swanson, USGS(Public domain.)

OpenPDAC: code overview (2)

- Code written in C++, making intensive use of **object orientated features**, as inheritance, template classes, virtual functions and operator overloading.
- In OpenFOAM, the classes are designed to define, discretize and solve PDE's through a finite-volume discretization scheme on unstructured 3D grid.
- The method of **parallel computing** used by OpenFOAM is known as **domain decomposition**, in which the geometry and associated fields are broken into pieces and allocated to separate processors for solution.
- The parallel running uses the public domain **openMPI implementation** of the standard message passing interface (MPI) by default, although other libraries can be used.

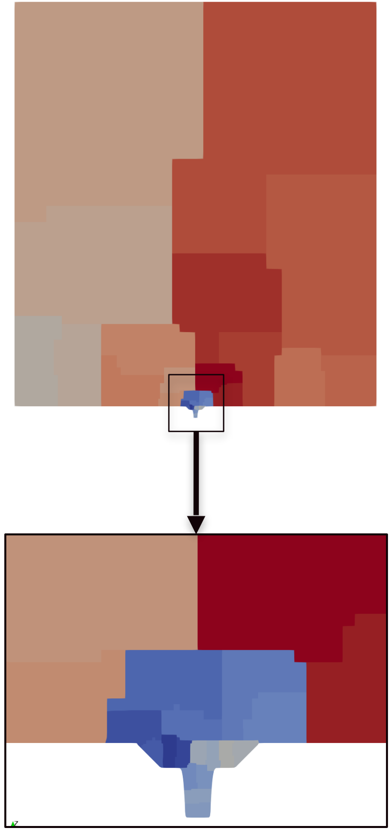

Example of OpenFOAM Scotch decomposition which requires no geometric input from the user and attempts to minimise the number of processor boundaries

OpenPDAC: ChEESE-1P legacy

A different OpenFOAM solver (ASHEE) was one of the flagship codes of ChEESE-1p, leading to a series of optimizations.

- Efficiency optimization based on mixed-precision
- Asynchronous I/O
- Ongoing experiments on GPU porting (exaFOAM partnership)

Up to:

$130*10^6$ cells

12,000 cores

HLRS
High Performance Computing Center | Stuttgart

CINECA

INGV

ChEESE

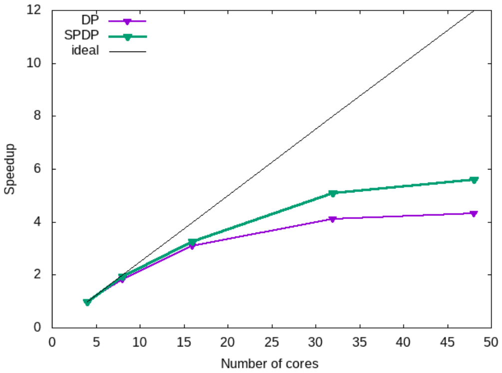

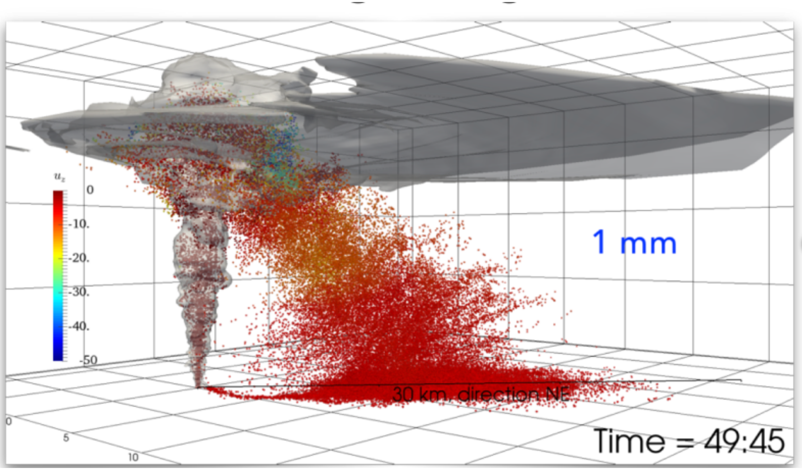

OpenPDAC: code repositories

https://github.com/demichie/OpenPDAC-11

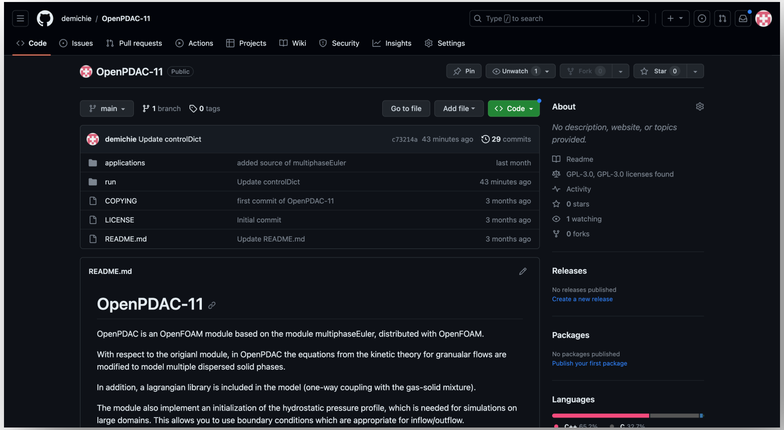

EOSC-SYNERGY SOFTWARE COMPLIANT

sqaaas software | silver

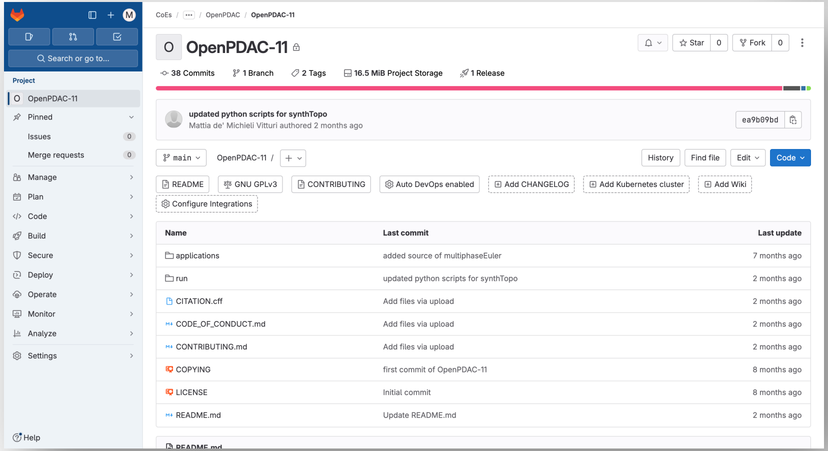

https://codehub.hlrs.de/coes/cheese-2p/o
OpenPDAC-11

OpenPDAC: CI/CD

CI/CD Github/Gitlab Actions (temporary)

Action	Test	Build	Publish
Push Master	**Run tests**	**Compile the code and Build a Docker container**	**Publish the Docker container**

- At present OpenPDAC is installed on LUMI (CPU partition, AMD EPYC@64 core), and soon on Leonardo (CPU partition, Intel Sapphire Rapids). We plan to implement CI/CD actions to update the installations of the code when new major releases are pushed to the repository.
- Mare nostrum

OpenPDAC: performance @ LUMI

Scalability

Mesh: ~**16M** cells
Max nodes: **64 (8192 CPU cores)**
Efficiency: **~ 0.7 with 2k cells/core** (256k cells/node)
Computational Speed: 28M ~cells/s

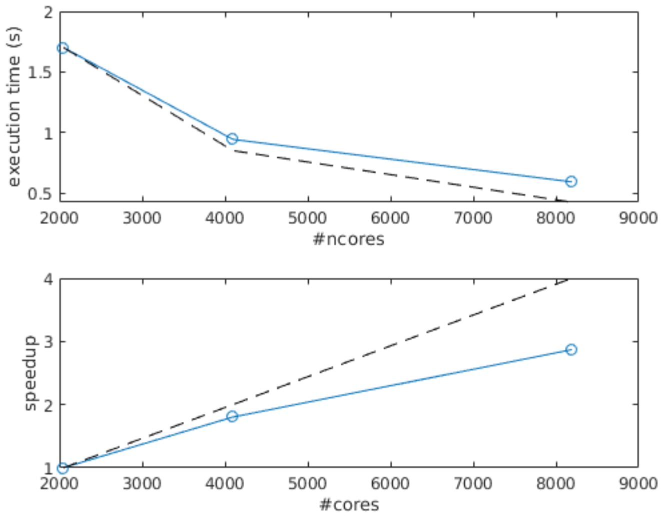

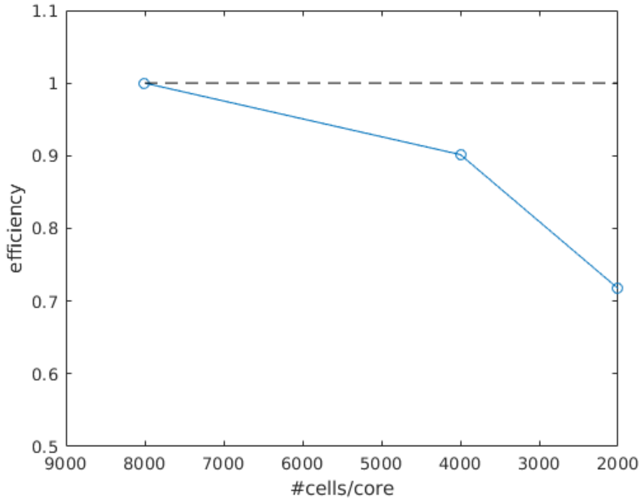

OpenPDAC: performance @ LUMI

Application profiling

- Profiling has been done with a small number of cores (512). The trace for 512 MPI ranks is larger than 30GB.
- POP metrics highlight that the parallel efficiency is low and need to be improved

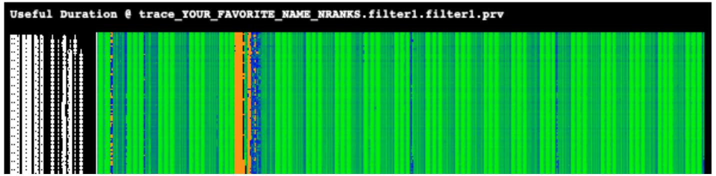

	MPI Processes
Global Efficiency	**36%**
Parallel Efficiency	**36%**
Load balance	**74%**
Communication effifciency	**49%**
Serialization efficiency	**69%**
Transfer efficiency	**71%**
Computational Scalability	**100%**

Parallel efficiency is only 36% which means that the code spends only roughly a third of its time doing actual computations.

OpenPDAC: Simulation Cases

SC6.1

Urgent high-resolution, 3D multiphase flow simulation of phreatic eruptions at Vulcano

(INGV, CIN; **Capability**)

TARGET: perform a single, 3D phreatic eruption scenario for the whole island of Vulcano with a prescribed resolution at ground (less than 1 m) within X hours (to be decided with the stakeholders).

SC6.2

Long-term probabilistic hazard maps for phreatic eruptions at Vulcano island

(INGV, CIN; **Capacity**)

TARGET: Perform a large ensemble of 3D numerical simulations to build a **Probabilistic Hazard Map** of hazardous actions (Pyroclastic Currents and Ballistics) with variable

- Vent location and geometry
- Temperature conditions
- Pressure conditions

OpenPDAC: workflow optimization

- Work in collaboration with HLRS High-Performance Computing Center, Stuttgart

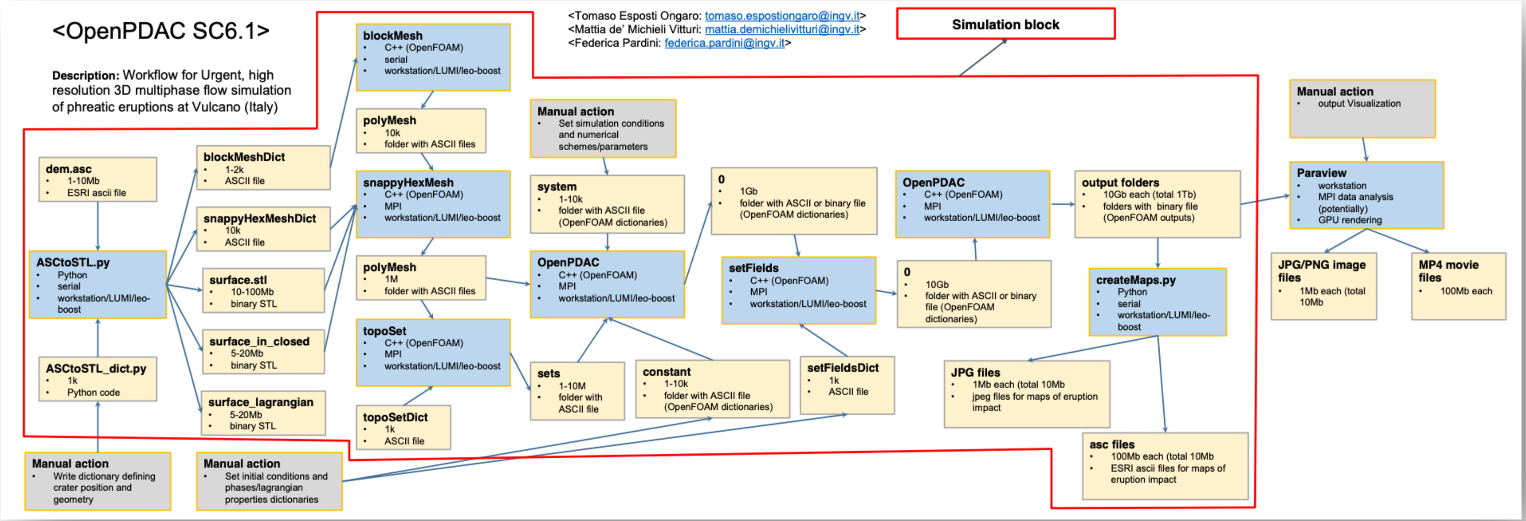

High Performance Computing Center | Stuttgart

Small scale scenario at Spiaggia di Levante

Vent Geometry (R/D)	5/50
Volume DRE [m3]	750
α_p	0.3
n_g	0.026
P [MPa]	5.0
T [°C]	200
Specific Energy [kJ/kg]	14.0
Rmax [km]	0.50
PDC Runout [km]	no

INPUT

OUTPUT

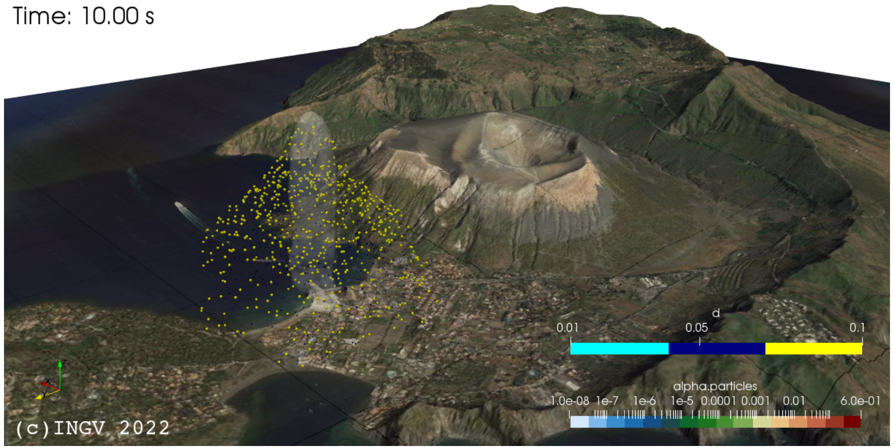

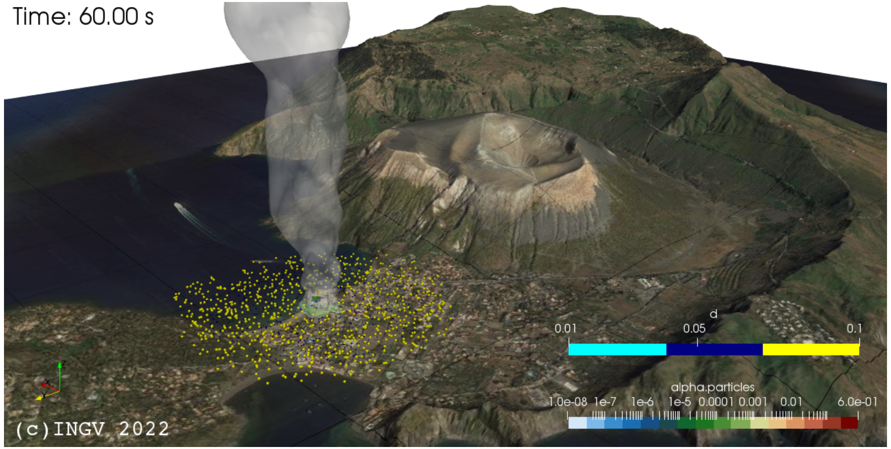

Small scale scenario at Spiaggia di Levante

Vent Geometry (R/D)	8/13
Volume DRE [m3]	1660
α_p	0.6
n_g	0.008
P [MPa]	5.0
T [°C]	200
Specific Energy [kJ/kg]	5.4
Rmax [km]	0.30
PDC Runout [km]	0.70

INPUT

OUTPUT

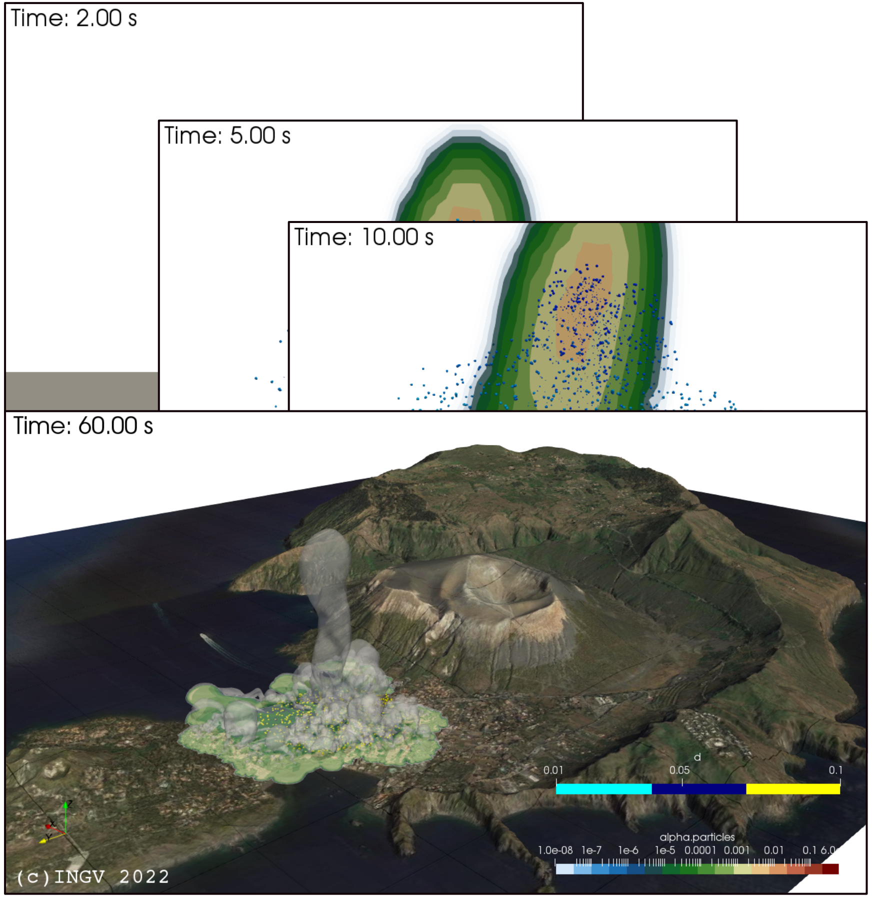

Small scale scenario at Spiaggia di Levante

Thank you!

@cheese-coe

http://cheese2.eu

@cheese-coe

@cheese_coe@techhub.social