D5.4 Service prototype and enabling

Version 1.4

Document Information

Document date	31/01/2022
Contract Number	823844
Project Website	https://cheese-coe.eu/
Contractual Deadline	31/01/2022 (M38)
Dissemination Level	CONFIDENTIAL
Nature	Report
Coordination and editing	Tomaso Esposti Ongaro (tomaso.espostiongaro@ingv.it)
Authors	Tomaso Esposti Ongaro, Marta Pienkowska (PD1), Jorge Macias (PD2), Alice Gabriel, Bo Li and Sara Wirp (PD5), Laura Sandri, Beatriz Martinez and Antonio Costa (PD6), Stefano Lorito (PD8), Arnau Folch (PD12)
Reviewers	Arnau Folch (afolch@bsc.es) Joan Farnos (joan.farnos@bsc.es) Marta Pienkowska (marta.pienkowska@erdw.ethz.ch)

The ChEESE project has received funding from the European Union's Horizon 2020 research and innovation programme under the Grant Agreement No 823844

Change Log

Version	Description of Change
v1.0	Initial draft for internal working
v1.1	Reviewed by the Project Coordinator
v1.2	Reviewed by Pilot leaders
v1.3	Reviewed by WP5 coordinator and sent out to internal reviewers
v1.4	Final review by WP5 coordinator, passed to project coordinator and project manager

Index

1. Introduction

One of the main objectives of WP5 is to coordinate the effort to **make Pilot Demonstrators** (PDs, developed in WP4), **available as Services** exploitable by a broader user community. Potential users in Task 5.4 are represented by the ChEESE Industrial and User Board (IUB), but a wider outreach is expected. As a measure of readiness, we use the Technology Readiness Level (TRL) as defined in Deliverable D5.1. 8 of the 12 PDs in ChEESE targeted at TRL higher than 5 and are considered as potential service candidates.

To achieve this goal, we have involved (in collaboration with WP6) the geophysical community (in coordination and synergy with other pan-European initiatives) and the non-academic stakeholders (industrial partners, observatories, civil protection authorities belonging to the IUB), in the definition of the Services and their validation. In particular, the end-users have cooperated in defining the service objectives, the gaps to fill, the relevant scientific/societal questions addressed, the technological bottlenecks.

This Deliverable D5.4 reports about ChEESE Task 5.4, in which the main scientific and technological outcomes of the Pilot Demonstrators and their related services are deployed in **Relevant (TRL 6) or Operational (TRL > 6) Environments.** The document is organised as follows**:**

- Section 2 reports about the different **operational service** prototypes enabled.
- Section 3 summarises the progress towards the **involvement of the Stakeholders**.
- Section 4 presents a **summary** of the main achievements in Task 5.4.
- Section 5 presents in more detail the **individual PD**’s results.

List of acronyms

TRL: Technology Readiness Level
EPOS: European Plate Observing System
ERCC: Emergency Response Coordination Centre
DG-ECHO: DG for European Civil Protection and Humanitarian Aid Operations
NEAMTWS: North-eastern Atlantic and Mediterranean Tsunamic Warning System
VAAC: Volcanic Ash Advisory Centre
ICAO: International Civil Aviation Organisation
SIAM: Sistema Italiano Allerta Maremoti (Italian Tsunami Alerting System)
UNDRR: United Nation Office for Disaster Risk Reduction
UC: Urgent Computing
P(V/S/T)HA: Probabilistic (Volcanic/Seismic/Tsunami) Hazard Assessment
PTF: Probabilistic Tsunami Forecast
T(E)WS: Tsunami (Early) Warning System

2. Operational service prototypes

In the framework of ChEESE, **operational services** are defined as those services giving access to HPC models, workflows and/or products **for making decisions**. The **operational environment** identifies the users, the access mode, the digital landscape (data, models and workflows), and the information flow involved in the decision-making process.

While Task T5.3 was devoted to scientific/technological challenges, Task T5.4 demonstrates that HPC applications and workflows developed in ChEESE can be implemented and **enabled in operational environments** to improve decision-making related to natural (solid-Earth) catastrophes. To this aim, the end-users have been directly and actively involved in this process.

2.1. Pilots focusing on urgent computing and workflows for early warning

Six pilots have demonstrated potential service exploitation during **"live" exercises**, in which target use-cases have been numerically addressed and service functional requirements (e.g., required time-to-solution, spatial/temporal resolution, minimum ensemble size), tested in an operational environment. Computational resources have been pre-allocated on HPC centres on the expected time-slots. Table 1 summarises the main objectives of the 6 live exercises performed in ChEESE.

PD	run date	Service delivered	Stakeholders	Time-to-solution	HPC resources
PD1	18/01/2022	Synthetic hake maps of ground motion proxies (Peak Ground Acceleration/Velocity, and others) provided in an Urgent Computing mode (1.5Hz live and 5Hz offline).	Global Parametrics, ARISTOTLE	30 minutes	10 nodes (480 CPU cores) on BSC's Mare Nostrum 4
PD2	22/11/2021	Tsunami impact map based on multi-scenario simulations (variable seismic sources) for the Spanish Tsunami Early Warning System	Instituto Géografico Nacional - Civil Protection (Spain)	1 minute	72 nodes (288 NVIDIA V100 GPUs) on CINECA's Marconi 100
PD5	05/11/2021	3D fully coupled (physics-based) high-resolution (1.5-5 Hz) earthquake and tsunami dynamics	As for PD1 and PD8	< 6 hours	16 nodes (4096 CPU cores) on SuperMuc2
PD6	04/11/2021	Short-term probabilistic assessment of volcanic tephra hazard/impact over a large-scale domain.	PLINIVS (Italian Civil Protection competence centre) ARISTOTLE (ERCC)	3 hours	24 nodes (1152 CPU cores) on BSC's Mare Nostrum 4
PD8	05/11/2021 (World Tsunami Awareness Day)	Rapid probabilistic Post-Event Tsunami Assessment following the 2020, October the 30th, Mw 7.0 Samos-Izmir earthquake.	NEAMTWS Tsunami Service Providers	45 minutes	800 nodes (3200 NVIDIA V100 GPUs) on CINECA's Marconi100

PD12	10/12/2021	Early probabilistic ash dispersal forecast (“Quantitative Volcanic Ash” service).	Buenos Aires VAAC	25 minutes	48 nodes (2304 CPU cores) on BSC’s Mare Nostrum 4

Table 1. List of live exercises carried out in Task 5.4.

All live exercises have followed a similar structure, represented in Figure 1.

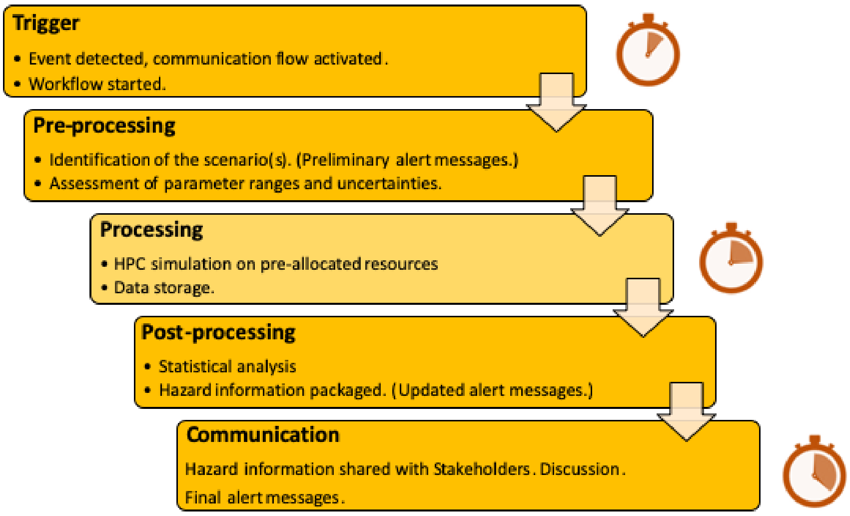

Figure 1. General flow-chart of T5.4 “live” exercises (see Section 5 for details on individual PD exercises).

Owing to the enduring COVID-19 pandemic, all “live” exercises have been carried out remotely using video-conference and mail.

2.2. Pilots focussing on probabilistic hazard assessment

Pilots focusing on workflows for probabilistic hazard assessment, based on a large ensemble of HPC simulations but without strict time constraints have been mostly described in Deliverable D5.3.

PD	Product delivered	Stakeholders	HPC resources
PD7	Probabilistic Tsunami Hazard Model	INGV-CAT Tsunami Warning Centre , SiAM	up to 256 nodes (1024 NVIDIA V100 GPUs) on CINECA’s Marconi100

Table 2. Operational services for Probabilistic Hazard Assessment not included in this report as “live” exercise.

2.3. Advancement of Technology Readiness Levels of ChEESE Services

According to the definition of the ChEESE TRL scale proposed in Deliverable D5.1, we report here the advancement of TRL throughout the project development. Note that some of the initial TRL values originally stated in the DoA have been revised in Deliverable D5.1.

- All the PDs foreseen in Task T5.4 have reached the target TRL=6, i.e. component integration and interoperability and **individual components are demonstrated in a relevant or in an operational environment**, with service functional requirements set by end-users. In some cases, part of the workflow (typically, the post-processing blocks) has been run "offline". This includes exercises with limited interaction with the stakeholders. Two additional PDs have challenged a "live" exercise to test their workflow in an operational environment: PD1 demonstrated the capability of Urgent Seismic Computing, although its TRL remains lower (around 6); PD5 (complemented by PD4) tested model capability in parallel with PD1 exercise (but not in "live" mode) and demonstrated the potential for integrating multi-physics-based simulation of earthquake-induced tsunami in rapid tsunami assessments (PD8 exercise); its TRL also is around 6 (demonstration in relevant environment).
- PD2 and PD12 have reached the target TRL=8, i.e. **integrated workflows are demonstrated in an operational environment**.
- Finally, two Pilot Demonstrators (PD2 and PD12) have already been used in *assisted operational mode* (TRL=9), i.e., **the operational service is deployed and operation-ready for different scenarios**. In particular, the PD2 flagship code Tsunami-HySEA is running faster-than-real-time tsunami simulations for ARISTOTLE service to ERCC, and PD12 full workflow, based on the FALL3D flagship code, has been run daily in operational mode on 1536 cores at MareNostrum-4 to advise regional airport management and air traffic during La Palma eruption (October-December 2021). These achievements will be documented in Deliverable D5.5 ("Access to deployed Pilot services").

Pilot	**PD1**	**PD2**	**PD5**	**PD6**	**PD7**	**PD8**	**PD12**
Initial TRL	3	3	3	3	3	3	3
Target TRL	5-6	6-7	6-7	7	5-7	6-8	6
Achieved TRL	**6**	**8-9**	**6**	**7**	**7**	**7-8**	**8-9**

Table 3. Summary of Pilot Demonstrators' TRL achieved with Task 5.4.

3. Role of the Stakeholders

ChEESE services are based on specific workflows (see WP3) customised to solve the use-cases and optimised to run with the ChEESE Flagship codes (see WP2) **SALVUS** (PD1), **Tsunami-HySEA** (PD2, PD7, PD8), **SeisSol** (PD5), and **FALL3D** (PD6, PD12). Some of the IUB members have taken an active role in the deployment of live exercises and in the testing of PD workflows in an operational environment (Table 4).

Stakeholders' institution	**Mission**	**Role played in Task 5.4 exercises**
ARISTOTLE	ARISTOTLE is a long-term operational, research, and cooperation project financed by the Directorate-General for European Civil Protection and Humanitarian Aid Operations (DG-ECHO).	PD6. Simulates reporting to ERCC. PD1. Evaluates rapid post-EQ assessment by using numerical simulations

	ARISTOTLE-eENHSP aims to continue strengthening the monitoring and analysis functions of the Emergency Response Coordination Centre (ERCC) by delivering a unique multi-hazard advice service at global level and on a 24/7 operational basis.	PD8. Potential end-user
PLINIVS	PLINIVS (Study Centre for Hydrogeological, Volcanic and Seismic Engineering) is a National Competence Centre on Volcanic Risk for the Italian Civil Protection.	PD6. Receives hazard information and elaborates risk/impact maps and information on its basis.
IGN	IGN (Spain) has the mandate to plan and manage systems for detecting seismic movements countrywide and their possible effects on the coasts, and to notify them to the institutions.	PD2. Provides trigger and source information for simulated EQ and associated fault characteristics. Elaborates results and issues Alert messages.
Global Parametrics	Global Parametrics' mission is to leverage advances in climate science, data modelling and financial engineering to build the tools needed to understand, manage and mitigate the risks of extreme weather and natural hazards anywhere in the world.	PD1. Interested in exploring numerical simulations for parametric risk assessment and index-based payouts.
Buenos Aires VAAC	The Buenos Aires VAAC is one of the 9 Volcanic Ash Advisory Centers established by the ICAO with the mission of monitoring and forecasting the location and trajectories of volcanic clouds occurring under their respective areas of responsibility. In the event of an eruption, the international civil aviation arrangements state that the affected VAAC must issue periodic Volcanic Ash Advisories (VAA), consisting of text messages including the forecasted ash polygons delineating unsafe flight areas.	PD12 exercise owner. Triggers the simulated events, issues the VAA messages. Coordinates the exercise.
CAT Tsunami Warning Centre	Centro Allerta Tsunami (CAT) of INGV is a Tsunami Service Provider (TSP) of the ICG / NEAMTWS (Intergovernmental Coordination Group for the Tsunami Early Warning and Mitigation System in the North-eastern Atlantic, the Mediterranean and connected seas), which is an integral part of the global warning and risk mitigation system tsunami, established and coordinated by UNESCO's Intergovernmental Oceanographic Commission (IOC).	PD8 exercise coordinator. Triggers the simulated events, gathers information from the NEAMTWS Tsunami Service Providers, elaborates data and issues Tsunami warnings.

Table 4. Involvement of the IUB members and stakeholders in D5.4.

4. Summary of the main achievements

HPC targets

- First operational Tsunami Probabilistic Forecast workflow executed on 800 nodes (3200 NVIDIA V100 GPUs).
- First Tsunami Early Warning workflow with uncertainty estimation executed on 72 nodes (288 GPUs) on CINECA's Marconi100 (Italy).
- First Volcanic Short-Term Probabilistic Hazard Assessment and Probabilistic Forecast workflows run on 24 to 48 nodes on BSC's Mare Nostrum 4.

Service-oriented workflow developments

- Development of the UCIS4EQ User Portal, a front-end solution for the Seismic Urgent Computing Service.

Operational service prototypes

- Prototype probabilistic km-resolution forecast products have been delivered to Volcanic Ash Advisory Centres for Volcanic Ash dispersal, with uncertainty estimated on hundreds of ensemble members, in less than 1 hour.

- Operational faster-than-real-time tsunami warning (PD2) has increased its capability by hundreds of times with respect to current state-of-the-art, including an early assessment of tsunami coastal hazard and innovative uncertainty evaluation.

- Operational volcanological services (PD6 and PD12) have been scaled by a factor $>10^3$ (in terms of number of cells, time steps, and ensemble numerosity/sample size) with respect to current state-of-the-art.

- Operational Probabilistic Tsunami Forecast (PD8) is a game changer: it allows increased accuracy (up to 10 m) with respect to precalculated scenarios requiring linearity; 2) it demonstrates statistical compatibility with data; 3) it allows for a deeper uncertainty exploration (larger ensemble), or for faster calculations. Down-scaling based on different sub-sampling strategies can be considered for operational Early-Warning.

- Two services (PD2 and PD12) based on ChEESE Flagship codes Tsunami HySEA and Fall3D have been deployed as fully Operational Services (TRL 9). Their codes, workflows and results are routinely used by the stakeholders to support decision-making processes.

5. PD sheets

The next sections report, for each PD, about the development and testing of the prototypal operational service and the results of the exercises.

PD1. Urgent Seismic Simulations

PD1	**Urgent Seismic Simulations**
Leader	Marta Pienkowska (ETH)
Participants	Juan Esteban Rodriguez (BSC), Marisol Monterrubio Velasco (BSC), Josep de la Puente (BSC), Otilio Rojas (BSC), Andreas Fichtner (ETH), Alice Agnes Gabriel (LMU)
Workflow	UCIS4EQ
Numerical Engine	Salvus
TRL initial	3
TRL target	5-6
TRL achieved	6

Description of the live exercise

- **Summary and objectives**

On 18 January 2022, ChEESE organised a live demonstration of the workflow developed in PD1: the Urgent Computing Integrated Services for EarthQuakes (UCIS4QE) integrated with Salvus. To this end, we have developed the front-end solution, the UCIS4EQ User Portal, that provides an interface to monitor the status of the service, examine the final results, and launch user-defined events. In preparation for the exercise we have progressed from TRL 5, where individual components of the prototype service had been validated in relevant environments, to TRL 6, with all components integrated and tested in a relevant environment prior to the live exercise.

The objective of the exercise was to introduce the stakeholders to the tool and its fully automatic execution process, detailing the steps of a sample complete end-to-end run, and monitoring those steps in real time. The feasibility of fully automatic urgent seismic simulations with UCIS4EQ has been demonstrated to end-users for the first time in a relevant environment.

- **Operational Challenges**

- Computational resources: in order to monitor the workflow in real-time, we have limited the live demonstration to include only a single low-frequency Salvus simulation instead of a usual suite of simulations generated by UCIS4EQ. Such a single full end-to-end run completed in approximately 30 minutes (including pre-processing and post-processing), allowing us to explore the status of the system live.
- Emulating an Urgent Computing environment: a dedicated reservation on a HPC cluster is required to reproduce the fact that jobs are not meant to be queuing in a regular queue. 10 nodes have been reserved on MareNostrum4 for this live exercise.

- User-modified event: in addition to the pre-defined and previously tested alert for the purpose of the exercise (mimicking the automatically generated alerts), we have launched a simulation with a set of parameters modified on the fly via the User Portal. We have shown the capacity of the User Portal to monitor multiple jobs, as well as the capability of UCIS4EQ to manage non-urgent ("offline") user-generated jobs.
- Data transfers: we have for the first time used automatic data transfers to and from B2Drop in real-time: first to upload the simulation results from the HPC Cluster to B2Drop, then to download them to view in the User Portal.

- **Scientific Challenges**

The main scientific challenge of our service is to prove that urgent seismic simulations can bring added value to post-event evaluations of the hazard. This was not the objective of the live exercise, as we cannot expect a single low-frequency simulation to reproduce the right order of magnitude of ground motion proxies. We have, however, presented results from previous executions that demonstrate that suites of high-frequency simulations from UCIS4EQ with Salvus can reproduce the right order of magnitude for a range of observed proxies, such as PGA, PGV, Arias Intensity or significant duration.

- **Workflow split into Tasks, Milestones, KPI**

Alert Service: The alert service passively listens to information coming in from FDSN databases in real time and decides which (if any) information needs to be calculated. It triggers the workflow should the event be classified as an urgent computing candidate.

EventDomains: The domain around the earthquake is defined for the simulation

ComputeResources: The computational resources are estimated.

Source parameters: The computation of the source parameters is subdivided into a few steps. First, the workflow rapidly computes a set of CMT parameters in the region given information databases from past earthquakes in that same area (**CMTCalculation**). Then, it estimates the fault size and asperities size (**SourceType**) and populates the fault with a kinematic description of the slip via the Graves Pitarka rupture generator (**SlipGenGP**).

InputParametersBuilder: All information is collected to generate a single YAML file with input parameters.

SalvusPrepare: The inputs generated by the workflow are processed for a simulation with Salvus, locating sources and receivers on the specified mesh, attaching absorbing boundaries, and defining output types, among others.

SalvusRun: The HPC wave propagation simulations are performed on the cluster.

SalvusPost: The post-processing module collates results, generates netCDF files with ground motion proxies, and then automatically generates plots on the cluster. The plots are then compressed and immediately uploaded to B2Drop at the BSC and can be viewed in the User Portal.

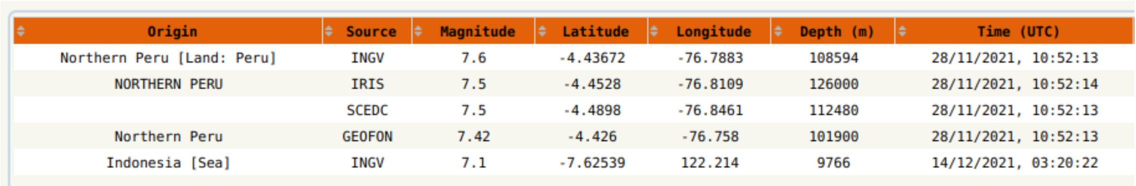

Origin	Source	Magnitude	Latitude	Longitude	Depth (m)	Time (UTC)
Northern Peru [Land: Peru]	INGV	7.6	-4.43672	-76.7883	108594	28/11/2021, 10:52:13
NORTHERN PERU	IRIS	7.5	-4.4528	-76.8109	126000	28/11/2021, 10:52:14
	SCEDC	7.5	-4.4898	-76.8461	112480	28/11/2021, 10:52:13
Northern Peru	GEOFON	7.42	-4.426	-76.758	101900	28/11/2021, 10:52:13
Indonesia [Sea]	INGV	7.1	-7.62539	122.214	9766	14/12/2021, 03:20:22

Figure 1.1. A snapshot of the Alert Service from the UCIS4EQ User Portal. The same earthquake is usually reported by multiple agencies and the workflow needs to unify that information prior to further execution.

Service	Status	InitTime	EndTime
EventDomains	SUCCESS	2022-01-17 09:44:55	2022-01-17 09:44:55
CMTInputs	SUCCESS	2022-01-17 09:44:55	2022-01-17 09:44:55
ComputeResources	SUCCESS	2022-01-17 09:44:55	2022-01-17 09:44:55
CMTCalculation	SUCCESS	2022-01-17 09:44:55	2022-01-17 09:45:06
SourceType	SUCCESS	2022-01-17 09:45:06	2022-01-17 09:45:06
SlipGenGP	SUCCESS	2022-01-17 09:45:06	2022-01-17 09:46:44
InputParametersBuilder	SUCCESS	2022-01-17 09:46:44	2022-01-17 09:46:44
SalvusPrepare	SUCCESS	2022-01-17 09:46:44	2022-01-17 09:49:02
SalvusRun	RUNNING	2022-01-17 09:49:02	

Figure 1.2. A snapshot from the UCIS4EQ User Portal monitoring the status of each task of the workflow after an alert has been generated via the Alert Service. On the snapshot, the SalvusRun is nearly finished and the workflow will move to the final stage of post-processing, SalvusPost.

Identification of the Stakeholders and their roles

Stakeholders' institution	Role played
ARISTOTLE	IUB member.
Global Parametrics (GP)	IUB member.
Swiss Seismological Service (SED)	Invited observer.
Institut Cartogràfic i Geològic de Catlalunya (ICGC)	Invited observer.
VESTEC	Invited observer.
eFlows4HPC	Invited observer.
DT-GEO	Invited observer.

Table 1.1. Stakeholders' involvement

- **Communication and information flow**

Direct contact has been established with the participants via e-mail prior to the event. The event was held virtually and the live demo has been performed by UCIS4EQ developers with the participants as passive observers - at TRL 6, the field specialists remain the target end user and the prototype is not ready for deployment to end users directly. The following timetable has been communicated:

Schedule (time in CET)	
14:30 – 15:00	**Presentation:** An introduction to the Urgent Computing Integrated Services for Earthquakes (UCIS4EQ) with Salvus.
15:00 – 15:30	**Live demo:** emulating a trigger of the Samos-Izmir October 30th 2020 event.
15:10 – 15:25	**Presentation:** SeisSol and ExaHyPE in PD1.
15:30 – 16:00	**Live demo:** a user-triggered event.
16:00 –	Q&A, discussion and feedback.

Table 1.2. Schedule of the event.

A short video summarising both the front-end and the back-end of UCIS4EQ is currently in preparation and is going to be distributed to the participants and shared with a wider audience via social media.

- **Ways to collect feedback from the Stakeholders**

The private event format was designed to encourage a discussion and direct feedback during the event, with dedicated time allocated for that purpose. The participants shared their comments and asked questions, and direct e-mail communication followed with some.

Target Operational Environment

- **Description of the operational context (main actors, procedures, etc)**

ChEESE-PD1 organised and coordinated the exercise, emulating the alert trigger manually for the Samos-Izmir October 30th 2020 earthquake. The simulations have been assigned to a specific reservation on the MareNostrum4 supercomputer and executed immediately.

For a real event, the alert would be automatically generated by the Alert Service that analyses information coming from FDSN and filters for events that fit the criteria for an urgent simulation. However, there is currently no ecosystem for urgent computing requests in HPC centres and thus no standardised way for an urgent allocation of resources.

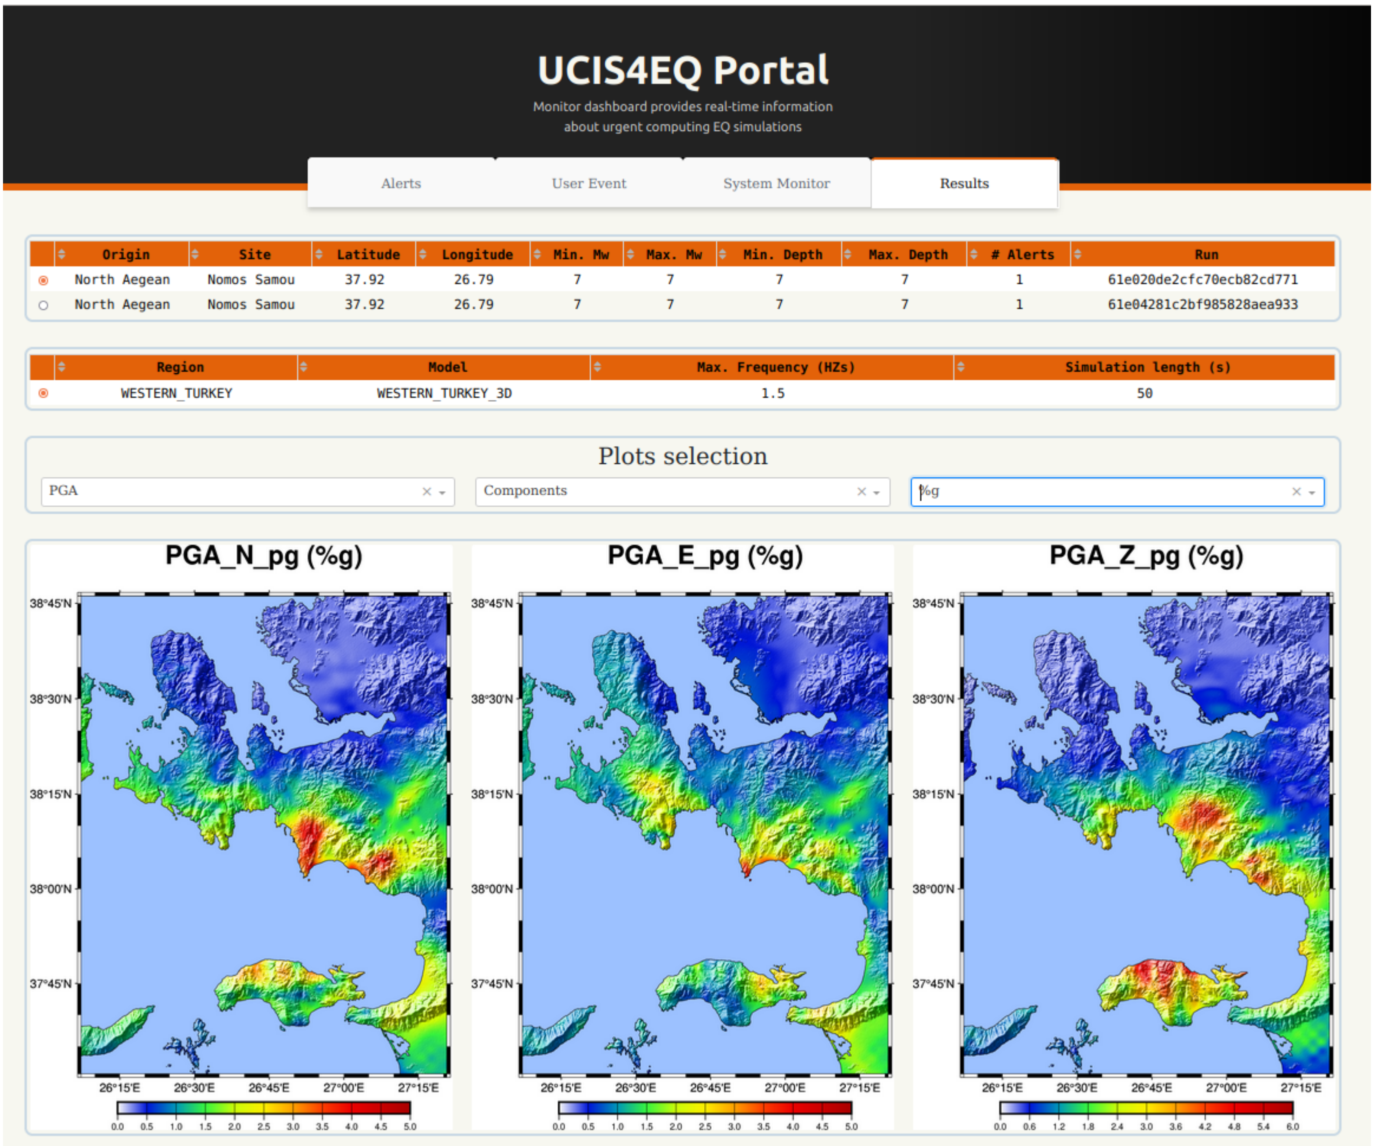

Figure 1.3. A snapshot from the UCIS4EQ User Portal where synthetic maps of ground motion proxies can be explored.

- **Expected added value with respect to existing procedures**

Wave propagation simulations complement data in a unique way by providing full time-histories that can be densely sampled in space. Such synthetic shake-maps from 3-D physics based simulations can incorporate a wide range of effects, such as those of the topography or of the bathymetry and the ocean, in order to best reproduce the ground shaking. Moreover, the numerical approach is sensitive to uncertainties in different ways than the current data-driven GMPE-based approaches. We therefore believe that the densely sampled synthetic shake maps (see Figure 1.3) are complementary to the existing shake maps and can provide additional valuable insights into post-event hazard assessment.

HPC Challenge

- **Computational resources allocated**

  A reservation of 10 nodes (a total of 480 CPUs) on MareNostrum4 supercomputer in order to emulate an Urgent Computing environment.

- **HPC Challenges**

  The exercise has been designed to showcase the prototype of the UCIS4EQ system for the first time. We aimed to monitor progress in real time - the simulation, therefore, was small and the exercise did not pose a particular HPC challenge. PD1 addressed relevant HPC challenges in the scope of the exascale testbed demonstrator of D3.6.

	Node specs	Total number of allocated nodes per exercise	Node per run	Number of runs	Duration of a single run	Data storage needs
Mare Nostrum 4 (BSC, Spain)	2x Intel Xeon Platinum 8160 24C (2.1 GHz)	10	10	2	20 minutes	700 MB

Table 1.3. HPC resources.

Exercise Run (only for live exercises)

- **Exercise report (with time indications)**

Scheduled time	Actual time	
14:30 – 15:00	14:35 - 15:05	**Presentation:** An introduction to the Urgent Computing Integrated Services for Earthquakes (UCIS4EQ) with Salvus.
15:00 – 15:30	15:05 - 15:22	**Live demo:** Introduce the UCIS4EQ User Portal, introduce the back-end, trigger the alert and explore the alert in the User Portal and the HPC job statistics on the BSC User Portal; introduce the B2Drop data repository.
	15:37 - 15:43	**Live demo:** Follow the final stages of the HPC run, then explore the generation of the user-defined event.
15:10 – 15:25	15:22 - 15:37	**Presentation:** SeisSol and ExaHyPE in PD1.
15:30 – 16:00	15:43 - 15:57	**Live demo:** Trigger another simulation via the user-triggered event; while it is running explore results of

		the first simulation, then discuss future developments needed for the service
16:00 –	15:57 - 16:25	Q&A, discussion and feedback. Explore briefly the results of the 2nd run in the User Portal.

Table 1.4. Timing of the exercise with respect to the schedule.

- **Deviations from original plan**

  The original date of the exercise had to be rescheduled for a later date. On the day of the exercise, however, there were no significant deviations from the original plan. The workflow has been thoroughly tested prior to the live demonstration, including the timing of the HPC jobs.

- **Bottlenecks, obstacles, unexpected difficulties**

  As above, there were no unexpected difficulties and both workflow executions completed as expected on the reserved resources.

- **Lessons learned and recommendations**

  The main take-away messages from the discussion with the stakeholders are as follows:

  - Further work is necessary to better understand how uncertainties in the source parameters translate to uncertainties in the final shake maps, and more tests need to be carried out to ensure the reliability of the results.
  - Further work is also necessary to improve the estimation of an ensemble of source parameters.

Potential target service

- **Key missing steps**

  Technological improvements: Integrate multiple HPC clusters and multiple data centres. Integrate a Workflow Manager.

  Parametrization: Improve the estimation of the source parameters and the simulated ensembles. Integrate multiple regions with appropriate velocity models.

  Data assimilation: Update the state of the workflow given updated incoming information, including assimilating data to calibrate final shake maps.

- **Computational resources needed if operational**

  Remains to be defined. Required computational resources are frequency dependent, while the frequency resolution is problem dependent. The current requirements for 5 Hz simulations that appear to reproduce the right order of magnitude of ground motion proxies are on the order of 90 GPUs for ~1.5h of wallclock time per simulation, and tens of simulations are required. This time should be further reduced on upcoming machines and with some improvements in the software. Moreover, some use cases may require much lower frequency resolution (for example for index-based insurance case study proposed by GP, the IUB member).

- **Estimated costs of the operational service**

  The cost of the urgent access to HPC resources. The events that qualify as urgent computing candidates are rare, so resources will not be needed often. However, the access will be required immediately.

- **Required policies and agreements to develop an operational service**

  Urgent Computing policies for HPC clusters that guarantee access to resources following a potentially destructive earthquake.

PD2. FASTER THAN REAL TIME TSUNAMI SIMULATIONS

PD2	FASTER THAN REAL TIME TSUNAMI SIMULATIONS
Leader	Jorge Macías (UMA)
Participants	Carlos Sánchez Linares (UMA) / Beatriz Gaite (IGN)
Workflow	Faster Than Real Time Tsunami Simulations
Numerical Engine	Tsunami-HySEA
TRL initial	2
TRL target	6-7
TRL achieved	8-9

Description of the live exercise

- **Summary**
  On 22 November 2021, within the framework of the ChEESE project (https://cheese-coe.eu/) a Live Demo of the Pilot Demonstrator 2, FTRT (Faster Than Real Time) Tsunami Simulations, with an application to Early Warning was organised. The problem addressed in this Live Demo was proposed by IGN, the Spanish Tsunami Warning Center and co-designed with the UMA team. It consisted in the design of the set-up of the computational component of the TEWS for the Spanish TWS for the Atlantic. The Live Demo consisted of a blind exercise, where the IGN provided the set of parameters for the 135 scenarios to be simulated. Indeed, the aim was to simulate the supposed actual scenario taking place but including some variability associated to the uncertainty in source definition as we describe in our presentation. Finally, alert levels were provided within a short time frame.

- **Objective**
  The main objective of this exercise was to provide a quick response for tsunami action, incorporating not only deterministic simulations of a single event but also including variability in the parameters that define the source of the tsunami, showing that the use of supercomputing resources allows the realisation of hundreds of tsunami simulation scenarios in a few minutes, allowing the analysis of the results to be incorporated into the Tsunami Early Warning System. More concretely, in this Live Demo, 135 simultaneous simulations were performed on four different computational domains in the North-Eastern Atlantic Ocean for tsunamis generated in the Gulf of Cadiz and surrounding area and with impact in the Western Iberian Peninsula and the Canary Islands in order to produce alert level along all the coast.

- **Operational Challenges**
The main operational challenge lies in having access to enough and reserved computing resources (GPU nodes) for these types of calculations. To this purpose, at least 34 nodes (4 GPUs each) are necessary to perform 135 simulations simultaneously and to be able to give the earliest alert message throughout the coast in less than 1 minute. During the exercise, CINECA made 68 nodes available with dedicated access to 272 NVIDIA V100 GPUs, which allowed dividing the overall computing time by two.

- **Scientific Challenges**
To address the uncertainty in the seismic tsunami source a set of equally probable scenarios was simulated (instead of a single deterministic scenario), taking into account the unknowns and uncertainties from earthquake source parameters. Thus, the first challenge lies in being able to give a list of simulation scenarios just after the earthquake occurrence, and perform all the simulations. The second challenge is then to make a quick analysis of the outputs of these simulations in order to give an alert message that reflects the uncertainty on source conditions.

- **Workflow split into Tasks, Milestones, KPI**

- *Occurrence and detection of an earthquake.* The Spanish National Geographic Institute (IGN) detects the occurrence of a hypothetical earthquake located in the Gulf of Cádiz and activates its tsunami warning protocol. The first step in this protocol is the characterization of the earthquake. In a few seconds, the IGN provided this information: the reference scenario was a magnitude Mw 7.8 earthquake with the epicentre located at 9.772º W, 35.442º N and at 50 km depth. Its fault parameters were selected from the closest mapped fault from the IGN fault database. In this particular case, the strike was 39° and the dip was 35°. These (and other) parameters define the reference scenario.

- *Including uncertainty in the parameters - generation of a list of equally-probable scenarios.* The range of parameter variation was selected from typical uncertainties in earthquake magnitude (± 0.2). Strike and dip errors were chosen greater than common values for the Atlantic active faults compilation in QAFI v.3 database, as ± 15º and ±20º, respectively. Also, the reference scenario generated four scenarios by shifting the location of the hypocenter half the width and half the length of the rupture within the fault-associated plane. Considering these quantities, a list of 135 (5×3×3×3) scenarios was generated as the input list of events to simulate.

- *Sending the list of events to be simulated to UMA.* The complete list with the labelled scenarios was sent to the UMA team in charge of carrying out the simulations. These previous steps were carried out automatically from the IGN system when detecting the earthquake

- *Massive launching of the simulations.* The UMA team was in charge of receiving and processing the data. Given the location of the earthquake, the computation domains and simulation time to be taken into account were previously considered in coordination with the IGN. Once connected to the cluster, the file provided by IGN with the list of scenarios was uploaded in the folder dedicated to this purpose. Finally, a script that automates this entire process is run, simultaneously launching the simulation of the 135 events in each of the considered domains.

- *Automatically post-processing of the results.* As each batch of 135 simulations concludes, the execution script sends a message to the user, in this case UMA, indicating that the job has finished, and that the post-processing of the results is carried out. A few seconds later, the summary of the analysis of the results appeared on the screen, and tables and maps were generated for their sending back to the IGN, who will include them in the Tsunami Warning System. All of the above processes ended in less than 7 minutes from the reception of the list of scenarios. The earliest alert message was provided in less than 1 minute and it has been updated after the conclusion of each batch of simulations.

Milestones
This was, to our knowledge, the **very first time that an on-the-fly numerical simulation of a large number of scenarios in just a few minutes was performed worldwide**. Besides, all this has been performed in the framework of a real TEWS (Tsunami Early Warning System) environment and the final result has been providing alerts levels in coastal segments and automatic message generation. Currently, no TEWS in the world uses computational resources as the ones that have been used here to produce numerical simulations in such a short time.

KPI (Key Performance Indicators)
Number of institutions presented during the Live Demo: 44
Number of attendees: 99
Number of visualisations of the recorded video: 77 (17 January 2021)

Identification of the Stakeholders and their roles

Stakeholders' institution	Role played
Instituto Geográfico Nacional (IGN)	Proposer

ARISTOTLE-eENHSP	Invited observer
NEAMTWS Tsunami Service Providers (CAT-INGV, KOERI, IPMA, NOA, CEA/CENALT)	Invited observers
PMEL/NOAA (USA)	Invited observer
SHOA (Chile)	Invited observer
CIGIDEN (Chile)	Invited observer
INETER (Nicaragua)	Invited observer
Univ. of Malta (Malta)	Invited observer
SINAMOT-UNA (Costa Rica)	Invited observer
UPRM/PR Seismic Network (Puerto Rico)	Invited observer
DIMAR (Colombia)	Invited observer
UNAM (Mexico)	Invited observer
FUNVISIS (Venezuela)	Invited observer
CNG-CNRS (Lebanon)	Invited observer
NSCEP (Israel)	Invited observer
JRC (EC)	Invited observer

Table 2.1. Stakeholders' involvement. Other attending institutions were: Puertos del Estado, BSC-National Center for Supercomputing, Instituto de Hidráulica Ambiental (IHC), Instituto Geológico y Minero (IGME-CSIC)), Instituto Español de Oceanografía (IEO-CSIC) and GREA (Grupo de Emergencia de Andalucía, Junta de Andalucía) from Spain, and other research institutions and universities as GFZ (Germany), IMO (Iceland), HLRS (Germany), Bull Atos (company), IDL Lisbon (Portugal), NIEP (Romania), ESPOL (Ecuador), U. Trieste (Italy), METU (Turkey), NGI (Norway), NIGGG (Bulgaria), AWI (Germany), UCL (UK), Sofia Univ. (Bulgaria), Univ. of Naples (Italy), UiB (Norway), UTFSM (Chile), Univ. of Naples (Italy), TUM (Germany), U. of Nevada (USA), among others. The total number of individual attendees was 99.

- **Communication and information flow**. The Live Demo was publicised through various media, including social media such as Twitter, web pages (ChEESE and EDANYA-UMA). Direct contact with all stakeholders was established through individual e-mails and also e-mailing to distribution lists for a wider dissemination. The organisation of the event was also presented in several scientific events and institutional meetings, such as the PTF ChEESE exercise the 5th of November, the Aristotle project, the ChEESE WP5 meeting, the Civil Protection meeting at Madrid and the ICG/NEAMTWS-XVII meeting. In the recording of a UMA dissemination material for secondary school students, the description of the work performed in the Demo was included (Cientificando, in Spanish, at YouTube: https://youtu.be/rHVNxYrfYe0 in the last part of the video). Finally the complete recording of the Live Demo has been uploaded to EDANYA YouTube channel and it is publicly available at https://www.youtube.com/watch?v=rkruUHAaleA with 77 visualisations (the 17th January 2021).

- **Collection of feedback from stakeholders**. The present Live Demo has been co-designed in straight collaboration and communication with our partners at IGN, the Spanish institution in charge of the National Tsunami Early Warning System. Therefore, it has been in the process of designing and preparing the Live Demo that we have strongly interacted with the stakeholder proposing the specific problem to be tackled. Currently we are analysing the results of the simulations performed during the Live Demo, in collaboration with our partners at IGN, and the results will be sent for publication soon. Besides this, at the end of the Live Demo we opened a time for comments and discussion to get some feedback from all the invited observers.

Target Operational Environment

- **Description of the operational context** (main actors, procedures, etc). In case of a real earthquake with the potential of generating a tsunami, the former is detected, located and its magnitude is estimated by the Seismological Service at IGN. According to their dataset, a catalogued fault is associated with the undergoing event and the corresponding Okada parameters for the reference scenario are set. Then the variability on the parameters is automatically generated and a standardised list of 135 scenarios is generated. This first part is a task of the IGN. Then this list is sent to the UMA team and, in a few seconds, a script for the automatic computation of all the scenarios in a pre-established order is run at the supercomputer infrastructure (in our particular case at CINECA). In a fully operational environment this should be done automatically without the intervention of the human hand, in this case UMA personnel, to launch the numerical simulations. This could be done by a rabbitMQ protocol as is currently done in the framework of the Aristotle project and the SPADA system, in which seismic data are generated at INGV Rome and sent automatically to the UMA computing cluster. Finally, when the simulations are progressively finishing, the results are output and automatic alert messages are generated, ready for dissemination.

- **Expected added value with respect to existing procedures**. Existing TEWS in the NEAM region (North Eastern Atlantic and Mediterranean) mostly rely on decision matrices not including a computational component. These systems also use precomputed dataset of scenarios, which are fast to access but have as main disadvantage the occurrence of unexpected events not computed in the database. When real time on-the-fly computing solutions are used these are restricted to a single deterministic scenario, which means to compute a solution with a high level of uncertainty. Here we propose to produce real time simulations including a certain amount of variability in order to assess the uncertainty in the seismic source.

HPC Challenge

- **Computational resources allocated**. A reservation for urgent computing of 72 nodes of 4x NVIDIA Volta V100 GPUs at CINECA HPC - Marconi100 partition. The reservation started at 2021-11-22T16:00:00 and ended at 2021-11-22T18:30:00.

Architecture	Node specs	Total number of allocated nodes per exercise	Node per run	Number of runs	Expected duration of a single run	Expected data storage needs
Marconi 100 (CINECA, Italy)	4x NVIDIA Volta V100 GPUs per node, NVlink 2.0 16 GB	72	34	4 (each run launches 135 simulations)	From 1 min to 7 min, depending on the resolution of the computational domain.	~ 2.3 GB

Table 2.2. Computational resources allocated for PD2.

Exercise Run

- Exercise report (with time indications). Here we include the proposed schedule for the Live Demo that was sent to all the attendees, besides the indicated times correspond to the timing in the recording we made of the event. The edited recording of the event can be found at PD2 ChEESE Live Demo on FTRT Tsunami Simulations for Early Warning

  Introduction
  0:00 Welcome and brief description
  13:53 Description by IGN of the proposed problem
  20:48 Description of the numerical setup (UMA)
  Live Demo
  34:54 IGN sends UMA the file with the data for the seismic sources
  38:27 UMA launches the simulations
  49:12 Presentation of the results as they are generated
  58:35 Discussion, Comments and Questions

- **Deviations from original plan**. There were no deviations from the original plan, the Live Demo developed as planned. It must be noted that we tested several times the complete workflow, even using different configurations and computational resources. Everything was tested in advance, including different clusters at two different tier-0 supercomputing centers (CNS-BSC and CINECA).

- **Bottlenecks, obstacles, unexpected difficulties**. The supercomputing resources were provided by CINECA, with the reservation of 72 nodes, each with 4 GPUs so that two rounds of 135 simulations each could be carried out simultaneously. At the time of the Live Demo execution, due to a confusion with the account destined for the reservation, the launched jobs remained in the queue of the system, not going into the running state. Due to this, the people in charge of CINECA support were contacted internally, who provided the solution to this issue in a very short time and the demo could be completed successfully after relaunching the jobs. This experience reinforces the need to carry out simulations of real scenarios (drills) under the most realistic conditions possible to detect all possible sources of failure or malfunction.

- **Lessons learned and recommendations**. From a scientific point of view, taking on the challenge of carrying out hundreds of tsunami simulations in real time gives an insight into the reliability and good accuracy that mathematical models behind these simulations have. The post-processing of the results, also in real time, represents a great advance that NTWCs must incorporate in order to make better and more supported decisions. From the technical point of view, it is very relevant to send the message that when "doing things live" some issues may arise and that the issue is part of the demo. Using computing resources "live" is different from doing something off-line. When off-line you can use a trial-and-error strategy and refine the results in order to match the desiderata, in a "live" demo this is not the case.

Potential target service

- **Missing steps.** The automatic connection, in case of emergency, between the National TEW Center and the tier-0 supercomputing facility.
- **Computational resources needed if operational.** The same resources as the one used in the present Live Demo.
- **Estimated costs of the operational service.** Once the service is deployed the sole cost corresponds to the access to the HPC supercomputing resources in the very rare event of a potentially tsunamigenic submarine earthquake.
- **Required policies and agreements to develop an operational service.** those related to the access to the supercomputing resources in case of an emergency event.

	Operational at IGN	ChEESE Live Demo
Model	Tsunami-HySEA v-3.6.1	Tsunami-HySEA v-3.8.1 MC
Target area	Area of the Gulf of Cádiz (Atlantic Ocean): from -18.5W to -1W in longitude and from 27N to 45N in latitude.	
Computational domain	1 computational domain with one numerical resolution:	4 computational domains with 3 numerical resolutions

	ATL (-21,26; 1,45)	1. ATL ½ arcmin (-18.5,27;-1,45) 2. GCN ¼ arcmin (-13,33.5;-1,45) 3. GCIC ¼ arcmin (-18.5,27;-5,38) 4. GC ⅛ arcmin (-13,33.5;-4,38) (left-bottom;right-top corners coordinates in longitude,latitude)
Model resolution	½ arcmin	½ - ¼ - ⅛ arcmin
Size of the problem (number of volumes)	5,472,000	• ATL - 4,536,000 • GCN - 7,948,800 • GCIC - 8,553,600 • GC - 9,331,200
Simulated time	8 hours	4 hours
Storage frequency for time series	30 sec	15 sec
Number of simulations	1	135
Computational resources	2 nodes / 2 NVIDIA V100 GPUs	68 nodes / 272 NVIDIA V100 GPUs
Time to solution	2.5 min	1 min - ATL 3 min - GCN & GCIC 7 min - GC Total time: ≈ 7 min (parallel)
Deterministic products	Tsunami arrival time Maximum tsunami height Alert levels	Tsunami arrival time Maximum tsunami height Alert levels
Include uncertainty	No	**Yes**

Table 2.3. Comparison between the computing component in the operational Early Warning System at IGN and the tested service in PD2 Live Demo.

PD5. Physics-Based Probabilistic Seismic Hazard Assessment

PD5	Physics-based Probabilistic Seismic Hazard Assessment
Leader	Alice-Agnes Gabriel (LMU)
Participants	D5.4: Bo Li, Sara Aniko Wirp, Thomas Ulrich (LMU), Lukas Krenz, Michael Bader (TUM) PD5 participants: LMU, TUM, IMO, BSC
Codes	SeisSol, ExaHyPe
TRL initial	4
TRL target	6-7
TRL achieved	6

Description of the live exercise / benchmark

Summary and Objectives

The PD5 team was represented by **LMU Munich** during both the November 5, 2021, **ChEESE PD8 live demonstration** and the January 18, 2022, **ChEESE PD1 live demonstration** (see Figure 5.1). In the few days before the demonstrations we performed high-resolution multi-physics simulations with the ChEESE flagship code SeisSol. During the events, we summarised these efforts in short presentations of about 25 minutes and engaged with stakeholders of both events.

Our overarching objectives were three-fold:

1) Illustrating the flexibility of PD1 and PD8 frameworks to integrate with other ChEESE flagship codes and pilot demonstrators.

2) Illustrating advantages and associated requirements of more realistic, multi-physics simulations that may be implemented in future workflows.

3) Mimicking a **rapid scientific response setting**, which spans hours to days after significant geohazardous events, during which advanced geophysical and field observations become readily available (e.g., Lacassin et al., 2020)[1].

Earthquake seismology is increasingly data-rich, a development driven by e.g. machine-learning enhanced seismic catalogues, space geodesy, seismic arrays, and distributed acoustic sensing (DAS). Data-driven methods reveal striking earthquake variability but are challenged by their inherent non-uniqueness. Physics-based

[1] 2020, R. Lacassin, M. Deves, S. Hicks, J.-P. Ampuero, R. Bossu, L. Bruhat, D. Wibisono, L. Fallou, E. Fielding, A.-A. Gabriel, J. Gurney, J. Krippner, A. Lomax, M. Ma'rufin Sudibyo, A. Pamumpuni, J. Patton, H. Robinson, M. Tingay, S. Valkaniotis, "Rapid collaborative knowledge building via Twitter after significant geohazard events", Geosc. Comm. Disc., 3, 129–146, doi:10.5194/gc-2019-23.

earthquake modelling can help to overcome these crucial challenges in order to harness simulations for meaningful hazard assessment.

In D5.4 we focus on our open-source Arbitrary high-order DERivative Discontinuous Galerkin (ADER-DG) software package **SeisSol** (https://github.com/SeisSol)**.** SeisSol solves the seismic wave equations in elastic, viscoelastic, anisotropic & poroelastic and viscoplastic material on unstructured tetrahedral meshes, and by design, permits: representing complex geometries - by discretizing the volume via a tetrahedral mesh; modelling heterogeneous media; its flux based formulation is natural for representing physics defined on interfaces; high accuracy since its modal flux based formulation allows us to suppress spurious (unresolved) high frequencies and high resolution since it is very suitable for optimal performance on parallel computing environments.

For probabilistic applications in PD5 we design databases that can account for physics-based earthquake models. Heterogeneous ground shakings show clear rupture directivity, fault geometry and topography effects and modelled ground motions complement empirical ground motion models.

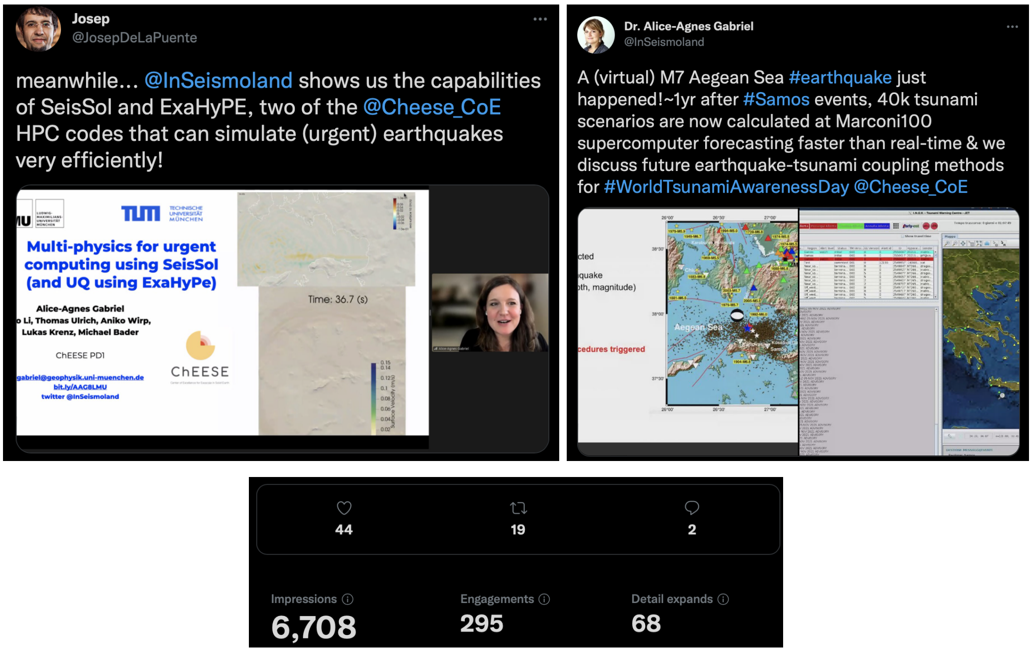

Figure 5.1: Top: social media (twitter) live coverage during the ChEESE PD1 (left) and PD8 (right) live demonstrations. Bottom: Exemplary impact for the PD8 related tweet (>6700 impressions).

PD8 specific contributions: We link PD5 with PD4, the first, physics-based tsunami-earthquake interaction pilot demonstrator led by TUM and LMU, that models in unprecedented detail the tsunami source[2]. The impact of a tsunami depends upon how

[2] L. Krenz, C. Uphoff, T. Ulrich, A.-A. Gabriel, L. Abrahams, E. Dunham, M. Bader (2021), “3D Acoustic-Elastic Coupling with Gravity: The Dynamics of the 2018 Palu, Sulawesi Earthquake and Tsunami", SC’21: Proceedings of the international conference for high performance computing,

the earthquake slip varies in both space and time. Some large undersea earthquakes result in small tsunamis. Conversely, significant tsunamis have been generated by smaller earthquakes. Ultimately, we want to base a forecast of the tsunami impact on the most accurate possible model of the earthquake rupture. The seismic waves travel faster than the tsunami waves and may give us valuable information about the 3-dimensional and time-dependent rupture process. In the ChEESE project, the difference in tsunami inundation resulting from different subduction earthquake scenarios has been explored using advanced source models which couple earthquake physics and the generation of seismic, acoustic, and water waves[3].

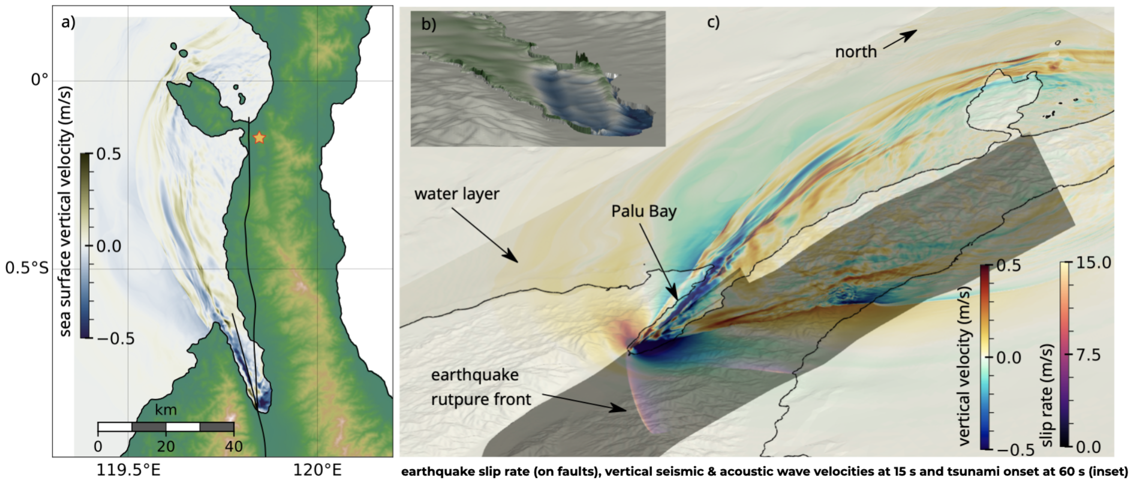

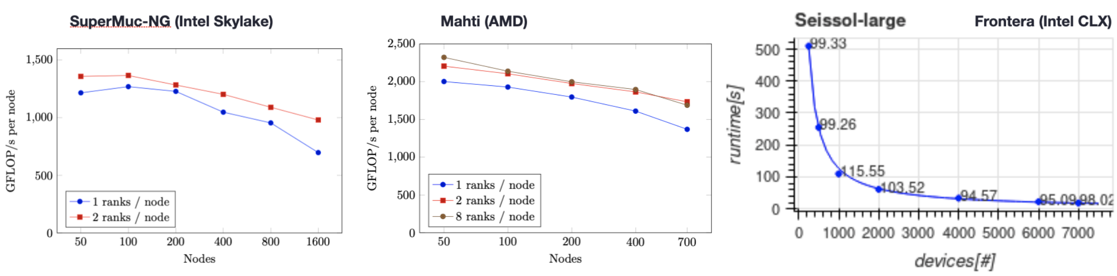

Fig. 5.2: 3D fully coupled earthquake-tsunami modelling of the 2018 Palu, Sulawesi events (Krenz et al., SC'21). SeisSol solves the elastic wave equation coupled to non-linear frictional sliding in a complex fault network & the acoustic wave equation, describing perturbations about an equilibrium hydrostatic state in a compressible, inviscid ocean of variable depth & the effects of gravitational restoring forces through a modification of the standard free surface boundary condition. We are resolving wave excitation of up to 30 Hz in the Fourier spectra of the recorded acoustic waves and 10 Hz seismics and excellent parallel efficiency, >70% on

networking, storage and analysis, Association for Computing Machinery, 63, 1–14, doi.org/10.1145/3458817.3476173.

[3] S.A. Wirp, A.-A. Gabriel, E.H. Madden, M. Schmeller, I. van Zelst, L. Krenz, Y. van Dinther, and L. Rannabauer (2020), "3D linked subduction, dynamic rupture, tsunami and inundation modelling: dynamic effects of supershear and tsunami earthquakes, hypocenter location and shallow fault slip", Frontiers in Earth Science, Geohazards and Georisks, 9, 177, doi:10.3389/feart.2021.626844.

SuperMUC-NG (Germany) and Mahti (Finland) and >95% efficiency on Frontera (preliminary results as a potential NSF leadership class computing facility application partner).

During the live demonstration we showcase that interference of seismic and acoustic waves may be dominant in data recorded by offshore instruments, which motivates fully coupled acoustic-elastic coupling with gravity solving an entirely new class of earthquake-tsunami problems. One-way linking of earthquake models to shallow water equations may omit tsunami dispersion, acoustic waves, horizontal momentum transfer and tsunami generation complexity. SeisSol allows for 3D acoustic-elastic coupling with gravity implemented via free surface tracking (gravitational effects) by linearised free surface boundary condition. This implementation compares well to one-way linking in simple setups and scales well: A multi-petaflop simulation of the 2018 Palu, Sulawesi events included 518 mio. elements = 261 billion degrees of freedom and ran for 5.5 hours on 3,072 nodes SuperMUC-NG at sustained performance of 3.1 PFLOPS. Excellent parallel efficiency was achieved on several machines >70% on SuperMUC-NG (Germany) and Mahti (Finland), >95% efficiency in preliminary results on Frontera (USA).

During the PD8 demo, we adapt a kinematic north-dipping earthquake model of the Samos earthquake consisting of 147 sub-faults provided by Ryo Okuwaki (Heidarzadeh et al., 2021) similar to kinematic inversion data products typically available within hours after large earthquakes.

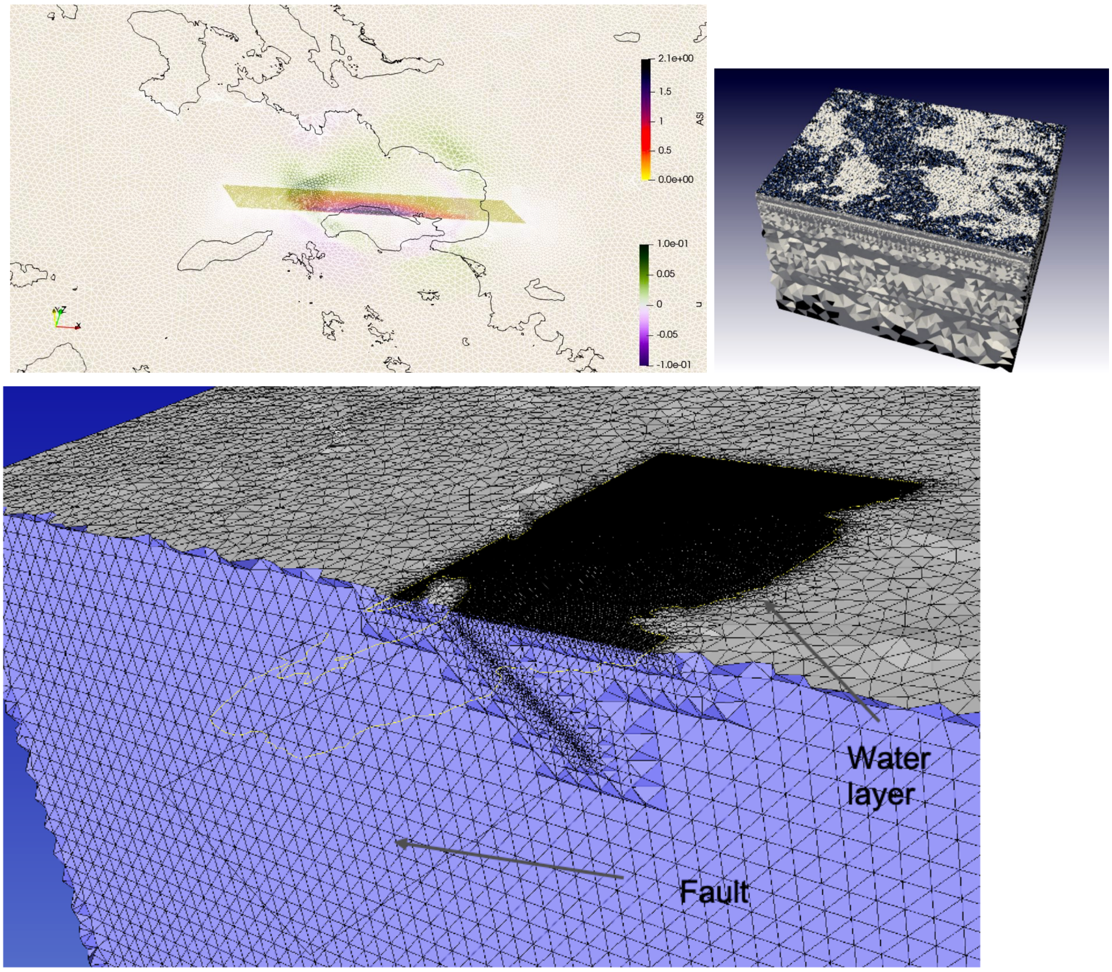

Fig. 5.3: SeisSol's rapidly constructed unstructured tetrahedral mesh for fully coupled 3D simulations of the 2020 Samos data-driven earthquake model and tsunami generation.

We are using SeisSol's multi-physics and geometrical capabilities to fully couple the earthquake model to seismic, acoustic and tsunami wave generation - modelled accurately and without aliasing.

To facilitate the rapid response mode, our 3D fully coupled earthquake-tsunami scenario accounts for a geometric bathymetry resolution of 400 m, which is rather coarse. We required **only 256 CPU hours** on SuperMUC-NG for a PD8 3D fully coupled earthquake-tsunami coupled simulation, generating synthetic seismic, acoustic and tsunami waves as well as dynamic seafloor deformation for 50 s - which is very fast. Our model revealed interesting dynamic effects such as the difference in dynamic complexity cf. considering only static seafloor uplift as tsunami source.

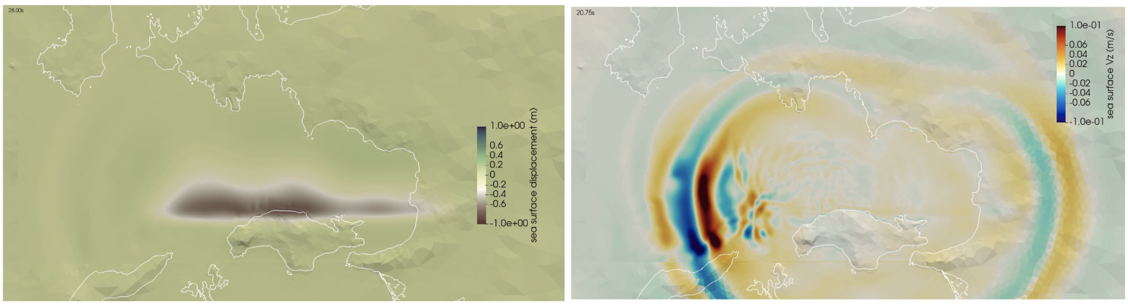

Fig. 5.4: Snapshots of seafloor displacement and seismic,acoustic and gravity wave propagation during the PD8 SeisSol Samos earthquake and tsunami generation simulation.

PD1 specific contributions: We show that the generic earthquake description developed in PD1 can be plugged in to alternative flagship codes **(serving as open-source/multi-physics or uncertainty quantification drop-in's)**; In this manner, the PD1 workflow can be adapted linking to (probabilistic) hazard and multi-physics studies.

We use SeisSol to verify crucial components of the PD1 workflow, independently of the forward solver. This includes the physical description of the earthquake source model, the simulation of seismic wave propagation and the calculation and analysis of ground shaking measures such as peak ground velocity (PGV) shake maps. We advance SeisSol's modelling workflow with a specific focus on developing **a fully automated mesh generation workflow**. This englobes developing tools to automatically translate simulation parameters from the PD1 workflow, currently based on Salvus, automatically into a SeisSol compatible format.

Generating all required simulation input files from the event .yaml file is trivial. Thus, we tackled the challenge of automatic, flexible and adaptive computational mesh generation for SeisSol's adoption in PD1.

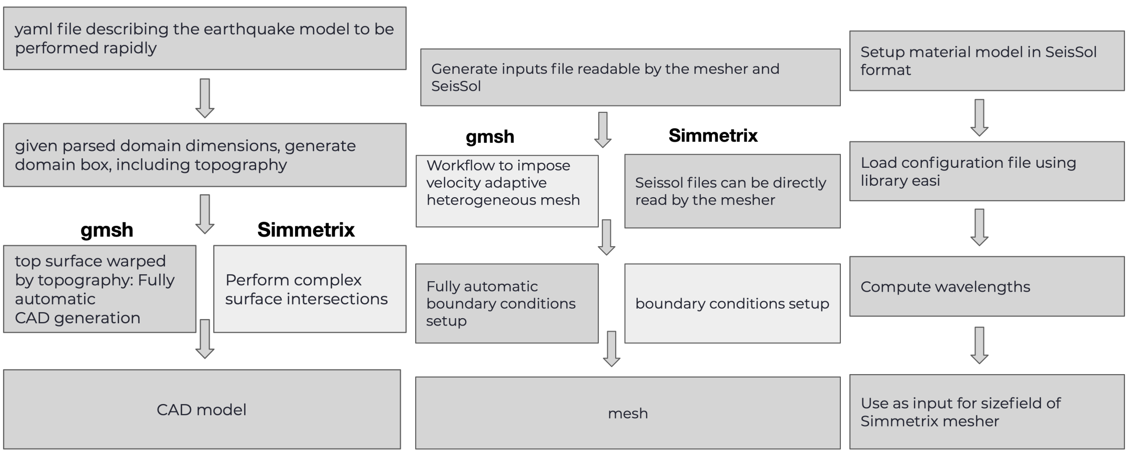

Fig. 5.5: Three SeisSol's workflows adapted for PD1 and automatic and flexible generation of complex and spatially adaptive computational meshes. Right: velocity-aware automatic and flexible computational mesh generation.

Specifically, we developed velocity-aware automatic and flexible computational mesh generation. Based on empirical analysis, we need approx. 2 elements per wavelength for accurate results using polynomial order 7 of SeisSol's basis functions. For example, assuming shear wave speed of 5 km/s and a target frequency 10 Hz required to resolve a wavelength of 500m and a target geometrical mesh resolution of 250m. Wave speeds in our models can vary a lot (e.g., for elasto-acoustic simulations the speed of sound in air is only 343 m/s!). To vary mesh size with wave speed requires an interface that reads material properties. The advantages of this new workflow include: (1) we only need to setup our computational model once for meshing and the SeisSol simulation, (2) fully automatic computation of mesh sizes, and (3) we can directly state the resolved maximum frequency.

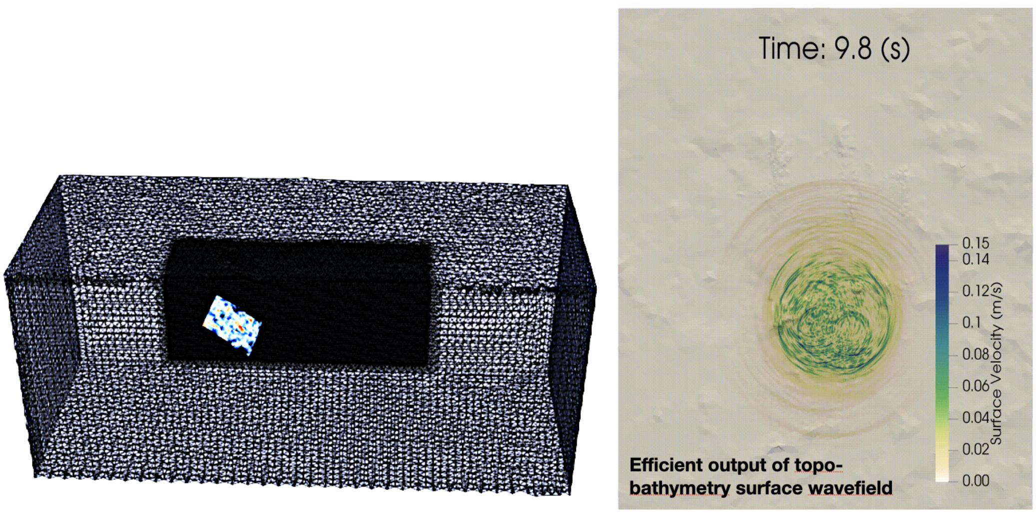

Fig. 5.6: SeiSol's PD1 Samos earthquake demonstration scenario.

For the PD1 live demonstration, we modelled one of the stochastic earthquake models provided by the PD1 workflow. The fault dimensions are 32.84 km length × 27.8 km width and we consider 1116 sub-faults (36 × 31). The SeisSol mesh we generated consists of 6.25 million elements to resolve 1.5 Hz in the seismic wavefield. We find large variability of ground shaking in the source region.

Operational Challenges

To facilitate our multi-physics simulations, it is required to have access to computing resources (CPU or GPU nodes).

We required **256 CPU hours on SuperMUC-NG** for a PD8 3D fully coupled earthquake-tsunami coupled simulation, generating synthetic seismic, acoustic and tsunami waves as well as dynamic seafloor deformation for 50 s.

We required **690 CPU hours on SuperMUC-NG** per simulation for PD1 high-resolution earthquake simulation, utilising a computational mesh consisting of 6.25 million elements and generating synthetic seismograms of 100 s in length.

Scientific Challenges

We used the fully-coupled approach to model the earthquake and the tsunami evolution simultaneously. This approach takes into account the unknowns of tsunami generation and provides new insights into the whole generation process. The modified free surface boundary condition of the water layer accounts for gravity restoring forces.
A first challenge lies in the meshing of the complex 3D simulation domain, including the detailed topography with the water layer on top. The input parameters of the kinematic source are easily constrained and the output is fastly elvaluated, such that varying sources can be efficiently compared. The analysis of the fully coupling process between earthquake source and tsunami generation fundamentally improves our understanding of earthquake-tsunami interaction in its full complexity and helps building blocks towards using fully-coupled models to complement commonly used tsunami models.

In addition, there are two major operational challenges for seismic ground shaking analysis. One is the uncertainty of the input kinematic source models, which show non-unique solutions and can differ significantly for different observation data and model assumptions. The other one is the constraints of the velocity structures, especially the near-surface structures.

- **Workflow split into Tasks, Milestones, KPI**

1. Constraining the input parameters Downloading the kinematic earthquake data and post-processing to fit the input files structure of SeisSol.
2. Creating a computational mesh of the 3-dimensional simulation domain The earthquake and tsunami domain is defined and a non-uniform, adaptive, unstructured tetrahedral mesh is created with the meshing software Simmodeler (provided by Simmetrix Inc., 2017). It enables the usage of geometrically complex and curved, intersected or segmented faults, as well as detailed topography. The mesh resolution is adapted, such that the fault, the waterlayer and the topography are highly resolved.
3. Simulating the earthquake and tsunami simultaneously Simulation of the 3-dimensional and time-dependent rupture process together with the tsunami generation and propagation in SeisSol. SeisSol solves the seismic wave equations with high-order accuracy in space and time. The fully-coupled earthquake-tsunami model is able to capture the direct impact of rupture variability on tsunami generation.
4. Post-processing and analysing the results The data are post-processed, visualised in Paraview and analysed. Videos of the

rupture evolution, the surface wavefield and the tsunami provide insights into the complex rupture-tsunami interaction. The time-dependent and spatial variable seafloor displacement translates into the sea surface displacement, which directly influences tsunami generation.

Identification of the Stakeholders and their roles

The stakeholders involved in the demonstration are those of PD1 and PD8.

Target Operational Environment

- **Description of the operational context (main actors, procedures, etc)**

To generate the ground shaking maps and coupled earthquake-tsunami for the Samos-Izmir October 30th, 2020 earthquake, we use the kinematic rupture models from different research groups, as the input rupture sources input in SeisSol. The models are either based on observational data (Heidarzadeh et al., 2021, used in PD8 demo)) or based on the CMT parameters (generated by UCIS4EQ in PD1). We first convert the source model to standard rupture format (SRF) and then use python scripts (https://github.com/SeisSol/Meshing/tree/thomas/improve_inp_converter), by interpolating more point sources, to smooth the original coarse point sources. We use Simmodeler (provided by Simmetrix Inc., 2017) to generate the mesh with topography/bathymetry data (https://download.gebco.net/). The simulation in SeisSol utilises either 1-D or 3-D velocity models. After the simulation is finished, we use post-processing scripts (https://github.com/SeisSol/SeisSol/tree/master/postprocessing/science/GroundMotionParametersMaps) to generate the ground shaking maps.

HPC Challenge

Architecture	Node specs	Total number of allocated nodes per exercise	Node per run	Number of runs	Expected duration of a single run	Expected data storage needs
Supermuc NG (LRZ, Germany)	Intel Xeon Skylake	17	17	2	~48 minutes for 100s simulation time	32 Gb

Table 5.1. Computational resources committed for PD5 exercise.

Potential target service

PD8: Multi-physics simulation capabilities open the possibility to fundamentally improve our understanding of earthquake-tsunami interaction in its full complexity, which includes identifying when 3D-2D one-way linked models are justified and when fully coupled modelling may be required.

PD1: We can open up new avenues for applications into urgent physics-based modelling. Not only is this now possible; it is also efficient. Our new simulation capabilities combined with PD1 workflows hold the possibility to fundamentally improve our understanding of earthquake hazard, by PD1 workflow linking to (probabilistic) hazard and multi-physics studies.

Current challenge and on-going efforts for advancing urgent-response applications:

1) Automated, efficient and adaptive geo-data merging into geometric (CAD) models to allow for automated meshing workflows.
2) Increasing computational cost if higher spatial resolution and longer simulations are required.
3) Quantification of uncertainties using ExaHype (see e.g. Seelinger et al., SC'21[4])

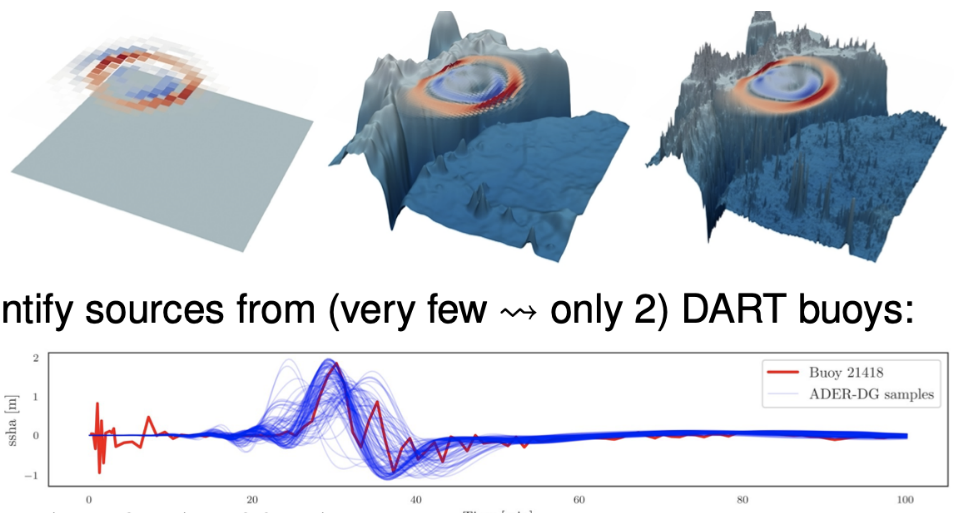

Fig. 5.7: ExahyPE's UQ application to tsunami modelling with adaptively refined bathymetry resolution available (Seelinger et al., SC'21).

[4] Linus Seelinger, Anne Reinarz, Leonhard Rannabauer, Michael Bader, Peter Bastian, and Robert Scheichl. 2021. High performance uncertainty quantification with parallelized multilevel Markov chain Monte Carlo. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC '21). Association for Computing Machinery, New York, NY, USA, Article 75, 1–15. DOI:https://doi.org/10.1145/3458817.3476150

PD6. Probabilistic Volcanic Hazard Assessment

PD6	Probabilistic Volcanic Hazard Assessment
Leader	Laura Sandri
Participants	INGV, BSC, IMO
Workflow	ST_PVHA_CF
Numerical Engine	Fall3D
TRL initial	3(*)
TRL target	7
TRL achieved	7

(*) The TRL values originally stated in the DoA, i.e. from initial 4 to final 6-7 in the particular case of PD6, have been revised after the redefinition of TRL in the context of ChEESE. See deliverable D5.1; Pilot analysis and development strategy; TRL Table on page 59.

Description of the live exercise / benchmark

- **Objective**. The objective of PD6 exercise was to show the feasibility and usefulness of short-term probabilistic assessment of volcanic tephra hazard in an operational environment, with the added value (enabled by HPC) of an assessment over a large-scale domain and considering the uncertainty associated with the eruption size parameters.
- **Operational Challenges**. The operational challenges were represented by dealing with real time performance of the PD6 distributed workflow, in particular:
    - Fetching real-time monitoring data from Osservatorio Vesuviano surveillance system, and processing them at INGV Bologna on the computer cluster ADA,
    - Fetching weather forecast from GFS and processing them at INGV Bologna,
    - Running large-scale and high-resolution tephra dispersal simulations with the Flagship code Fall3D on MareNostrum at BSC, and
    - Handling the output from the above on ADA cluster at INGV-Bologna to produce hazard assessment for the end-users. For redundancy reasons, the same workflow would have been operational, in case of necessity, at INGV Osservatorio Vesuviano on the computer cluster Ceneri.

- **Scientific Challenges**. The scientific challenges were related to prove the feasibility of using such hazard evaluations to produce useful impact assessment of tephra ground load at the scale of a country, in particular over mobility and electrical networks. Currently, a prototype operational tool (BET@OV) is implemented at INGV Osservatorio Vesuviano to evaluate daily the short-term probabilistic tephra-load hazard from Campi Flegrei, but it has two main limitations: (*i*) it is able to simulate only one eruptive scenario per size, so it neglects the uncertainty related to the eruption size and the related eruptive source parameters (ESP), and (*ii*) it covers a geographic domain restricted to the Campania administrative region (scale of about 200 km). The ChEESE's ST_PVHA_CF tested here overcomes these limitations.

	Current operational workflow BET_OV	**ChEESE workflow ST_PVHA_CF**
Computational domain	longitude from 13.75ºE to 15.36ºE latitude from 40.49ºE to 41.25ºE	longitude from 8.62ºE to 20.12ºE latitude from 36.60ºN to 45.22ºN
Horizontal resolution	0.005º (46137 grid points)	0.02º (248832 grid points)
Vertical resolution	20/30 layers	9/50 layers
Tephra dispersal model	FALL3D-v7.2	FALL3D-v8.1 (developed during ChEESE)
Number of scenarios	3	300
Cluster	Ceneri	Distributed workflow: the different tasks can be sent to different clusters and are managed by the ChEESE's Work Flow manager System WMS-L

Table 6.1. Comparison between BET_OV (prototype tool) and ST_PVHA_CF (tested) models setup.

- Workflow split into Tasks, Milestones, KPI

Task 1. a) Download of INGV-OV monitoring data (simulated – we took the data from 26th April 2020, when a seismic swarm with Mmax=3.3 was recorded) b) Run of BET_EF module
Task 2. a) Download of weather forecast (actual on 4th November 2021) b) Generate Fall3D input files for 3 sizes, about 100 runs for each size → about 300 different scenarios (different ESPs) c) Run 300 Fall3D simulations, with eruptions starting at 00 on 5th November each 72-hour long (5th, 6th and 7th November)
Task 3. a) Probabilistic combination of the simulations b) Production of output: ● Plot of time evolution of the probability of unrest and eruption ● Map of the spatial probability of vent opening ● Probability maps to overcome a selected ground load threshold ● Hazard maps to overcome a selected probability threshold in ground load ● Tables with mean probability to overcome all the ground load thresholds

Identification of the Stakeholders and their roles

Stakeholders' institution	Role played
PLINIVS	IUB member and Centre of Competence of the Italian Civil Protection. Use of PD6 results for short-term impact assessment of ground ash load on mobility networks (road, railways, seaports and airports) and electric network.
ARISTOTLE	IUB member. Use of PD6 results to open a report to ERCC. Update ERCC with PLINIVS impact assessment.

Table 6.2. Stakeholders' involvement

- **Communication and information flow.** The participation has been remote due to COVID-19 pandemics. However, it resembled a real case, in which in case of unrest at Campi Flegrei personnel from Bologna would very likely run the workflow remotely and communicate with end-users remotely. For the exercise, 5 mailing lists were created to reach participants in a proper communication flow: ChEESE-PD6 team, PLINIVS team, Aristotle team, Osservatorio Vesuviano directorship, and ChEESE observers who had asked to be in the loop. The following timetable with activity at the different steps were defined and accomplished:

h. 8:30-9:30 CET:

Videoconference on Google Meet to:
- test remote communication flow;
- shortly illustrate the ChEESE project,
- illustrate PD6: the goal of the exercise;
- illustrate the (simulated) volcanological situation at Campi Flegrei as recorded by Osservatorio Vesuviano surveillance system, and release of a bulletin via email (from ChEESE-PD6 mailing list specifically created for the exercise) from the latest simulated monitoring data

h. 13 CET:

Transmission of data from INGV to end-users (PLINIVS and ARISTOTLE) by email:
- ChEESE-PD6 sent to PLINIVS mailing list the tables with exceeding probability;
- ChEESE-PD6 sent a report containing PD6 results to PLINIVS, ARISTOTLE and Osservatorio Vesuviano mailing lists.

h. 15 CET:

Transmission of data from PLINIVS:
- PLINIVS loaded a report containing impact maps onto their ftp repository and notified by email ChEESE-PD6, ARISTOTLE and Osservatorio Vesuviano mailing lists.

h. 16-17 CET:

Feedback from end-users. Debriefing.

- **Ways to collect feedback from the Stakeholders**

We had a debrief phase at the end of the exercise day, discussing for an hour the potential in the workflow and its results.

Target Operational Environment

- **Description of the operational context** (main actors, procedures, etc). ChEESE-PD6 organised and coordinated the exercise, specifically producing a simulated bulletin of the monitoring data (see Appendix), downloaded the meteo data and the monitoring data, run the Fall3D simulations on MareNostrum4, run the workflow and sent the hazard products to end-users. PLINIVS received the hazard products and processed them to quantify the impact of tephra on the Italian mobility network (roads, railways, airports and seaports) and on the electric network. They also produced a report (see Appendix). ARISTOTLE responded to Osservatorio Vesuviano information on volcanic unrest by opening a report for the ERCC. Later, they produced updated reports to ERCC containing the hazard products from ChEESE and the impact assessment from PLINIVS.
- **Expected added value with respect to existing procedures.** We see two main added values. First, the operational run of such short-term hazard enables short-term impact assessment on a country-scale and high-resolution domain is unprecedented, demonstrating what it can be done daily in case HPC resources are available and dedicated. Secondly, we also showed that both short-term hazard and impact assessment can be used and coordinated at a European level by involving Aristotle and, consequently, ERCC.

HPC Challenge

- Computational resources allocated

MareNostrum at BSC:

- 24 nodes per 1 hour, approximately, for Tasks 2-b,c and 3-a. Summary:
- 2104 CPU hours for simulations size Low (26296 seconds real time x 6 nodes x 48 processors/node)
- 7752 CPU hours for simulations size Medium (96892 seconds real time x 6 nodes x 48 processors/node)
- 15000 CPU hours for simulations size High (93765 seconds real time x 12 nodes x 48 processors/node)
- 16 CPU hours for generation of scenarios and analysis (20 minutes real time x 1 nodes x 48 processors/node)
- Total: 24872 hours CPU time.

ADA at INGV-Bologna:

- 12 nodes per 2 hours (for downloading meteo data and generating outputs)

Ceneri at INGV-OV:

- 1 node per 10 minutes (for calculation of probabilities related to eruption occurrence and position of the vent)

	Architecture	Node specs	Total number of allocated nodes per exercise	Node per run	Number of runs	Expected duration of a single run (average)	Expected data storage needs
Mare Nostrum 4 (BSC, Spain)	x86_64		24	6 6 12 1	Task 2-c: 100 runs size L 100 runs size M 100 runs size H 1 run Task 2-b 1 run Task 3-a	5 min 20 min 20 min 20 min 20 min	2.6 Gb 4.6 Gb 4.6 Gb 31 Mb 100 Mb
INGV cluster "ADA"	x86_64		12	12	1 run Task 2-a 1 run Task 3-b	30 min 90 min	100 Mb 1.8 Gb
INGV cluster "ceneri"	x86_64		1	1	1 run Task 1	10 min	20 Kb

Exercise Run (for live exercises)

- **Exercise report (with time indications).** The exercise followed the scheduling reported above. In the appendix we provide the available material developed for and during the exercise. In particular we provide:
  - the simulated bulletin (Appendix PD6.1)
  - the report with PD6 results provided to PLINIVS and Aristotle (Appendix PD6.2)
  - the report with the impact assessment provided by PLINIVS (Appendix PD6.3)

- **Deviations from original plan.** N/A
- **Bottlenecks, obstacles, unexpected difficulties.** We encountered a problem in the input of some of the monitoring data in the workflow. This is due to the workflow being set to read and process automatically the monitoring data stored in Osservatorio Vesuviano databases. However, the geochemistry data are not currently stored in any database real-time, so we had to input them manually. This problem will be fixed in the near future.
- **Lessons learned and recommendations**. End-users appreciated the potential of the workflow's hazard results, so we believe that a service able to provide these results operationally could be exploited.

Potential target service

- **Missing steps**

In the workflow, we need to insert ad hoc input options for those monitoring data that are not operationally stored into databases.

- **Computational resources needed if operational**

In case of an operational service, similar resources to those used in the exercise would be needed and reserved for the service.

Appendix PD6.1

The online file PD6.1a.pdf is the simulated bulletin released in the morning (9:30 am CET) of November 4th by Osservatorio Vesuviano, illustrating the last month of monitoring and the updates since the last days.

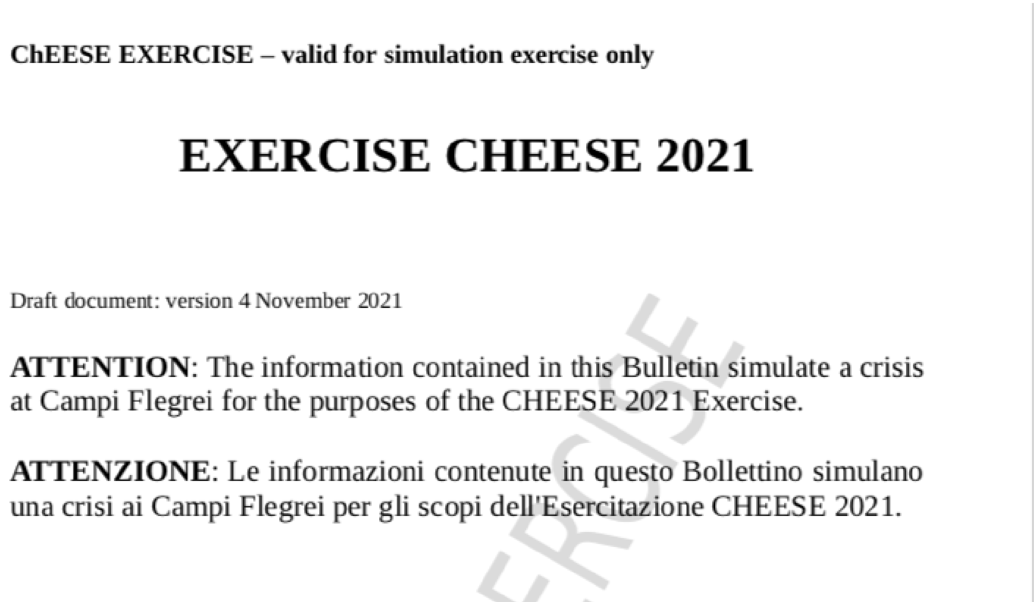

ChEESE EXERCISE – valid for simulation exercise only

EXERCISE CHEESE 2021

Draft document: version 4 November 2021

ATTENTION: The information contained in this Bulletin simulate a crisis at Campi Flegrei for the purposes of the CHEESE 2021 Exercise.

ATTENZIONE: Le informazioni contenute in questo Bollettino simulano una crisi ai Campi Flegrei per gli scopi dell'Esercitazione CHEESE 2021.

Figure 6.1. Extract from the Exercise Report PD6.1a. Simulation of unrest conditions at Campi Flegrei (Italy).

The online file PD6.1b.pdf is a simulated update to the bulletin released at 9:40 am CET of November 4th by Osservatorio Vesuviano, illustrating the most recent updates from November 4th in the geochemistry measurements.

Appendix PD6.2

The online file PD6.2.pdf is the report that the PD6-ChEESE group generated by 3pm on November 4th, containing the hazard results from PD6 workflow.

00-24 h	1%HM - 2.5th percentile 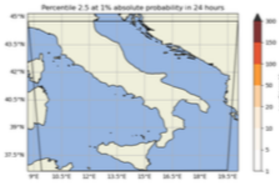	1%HM - mean 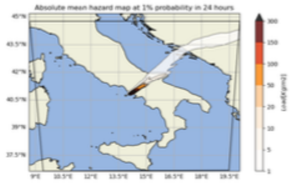	1%HM 97.5th - percentile 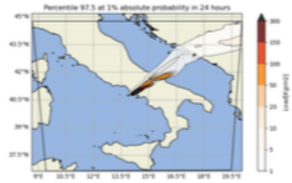
00-48 h	1%HM - 2.5th percentile 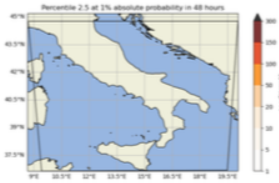	1%HM - mean 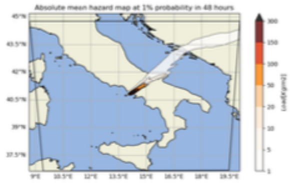	1%HM - 97.5th percentile 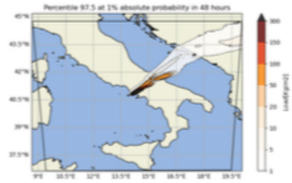

Figure 6.2. Extract from PD6.2 exercise report. Absolute mean hazard map at 1% probability (with percentiles 2.5th and 97.5th to show the uncertainty) for the first 24 and 48 hours since the beginning of the eruption

Appendix PD6.3

The online file [PD6.3.pdf]() is the report that PLINIVS generated by 3 pm on November 4th containing the impact results based on PD6 hazard assessment.

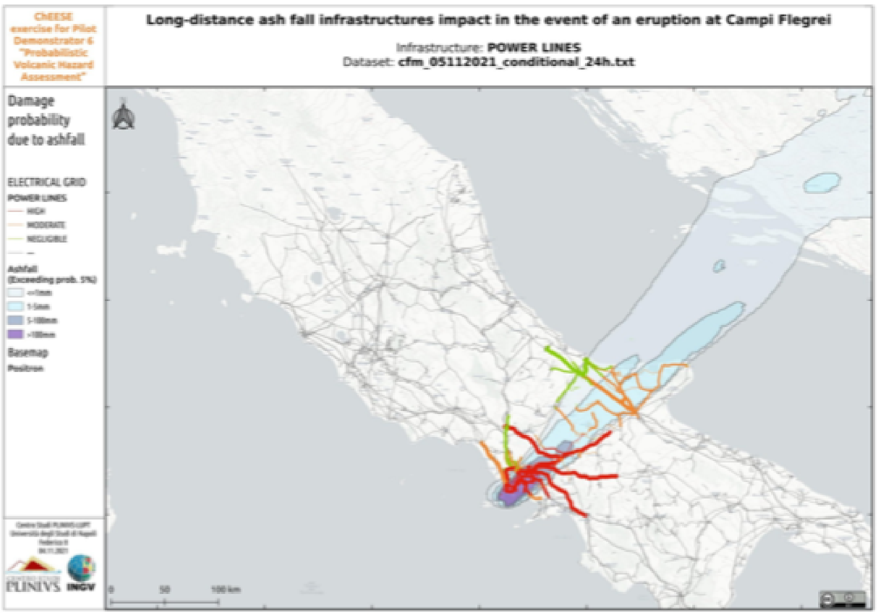

Figure 6.3. Extract from PD6.3 report. Medium size eruption impact on Electrical Grid Power Lines.

PD8. Probabilistic Tsunami Forecasting for Early Warning and Rapid Post Event Assessment

PD8	Probabilistic Tsunami Forecasting for Early Warning and Rapid Post Event Assessment
Leader	Stefano Lorito (INGV)
Participants	• INGV (Manuela Volpe, Jacopo Selva, Fabrizio Romano, Roberto Tonini, Fabrizio Bernardi, Maria Concetta Lorenzino) • UMA (Carlos Sanches-Linares, Jorge Macias, Marc de la Asuncion, Jose Manuel Gonzalez Vida, Manuel J. Castro) • NGI (Steven Gibbons, Finn Løvholt, Malte Vöge, Sylfest Glimsdal) • CINECA (Silvia Giuliani, Isabella Baccarelli, Piero Lanucara) • HLRS (Alexey Cheptsov) • UniNA (Antonio Scala) • GFZ (Andrey Babeyko)
Workflow	Suite of bash, python/Matlab, C, GMT codes for workflow execution, pre- and post-processing (https://gitlab.rm.ingv.it/; on the intranet and users need authorization)
Numerical Engine	Tsunami-HySEA
TRL initial	3
TRL target	6-8
TRL achieved	7-8 (9)*

*The version used for the exercise, which includes running of large ensembles of tsunami simulations can be considered at TRL 7, and it is expected to be raised to TRL 8 within the project duration, since both engineering for usage within an Urgent Computing experiment and integration with WMS-light workflow manager for orchestration across distributed resources are being achieved. The version based on pre-calculated ensemble scenarios has achieved TRL 9, since it is up and running in the operational environment of the Italian Tsunami Service Provider (CAT-INGV).

Description of the live exercise

ChEESE Urgent Computing Samos Earthquake and Tsunami Exercise - World Tsunami Awareness Day (WTAD).

On the 5th of November, on the occasion of the World Tsunami Awareness Day (WTAD, https://tsunamiday.undrr.org/) organized by UNDRR, we conducted a table-top Urgent Computing exercise and Rapid Post-Event Tsunami Assessment.

Objective

The exercise had several purposes:

1. to show the potentiality of High Performance Computing for Probabilistic Hazard Forecasting, in particular as far as Uncertainty Quantification is concerned;
2. to test the feasibility and to analyze the advantages offered by Urgent Computing, at least for Rapid Post Event Assessment, and in perspective for Early Warning;

3. to interact with the real stakeholders of this PD, such as the tsunami warning centers, the national civil protection agencies, the observatories and scientific institutions providing advice to decision-makers during the post-event period, and the European Civil Protection in charge of organizing first and second responders.

As illustrated schematically by Figure 8.1 and described in previous Deliverables (e.g. D4.8 and D4.3), PD8 PTF provides a rapid probabilistic forecast of tsunami inundation following an earthquake offshore. We reiterate that **there exist two versions of the PTF**.:

1. The first version is based on pre-calculated numerical simulations of tsunami scenarios. It works in near-real time on the premises of the CAT-INGV Tsunami Warning Centre (https://www.ingv.it/cat/en/) in the so-called Early-warning mode, which is based on pre-calculated ensembles of tsunami scenarios and their linear combinations according to the earthquake parameters. (TRL 9)
2. The second version is based on simulation ensembles to be run from scratch on large enough HPC clusters in the Urgent Computing mode. This is the version showcased in this live exercise. (TRL 7-8)

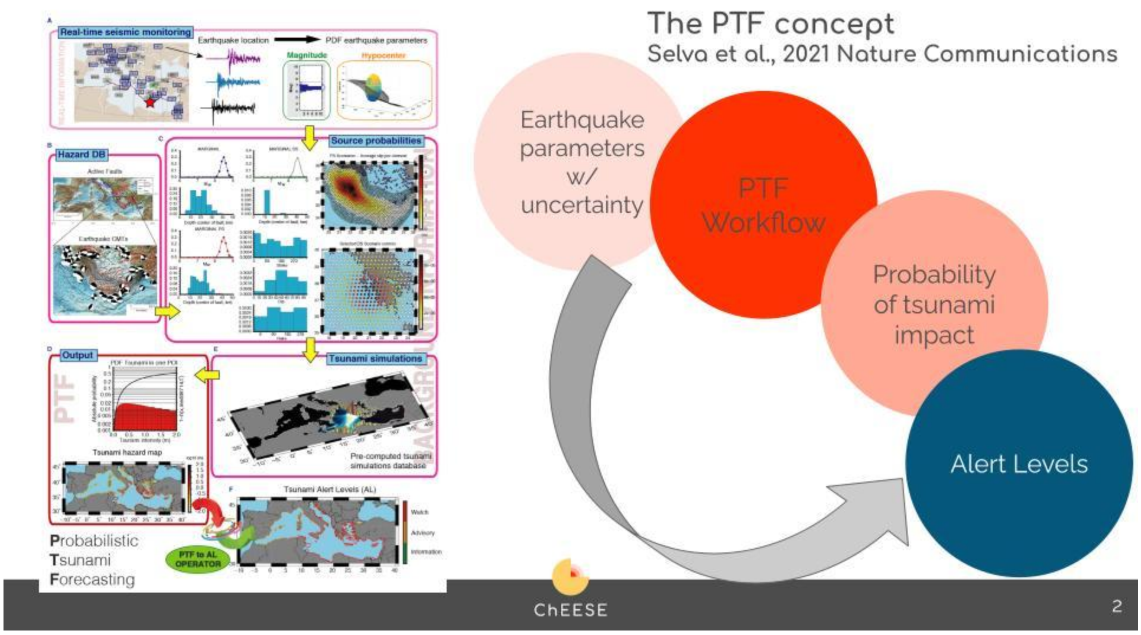

Figure 8.1. The PD8 PTF concept. The PTF Workflow block includes large ensembles of numerical simulations for uncertainty quantification.

Based on real-time seismic parameters, the workflow initializes the run of a large ensemble of numerical simulations with the Tsunami-HySEA Monte Carlo version, provided by UMA. It then produces the tsunami forecast at specific points within the Mediterranean Sea by combining the source probabilities with the ensemble simulation output (either from pre-calculated scenarios or from those computed on the fly). This service provides a rapid probabilistic forecast of tsunami inundation before tsunami observations are available. The large uncertainties about earthquake location and magnitude, as available in the first minutes, are reflected in forecasting uncertainty.

The overarching objective of the exercise was to demonstrate the functioning and characteristics of the PD8 PTF workflow, while testing it in an almost realistic Emergency Computing configuration, through a hindcasting experiment exploiting computational infrastructure at CAT-INGV for non-HPC pre- and post-processing and a large dedicated portion of the Tier-0 MARCONI100 supercomputer.

Description of the live exercise

On the 5th of November, 2021, we reconstructed a major event that occurred almost one year before: the tsunami following the October 30, 2020, Mw 7.0 Samos-Izmir earthquake. The earthquake caused building collapses and about 100 fatalities at Izmir, and a moderate tsunami which hit the close-by island of Samos to the South and the coast of Turkey to the North, causing one fatality.

In particular, we recreated the event and the response to it by several actors by setting up a realistic chain of events following the earthquake detection within a dedicated virtual meeting.

The exercise was led by INGV and conducted in several steps between about 9-14 CET. The number of participants peaked at about 70 people.

In the beginning, there was an introduction by INGV illustrating the program of the exercise, the basics of the ChEESE project, and of the ChEESE PD8. This was followed by a presentation by CINECA illustrating the Marconi100 supercomputer which would be used for the live Urgent Computing segment of the exercise. As detailed in the following, an ensemble of tsunami simulations has been run to evaluate the Probabilistic Tsunami Forecast (PTF, https://www.nature.com/articles/s41467-021-25815-w).

Then, the simulation of the earthquake and tsunami, and of the events and actions in response was simulated as follows.

Exercise Report with time indications

The seismic event occurred at about 9:40 CET. The earthquake was first **detected at 9:43 CET** by NOA, KOERI, and INGV, which are Tsunami Service Providers (TSPs) of the NEAMTWS (http://www.ioc-tsunami.org/) that issued alert messages for this tsunami in 2020.

The three TSPs (NOA, KOERI, CAT-INGV) issued their warning messages **in the following 10 minutes** and illustrated them to the audience, regarding the first Tsunami Alerts produced with the Decision Matrix (DM), which is a lookup table that associates the earthquake parameters with the alert levels. This part of the procedure was executed in an identical manner to the procedure carried out in 2020. Additionally, the TSPs demonstrated some enhanced alert messages, enriched with maps and other graphic information which will be implemented in their operations at some point in the future. Figure 8.2, for example, shows a prototype of the maps that CAT-INGV will distribute in the future along with the usual text messages that are currently distributed during an alert.

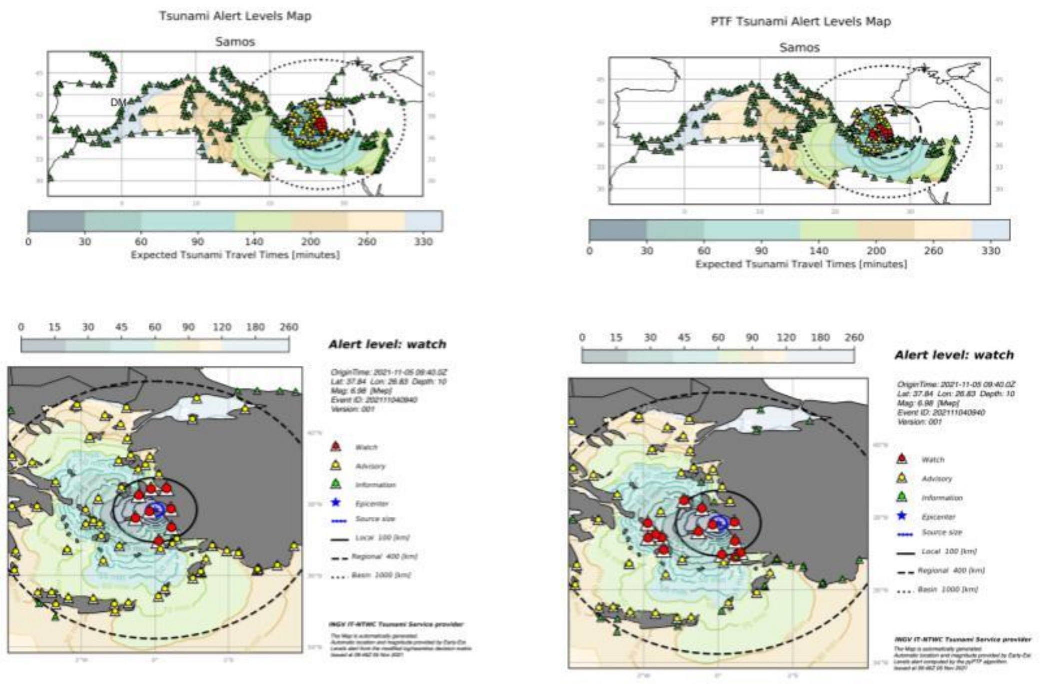

Figure 8.2. Future CAT-INGV enhanced alert products. Comparison between the alert levels obtained: with the Decision Matrix (DM), panels to the left; with the Probabilistic Tsunami Forecasting (PTF), panels to the right. In the near-field of the earthquake source (bottom panels), it is visible how PTF reduces overestimations and has a more realistic pattern driven by tsunami propagation over a complex bathymetry and taking into account topographic obstacles, which are not accounted for by the DM.

For the Samos event hind-casting experiment we considered an uncertainty exploration for the earthquake parameters up to 2 sigma. This resulted in ~38000 tsunami scenarios.

During the exercise, INGV first ran the PTF in the Early Warning mode. This run started in parallel with the determination using the DM, as soon as the earthquake parameter estimates were available, a few minutes after the earthquake occurrence. The results of the PTF in Early-warning mode were immediately compared with the alerts released with the DM. It was possible to point out unambiguously the first order advantages of the method based on uncertainty quantification through the combination of the results from large ensembles of pre-calculated numerical simulations (Figure 8.1).

At about **10:00 CET**, the ensemble of runs necessary for the **PTF in Urgent Computing Mode was then submitted to the Marconi100 supercomputer @CINECA.** The entire workflow execution, including pre-processing of the input, simulation staging, run, post-processing and transfer of the results took less than one hour (~ 40 min).

A total of 805 nodes were available, each one equipped with 4 NVIDIA V100 GPUs, reserved by CINECA for this exercise. The simulations were run through a job array saturating all the available resources after a few tens of seconds. Each element of the array run on one node, preparing the simulation input for 4 simulations, executing the simulations on a single GPU, and performing the post-processing on the HySEA outputs.

The computation domain consisted of 1981x1321 cells (~ 5x5 deg domain at 30 arcsec), with each run taking ~30 seconds on a single V100 for simulating 4 hours of propagation with Tsunami-HySea which solves nonlinear shallow water equations. The tsunami time-series were recorded offshore both at 50 and at 10 m depth, approximately every 2 km along the coast. The amplitudes were extrapolated to 1m depth using Green's law (1D energy conservation). The results obtained at the 50 m isobath are directly comparable to those obtained with the pre-calculated scenarios. This allows an assessment of the gain of accuracy provided by performing on-the-fly individual runs for each ensemble element as opposed to the linear combinations of pre-calculated scenarios. The pre-calculated scenarios are limited to the 50 m isobath from the requirement of linearity. The Tsunami HySEA simulations are not limited in this way and we can extend the simulations far closer to the shore: at the 10 m isobath. The improvement in accuracy of the amplitudes extrapolated to 1m depth can then be assessed by comparing the output extrapolated from the 50m time-series and the (more accurate) output extrapolated from the 10m time-series.

We point out that we have conducted the urgent computing exercise within the typical timeframe of operational advice in support of first and second disaster response, as in ARISTOTLE (http://aristotle.ingv.it/tiki-index.php) reporting to the DG-ECHO Emergency Response and Coordination Centre (ERCC). ARISTOTLE indeed reported and provided expert advice to ERCC on the occasion of the 2020 Samos-Izmir event. For this reason, the forecasting that ARISTOTLE would obtain with the tools available today was also presented by the Tsunami Hazard Group deputy chair on shift on 5th of November 2022 (NOA).

Ways to collect feedback from the stakeholders

In the following hours, a thorough presentation and discussion of the comparison among the different forecasts (DM, single simulations, PTF ensemble in early warning and in urgent computing mode) and with observations, including run-ups (Figure 8.3), took place.

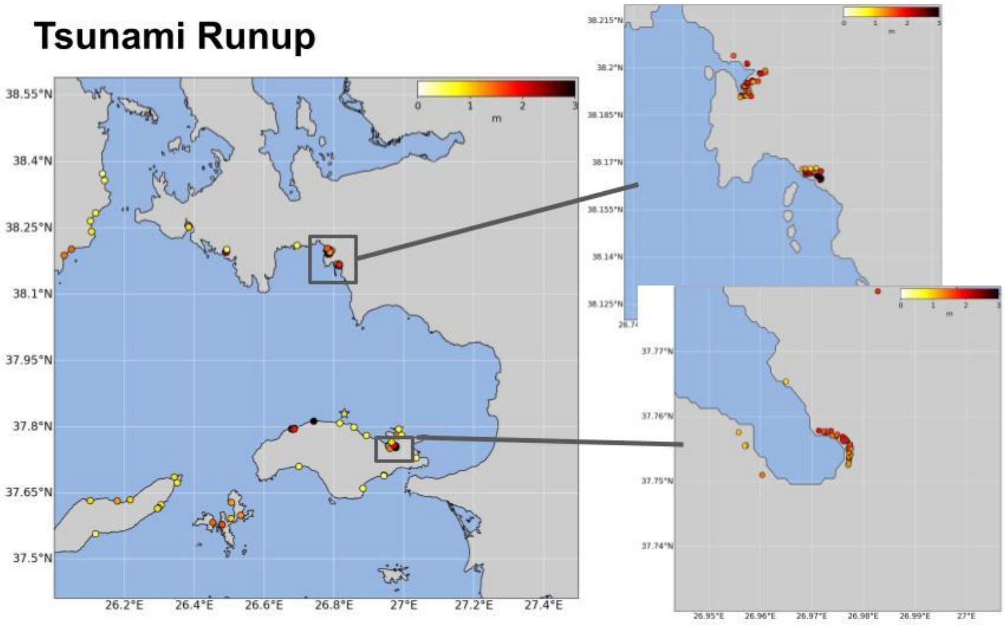

Figure 8.3. Summary of the tsunami runup measurements as a result of several surveys conducted in the aftermath of the 2020 Samos-Izmir earthquake, namely those reported in Triantafyllou et al., 2020, https://doi.org/10.3390/jmse9010068; Dogan et al., 2021, https://doi.org/10.1007/s00024-021-02693-3 ; Kalligeris et al., 2021, https://doi.org/10.1007/s10518-021-01250-6.

As an example, we show here the comparison between the results obtained from the PTF in the Urgent computing mode and the simulations performed by ARISTOTLE for the two nodal planes available from the CMT moment tensor (Figure 8.4). It is evident that depending on the nodal plane chosen as the earthquake faulting mechanism, the two simulations either over- or under-estimated the observed runups from direct propagation of the tsunami both to the southern coast of Turkey and to the northern coast of Samos and Ikaria Greek Islands. Conversely, on Fourni Island, both simulations underestimate the results. **The runup values are instead always inside the PTF range.** The PTF range of course embeds the results of the individual simulations, becauses they are included within the variability of the ensemble, which depends in turn on the uncertainty of the input seismic parameters. A thorough discussion of these results will be included in a scientific publication related to the urgent computing exercise.

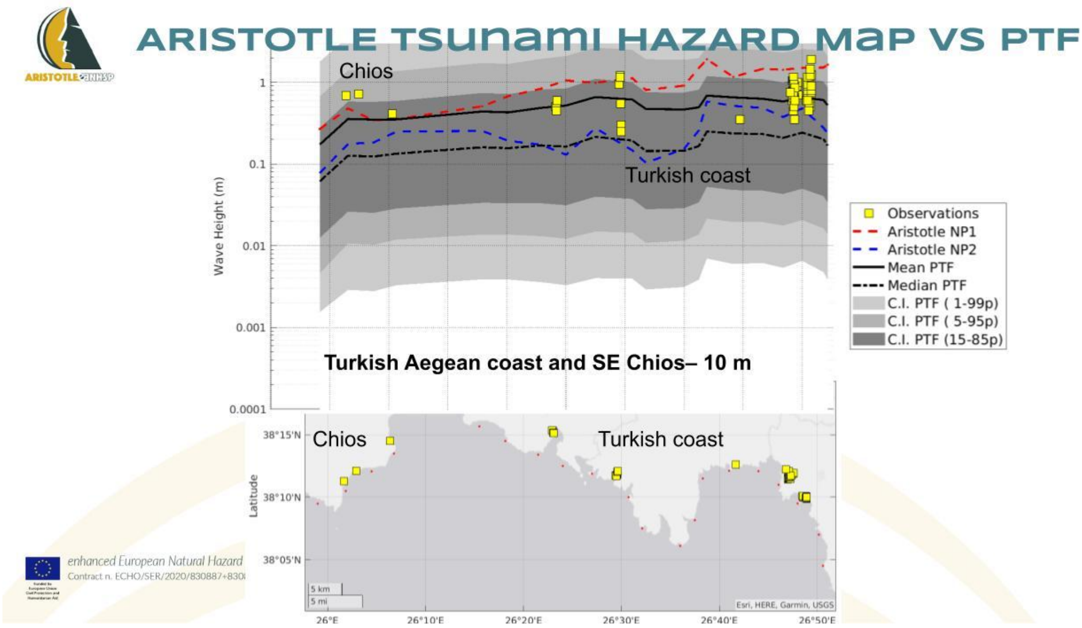

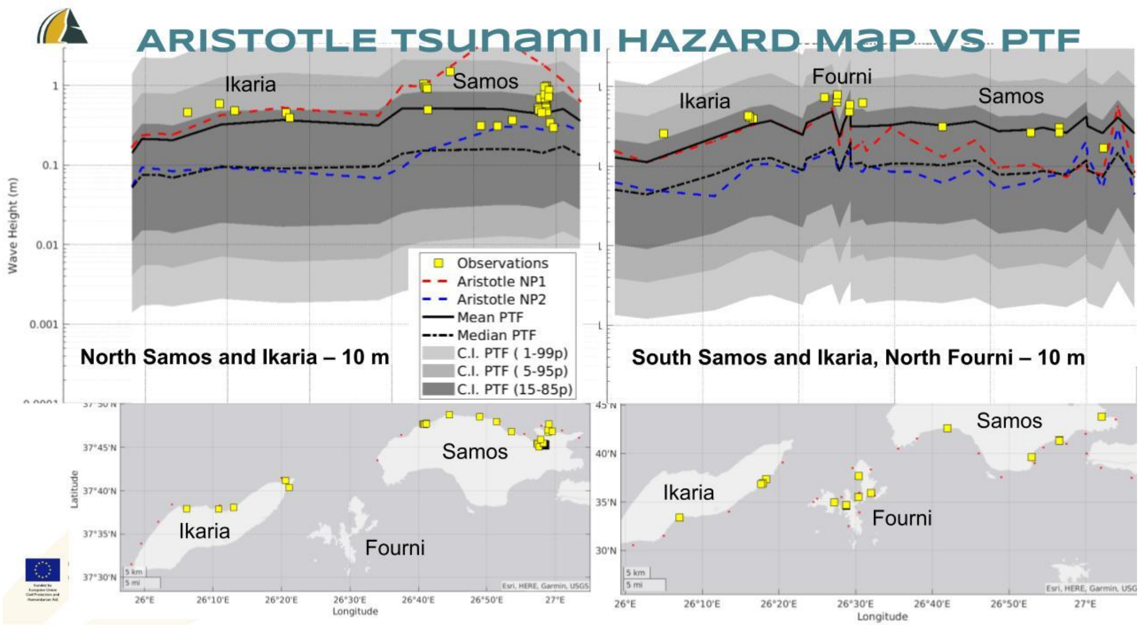

Figure 8.4. Comparison with the Samos 2020 runup observations of the forecasts obtained with the PTF and with the ARISTOTLE simulations.

It is important to remark that the participants involved in the discussion were not only from the scientific community. For example, an important contribution to the discussion and general comments or remarks came from the representative of the NEAMTWS TSPs (KOERI, NOA, and INGV), in charge for tsunami warning in the region, from the ARISTOTLE officer (NOA), and from national and European Civil Protection representatives. The potentiality and applicability of the PTF for early warning operations and post-disaster assessment were then discussed from different angles thanks to the presence of a variety of stakeholders. The ARISTOTLE officer and the Italian Civil Protection representative explicitly recognized the added value offered by the probabilistic approach for informed tsunami risk assessment and management.

The workshop concluded with four short presentations, each of length 15 minutes, describing four other elements of the ChEESE project, each aimed at strengthening a different aspect of operational tsunami science. The first, “Towards HPC-based Early Warning at the Spanish National Warning Center”, was presented by UMA based on ChEESE Pilot Demonstrator 2: Faster-Than-Real-Time tsunami simulation. This presentation described the advancement and optimization of code with results of timings for parallel Urgent Tsunami Simulations. The second, “Earthquake-Tsunami Coupled Simulation”, was presented by LMU/TUM based on ChEESE PD5 “Physics-based probabilistic seismic hazard assessment”. This presentation described the results of simulations performed using the SeisSol software in which the rupture process and tsunamigenesis is modeled in high-resolution. Such calculations are currently too expensive for urgent computing and Tsunami Forecast but will provide valuable insights into the physics of tsunamigenesis which will increase the accuracy of Urgent Tsunami Computations in future years. The third, “ChEESE Workflow Management System”, was presented by HLRS based ChEESE Work Package 3 (“HPC modelling workflows and tools”). Here the WMS-light system was presented which facilitates the operation of complex workflows in geoscience on HPC. The fourth, “High Performance Computing for Probabilistic Tsunami Hazard Analysis”, was presented by NGI based on ChEESE Pilot Demonstrator 7: PTHA (see Figure 8.5). Here the PTHA workflow was presented together with an explanation of why high-resolution long term local hazard assessment is necessary, and how the ChEESE project has allowed us to progress from regional to local hazard assessment through the exploitation of HPC.

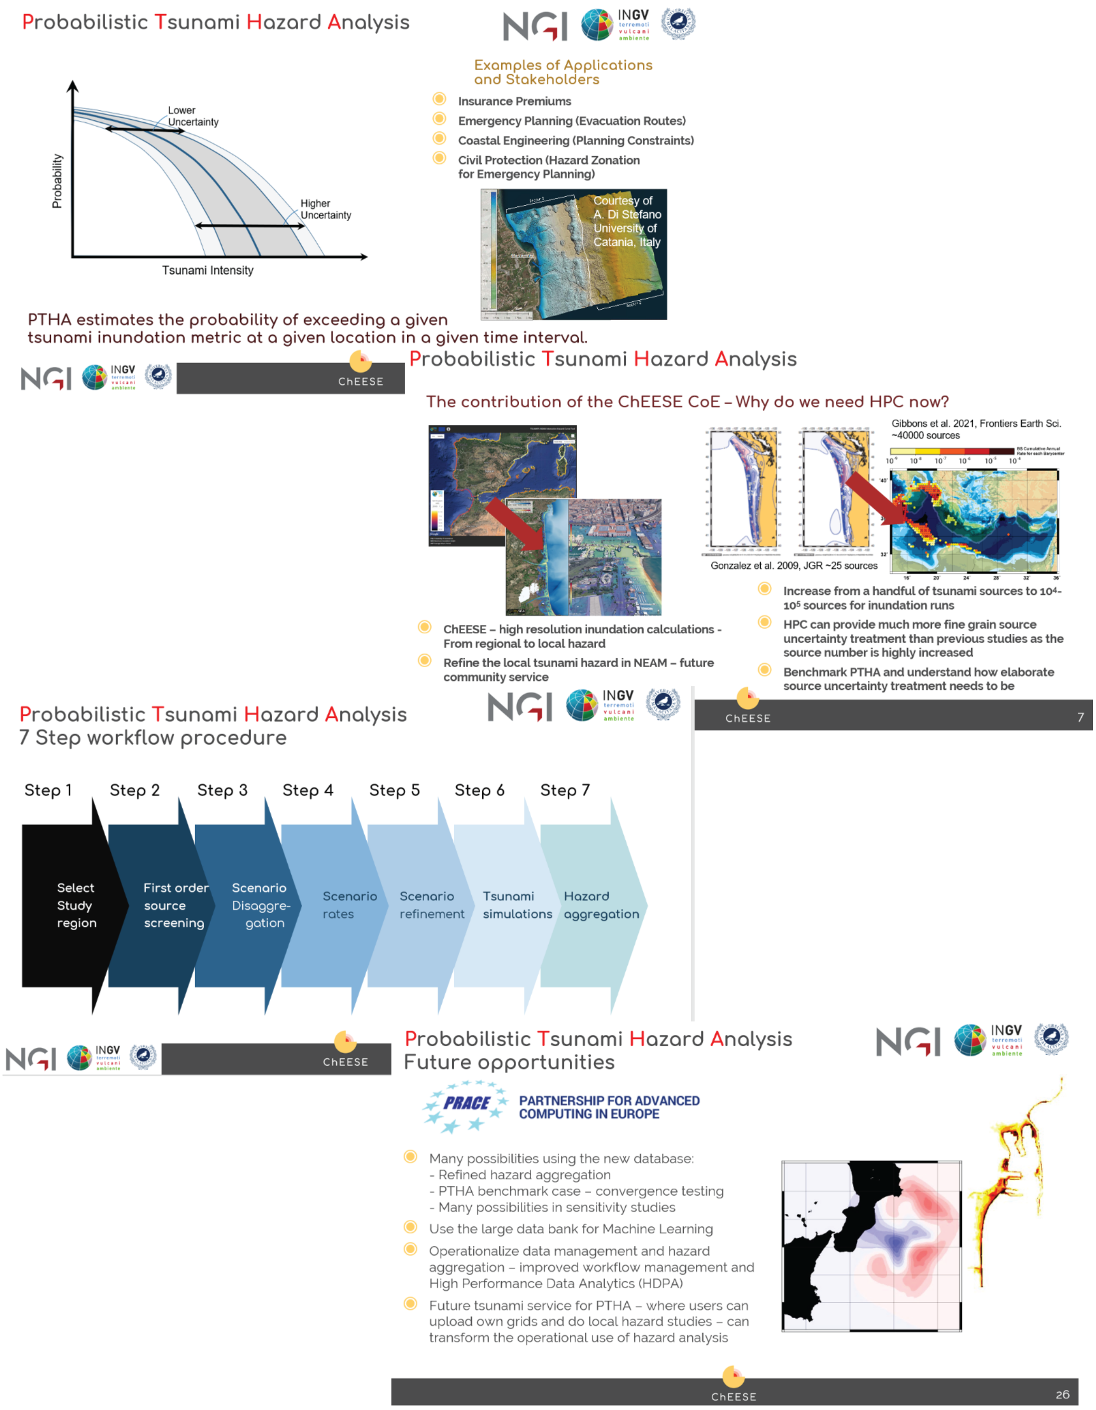

Figure 8.5. Excerpts from the presentation held on ChEESE PD7: Probabilistic Tsunami Hazard Analysis. It was presented how local tsunami hazard analysis benefits society and how high-resolution local PTHA is only possible through the use of HPC. The workflow was presented together with results for a local PTHA for Eastern Sicily and the prospects for future service provision in local PTHA. Details of the PTHA studies are provided in https://doi.org/10.3389/feart.2020.591549 and https://doi.org/10.3389/feart.2021.628061.

Identification of the Stakeholders and their roles

Stakeholders' institution	Interest	Role played
ARISTOTLE ENSHP	In charge of rapid Post-event assessment; may include the results in the reporting to ERCC	Active participant: the officer on shift presented their own assessment and view on the potential added value from PTF and from Urgent Computing
CAT-INGV, NEAMTWS TSP, ITALY	Runs the Italian National Tsunami Warning Centre and is a NEAMTWS Tsunami Service Provider in charge of monitoring the Mediterranean Sea. Will adopt the PTF as an operational tool for tsunami forecasting.	Hosted the Urgent Computing Exercise. Provided Seismic solution as input data to the workflow and the computational infrastructure (non-HPC pre- and post-processing), and the alert messages for the event based on traditional methods.
NOA, KOERI, NEAMTWS TSPs of Turkey and Greece	Run the Greek and Turkish Tsunami Warning Centres and are NEAMTWS Tsunami Service Providers in charge of monitoring the Mediterranean Sea	Provided the alert messages for the event based on traditional methods and participated actively in the discussion.
Other TSPs and members of the NEAMTWS SC	Governance and operation of the NEAMTWS	Observer
Italian Department of Civil Protection	Tsunami Alert dissemination on the Italian territory and coordination of the Italian Tsunami Alert System. Will adopt PTF as an operational approach for tsunami alert in Italy.	Observer
EPOS Candidate TCS Tsunami	Get to know about potential future services to be distributed and potentially integrated into ICS-D.	Observer
IGN - National Geographic Institute, Spain	Runs the Spanish National Tsunami Warning Centre, which is considering using scenario ensembles to be run on the fly for operational tsunami alert.	Observer

Table 8.1. Stakeholders' involvement

We report here for completeness the agenda of the day that was distributed in advance to invite the participants. Overall, the meeting went as planned.

Exercise promoted by the ChEESE Project (https://cheese-coe.eu/) to show the potentiality of Urgent Computing for Rapid Post Event Assessment: the 2020 Mw 7 Samos earthquake With the participation of the NEAMTWS Tsunami Service Providers NOA, KOERI, CAT-INGV; the ARISTOTLE Tsunami Hazard Group; National and European Civil Protection representatives; EPOS representatives	
9:00	Introduction 1) Welcome and Exercise agenda (INGV) 2) The ChEESE project (BSC) 3) The ChEESE PD8, Probabilistic Tsunami Forecasting (PTF) for early warning and rapid post event assessment (INGV) 4) HPC resources (CINECA)
9:30	Simulation of the 2020 Samos Event: 1) Earthquake detected (9:33) 2) Alarm sounds - First Earthquake parameters available (9:33-9:35) - all procedures start: a) Tsunami Alert produced with Decision Matrix (DM) by NOA, KOERI, and CAT-INGV b) PTF in Early Warning mode (based on pre-calculated scenarios) c) PTF in Urgent Computing Mode: submission of massive simulations to the Marconi100 supercomputer @CINECA
9:45	Tsunami Warning Messages by NEAMTWS TSPs (CAT-INGV, NOA, KOERI), based on DM, short presentations by 1) NOA (upgraded operations: updated DM, enhanced mapping product, use of Mw, national messages) 2) KOERI (enhanced products, potential use of numerical simulations) 3) INGV (DM and preliminary enhanced products)
10:15	Tsunami description (NOA, KOERI/GTU): Marigrams, Ongoing Messages, eye-witness and footages during the tsunami
10:35	PTF Results - comparison with DM and data (INGV) 1) Early warning mode 2) Urgent Computing mode 3) Role of the thresholds / conservatism 4) Comparison of the different forecasts (DM, single simulations, PTF ensemble in early warning and in urgent computing mode) with observations, including run-ups
11:05	Discussion (Chair NOA)
11:45	Added Value by Urgent Computing to ARISTOTLE Emergency Reporting (ARISTOTLE Tsunami Hazard Group deputy chair)
12:10	Other applications and Future developments: 1) Towards HPC-based Early Warning in Spanish National Warning Centre - ChEESE PD2 (UMA) 2) Earthquake-tsunami coupled simulation - ChEESE PD4 (LMU/TUM) 3) ChEESE Workflow management system (HLRS) 4) ChEESE PD7 (NGI)
13:20	Closing remarks - End of the exercise

Figure 8.6. Agenda of the Urgent Computing Workshop.

Operational Challenges and Potential target service

There are two distinct situations as far as the operational challenges are concerned.

The Early-warning mode version requires minor code optimization to reduce the execution time as much as possible. Otherwise, it is up and running and already fed by seismic solutions in real time. The production of messages can be activated on demand.

The situation is different when the Urgent Computing mode is concerned. During the exercise, the reservation of the MARCONI100 nodes was decided well in advance, as an in-kind contribution to the ChEESE project on the CINECA side. There is no guarantee nor a plan presently in place to make it happen in response to the next potentially catastrophic event. Hence, the principal operational challenge is that there need to be agreements in place to make it possible. Apart from this, it still appears necessary to dig a bit more into the problem of freeing the resources in a reasonable time, and to reduce some bottlenecks in the post-processing and visualization phase, as this exercise was the first experience of his kind and at this scale that we know about.

Apart from that, both the communication and the information flows do not represent a challenge, as long as either the database of pre-calculated scenarios, or the computational resources are available to the monitoring and warning centres in charge. The input and output information is indeed manageable.

Scientific Challenges and Target Operational Environment

Refinement of uncertainty modelling and calibration are required: more testing with real tsunamis is ongoing, also to provide decision-makers with more elements to set the thresholds to convert the PTF into alert levels, which is a political choice that should be supported by full uncertainty communication from the scientists.

Another challenge being dealt with is to build the evolutionary version of the PTF, to update the ensemble with incoming data (seismic & tsunami), in order to progressively reduce the uncertainty. This is likely easier for distant tsunamis than for local tsunamis, and for rapid post-event assessment rather than for early warning, which makes it a potentially typical Urgent Computing case involving data fusion.

The third challenge regards the Ensemble (sub-)sampling optimization to decrease computational demand. This would make it easier to allow for applications using smaller scale HPC clusters and would foster the diffusion of the approach to more centres in a more sustainable way.

Description of the operational context and of the different phases and actors for this exercise

- The INGV team run the workflow from the non-HPC computational environment of CAT-INGV (Centro Allerta Tsunami, Tsunami Warning Centre). The INGV team worked either remotely or on the premises of CAT-INGV;
- The seismic information was provided by running the Early-Est software in the CAT-INGV operational environment, using data streams downloaded by IRIS for the 2020 Earthquake. The seismic solutions triggered the workflow that starts the preprocessing on the local computational infrastructure;
- The TSPs NOA and KOERI issued the alert messages following the same procedures they used in 2020, when the event occurred;
- The CINECA team supervised the MARCONI100 supercomputer and made the node reservation.

- The tsunami simulation parameters were sent automatically to MARCONI100. The simulation post-processing steps were carried out on MARCONI100 and then the output was automatically transferred back to the CAT-INGV, where the forecast was assembled and the visualization was provided and made available to the stakeholders;
- NGI and UMA teams contributed to the preparation of the exercise in the problem definition and in the preliminary tests.

Expected added value with respect to existing procedures

It was shown that all the different PTF models (Early-warning, Urgent Computing at 50 and 10 m) are statistically compatible with data. Its adoption in the operational environments then would represent an added value over some of the existing procedures. It allows in fact for the quantification of uncertainties in real-time tsunami forecasting, which in turns favours a clear separation of the roles between scientists and decision makers.

It was also demonstrated that Urgent Computing simulations allow increasing accuracy, for example with respect to precalculated scenarios requiring linearity. Exascale HPC resources may allow to use denser target points or even to model inundation at specific locations via high-resolution simulations. On the other hand, they would allow for a deeper uncertainty exploration (larger ensemble), or for faster calculations if resolution and ensemble cardinality are fixed. Down-scaling based on different sub-sampling strategies can be considered for early-warning.

HPC Challenge

- **Computational resources allocated:** 805/980 compute nodes on the Tier-0 MARCONI100 Cluster @CINECA, each equipped with 4 x NVIDIA Volta V100 GPUs
- **HPC Challenges:** full exploitation of the dedicated resources; minimization of the latency of the Workload Manager, especially in case of large ensembles of simulations; parallel execution of post-processing on the nodes;

Architecture	Node specs	Total number of allocated nodes per exercise	Run per Node	Number of runs	Expected duration of a single run	Expected data storage needs
MARCONI100 @CINECA, Italy	4x NVIDIA Volta V100 GPUs per node	805	4	38736	~30 sec	20M per run

Table 8.2. Computational resources committed for PD8 live exercise.

- **Bottlenecks, obstacles, unexpected difficulties:** during the preparation of the exercise, several tests were performed with the support of the CINECA staff
  - The main effort was dedicated to fully and efficiently exploiting the available dedicated resources. In the first version of the workflow, the

preparation of the simulation input files was done before submitting the run to the nodes. Due to the very short duration of the simulations in the exercise (because of the use of a small local domain), this strategy did not allow to exploit all of the dedicated resources, as the scheduler saturated at ~250 nodes;
    - Another effort was put for fractionating the ensemble of simulations in big jobs to be executed on big groups of nodes (i.e. 64 or 256 nodes) exposed to the risk of loss of big chunks of simulations in case of a crash of one node;
    - The output post-processing needed to be optimized not to waste time.
- **Lessons learned and recommendations**
    - A job array was the better strategy to fully and simultaneously exploit the 805 nodes available through the specific reservation;
    - Each element of the array was executed on a single node to minimize the failures;
    - The pre-processing operations performed outside of the nodes were been minimized to optimize the performances of the SLURM Workload Manager;
    - The post-processing operations were parallelized within the node through the use of the *srun* Slurm command, allowing to define multiple jobs in a co-scheduled heterogeneous job.

PD12. Buenos Aires VAAC volcanic ash exercise

PD12	High-resolution volcanic ash dispersal forecast
Leader	BSC
Participants	BSC / INGV / IMO
Workflow	PDAF + FALL3D
Numerical Engine	FALL3D (v8.1)
TRL initial	2
TRL target	5
TRL achieved	8-9

Description of the live exercise

Volcanic clouds jeopardise aerial navigation and operations. Volcanic aerosols and fine ash particles, angular in shape and highly abrasive, can damage turbine blades, airplane windscreens and fuselage, disrupt navigation instruments and, in the worst scenario, cause the clogging of cooling passages and potential engine stall. In 1993, the International Civil Aviation Organization (ICAO) established a global network of 9 Volcanic Ash Advisory Centers (VAACs) with the mission of monitoring and forecasting the location and trajectories of volcanic clouds occurring under their respective areas of responsibility. In the event of an eruption, the international civil aviation arrangements state that the affected VAAC must issue periodic Volcanic Ash Advisories (VAA), consisting of text messages including the forecasted ash polygons delineating unsafe flight areas. The VAAs are issued also in graphic form in the so called Volcanic Ash Graphics (VAG). In the past decade, this qualitative (ash/no-ash) approach, combined with the increase and congestion of global air traffic routes, has yield to some undesired overreaction situations causing millions in economic losses to airlines and stakeholders (see, e.g. the April-May 2010 aviation breakdown in Europe). At present, VAACs are under a profound transition (see the roadmap of the International Airways Volcano Watch -IAVW- for 2013-2037) to digital format for all volcanic ash information and provision of improved and more efficient volcanic hazard information services. This includes the development of the next generation of quantitative ash cloud forecasts and probabilistic (uncertainty) information as well as their integration into the System Wide Information Management (SWIM) environment. In fact, current (super)computers already allow solving this urgent computing scenario with the space-time model resolution requirements demanded by the aviation stakeholders.

VAACs perform periodic exercises in order to practise the conduct of volcanic activity and test new procedures before global implementation (see the ICAO Doc 9766, handbook on the International Airways Volcano Watch). In this regard, the Buenos Aires VAAC conducted an exercise on 10th December 2021 aimed at testing the new service “Quantitative Volcanic Ash” (expected to be operative by late 2023) among National Agencies, Meteorological Watch

Officers, air traffic Area Control Centers (ACCs), and the ICAO International NOTAM Office (the ICAO Office responsible for providing aeronautical information services). During the exercise, the ChEESE PD12 was run using in-house resources (MareNostrum-4 supercomputer) provided by the Barcelona Supercomputing Center (BSC). The objective of this exercise was to practice the response to reports of volcanic ash within the region of responsibility of the Buenos Aires VAAC in an operational environment and provide ensemble forecasts to the aviation community of ash cloud extent and movement. The routine Buenos Aires VAAC exercise in the South American airspace has provided the opportunity to validate the ChEESE volcanic ash PD12 as a service (Technology Readiness Level 7-8) and test its future operational deployment.

Identification of the Stakeholders and their roles

Stakeholders' institution	Role played
BSC	Run the PD12 @MN-4 during the live exercise
Argentinean Meteorological Service (Buenos Aires VAAC)	Run the exercise with their current operational setting (see Table 12.2)
Meteorological Watch Officers, air traffic Area Control Centers (ACCs), and the ICAO International NOTAM Office	End-users (decision makers)

Table 12.1. Stakeholders' involvement

Target Operational Environment

- **Description of the operational context (main actors, procedures, etc)**

  The operational urgent computing requirement was to deliver simulation results within a strict time window after the issue of each Volcanic Ash Advisory (VAA) or SIGMET message. Given the computational resources available for this exercise (48 dedicated nodes of the MN-4 supercomputer), the dispersal model was configured to ensure a timely execution with 48 ensemble members. The ensemble members were generated by perturbing eruption conditions notified in the VAAs and the underlying wind field. Table 12.2 compares the characteristics of the current Buenos Aires VAAC operational setting with those of the ChEESE PD12 as configured for this particular validation exercise.

- **Expected added value with respect to existing procedures**
  In terms of added value, two types of products fulfilling the IAVW requirements can be generated from PD12 ensemble-based model runs: (1) deterministic products giving a deterministic forecast based on some combination of the ensemble members (e.g. ensemble mean) and, (2) probabilistic products giving a probabilistic forecast based on the fraction of ensemble members that verify a certain condition, e.g. the probability of ash concentration exceeding a flight safety threshold.

HPC Challenge

	Current setup at BsAs VAAC	ChEESE PD setup (test)

Model	FALL3D-v7.2	FALL3D-v8.1 (ensemble-based, developed during ChEESE)
Computational domain	Area of the Buenos Aires VAAC: from 90ºW to 10ºW in longitude and from 60ºS to 10ºS in latitude (80ºx50º)	
Horizontal model resolution	0.25º	0.15º
Vertical model resolution	20 layers	50 layers
Number of ash bins	5	10
Ensemble members	1	48
Underlying meteorology	GFS forecasts (cycles 00, 06,12, 18)	
Hardware	Small cluster (10 cores)	48 nodes of MN-4 (2304 cores)
Time resolution (output)	Every 6 hours	Every 1 hour
Time-to-solution (target)	70 minutes (maximum) according to KPIs	The challenge was to deliver PD products in <25 min following each VAA
Deterministic products	- Column mass load - Concentration at Flight Levels (FL) from FL050 to FL350	- Ensemble-mean column mass load - Ensemble-mean concentration at Flight Levels (FL) from FL050 to FL350
Probabilistic products	No	- Probability to exceed column mass threshold (0.1 g/m^2) - Probability to exceed concentration thresholds (0.2, 2, 5 and 10 mg/m^3) at FLs from FL050 to FL350
Uncertainty estimation	No	- Based on the ensemble spread for column mass load and concentration

Table 12.2. Comparison between BsAs VAAC (operative) and ChEESE PD (tested) model setups.

Exercise Run (only for live exercises)

The exercise started 10th December 2021 at 11:00 UTC with the issue of a VAA reporting the start of an eruption. Successive VAAs were issued during the day to report changes on the eruption conditions (column height and emission) and, finally, the cease of eruptive activity. Each VAA triggered a forecast and the delivery of products containing the most updated information. Table 12.3 shows the exercise timeline for the current Buenos Aires VAAC setup and the tested ChEESE PD12.

Time (UTC)	Time (CET)	BsAs VAAC	ChEESE PD
11:00	**12:00**	Exercise starts	Exercise starts

11:02	**12:02**	The VAAC receives a message from local observatory stating that “According to GOES satellite observations, an eruption started at Cerro Negro volcano (Argentina) with an observed column height of 14 km a.s.l.	The BsAS VAAC communicates the eruption start message (via mail)
11:03	**12:03**		FALL3D ensemble run launched at MN-4 (using GFS cycle 00) with the eruption starting at 11:00 UTC
11:10	**12:10**	The VAAC posts the first VAA message and launches their dispersal model	
11:16	**12:16**		PD12 ensemble simulation ends and the post-process of results starts
11:21	**12:21**		Post-processed results (maps, 128MBy) successfully uploaded to a shared repository
			Total time to solution: 19 minutes
12:05	**13:05**	VAAC issues a second VAA and the first VAG with the forecasts (up to T+18)	
		Total time to solution: 55 minutes	
16:01	**17:01**	BsAs VAAC receives a new message saying that the eruption column height has decreased to 4 km a.s.l.	The BsAS VAAC communicates the changes in the eruption conditions (via mail)
16:02	**17:02**		A second FALL3D ensemble run launched at MN-4 (using GFS cycle 06) with eruption start time 11:00 UTC (14km) and a second pulse at 16:00 UTC (4km)
16:17	**17:17**		PD12 ensemble simulation ends and the post-process of results starts
16:22	**17:22**		Post-processed results (maps, 128MBy) successfully uploaded to a shared repository
			Total time to solution: 20 minutes
16:45	**17:45**	VAAC posts the third VAA and the second VAG with the updated forecasts	

		Total time to solution 44 minutes	
19:00	**20:00**	VAAC receives a new message saying that the eruption has ceased	The BsAS VAAC communicates than the eruption has ended (via mail)
19:01	**20:01**		A third FALL3D ensemble run launched at MN-4 (using GFS cycle 12) with eruption start time 11:00 UTC (14km) and a second pulse at 16:00 UTC (4km) lasting for 3h
19:16	**20:16**		PD12 ensemble simulation ends and the post-process of results starts
19:21	**20:21**		Post-processed results (maps, 128MBy) successfully uploaded to a shared repository
			Total time to solution: 20 minutes
19:45	**20:45**	VAAC posts the forth VAA and the third VAG with the updated forecasts	
		Total time to solution: 45 minutes	
23:00	**00:00**	VAAC issues and eruption end message	
		End of the exercise	End of the exercise

Table 12.3. Exercise timeline

The exercise consisted therefore on 3 different forecasts, all successfully delivered within the pre-defined time-to-solution limits (Table 12.4).

		BsAs VAAC time to solution (min)	ChEESE PD12 time to solution (min)
Forecast 1	First VAG posted	55	19
Forecast 2	Second VAG posted	44	20
Forecast 3	Third VAG posted	45	20

Table 12.4. Time to solution for each forecast.

For illustrative purposes, Figure 12.1 shows some ChEESE PD12 results (1a for ensemble-base deterministic and 1b for ensemble-based probabilistic results). Figures 12.2 to 12.4 show the comparison of ChEESE PD12 results with the first, second and third VAG (VAAC forecasts) respectively.

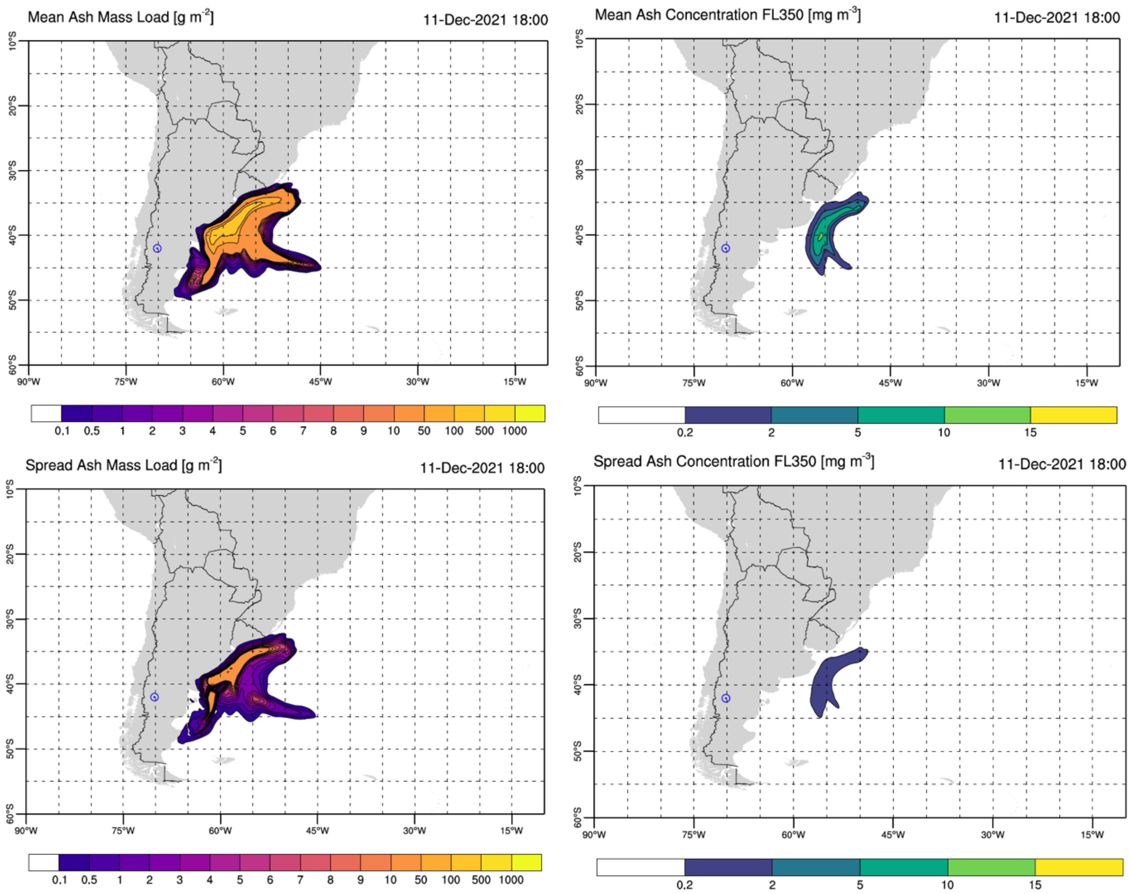

Figure 12.1a. Example of ChEESE PD12 ensemble-based deterministic results. Left: ensemble-mean column mass load in g/m^2 (top) and associated uncertainty (bottom). Right: ensemble-mean concentration in mg/m^3 (top) at flight level FL350 and associated uncertainty (bottom). Results forecasted for 11 Dec at 18:00 UTC.

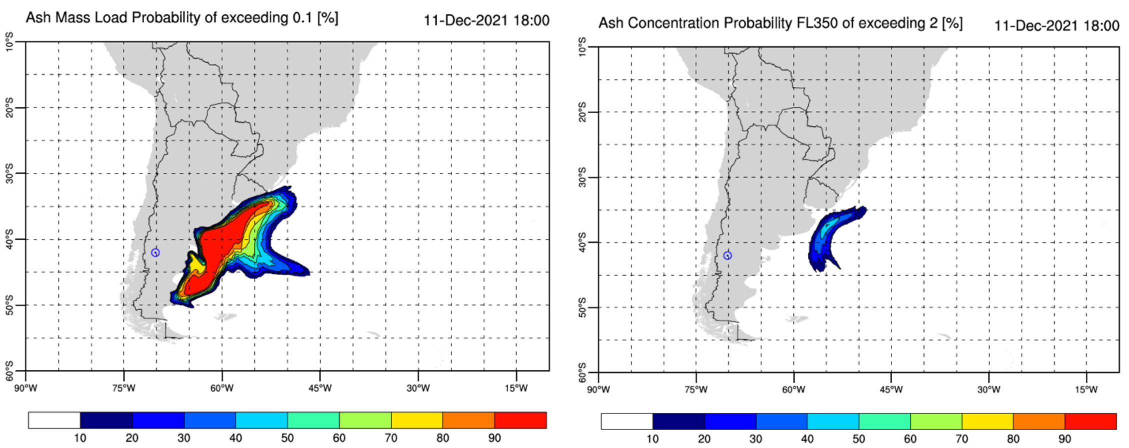

Figure 12.1b. Example of ChEESE PD12 ensemble-based probabilistic results. Left: probability of column mass exceeding 0.1 g/m^2. Right: probability of concentration exceeding 2 mg/m^3 at FL350. Results forecasted for 11 Dec at 18:00 UTC.

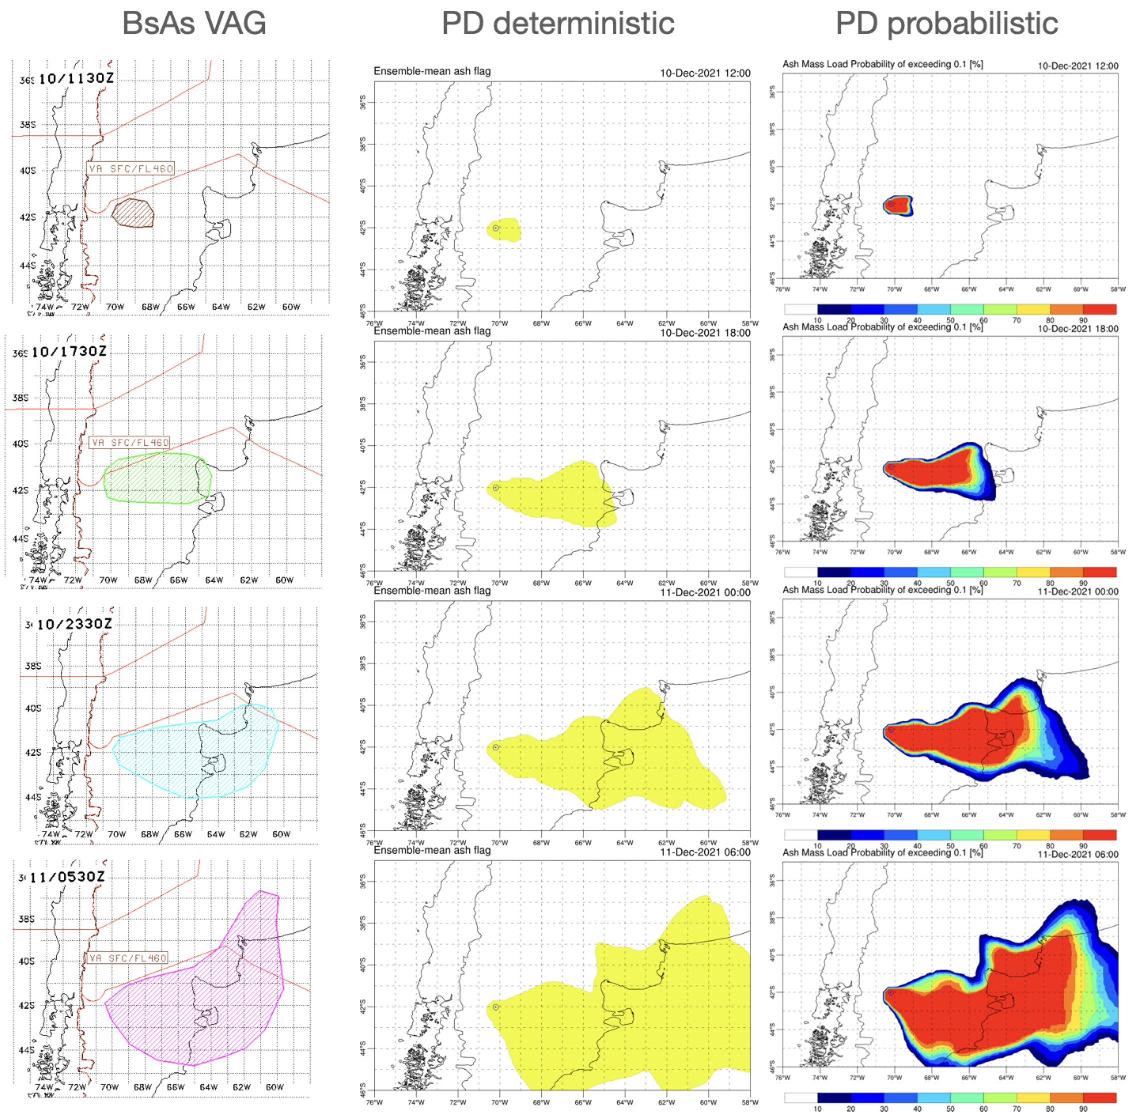

Figure 12.2. First simulation. Left: First BsAs Volcanic Ash Graphic (VAG) issued at 12:05 UTC. Forecasted times correspond to T (11:30 UTC), T+6 (17:30 UTC), T+12 (23:30 UTC) and T+18 (05:30 UTC). Contours delineate the qualitative ash/no ash polygons. Center: ChEESE PD12 deterministic product based on the column mass ensemble mean. The contour of 0.1 g/m^2. Results at similar time instants. Right: ChEESE PD12 probabilistic product based on the column mass ensemble mean. Contours show the probability (in %) to exceed 0.1 g/m^2.

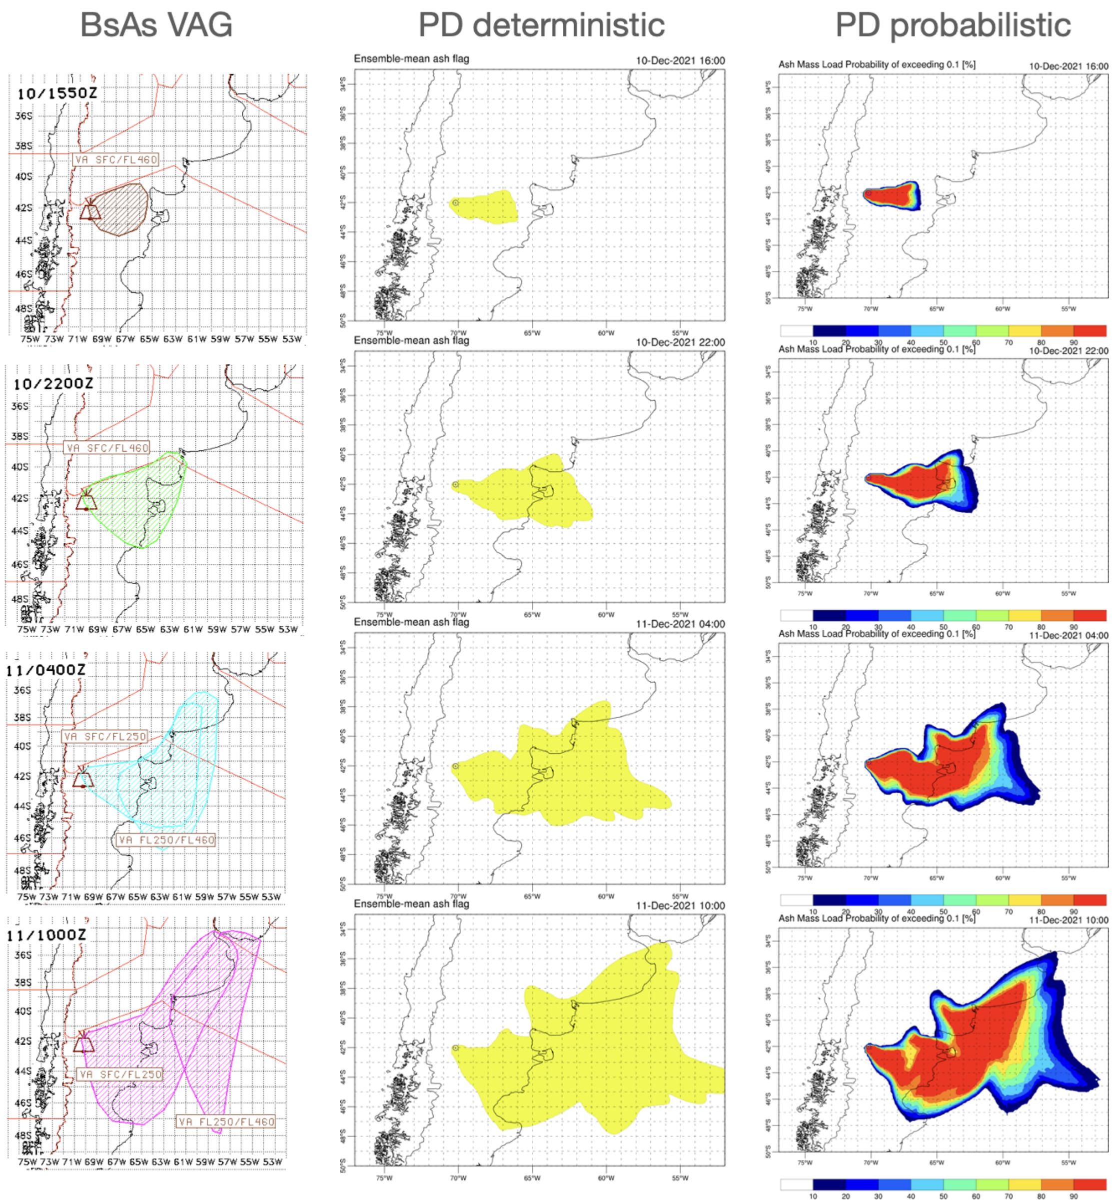

Figure 12.3. Second simulation. Left: Second BsAs Volcanic Ash Graphic (VAG) issued at 16:45 UTC. Forecasted times correspond to T (15:50 UTC), T+6 (22:00 UTC), T+12 (04:00 UTC) and T+18 (10:00 UTC). Contours delineate the qualitative ash/no ash polygons. Center: ChEESE PD12 deterministic product based on the column mass ensemble mean. The contour of 0.1 g/m^2. Results at similar time instants. Right: ChEESE PD12 probabilistic product based on the column mass ensemble mean. Contours show the probability (in %) to exceed 0.1 g/m^2.

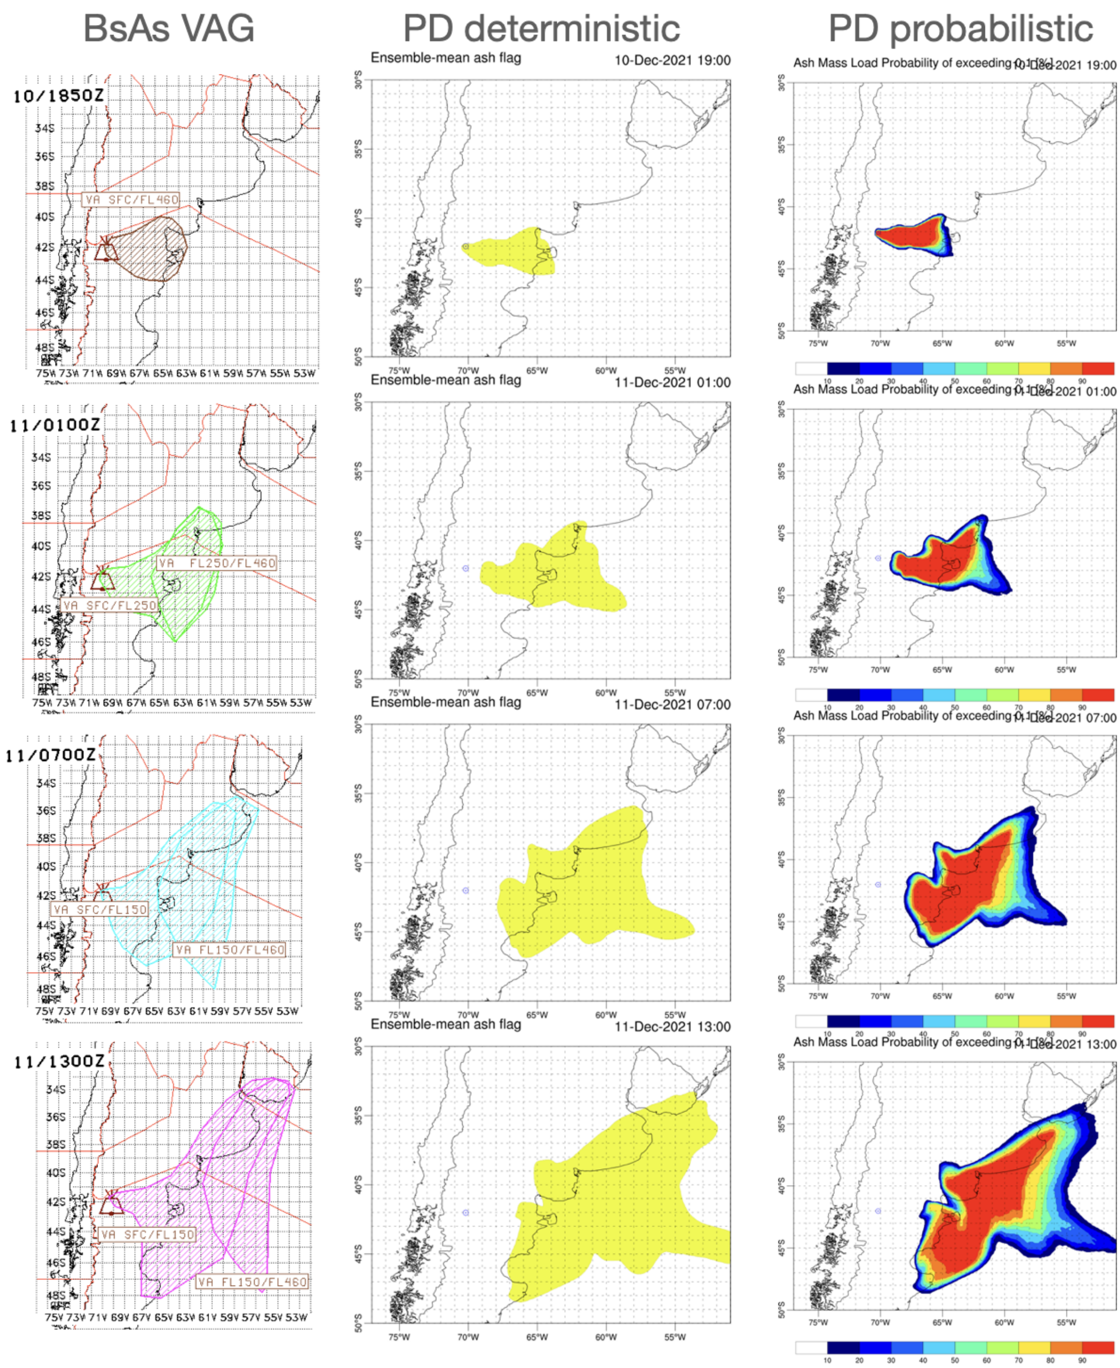

Figure 12.4. Third simulation. Left: Third BsAs Volcanic Ash Graphic (VAG) issued at 19:45 UTC. Forecasted times correspond to T (18:50 UTC), T+6 (01:00 UTC), T+12 (07:00 UTC) and T+18 (13:00 UTC). Contours delineate the qualitative ash/no ash polygons. Center: ChEESE PD12 deterministic product based on the column mass ensemble mean. The contour of 0.1 g/m^2. Results at similar time instants. Right: ChEESE PD12 probabilistic product based on the column mass ensemble mean. Contours show the probability (in %) to exceed 0.1 g/m^2.

Conclusions and potential target service

When compared with the current BsAs VAAC operational products (see Table 12.2), the ChEESE PD12 forecasts as set for this particular exercise have:

- Higher space resolution (0.15º vs 0.25º and 50 vs 20 vertical layers)
- Higher output frequency (1h vs 6h)
- Double number of ash bins (10 vs 5)
- Ensemble versus single-run forecast (48 vs 1 member)
- Better time to solution (20 vs 45-55 min)
- An array of quantitative products (including probabilistic forecasts)

In the context of pre-exascale, this urgent computing exercise has been done with "limited computational resources" (2304 MN-4 cores only). However, **results have shown that the setup of an urgent computing HPC access mode on current (and upcoming) supercomputers would allow delivering km-resolution forecast products with hundreds of ensemble members.** It is expected that around 256 members should be sufficient to account for the typical uncertainties accompanying ash forecasts.