ChEESE: The European Union prepares to enter the exascale era

"*Today, your cell phone has more computer power than all of NASA back in 1969 when it sent two astronauts to the moon. [...] The Sony Playstation of today, which costs $300, has the power of a military supercomputer of 1997, which cost millions of dollars.*"
[Michio Kaku, Theoretical Physicist]

As the American computer scientist and essayist Ray Kurzweil shows, technological progress is an exponential growth process. If we had to describe the growth curve of computing capacity from the first programmable electronic computer (the *Colossus*, used by the British during World War II to decipher Nazi messages) to today's supercomputers, we certainly could not use a linear relationship. In the recent decades we have witnessed an incredible technological acceleration in the field of **HPC** (**H**igh **P**erformance **C**omputing or supercomputing), which has allowed us to achieve computing capabilities once unthinkable. Let's take the computing performance, which is generally expressed in **FLOPS** (**FL**oating **P**oint **O**perations **P**er **S**econd), i.e. the number of floating-point calculations that a computer can perform in one second: if 2010 saw the advent of **petascale** computing (10^{15} FLOPS, or one million of billions FLOPs), 2020 marked the beginning of the **exascale** era, characterized by computers that are capable of reaching 10^{18} FLOPS (or one billion of billions FLOPs, i.e. a thousand times more). The technological efforts behind these incredible numbers lie not only in hardware development, but also in software development: a computing infrastructure perfectly capable of sustaining a very large number of calculations is in fact useless if the implemented algorithms are inefficient or, even worse, unable to handle the specific characteristics of HPC computing. But why do we need exascale computing? For two fundamental reasons: 1) current computational capacity is not sufficient to solve certain scientific problems; 2) current computation times are incompatible with certain specific applications. In the remainder of this article we will look at some examples of the scientific challenges for which we need exascale machines.

The European HPC system

While United States and China contend for the scepter of the fastest supercomputer in the world, Italy holds - as of March 2021 - a very respectable position with two supercomputers among the top 11 (TOP500 List): ENI's **HPC5**, in the 8th position (but 6th if we look at the Green500 List, being one of the most sustainable in the world), and CINECA's **Marconi-100,** in the 11th position. For its part, the European Union is deploying ambitious strategies and substantial budgets, and is determined to play an increasingly important role in the development of high-performance computing.
In 2018, **EuroHPC**, the **European High Performance Computing Joint Undertaking**, was established with the dual objective of developing a supercomputing infrastructure and funding research and innovation activities. The budget earmarked for this undertaking is **1400 million euros** for the 2019-2020 biennium alone, a large sum that testifies the European Union's interest in this technological challenge. The European supercomputing infrastructure, which currently includes 7 HPC systems, will be expanded to include 3 petascale machines by 2021 (one of them in Italy, at CINECA) and 2 exascale machines by 2023 (the first one is scheduled for 2022). Infrastructure and technology are not, however, the only pillars this European vision is based on: a third and fundamental pillar is that of applications, represented by **14 European Centers of Excellence**, that

are the real protagonists of technological progress and consolidation of European leadership in the field of supercomputing.

The development of the European HPC-ecosystem was fostered by **PRACE** (**P**artnership for **A**dvanced **C**omputing in **E**urope), supported by the PRACE Members and by the EU through a series of implementation projects. PRACE offers a persistent pan-European Research Infrastructure of world-class supercomputers in all major architectural classes, a comprehensive range of services, support activities, and programmes to foster the European HPC ecosystem.

Figure 1: The PRACE infrastructure in its current state (prace-ri.eu)

ChEESE

ChEESE (https://cheese-coe.eu/) is a consortium of 13 partners from 7 European countries (Italy participates with INGV and CINECA), coordinated by the Barcelona Supercomputing Centre (BSC). The goal of this European project is to create an European **Center of Excellence in the domain of Solid Earth**, capable of fully exploiting the new petascale (2020) and exascale (2022) technologies.

Specifically, ChEESE has the following objectives:

- To prepare 10 European codes for the emerging exascale architecture that will allow to reproduce a given physical process through numerical simulations;
- To address 15 scientific challenges in the domain of Solid Earth which require high performance computing;
- To provide 12 pilot demonstrators (PD), i.e. pilot examples of code applications to Solid Earth problems, in order to demonstrate the utility of supercomputing in building useful services for the society (e.g., seismic, or volcanic, or tsunami hazards);
- To integrate institutions, academia, research, and private industry into a common goal.

The 15 scientific challenges that researchers within ChEESE propose to address are not only very ambitious but can also have a great impact on society. They are in fact related to the study of earthquakes, tsunamis, volcanoes and geomagnetic field and are organized in the following services:

1) **Urgent computing** (simulations in emergency conditions, for which it is necessary to drastically reduce the computation time in order to provide timely services); 2) **Probabilistic hazard estimation** (estimation of the hazard in the short or longer term of a phenomenon in progress or which has a given probability of occurring); 3) **Early warning**; 4) **Seismic tomography for the characterization of the subsurface**.

The role of INGV

Within the ChEESE project, INGV plays a major role, with more than 20 researchers involved in the development of 2 numerical simulation codes and 7 applications to Solid Earth problems, in the creation of a web portal which serves as a repository for codes and datasets, in the definition of strategies for the implementation of ChEESE services and in the inclusion of several products, software and services in EPOS-ERIC (https://www.epos-eu.org/). In the long term, these developments are expected to significantly improve the seismic and volcanic surveillance, tsunami monitoring and warning services (http://www.ingv.it/cat/it/), as well as the short- and long-term hazard estimates that INGV provides to the Italian Civil Protection (see also https://cheese-coe.eu/media/news/research-societal-relevance-how-cheese-and-urgent-computing-may-enhance-ingvs-hazard).

Volcanoes

As for the applications to volcanoes within ChEESE, INGV is leading a PD aiming to develop the **ASHEE** code for the simulation of dynamics of explosive eruptions. The optimization of the computational code and workflows will increase by an order of magnitude the spatial accuracy of numerical simulations, yielding a better physical description of multiphase processes and a better multiscale analysis of turbulent dispersion processes. On the other hand, the decrease in computational time will also enable 3D models to be used probabilistically, using *ensembles* of simulations to assess the forecast uncertainty.
As for the volcanic hazard analysis, INGV is also leading a PD that illustrates the usefulness of exascale computation in workflows for volcanic ash (falling to the ground or suspended in the atmosphere) impact estimates simulated with the **FALL3D** code. The illustrative example is for the volcanoes Campi Flegrei in Naples, Italy, and Jan Mayen, a volcanic island located in the North Atlantic. The ashfall hazard estimates are calculated for the short term (i.e. by running a few hundred simulations that consider the meteorological wind prediction for the following few days) and for the long term (i.e., for the following few decades), by running thousands of simulations that consider the climatological wind statistics. In both cases, the exascale calculation allows the simulations to be run in a reasonable time and, in the case of the short term, to be run in a time that is short enough for risk mitigation purposes. In addition, exascale computing allows simulations to be performed over very large areas, reaching thousands of km, at a resolution of about 1 km (see, for example, Figure 2). This strongly improves the capability to simulate, and thus to assess, the hazard associated with low-probability high-impact events.

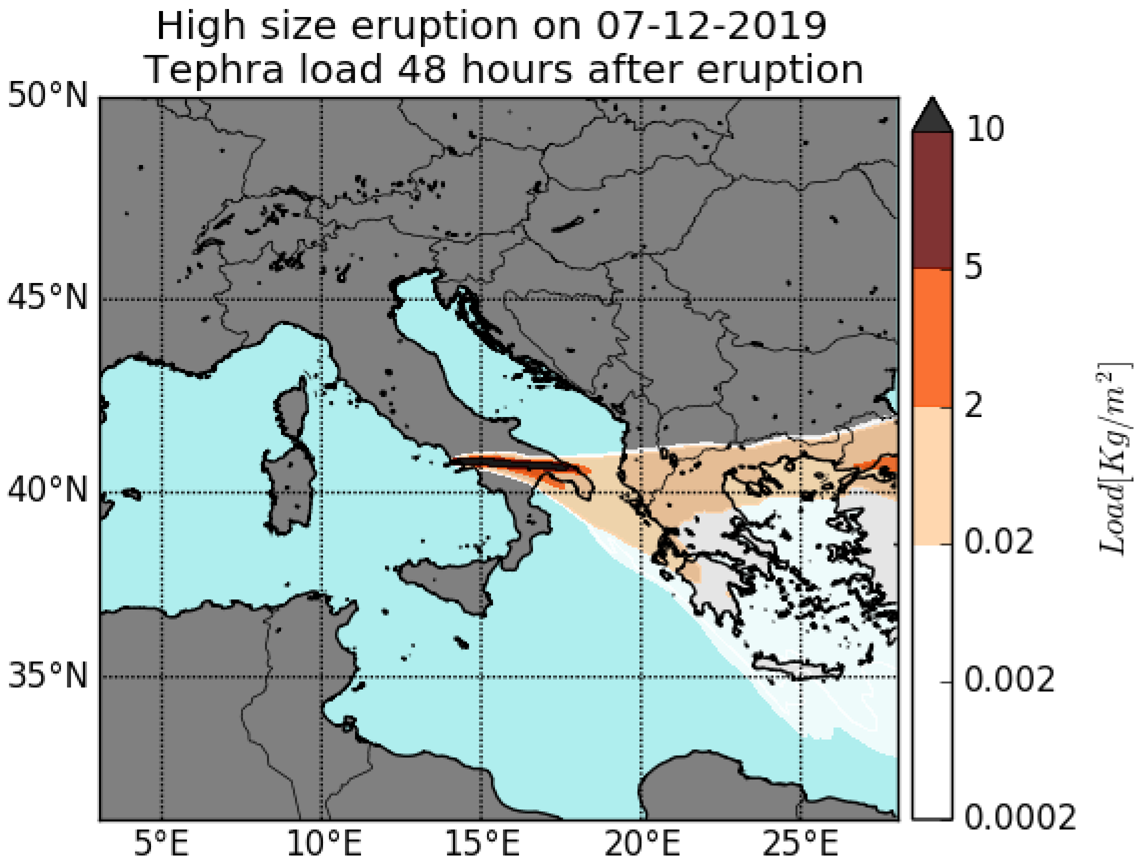

Figure 2: Example of hazard estimation of ground-accumulated ash from a highly explosive eruption at Campi Flegrei, calculated using wind forecasts from December 7, 2019 for the following 48 hours, over a 2000 km x 2000 km domain

Earthquakes

INGV researchers are engaged in the development of ChEESE applications related to urgent computing - in case of strong earthquakes - and seismic tomography. In the first case, the goal is to implement exascale computing resources to realize highly accurate shaking maps ("**shakemaps**") based on physical models within a very short time (hours) from the occurrence of a medium-to-large magnitude earthquake. Currently, these maps are based on empirical relationships that can be improved with numerical simulations, in order to better represent the complex uniqueness of each seismic event. Fostering the development of urgent computing applications, also by means of real-time data assimilation, is of crucial importance, since the timely and accurate estimation of the expected shaking distribution represents an essential support for rescue operations. A very effective tool related to the *shakemaps* is represented by the *shakemovies*, which allow to efficiently visualize the expected shaking in areas at different distances from the epicenter. Figure 3 shows an example, which is related to the recent seismic event in Croatia.

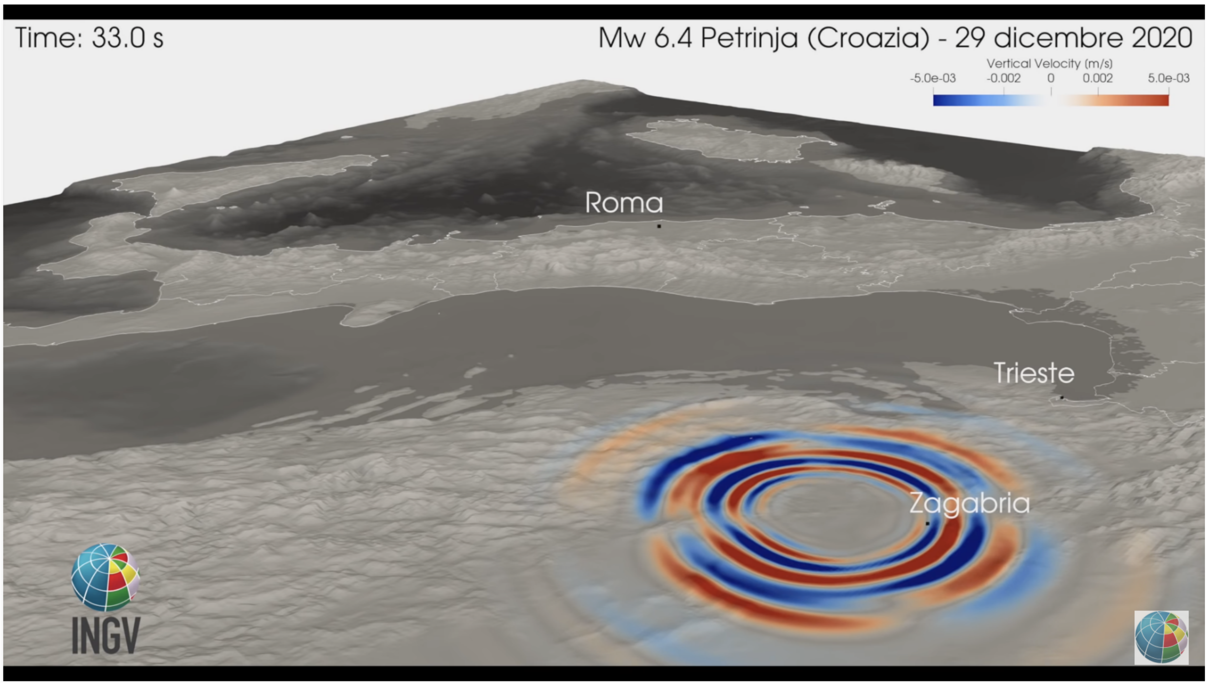

Figure 3: Excerpt from INGVterremoti video showing the simulation of propagation of seismic waves generated by the Mw 6.4 Croatia earthquake, occurred on December 29, 2020. The blue color indicates downward ground motions, while the red color indicates upward ground motions. The color intensity represents different values of vertical velocity (m/s): the more intense the blue (red), the faster the ground moves downward (upward).

A very important aspect that emerges from simulations like this is the inhomogeneous distribution of ground velocities: points equidistant from the epicenter do not experience the same shaking, due to local conditions (topography, soil type, ...) that can significantly affect the propagation of seismic waves. This is one of the main reasons why it is necessary to obtain shaking estimates that are based on realistic physical models.

To this end, INGV researchers are involved in another key application related to the study of earthquakes, namely seismic tomography. This **imaging** technique (which is very similar to X-rays or ultrasound techniques in the medical field) allows to produce realistic and physically plausible models of the subsurface. Higher quality subsurface models yield better simulations of seismic wave propagation and, in turn, more accurate estimates of ground shaking.

Tsunamis

Regarding the study of tsunamis within the ChEESE project, INGV is involved in two PDs for short-term - thus immediately after the occurrence of a seismic event - and long-term hazard assessment.

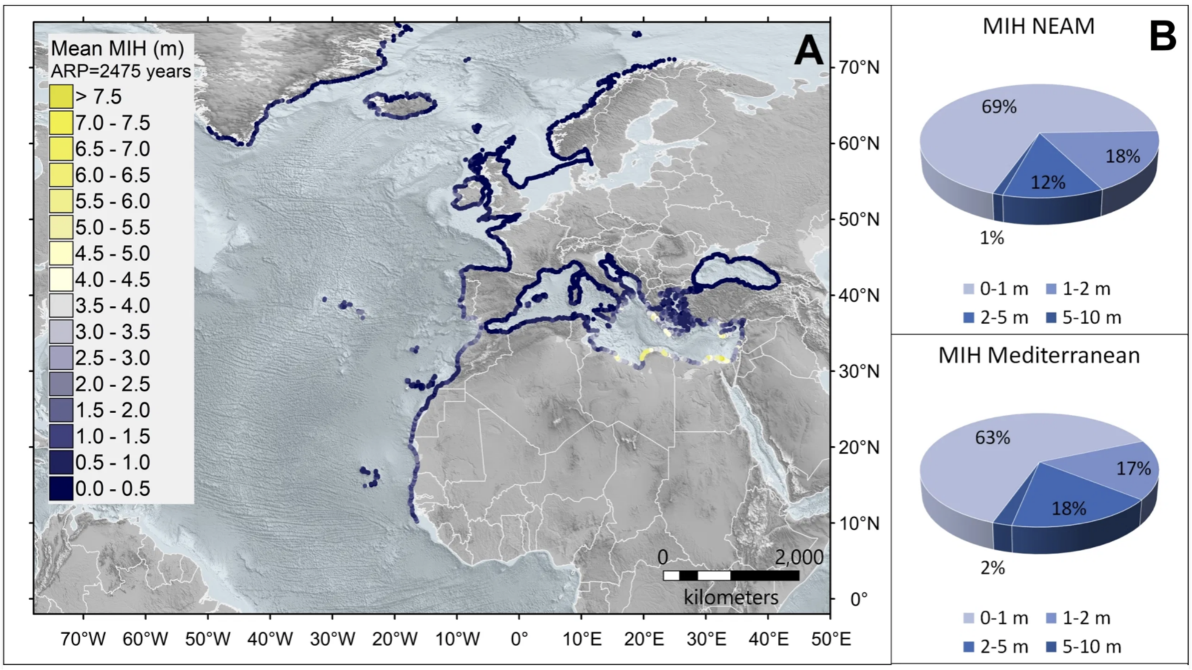

Figure 4: One of the outputs of NEAM Tsunami Hazard Model 2018 (NEAMTHM18), the PTHA model related to the Mediterranean, NE Atlantic and connected areas. The map shows the MIH ("Max Inundation Height") expected for an average return time of 2475 years.

The first PD concerns the **Probabilistic Tsunami Hazard Assessment (PTHA)**. PTHA methodologies allow the assessment of the probability that a given measure of tsunami intensity (e.g., maximum run-up height) will be exceeded at a particular location in a given time period. Behind this seemingly simple sentence lies an extremely complex problem. It is in fact necessary to combine a very large number of parameters and a variety of different types of data to estimate how their uncertainty propagates into hazard models. In fact, the variability of the input parameters - such as the existence of seismogenic faults and their typology, the frequency of earthquakes of a certain magnitude on a given fault, the uncertainties on bathymetry and topography - determines an extreme variability in the plausible tsunami scenarios and in their probability of occurrence. Currently, computational resources place limits on the ability to explore the total uncertainty and to model local-scale tsunami propagation, in which case the complexity of the process and the need for high-resolution modelling inherently require extremely high computational time and resources. Exascale computers would therefore allow researchers to overcome such limitats, thus improving their ability to return more physically plausible scenarios and, at the same time, to quantify the total uncertainty.

The second PD, coordinated by INGV, is related to **early warning** and **post-disaster assessment** (rapid coastal impact estimation). Since the propagation speed of tsunami waves is much lower than that of seismic waves, it is possible to use seismic recordings to estimate the impact of the expected tsunami on the coast within a reasonable time. To do this, however, it is necessary to take into account the uncertainty about the seismic source that triggered the tsunami: in fact, in the first minutes after an earthquake, the seismic source is only known approximately. Exascale computing would enable researchers to explore in a very short time a large number of seismic sources that are compatible with the first-minutes observations. Furthermore, it would yield a better, faster and more detailed estimation of the expected impact on the coast, while accurately quantifying the uncertainty.