ChEESE

www.cheese-coe-eu

Center of Excellence for Exascale in Solid Earth

# FTRT tsunami simulations: challenges & solutions towards HPC

Jorge Macías et al
**EDANYA Research Group**
**Universidad de Málaga**

EGU 2020 - NH5.1 - Tsunamis

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

# What is ChEESE?

**C**enter of **E**xcellence for **E**xascale in **S**olid **E**arth

## Project information

### ChEESE

**Grant agreement ID: 823844**

Status
**Ongoing project**

| Start date | End date |
|---|---|
| **1 November 2018** | **31 October 2021** |

Funded under:
**H2020-EU.1.4.1.3.**

Overall budget:
€ 7 683 241,25

**EU contribution**
**€ 7 683 241,25**

Coordinated by:
**BARCELONA SUPERCOMPUTING CENTER - CENTRO NACIONAL DE SUPERCOMPUTACION**
**Spain**

# Objectives

1. Prepare 10 Community flagship European codes for the upcoming pre-Exascale (2020) and Exascale (2022) supercomputers.
2. Address 15 scientific, technical, and socio-economic Exascale Computational Challenges (ECC) in the domain of Solid Earth.
3. Develop 12 Pilot Demonstrators and enable services oriented to society on critical aspects of geohazards like hazard assessment, urgent computing, and early warning forecast.
4. Integrate around HPC and HDA transversal European institutions in charge of operational geophysical monitoring networks, Tier-0 supercomputing centers, academia, hardware developers, and third-parties from SMEs, Industry and public governance bodies (civil protection).

# ChEESE Consortium

- 13 partners from 4 EU Members States (Spain, Italy, Germany, France) and 3 Associated Countries (Norway, Switzerland, Iceland).

Among them

- 3 Tier-0 supercomputing centers (BSC, CINECA, HLRS).
- 3 institutions in charge of Geophysical monitoring networks in 3 European countries (INGV, IMO, IPGP).

# x Industry and Users Board

GEO GROUP ON EARTH OBSERVATIONS

GSNL

EPOS EUROPEAN PLATE OBSERVING SYSTEM

PROTEZIONE CIVILE
Presidenza del Consiglio dei Ministri
Dipartimento della Protezione Civile

REPSOL

GW4 Isambard

GOBIERNO DE ESPAÑA
MINISTERIO DE FOMENTO
INSTITUTO GEOGRÁFICO NACIONAL

Schlumberger

IASPEI
International Association of Seismology and Physics of the Earth's Interior

RÍKISLÖGREGLUSTJÓRINN
ALMANNAVARNADEILD

IAVCEI
International Association of Volcanology and Chemistry of the Earth's Interior

MITIGA SOLUTIONS

# 10 Flagship Codes

FLAGSHIP CODES

| Code name | Area | Application |
|---|---|---|
| **ExaHyPE** | Computational Seismology | Earthquake simulations in highly heterogeneous media |
| **Salvus** | Computational Seismology | Wave propagation through complex unstructured domains, seismic tomography, medical ultrasound tomography |
| **SeisSol** | Computational Seismology | Earthquake simulations in various settings including induced earthquakes |
| **SPECFEM3D** | Computational Seismology | Earthquake simulations in various settings, imaging of complex geological objects |
| **PARODY_PDAF** | Magnetohydrodynamics | Geodynamo simulations and ensemble assimilation experiments |
| **XSHELLS** | Magnetohydrodynamics | Geodynamo simulations, liquid metal planetary core simulations |
| **ASHEE** | Physical Volcanology | Volcanic plume and PDC simulator |
| **FALL3D** | Volcanology | Forecast of volcanic ash clouds, volcanic ash fallout and volcanic ash resuspension, tephra dispersal and fallout hazard assessment |
| **T-HySEA** | Tsunamis | Faster-than-real-time (FTRT) simulations for Tsunami Early Warning Systems (TEWS), Probabilistic Tsunami Hazard Assessment (PTHA), tsunami inundation maps |
| **L-HySEA** | Tsunamis | Subaerial and submarine granular landslides, landslide generated tsunamis |

# 15 Exascale Computational Challenges (ECC)

**In Computational Seismology**

- Full-waveform inversion (FWI)
- High-resolution subsurface imaging
- Near real-time seismic scenarios
- Physics-based Probabilistic Seismic Hazard Analysis (PSHA)

**In Magnetohydrodynamics (MHD)**

- The dynamo model
- Earths magnetic field evolution

**In Physical Volcanology**

- Volcanic plumes and Pyroclastic Density Currents (PDCs)
- High-resolution volcanic ash dispersal
- Subsurface thermo-fluid dynamics of magmas
- Probabilistic Volcanic Hazard Analysis (PVHA)

**In Tsunami Modelling**

- Faster Than Real Time (FTRT) tsunami computations
- Near real-time tsunami source inversion
- Probabilistic Tsunami Forecast (PTF) for early warning and rapid post-event assessment
- Probabilistic Tsunami Hazard Analysis (PTHA)

**In Observational Seismology**

- Automated array-based statistical detection and restoration of seismic slow-earthquakes

# Motivation. FTRT and TEWS for the NEAM region

**Aim:** Focus in a very specific problem

Achieving much FTRT tsunami simulations in the context of TEWS for the NEAM

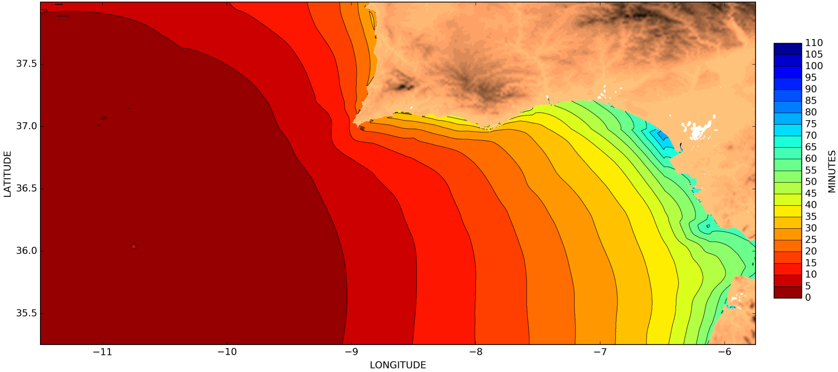


**NEAM** stands for North Eastern Atlantic and Mediterranean

# Motivation. FTRT and TEWS for the NEAM region

**Aim:** Focus in a very specific problem

Achieving much FTRT tsunami simulations in the context of TEWS for the NEAM

**As modelers / Numerical specialists**

- Developing efficient numerical tools to simulate tsunamis
- Extremely robust codes
- Need to compute extremely fast (if aim is saving lives)
- This (computational tools in TEWS) was UNTHINKABLE some years ago

# How we do it

## Two Ingredients

### 1. Numerical model: Tsunami-HySEA

- Robust
- Efficient
- Precise
- Validated


### 2. GPU and multi-GPU

- Extremely fast computing (and inexpensive)

# Motivation. The state of the art

## How TEWS in NEAM region do work

- Decision Matrices
- Precomputed Databases

# The result

## A novel approach

**Computing in Real Time**

- **2018 NVIDIA Global Impact Award**

### Global Impact Award Finalist Using GPUs with Aim to Spare Lives Ahead of Tsunamis

March 12, 2018 by TONIE HANSEN

---

**The University of Málaga team advances capabilities of tsunami early warning systems.**

## The rules of the game have changed

# 1. The Mediterranean challenge (by INGV)

## INGV. A TEWS for all the Mediterranean

- Computational domain: the whole Mediterranean
- Spatial resolution: 30 arc-sec.
- Size of the problem: $5,221 \times 1,921 = 10,029,541$ cells
- Simulation time: 8 hours

# 1. The Mediterranean challenge (by INGV)

## Output

- Times series at 17,000 predefined locations (PoIs)
- Maximum height in all the domain

# 1. The Mediterranean challenge - 2014

## The Challenge:

**Do it in less than 6 min!!!**
(when around ten hours were required)

# 1. The Mediterranean challenge - 2014

## The Challenge:

**Do it in less than 6 min!!!**

(when around ten hours were required)

## Using on single GPU (Titan Black)

**35 min 25 s!!!**

(really good but still far from away the challenge)

# 1. The Mediterranean challenge - 2014

## The Challenge:

**Do it in less than 6 min!!!**
(when around ten hours were required)

## Using on single GPU (Titan Black)

**35 min 25 s!!!**
(really good but still far from away the challenge)

## Multi-GPU implementation

**More than one single GPU can be used**
(2 GPUs →19 min)

# Times for the Mediterranean test case

Table: **2014** Computing times and speed-up - nVIDIA Titan Black GPUs (Kepler, 2012)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 2141.1 (35 min 41 s) | 1.00 |
| 2 | 1139.5 (18 min 59 s) | 1.88 |
| 4 | 601.3 (10 min 1 s) | 3.56 |
| 8 | 378.1 (6 min 18 s) | 5.66 |
| 10 | 352.0 (5 min 52s) | 6.08 |

**Requirement: computing time $<$ 6 min**

# Times for the Mediterranean test case

Table: **2014** Computing times and speed-up - nVIDIA Titan Black GPUs (Kepler, 2012)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 2141.1 (35 min 41 s) | 1.00 |
| 2 | 1139.5 (18 min 59 s) | 1.88 |
| 4 | 601.3 (10 min 1 s) | 3.56 |
| 8 | 378.1 (6 min 18 s) | 5.66 |
| 10 | 352.0 (5 min 52s) | 6.08 |

**Requirement: computing time $<$ 6 min**

## Continuous improvements (2015-2017)

- Static load balancing
- CFL adjustment
- Writing while computing
- Overlapping of processes

# Times for the Mediterranean test case

**2017 Computing times and speed-up***

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 1,764.0 (29 min 24 s) | 1.00 |
| 2 | 908.6 (15 min 9 s) | 1.94 |
| 4 | 507.8 (8 min 28 s) | 3.47 |
| 8 | 312.1 (5 min 12 s) | 5.65 |
| 12 | 259.0 (4 min 19 s) | 6.81 |

**Requirement: computing time $<$ 6 min**

* Times for nVIDIA Titan Black GPUs (Kepler, 2012). 1 Gb ethernet network

# Times for the Mediterranean test case

## Comparison Computing times 2014 vs 2017

| # GPUs | 2014 | 2017 |
|---|---|---|
| 1 | 35 min 41 s | 29 min 24 s |
| 2 | 18 min 59 s | 15 min 9 s |
| 4 | 10 min 1 s | 8 min 28 s |
| 8 | 6 min 18 s | 5 min 12 s |
| 10/12 | 5 min 52s | 4 min 19 s |

**Requirement: computing time $<$ 6 min**

* Times for nVIDIA Titan Black GPUs (Kepler, 2012). 1 Gb ethernet network

# Times for the Mediterranean test case

## 2017 Computing times and speed-up (Kepler, 2012)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 1,764.0 (29 min 24 s) | 1.00 |
| 2 | 908.6 (15 min 9 s) | 1.94 |
| 4 | 507.8 (8 min 28 s) | 3.47 |
| 8 | 312.1 (5 min 12 s) | 5.65 |
| 12 | 259.0 (4 min 19 s) | 6.81 |

**But also new architectures (2018)...**

**2 NVIDIA Tesla P100** ... (already "obsolete")

**(Pascal architecture 2016)**

# Times for the Mediterranean test case

## 2017 Computing times and speed-up (Kepler, 2012)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 1,764.0 (29 min 24 s) | 1.00 |
| 2 | 908.6 (15 min 9 s) | 1.94 |
| 4 | 507.8 (8 min 28 s) | 3.47 |
| 8 | 312.1 (5 min 12 s) | 5.65 |
| 12 | 259.0 (4 min 19 s) | 6.81 |

**But also new architectures...**

**2 NVIDIA Tesla P100** - 257 sec "obsolete"!!!

(Pascal architecture 2016)

# Times for the Mediterranean test case

**2017 Computing times and speed-up (Kepler, 2012)**

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 1,764.0 (29 min 24 s) | 1.00 |
| 2 | 908.6 (15 min 9 s) | 1.94 |
| 4 | 507.8 (8 min 28 s) | 3.47 |
| 8 | 312.1 (5 min 12 s) | 5.65 |
| 12 | 259.0 (4 min 19 s) | 6.81 |

**But also new architectures...**

**2 NVIDIA Tesla P100** - 257 sec "obsolete"!!!

**(Pascal architecture 2016)**

**And, if this one it is obsolete, then what if...**

# Times for the Mediterranean test case

## 2017 Computing times and speed-up (Kepler, 2012)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 1,764.0 (29 min 24 s) | 1.00 |
| 2 | 908.6 (15 min 9 s) | 1.94 |
| 4 | 507.8 (8 min 28 s) | 3.47 |
| 8 | 312.1 (5 min 12 s) | 5.65 |
| 12 | 259.0 (4 min 19 s) | 6.81 |

## Using one single NVIDIA Tesla V100 (2018)

# Times for the Mediterranean test case

**2017 Computing times and speed-up (Kepler, 2012)**

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 1,764.0 (29 min 24 s) | 1.00 |
| 2 | 908.6 (15 min 9 s) | 1.94 |
| 4 | 507.8 (8 min 28 s) | 3.47 |
| 8 | 312.1 (5 min 12 s) | 5.65 |
| 12 | 259.0 (4 min 19 s) | 6.81 |

## Using one single NVIDIA Tesla V100 (2018)

Newest NVIDIA Tesla V100 (2018) - **1 GPU - 284.14 sec**

**(Volta architecture 2018, released end 2017)**

# Times for the Mediterranean test case

## 2017 Computing times and speed-up (Kepler, 2012)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 1,764.0 (29 min 24 s) | 1.00 |
| 2 | 908.6 (15 min 9 s) | 1.94 |
| 4 | 507.8 (8 min 28 s) | 3.47 |
| 8 | 312.1 (5 min 12 s) | 5.65 |
| 12 | 259.0 (4 min 19 s) | 6.81 |

**Using one single NVIDIA Tesla V100 (2018)**

1 GPU - 284.14 sec

**But we can use more GPUs!!!**

# Times for the Mediterranean test case

**Using one single NVIDIA Tesla V100 (2018)**

**1 GPU - 284.14 sec**

**But we can use more GPUs!!!**

**4 GPU - 97.20 sec**

# Times for the Mediterranean test case

2018 Computing times and speed-up (Volta, 2018)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 284.14 (4 min 44 s) | 1.00 |
| 4 | 97.20 (1 min 37 s) | 2.92 |
| 8 | 66.54 s | 4.27 |
| 16 | 50.09 s | 5.67 |
| 32 | 38.24 s | 7.43 |

Requirement: computing time $<$ 6 min

# Times for the Mediterranean test case

## 2019 Computing times and speed-up (Volta, 2018)

| # GPUs | Computing times | Speed-up |
|---|---|---|
| 1 | 284.14 (4 min 44 s) | 1.00 |
| 4 | 97.20 (1 min 37 s) | 2.92 |
| 8 | 58.50 s | 4.86 |
| 16 | 40.07 s | 7.09 |
| 32 | 30.93 s | 9.19 |

**Requirement: computing time $<$ 6 min**

# Times for the Mediterranean test case

## 2018 vs 2019 Computing times and speed-up (Volta, 2018)

| | 2018 | | 2019 | |
|---|---|---|---|---|
| # GPUs | Comput. time | Speed-up | Comput. time | Speed-up |
| 1 | 4 min 44 s | 1.00 | 4 min 44 s | 1.00 |
| 4 | 1 min 37 s | 2.92 | 1 min 37 s | 2.92 |
| 8 | 66.54 s | 4.27 | 58.50 s | 4.86 |
| 16 | 50.09 s | 5.67 | 40.07 s | 7.09 |
| 32 | 38.24 s | 7.43 | 30.93 s | 9.19 |

**Requirement: computing time $<$ 6 min**

# Times for the Mediterranean test case

2019 Computing times and speed-up (Volta, 2018)

| | Before Audit | | After Audit | |
|---|---|---|---|---|
| # GPUs | Comput. time | Speed-up | Comput. time | Speed-up |
| 1 | 4 min 44 s | 1.00 | 4 min 46 s | 1.00 |
| 4 | 1 min 37 s | 2.92 | 1 min 23 s | 3.44 |
| 8 | 58.50 s | 4.86 | 48.62 s | 5.89 |
| 16 | 40.07 s | 7.09 | 31.49 s | 9.10 |
| 32 | 30.93 s | 9.19 | 23.24 s | 12.33 |

Requirement: computing time < 6 min

# 2. The Pacific scenario. Computing times from 2014 to 2018

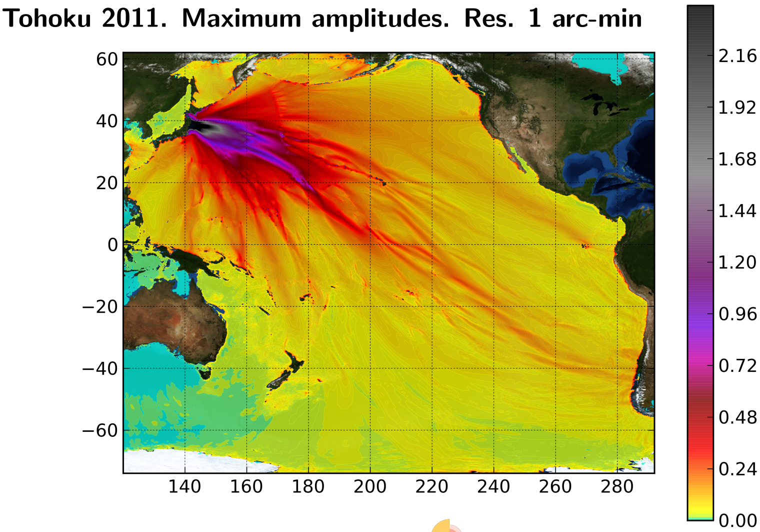

# Tohoku 2011. 2014 Computation times

## Propagation in global domain

**6 hours** (21,600 s) were simulated using three resolution levels:

- Original resolution (2 arc-min): 7,430,699 cells $(2,581 \times 2,879)$
- Resolution x2 (1 arc-min): 29,722,796 cells $(5,162 \times 5,758)$
- Resolution x4 (30 arc-sec): 118,891,184 cells $(10,324 \times 11,516)$

# Tohoku 2011. 2014 Computation times

## Propagation in global domain

**6 hours** (21,600 s) were simulated using three resolution levels:

- Original resolution (2 arc-min): 7,430,699 cells ($2,581 \times 2,879$)
- Resolution x2 (1 arc-min): 29,722,796 cells ($5,162 \times 5,758$)
- Resolution x4 (30 arc-sec): 118,891,184 cells ($10,324 \times 11,516$)

| | 2 arc-min | # FTRT | 1 arc-min | # FTRT | 30 arc-sec | # FTRT |
|---|---|---|---|---|---|---|
| 1 GPU | 7m 28.4s | 48.16 | 35m 36s | 10.11 | | |
| 2 GPUs | 4m 19.5s | 43.22 | 21m 41s | 17.37 | 2h 35m 21s | 2.32 |
| 4 GPUs | 2m 27s | 146.94 | 11m 38.8s | 30.91 | 1h 21m 49s | 4.4 |
| 8 GPUs | **77.74s** | 277.85 | 6m 22.3s | 56.5 | 43m 55s | 8.20 |

*** Times for nVIDIA Titan Black GPUs (Kepler, 2012). 1 Gb ethernet network (EDANYA cluster)**

# Tohoku 2011. 2018 Computation times

## Propagation in global domain

- Duration: **24 h** simulation / Numerical resolution: **2 arc-min** / Size: **21.0 M cells**

| # GPUs | Runtime (s) | # Runtime (min) | Speedup | Efficiency |
|---|---|---|---|---|
| 1 | 935.04 | **15.58** | 1 | 100% |
| 4 | 280.25 | **4.67** | 3.34 | 83% |
| 8 | 171.47 | **2.86** | 5.45 | 68% |
| 16 | 100.87 | **1.68** | 9.27 | 58 % |
| 24 | 79.06 | **1.32** | 11.83 | 49 % |
| 32 | 64.99 | **1.08** | 14.39 | 45 % |
| 40 | 59.83 | **1.00** | 15.63 | 39 % |
| 48 | 52.75 | $<$ **1** | 17.73 | 37 % |
| 56 | 48.95 | $<$ **1** | 19.10 | 34 % |
| 64 | 46.20 | $<$ **1** | 20.14 | 32 % |

**Computations performed at the CTE-POWER of the BSC (www.bsc.es)**

* Times for nVIDIA V100 GPUs (Volta, 2018).

# Tohoku 2011. 2019 Computation times

## Propagation in global domain

- Duration: **24 h** simulation / Numerical resolution: **1 arc-min** / Size: **84.2 M cells**

| # GPUs | Runtime (s) | # Runtime (min) | Speedup | Efficiency |
|---|---|---|---|---|
| 1 | 7343.52 | **122.39** | 1 | 100% |
| 4 | 1919.26 | **31.99** | 3.83 | 95.6% |
| 8 | 994.51 | **16.58** | 7.38 | 92.2% |
| 16 | 522.61 | **8.71** | 14.05 | 87.8 % |
| 24 | 362.93 | **6.05** | 20.23 | 84.3 % |
| 32 | 280.29 | **4.67** | 26.20 | 81.9 % |
| 40 | 233.87 | **3.90** | 31.40 | 78.5 % |
| 48 | 203.53 | **3.39** | 36.08 | 75.2 % |
| 56 | 179.20 | **2.99** | 40.98 | 73.1 % |
| 64 | 161.40 | **2.69** | 45.50 | 71.1 % |

**Computations performed at the CTE-POWER of the BSC (www.bsc.es)**

* Times for nVIDIA V100 GPUs (Volta, 2018).

# Computing times in the Pacific (2014 vs 2019)

## Propagation in global domain

- Duration: **24 h** / Numerical resolution: **1 arc-sec** / Size: **84.2 M cells** (original 7.4M)

| # GPUs | 2014 | 2019 |
|---|---|---|
| 2 | 6h 54 min 16 s | **1h 10 min** |
| 4 | 3h 38 min 11 s | **32 min** |
| 8 | 1h 57 min 7s | **16 min 26 s** |
| 16 | | **8 min 43 s** |
| 24 | | **6 min 3 s** |
| 32 | | **4 min 40 s** |
| 40 | | **3 min 54 s** |
| 48 | | **3 min 23 s** |
| 56 | | **3 min** |
| 64 | | **2 min 41 s** |

## Compute time

- Divided by 6 - 7

# 3. Future perspectives

## We are in a new era

- **FTRT simulations in TEWS are now possible**
- **Several scenarios (focal mechanisms) can be investigated**
- **It is the way to (computational) PTF?**
- **Local high-resolution inundation assessment (dedicated GPUs)**
- **More physics for propagation**

# 3. Future perspectives

## We are in a new era

- FTRT simulations in TEWS are now possible
- **Several scenarios (focal mechanisms) can be investigated**

## In the Mediterranean case - with 64 NVIDIA V100

- **8 scenarios in less than 1 minute (48 s)**

# 3. Future perspectives

## We are in a new era

- FTRT simulations in TEWS are now possible
- **Several scenarios (focal mechanisms) can be investigated**

## In the Mediterranean case - with 64 NVIDIA V100

- **8 scenarios in less than 1 minute (48 s)**
- **16 scenarios in 1 min 23 s**

# 3. Future perspectives

## We are in a new era

- FTRT simulations in TEWS are now possible
- **Several scenarios (focal mechanisms) can be investigated**

## In the Mediterranean case - with 64 NVIDIA V100

- **8 scenarios in less than 1 minute (48 s)**
- **16 scenarios in 1 min 23 s**
- **32 scenarios in less than 3 min**

# 3. Future perspectives

## We are in a new era

- FTRT simulations in TEWS are now possible
- **Several scenarios (focal mechanisms) can be investigated**

## In the Mediterranean case - with 64 NVIDIA V100

- **8 scenarios in less than 1 minute (48 s)**
- **16 scenarios in 1 min 23 s**
- **32 scenarios in less than 3 min**
- **64 scenarios in less than 5 min (4 min 46 s)**

# 3. Future perspectives

## We are in a new era

- FTRT simulations in TEWS are now possible
- **Several scenarios (focal mechanisms) can be investigated**

## In the Mediterranean case - with 64 NVIDIA V100

- **8 scenarios in less than 1 minute (48 s)**
- **16 scenarios in 1 min 23 s**
- **32 scenarios in less than 3 min**
- **64 scenarios in less than 5 min (4 min 46 s)**

## In the "less than 6 minutes" challenge

Now we can compute not just one single scenario **but 64 scenarios**

# 3. Future perspectives

## But let us think BIG!!!

- **What if when an event occurred ...**

# 3. Future perspectives

## But let us think BIG!!!

- **What if when an event occurred ...**
- **a large GPU infrastructure stopped for providing full computational resources**

# 3. Future perspectives

## But let us think BIG!!!

- **What if when an event occurred ...**
- **a large GPU infrastructure stopped for providing full computational resources**
- **What if HUNDREDS of GPUs ...,**

# 3. Future perspectives

## But let us think BIG!!!

- **What if when an event occurred ...**
- **a large GPU infrastructure stopped for providing full computational resources**
- **What if HUNDREDS of GPUs ...,**
- **THOUSANDS of GPUs were available for such computation?**

# 3. Future perspectives

## But let us think BIG!!!

- What if when an event occurred ...
- a large GPU infrastructure stopped for providing full computational resources
- What if HUNDREDS of GPUs ...,
- THOUSANDS of GPUs were available for such computation?
- What if the computational component were no more an issue?

# Thanks for your attention

## Faster and Faster Tsunami simulations with

CHEESE

Center of Excellence for Exascale in Solid Earth



https://www.uma.es/edanya
https://edanya.uma.es/hysea/
e-mail: jmacias@uma.es

+FB: https://www.facebook.com/edanya.uma/
+TW: @EdanyaUMA, @JorgeMACSAN
Subscribe to the EDANYA YouTube channel