Prospects of real time tsunami inundation estimates with TsunAWI - Studies in the LEXIS project

Natalja Rakowsky[1], Thierry Goubier[2], Sven Harig[1]

[1] Alfred Wegener Institute, Helmholtz Centre for Polar and Marine Research, Bremerhaven, Germany

[2] CEA List, DSCIN/LCYL, Gif sur Yvette Cedex, France

EGU General Assembly 2021, 19-30 April 2021, online

LEXIS builds an advanced engineering platform at the confluence of HPC, Cloud and Big Data. LEXIS develops infrastructure to enable workflows and demonstrates its abilities through three large-scale socio-economic pilots

- aeronautics
- weather & climate
- catastrophe alert systems: earthquake & tsunami

The LEXIS work flow of the earthquake and tsunami pilot with the tsunami inundation simulation.

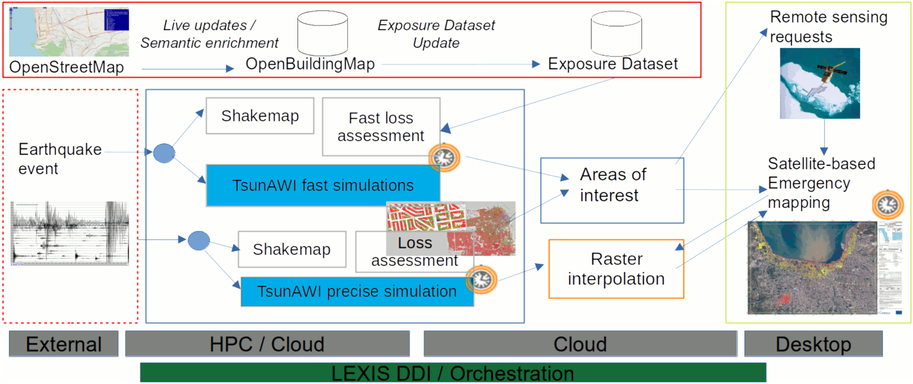

Former code optimization

- OpenMP parallelization, NUMA aware
- Mask dry areas, avoid to compute zeros
- Order mesh vertices along a space filling curve for good data locality
- Precompute auxiliary fields and data structures

Added last year

- Reduce floating point precision: double $\rightarrow$ single, $\approx$30% faster.
- **MPI parallelization**

With the restriction to one compute node, only regional simulations were possible in real time. Resolution was crucial, too.

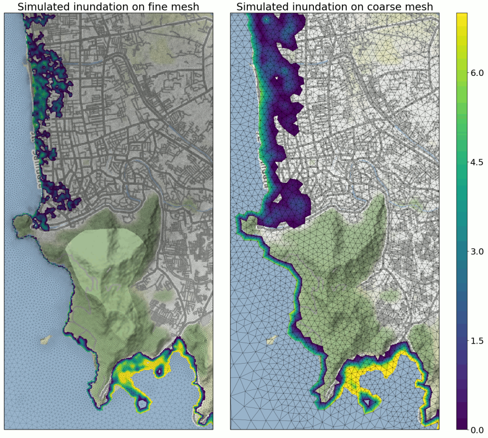

Simulated inundation [m] of a tsunami caused by a hypothetical earthquake, Mw=8.8, west off Padang, Sumatra, Indonesia.

The MPI parallel TsunAWI allows to simulate large ocean basins or e.g., the Indonesian Archipelago in real time.
Before, these scenarios had to be pre-computed, here one of 16,000 scenarios for InaTEWS, the Indonesia Tsunami Early Warning System.

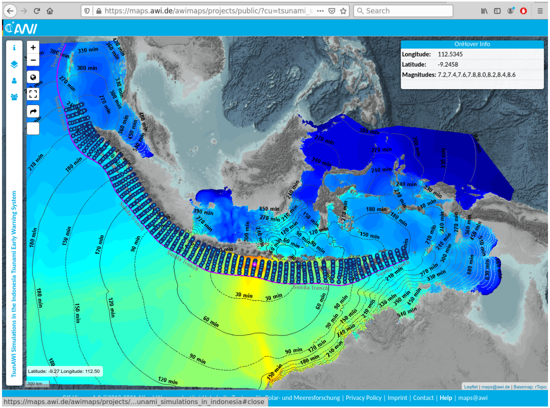

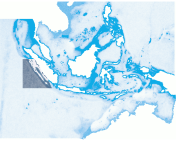

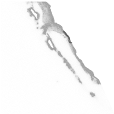

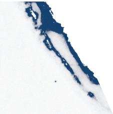

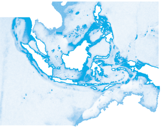

	regional coarse	regional fine	Indonesia
Number of vertices	230,000	1,240,000	11,110,000
Resolution	200m-15km	20m-5km	150m-20km
Time step dt	1.5s	0.15s	1.0s
Model time	2h	2h	12h

Compute time for time stepping: 2x Intel Xeon Cascade Lake 48core

	regional coarse	regional fine	Indonesia
1 node	4s	171s	707s
2 nodes	3s	96s	378s
4 nodes	–	58s	350s
10 nodes	–	34s	152s
20 nodes	–	–	89s
40 nodes	–	–	61s

Post processing and data products

The simulation data, for LEXIS the maximum inundation height, has to be interpolated from the 2D unstructured mesh to geotiff raster data.

Also, data from ensemble runs based on possible earthquake bottom movements have to be aggregated.

GPUs and FPGAs are well suited for this task. Porting is ongoing.

Integration in a real-time workflow

Fast ensemble simulations with tsunAWI are to be dispatched on multiple HPC centers and the cloud, with tentative over-submissions and aggregation of completed runs, in a 60 seconds time-from-event window.

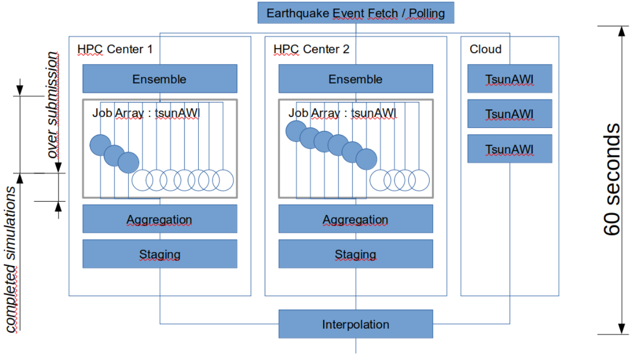

Outlook

- Add functionality to MPI version (tide gauge data, raster output, benchmarks)
- Optimise MPI
  - Simplify setup
  - split loops to better overlap computation and communication
  - One-sided MPI-2, e.g., MPI_Put only nonzero values
  - weighted partitioning, e.g. vertices on land vs. in the ocean
  - I/O asynchronous
- Optimise OpenMP/hybrid: So far, barriers are placed safety first

TsunAWI Materials

- Source code `https://gitlab.awi.de/tsunawi`
- Documentation `https://tsunami.awi.de/`

Thanks to all partners in LEXIS! https://lexis-project.eu

We also thank the North German Supercomputing Alliance (HLRN) and Atos/Bull for support, in particular Sebastian Krey (GWDG) for optimizing the compute node setup, the Scientific Board of HLRN for granting a consulting project, and Alesja Dammer and John Donners (Atos/Bull) for many hints on code optimization, in particular the hybrid parallelization.

This project has received funding from the European Union's Horizon 2020 research and innovation program under grant agreement 825532