Towards an exascale computing to recover the crustal Earth physical properties using the full seismic waveform inversion

Alizia Tarayoun[1], Romain Brossier[1], Ludovic Métivier[2,1]

[1] Univ. Grenoble Alpes - ISTerre, [2] Univ. Grenoble Alpes - CNRS - LJK

I- Full Waveform Inversion (FWI)

The Full Waveform Inversion (FWI) is a high-resolution seismic imaging technique allowing to reconstruct the physical properties of the Earth's subsurface. It is based on the iterative minimization of a misfit function related to the distance between observed and calculated seismic data, over a space of model parameters describing the medium. The iterative modification of the model parameters is obtained using the gradient of the misfit function. It is a challenging research topic combining geophysics, applied mathematics and high performance computing sides.

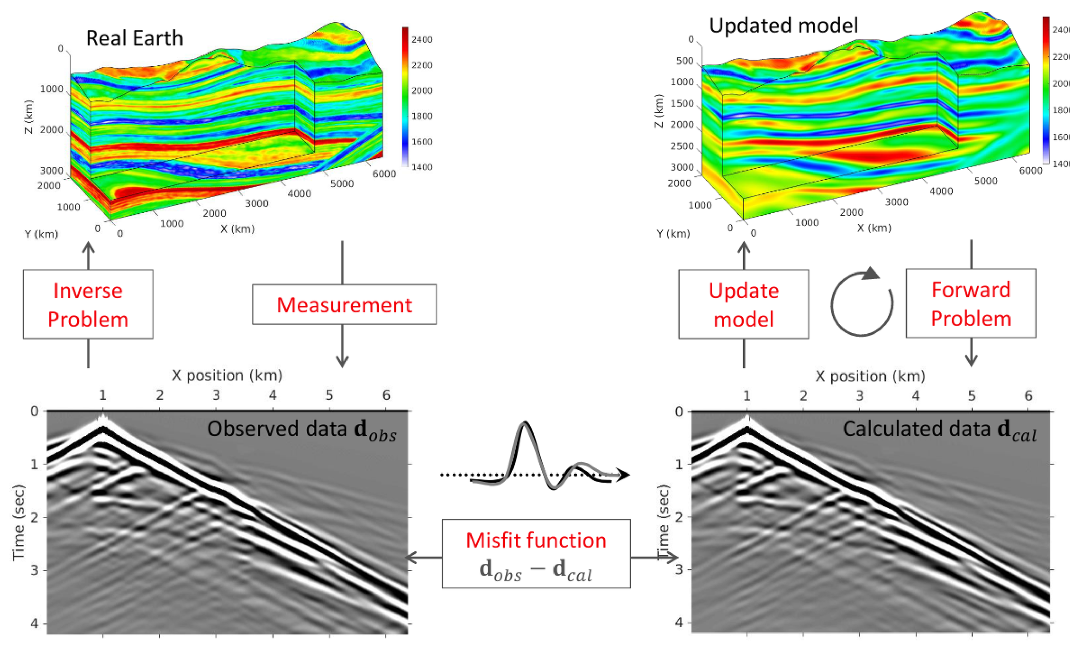

The FWI iterative workflow and steps are presented below:

1. Forward modeling →HPC intensive !
2. Compute misfit function
3. Adjoint modeling →HPC intensive !
4. Compute gradient →HPC intensive !
5. Compute model update
6. Update model

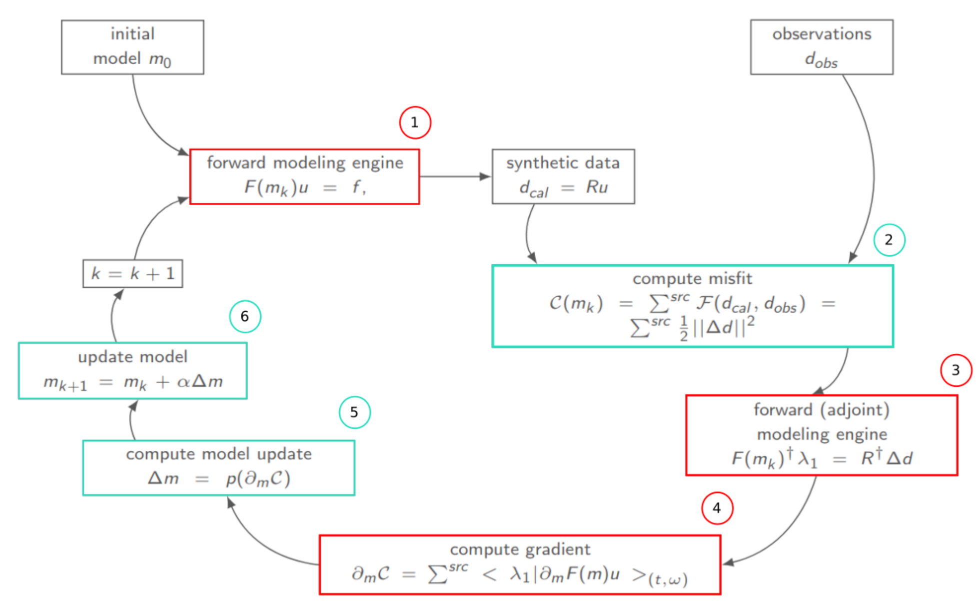

II - HPC designed SEM46 code

SEM46 is a seismic modeling and FWI 3D code for crustal exploration (Trinh et al., 2019).

→ Forward modeling and mesh description

The discretization is made through finite element (FE) scheme. A structured cartesian-based hexaedron mesh is implemented with a variable element-size strategy which can be used to reduce the numerical cost. The Cartesian-based mesh also allows to determine the spatial position and the neighbors of each element without any extra cost, which avoids a heavy searching operator over the global mesh as usual in standard FE methods.

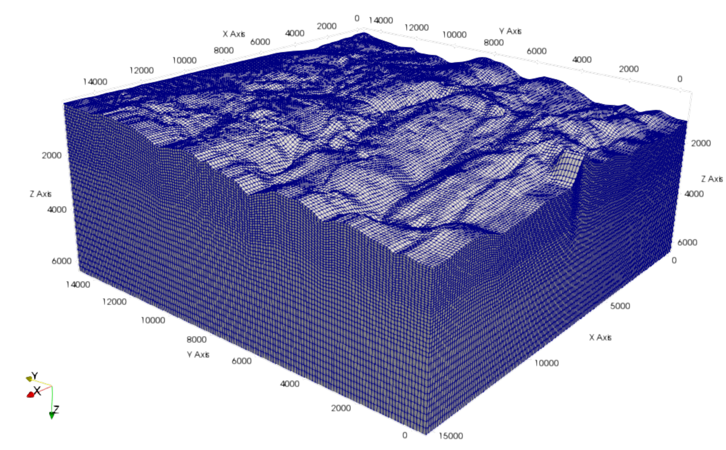

→ Parallel implementation

Two-level MPI-based parallelization is implemented:

- one level is designed on Cartesian-based domain decomposition, allowing an efficient load-balancing thanks to the Cartesian-based mesh
- one level over seismic sources

The parallel efficiency has been assessed through strong scalability tests on different HPC platforms: local (GRICAD), national (CINES/IDRIS/TGCC), European (BSC); showing good properties.

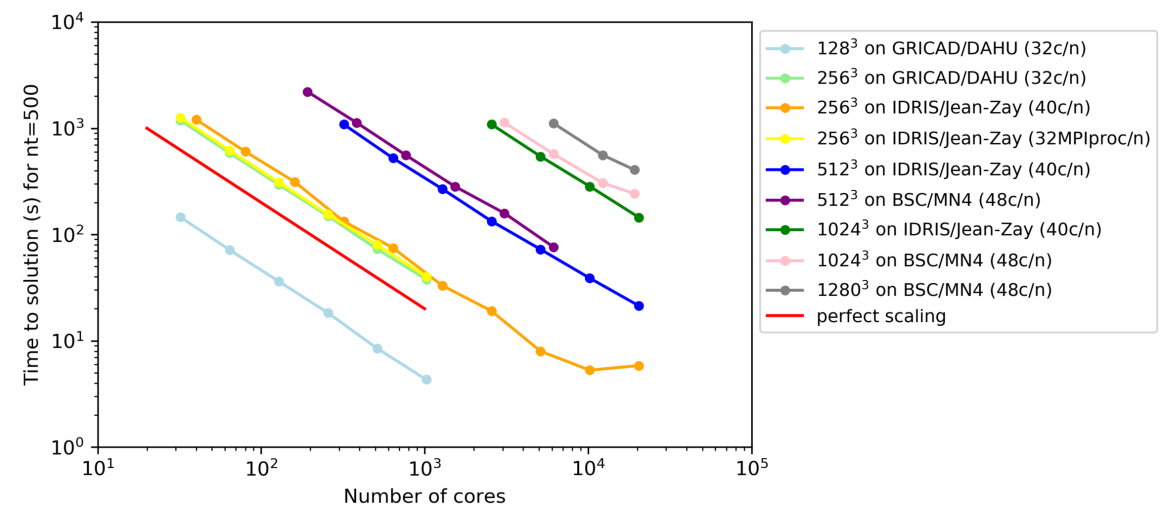

The largest scalability run has been made on BSC/MareNostrum4 machine using 19200 cores (400 nodes) with a mesh of 1280^3 elements (1.34^{11} dof).

→ Wavefield storage strategy

The gradient computation requires simultaneously to access the incident and an adjoint wavefields. To store/access the incident wavefield, two possibilities:

- The incident field is stored in memory or disk during the forward simulation, and read during the adjoint simulation
- The incident field is recomputed from the stored wavefield in the boundaries synchroneously with the adjoint field's propagation. A Checkpointing-assisted reverse-forward simulation method (CARFS) has been developed to achieve a better compromise between the memory requirement and the recomputation ratio (Yang et al., 2016).

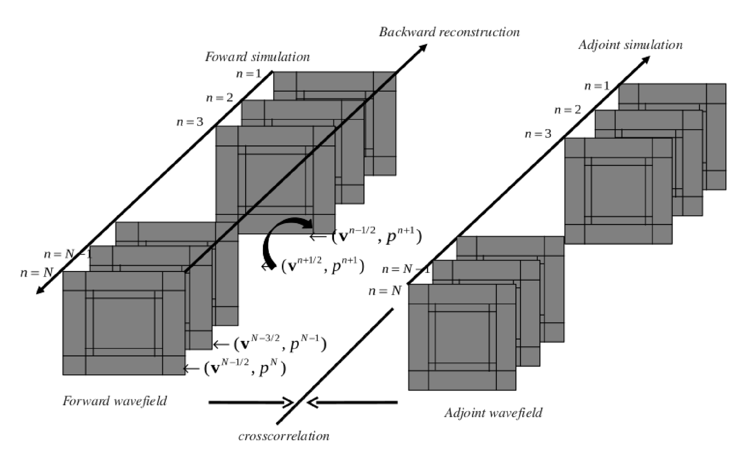

→ Application examples

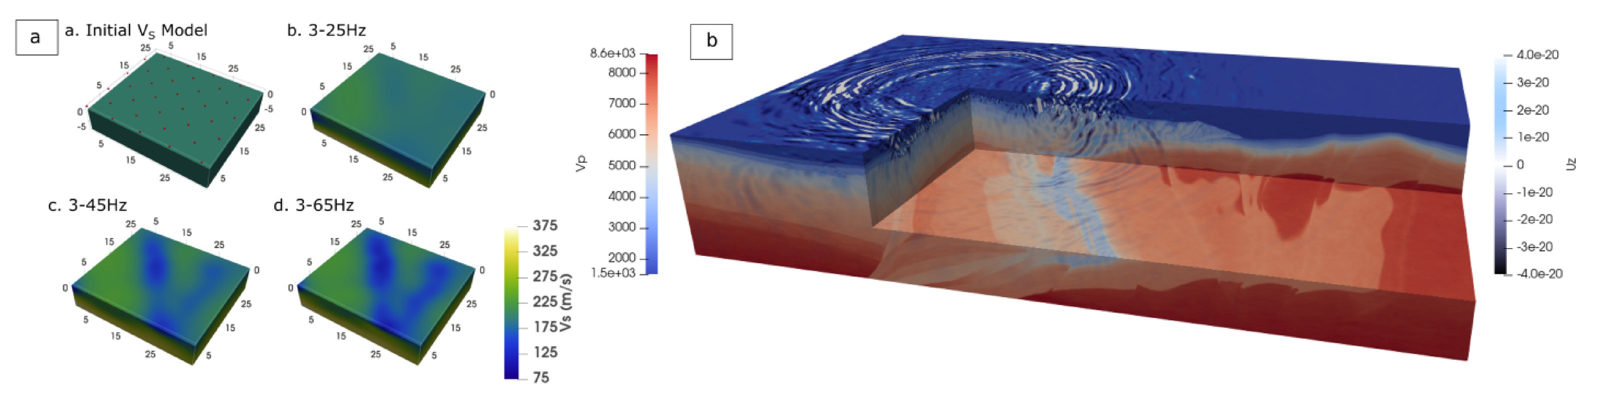

Figure: (a) FWI model reconstruction through frequency bands in a trench (Germany) (a) initial (d) final models. (b) 3D wavefield for deep crustal target (Nankaï subduction, Japan)

Application	Mesh size (z,x,y)	Model size (z,x,y)	Computing hours (sequential time)	Output volumetry (TB)
(a)	13×43×38	6×35×30 m	41415	3.5
(b)	94×827×477	30×175×100 km	152000	18

Applications (a) and (b) show a successful 3D FWI and a 3D forward modeling, respectively. For the latter, the FWI is currently challenging in terms of computional cost. Reaching 3D reconstruction of such kind of big model is one example of what we could benefit from exascale computing.

III - Exascale perspective

The goal is to adapt the SEM46 code for having a finer scale granularity for tackling larger problems and larger HPC systems while improving the fault resilience (in the same way of COFII[1] and COMPSs[2] projects).

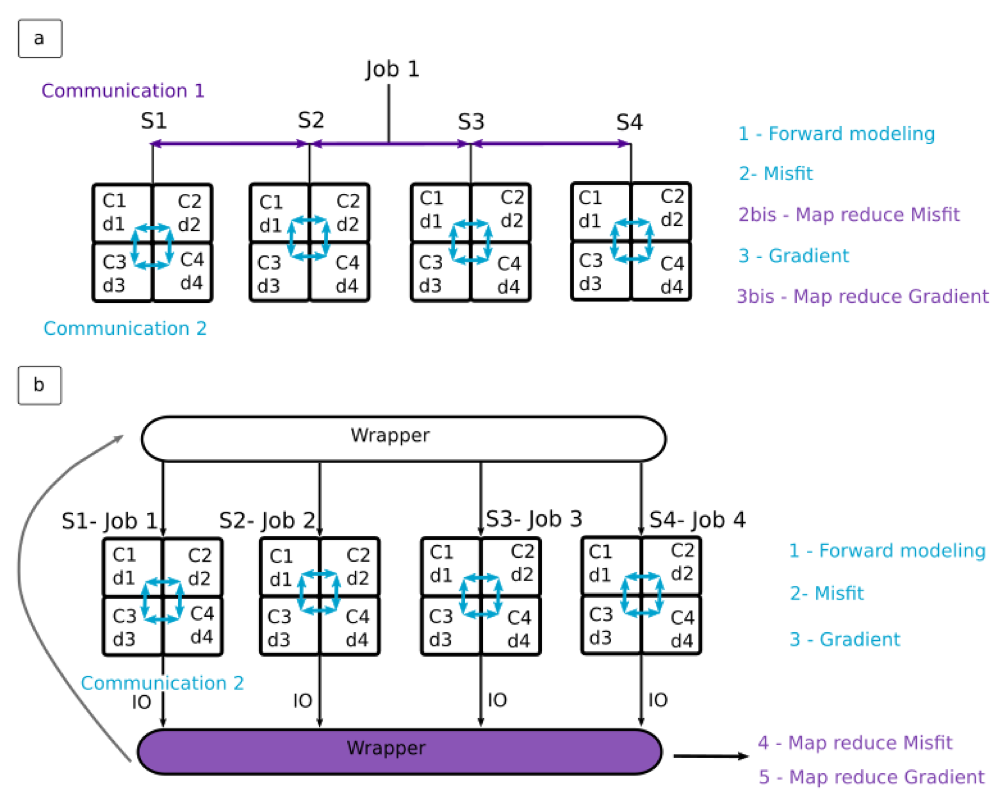

Figure: Simplified scheme of actual (a) and exascale-compatible (b) parallel tasks management. S1 is for Source 1, C1 is for Core 1 and d1 is for domain decomposition 1.

Current status of SEM46 code (a) 1 job controls the whole workflow (the 5 steps + iterations) for all seismic sources and domain decomposition for each source → highly sensitive to any hardware and software failure.

Proposed idea : redesign the workflow to reach a finer-scale job granularity (b). A wrapper would manage separated job launching for each seismic source. Each job will therefore only consider the heavy computational tasks of wavefield simulations, misfit function and gradient computation for one shot. The wrapper would manage the relaunching of failed jobs, the gathering of gradients, interaction with the non-linear optimization toolbox and launching of next batch of jobs for next iteration.

The advantages are:

- fault resilience management
- asynchronous jobs → can deal with different volume resources in time and more hardware flexibilities

Several points of vigilance: increase of IO level and errors management.

References

[1] https://github.com/ChevronETC/Schedulers.jl
[2] https://github.com/bsc-wdc/compss

Trinh, P. T., Brossier, R., Métivier, L., Tavard, L., and Virieux, J. (2019). Efficient 3D time-domain elastic and viscoelastic Full Waveform Inversion using a spectral-element method on flexible Cartesian-based mesh. *Geophysics*, 84(1):R75–R97.

Yang, P., Brossier, R., Métivier, L., and Virieux, J. (2016). Wavefield reconstruction in attenuating media: A checkpointing-assisted reverse-forward simulation method. *Geophysics*, 81(6):R349–R362.

Acknowledgements

We thank all the sponsors of the SEISCOPE consortium. This study was granted access to the HPC resources of GRICAD and MareNostrum4 infrastructures and CINES/IDRIS/TGCC under the allocation 046091 made by GENCI.

alizia.tarayoun@univ-grenoble-alpes.fr