Toward efficient elastodynamic modeling of Sequences of Earthquakes and Aseismic Slip on geometrically complex fault networks

Northwestern University

A divide-and-conquer strategy

Ciardo, F. and Romanet, P. (2026). A divide-and-conquer strategy for fast elastodynamic simulation of earthquakes and aseismic slip on fault networks*, Bulletin of the Seismological Society of America.* ISSN 0037-1106. doi: 10.1785/0120260073.

Roy, A., Ciardo, F. and Romanet, P. A GPU-Accelerated Non-Periodic Spectral Boundary Integral Equation Method for Three-Dimensional Earthquake-Cycle Simulations*, Under Review*.

Federico Ciardo, Ass. Professor @ Northwestern University

SCEC SEAS Workshop - Nov. 20, 2026

Ph.D. Pierre Romanet
Researcher at CEREMA and
Université Côte d'Azur (Geoazur),
Sophia-Antipolis , France



Ph.D. Anubhav Roy
Postdoctoral scholar in the Geoforce
group @ Northwestern University
Now assist. Prof. at IIT KGP



Introduction

- **The physical challenge**: SEAS simulations connect years for slow deformation to seconds of dynamic rupture. Quasy-Dynamic models reduce computational cost through radiation damping, but do not resolve wave propagation and the associated dynamic interaction between faults.

- **Volumetric methods**: (finite differences, finite, spectral or virtual elements) accommodate complex geometries, heterogeneous materials, and off-fault deformation. However, discretizing the surrounding medium makes long-duration, fully dynamic 3D simulations computationally demanding.

- **Space-time BIEM** restricts discretization to the faults in a homogeneous elastic medium and accommodates complex fault networks. Its fully dynamic formulation requires spatially coupled convolutions over past slip, creating substantial computational and memory costs.

- **Spectral BIEM** decouples spatial Fourier modes, leaving temporal convolutions evaluated independently for each mode. This is efficient for planar faults, but conventional periodic formulations require enlarged computational domains to limit artificial interactions between fault replicas.

- **Our hybrid strategy**: Combine non-periodic spectral self-interactions with hierarchically accelerated space–time interactions between faults.

2-Dimensional modeling

Divide-and-Conquer strategy

A divide-and-conquer strategy: simplest configuration

Displacement field given by 2D scalar wave equation:

$$\nabla^2 u = \frac{1}{c_s^2}\frac{\partial^2 u}{\partial t^2}, \quad u_3 = u(\mathbf{x},t), \quad u_1 = u_2 = 0, \quad \mathbf{x} \in \Gamma$$

$$\sigma_{31} = \mu\, \partial u/\partial x_1 \qquad \sigma_{32} = \mu\, \partial u/\partial x_2$$

Slip and shear stress on each fault:

$$\delta(\mathbf{x},t) = \Delta u(\mathbf{x},t) = u^+ - u^- \qquad \mathbf{x} \in \Gamma_i$$

$$\tau(\mathbf{x},t) = \sigma_{31}(\mathbf{x},t)n_1 + \sigma_{32}(\mathbf{x},t)n_2$$

Equilibrium:

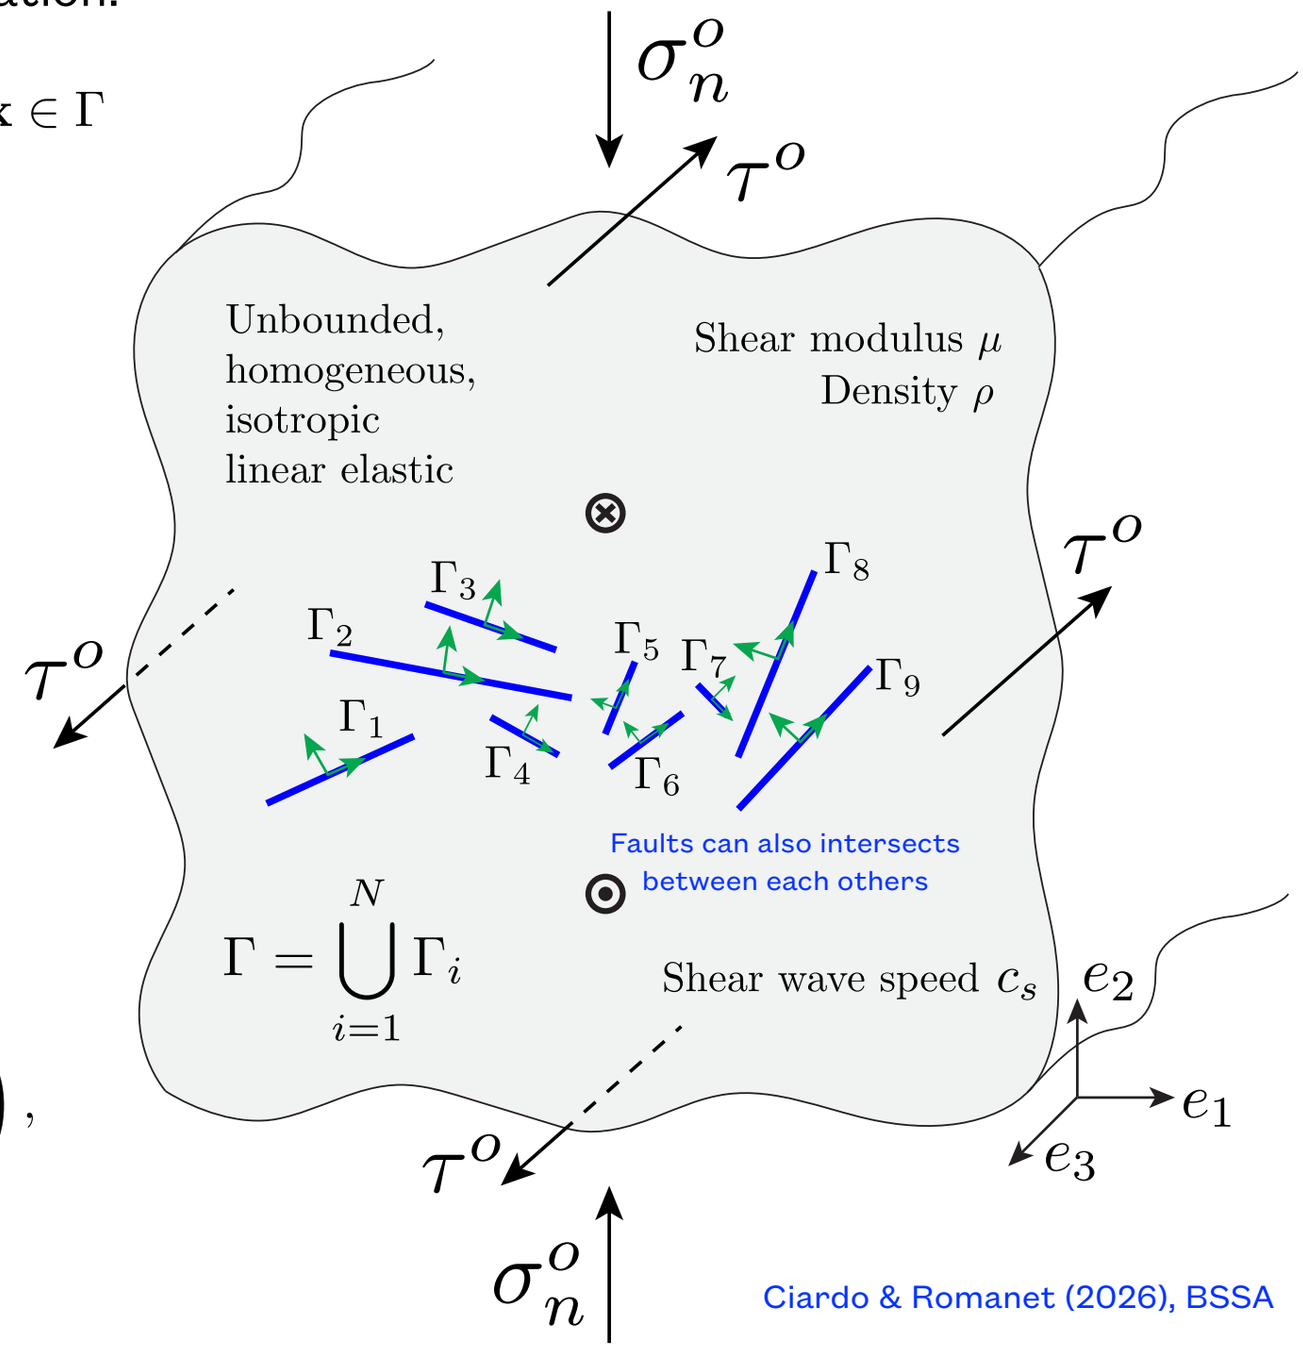
$$\tau^{\mathrm{el}}(\mathbf{x},t) + \tau^o(\mathbf{x},t) = f_c(\mathbf{x},t)\cdot\sigma_n^o(\mathbf{x},t)$$

Rate-and-state friction:

$$f(V,\theta) = f_* + a\ln\left(\frac{V(\mathbf{x},t)}{V_*}\right) + b\ln\left(\frac{V_*\theta(\mathbf{x},t)}{D_c}\right),$$

$$\frac{\partial\theta}{\partial t} = \frac{1}{\epsilon}\left[\left(\frac{V\theta}{D_c}\right)^{-\epsilon} - 1\right]\frac{V\theta}{D_c}$$

Ruina (1983), JGR



Ciardo & Romanet (2026), BSSA

A divide-and-conquer strategy: traction decomposition

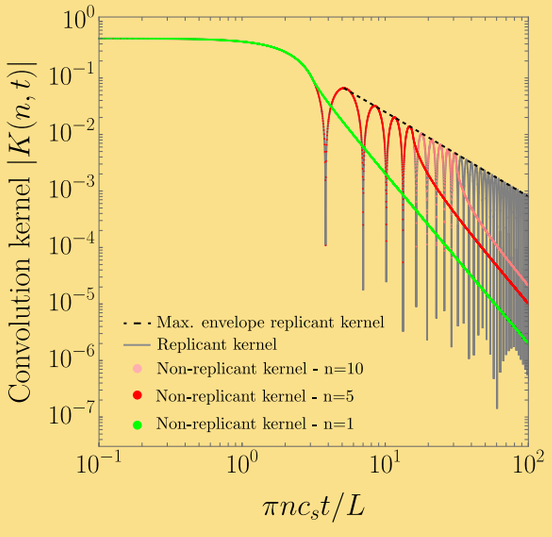
$$\tau^{\mathrm{el}}(\mathbf{x},t) = \underbrace{\left[-\frac{\mu}{2c_s}V_i(\mathbf{x},t) + f_{\Gamma_i\to\Gamma_i}(\mathbf{x},t)\right]}_{\text{Self-effects}} + \underbrace{\sum_{j\neq i} f_{\Gamma_j\to\Gamma_i}(\mathbf{x},t)}_{\text{Interaction effects}}.$$

Non-periodic spectral BIEM to eliminate spurious interactions from fault replicas

$$f_{\Gamma_i\to\Gamma_i}(\xi,t) = \sum_{n=-\infty}^{\infty} F_n(t)e^{ik_n\xi}, \qquad k_n = \frac{\pi n}{L_i}$$

$$F_n(t) = -\frac{\mu c_s k_n^2}{2}\int_{-\infty}^{t} K(n,t-\theta)D_n(\theta)\mathrm{d}\theta$$

In the spectral domain, the non-local operator is diagonal -> only time convolution!



Kernel decay allows very efficient spectral-mode dependent truncation! This allows to save memory and speed up time convolution.

Main computational cost is associated with FFT $\mathcal{O}(n_{el}\cdot\log(n_{el}))$ —> problem conquered!

Cochard & Rice (1997), JMPS Noda (2021), Earth, Planets and Space

Fully dynamic Space-Time BIEM accelerated using $\mathscr{H}$-matrices

Discretization of spatial-temporal convolution yields fully-populated, three-dimensional array. Hence:

Memory and operations $\mathcal{O}(N_{\mathrm{elts}}^2\cdot N_t)$ **BOTTLENECK, since** $N_t \propto N_{elts}$

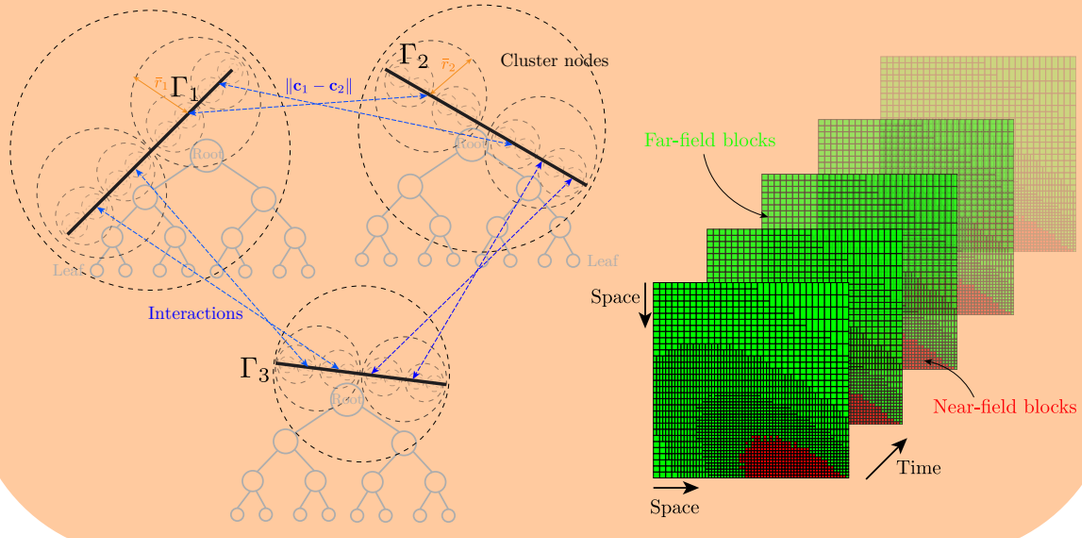
Lapusta et al. (2000), JGR

Our innovations:

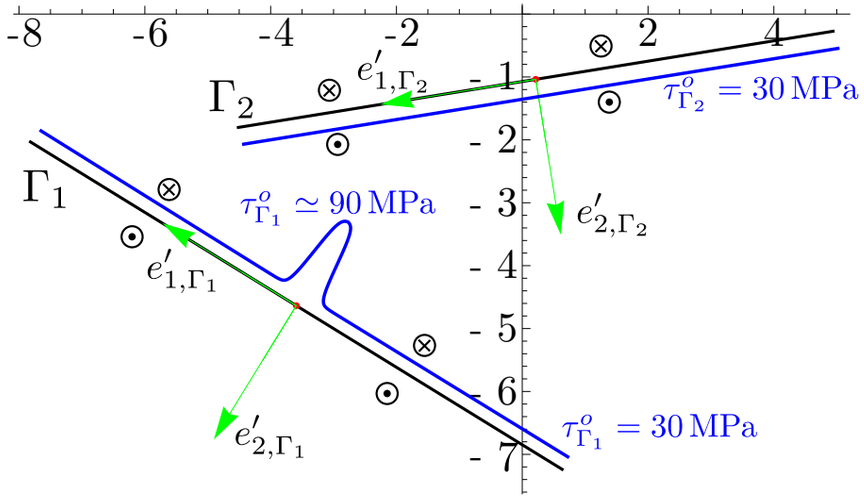
- Independent binary trees for each fault
- Far-field interactions compressed using low-rank approximations (Singular Value Decomposition)
- Near-field interactions retained in full to preserve accuracy
- Causality-based truncation reduces the elastodynamic history and speed-up space-time convolution

Hierarchical compression reduces memory and computational complexity —> problem conquered!

$$\mathcal{O}(N_{elts}^3) \to \mathcal{O}(N_{elts}^2\log(N_{elts}))$$



Numerical example: two-fault problem



Input parameters:

$b = (1.2/0.9) \cdot 10^{-2}$ $\quad a = 1.2 \cdot 10^{-2}$

$D_c = 10^{-2}\,\mathrm{m}$ $\quad f_* = 0.6$

$V_* = 10^{-9}\,\mathrm{m/s}$ $\quad \mu = 30\,\mathrm{GPa}$

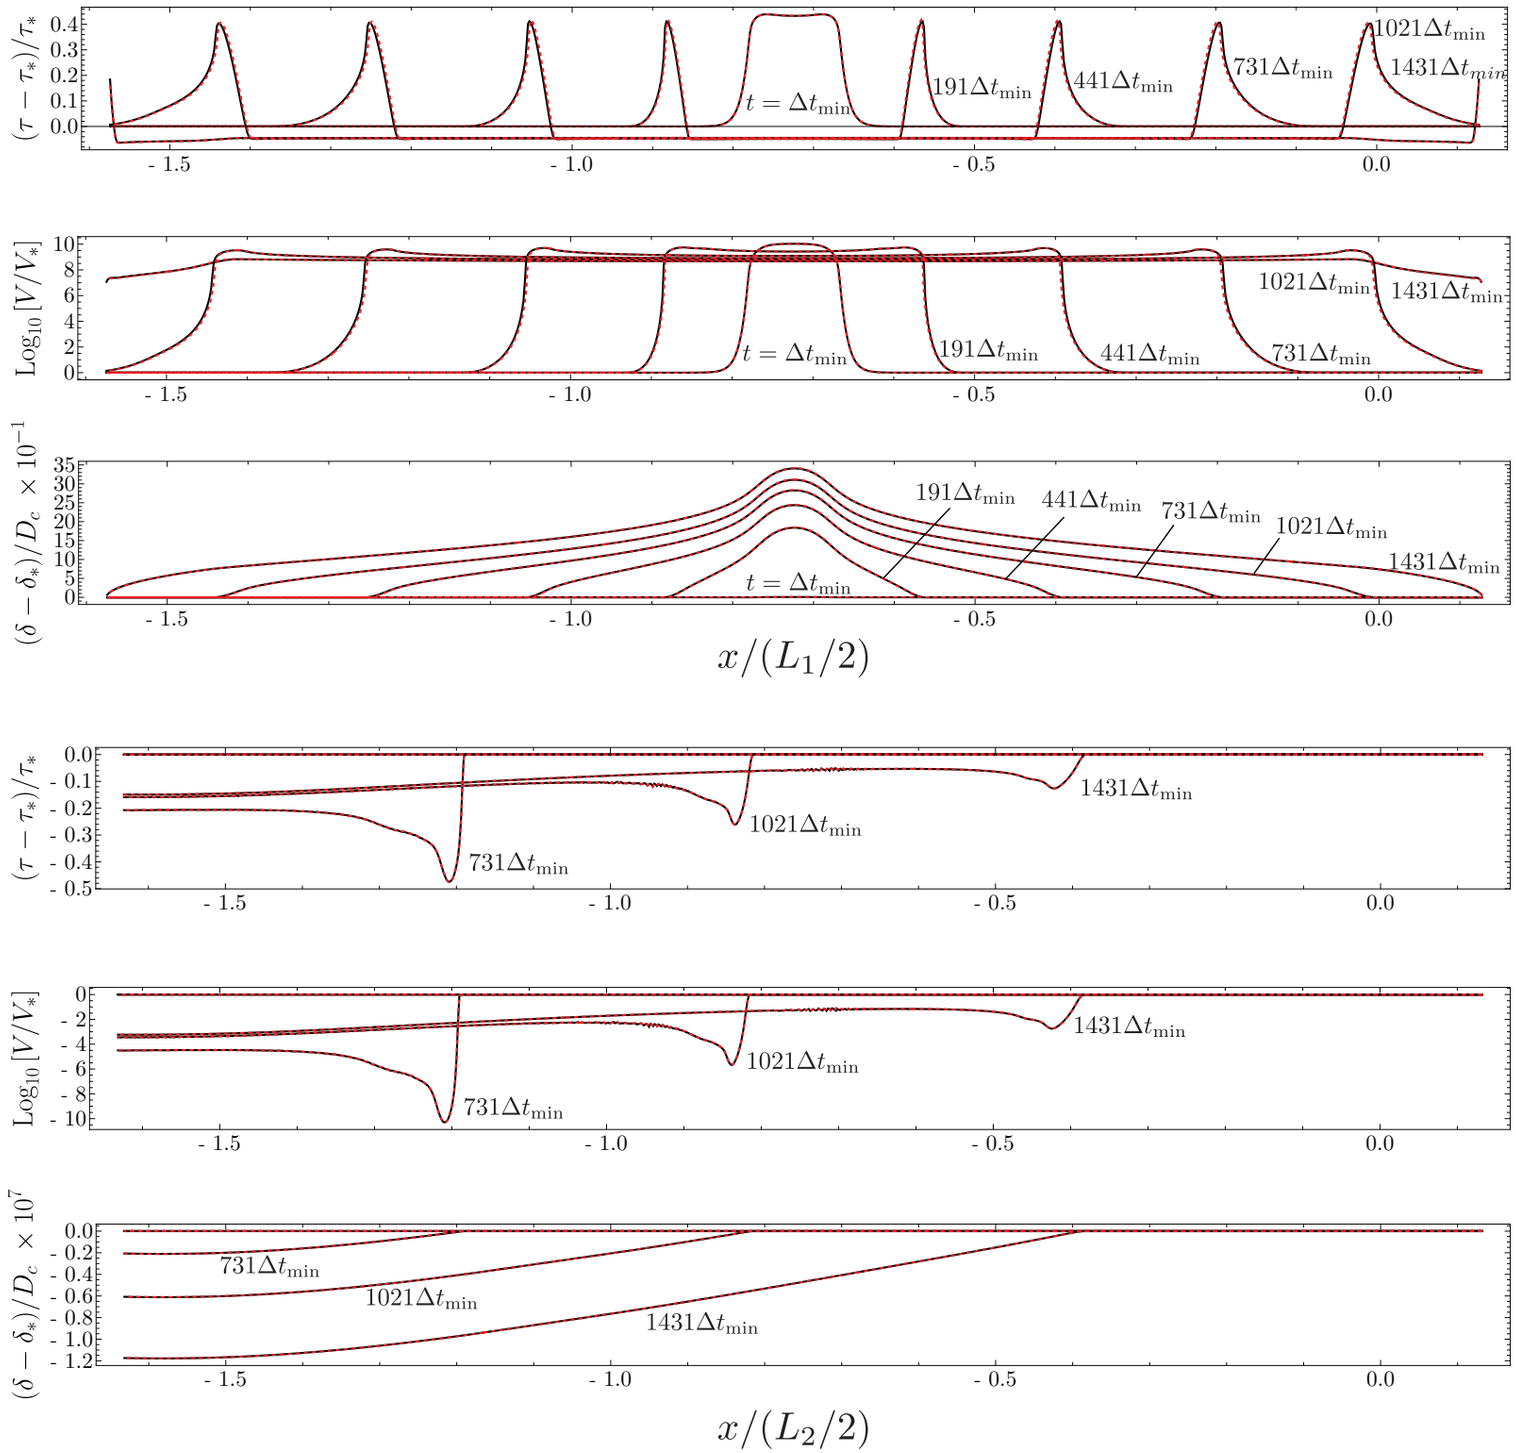
$c_s = 3464\,\mathrm{m/s}$ $\quad \sigma'_n = 50\,\mathrm{MPa}$

$\theta_{init} = D_c/V_*$ $\quad \epsilon_{SVD} = 10^{-6}$

$d_{\mathscr{H}} = 50$ $\quad \eta_{\mathscr{H}} = 10$

$$\Delta t_{\min} \simeq 1.38 \cdot 10^{-3}\,\mathrm{s}$$

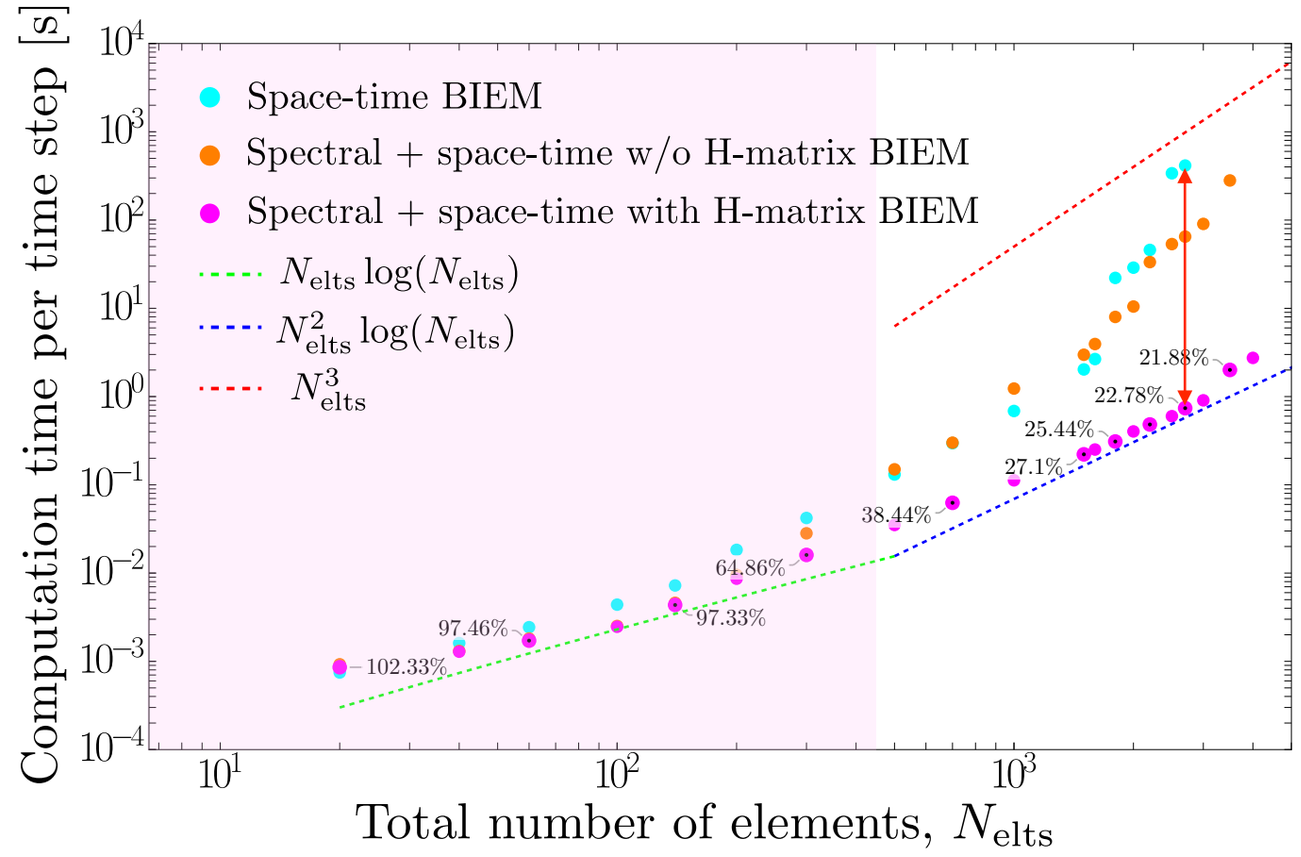
Compression ratio $c_r = 24.7\,\%$



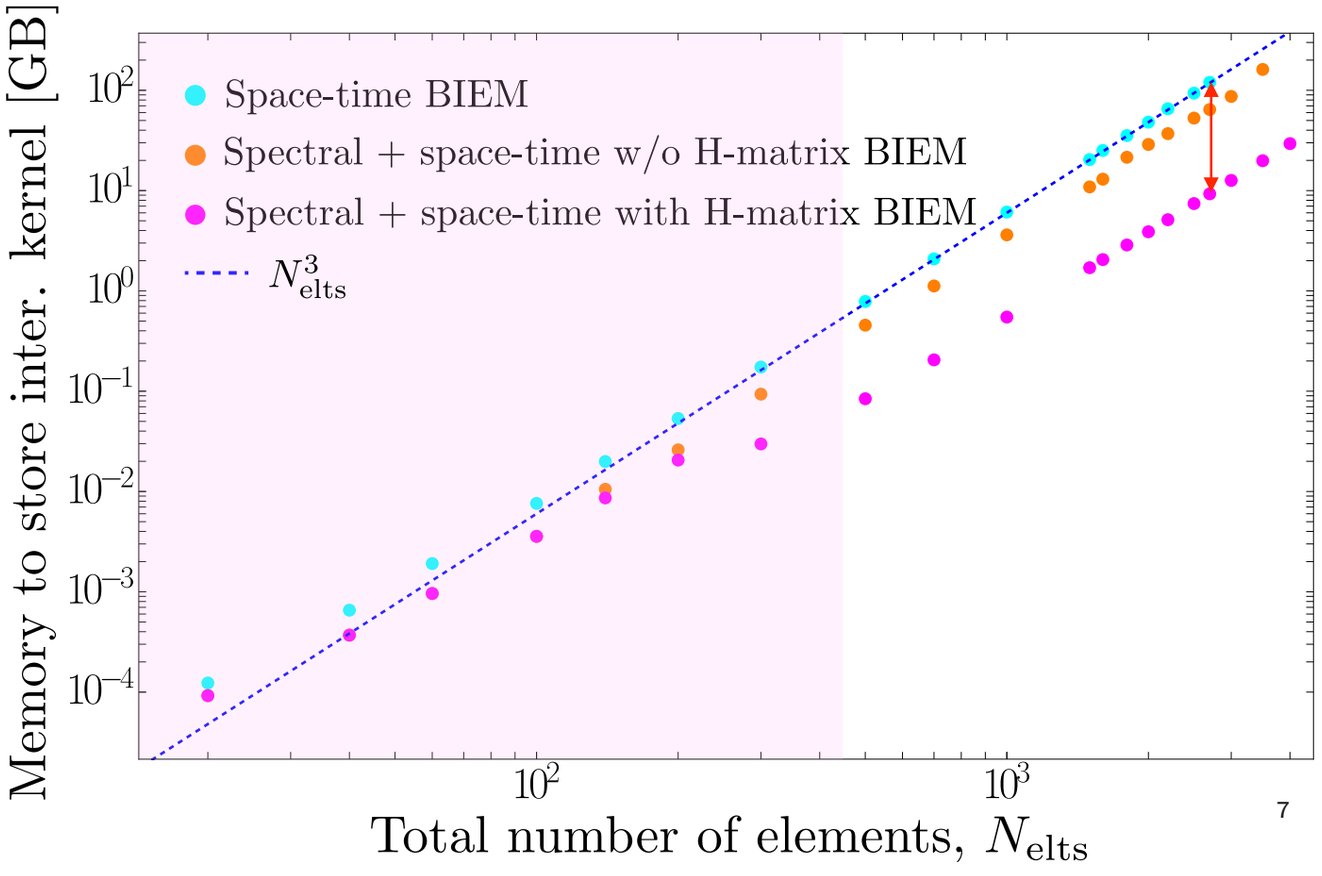
Space-Time BIEM — Hybrid ST/Spectral BIEM

Performance tests

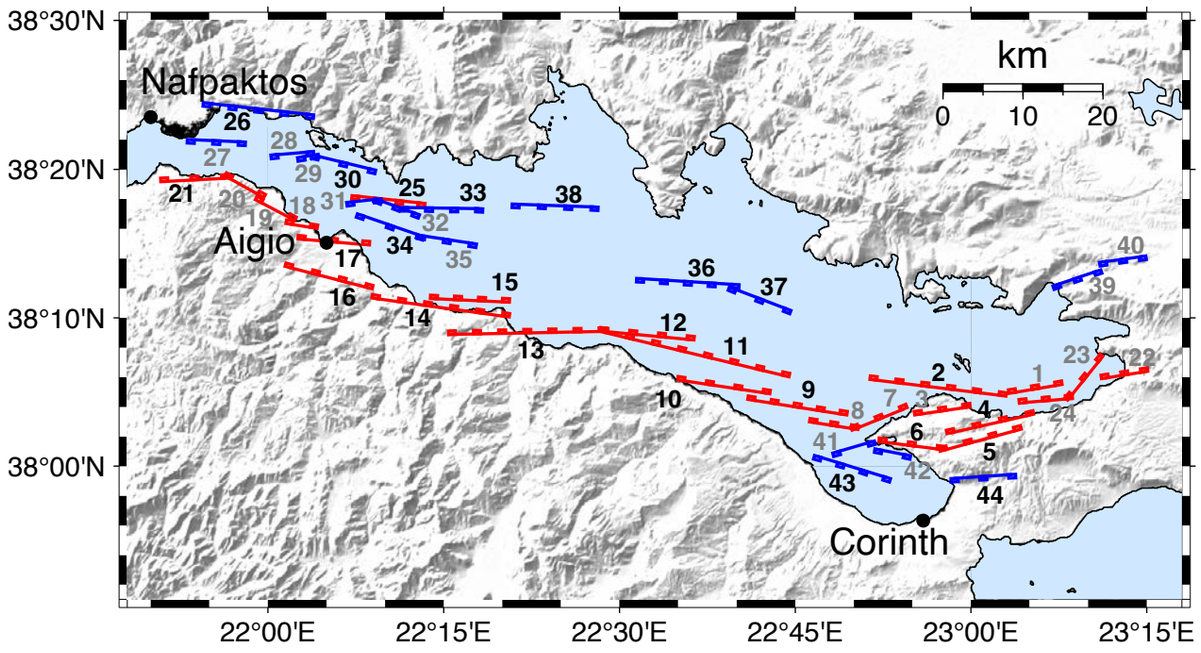
For sufficiently large problems, the proposed hybrid formulation reduces the computation time per time-step by more than two orders of magnitude, while achieving approximately one order of magnitude reduction in memory requirements, compared to classical Space-Time BIEM



In practical terms, a simulation that would require on the order of one year of wall-clock time using the traditional space–time formulation can be completed in approximately 10 hours with the proposed hybrid spectral / space–time approach.

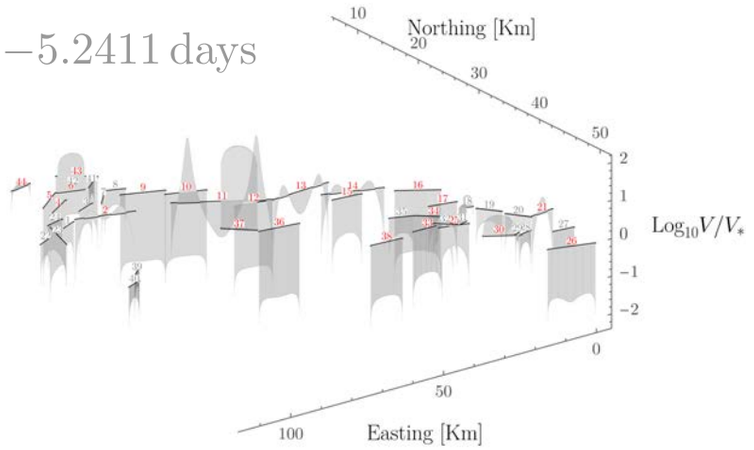


Slip dynamics in the extensional Corinth rift, Greece

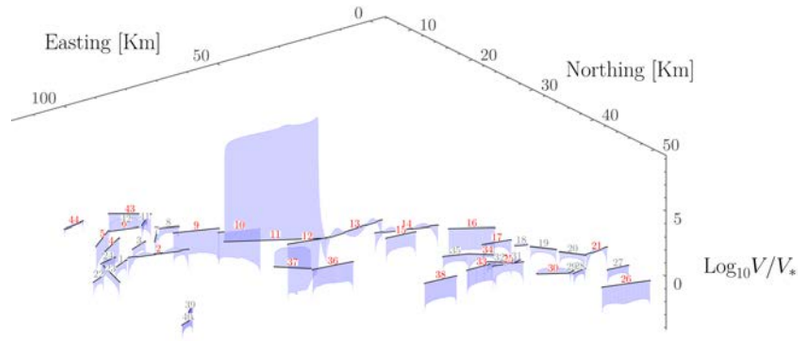


- 44 mapped, steeply dipping normal faults
- 32 years of simulated seismic and aseismic slip
- Spontaneous dynamic ruptures and wave-mediated fault interactions

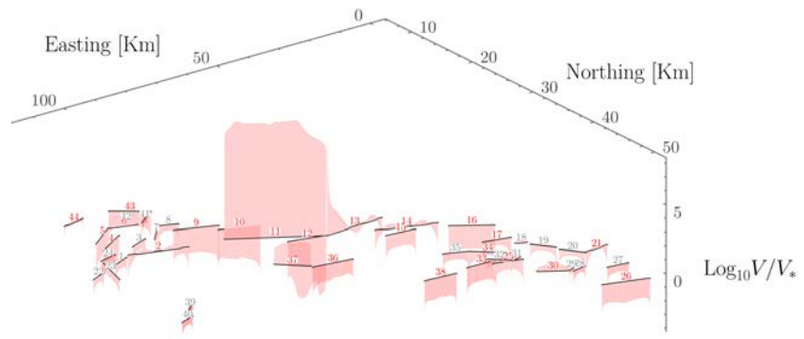
$(t - t_{\text{nuc}}) \simeq -5.2411$ days



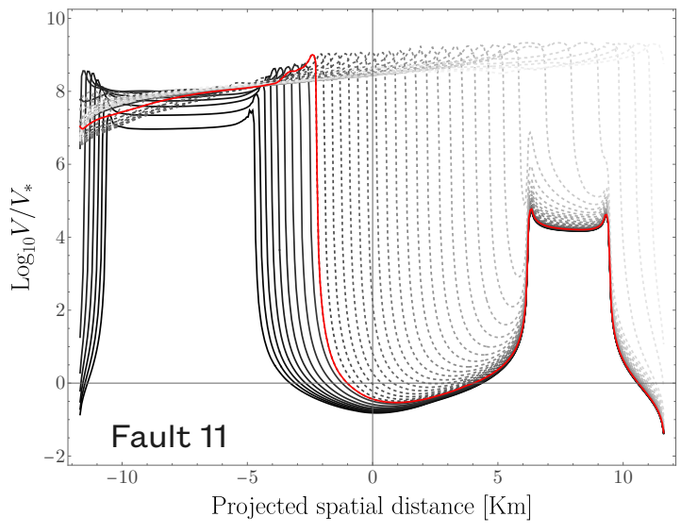
$(t - t_{\text{nuc}}) \simeq 9.79$ sec.

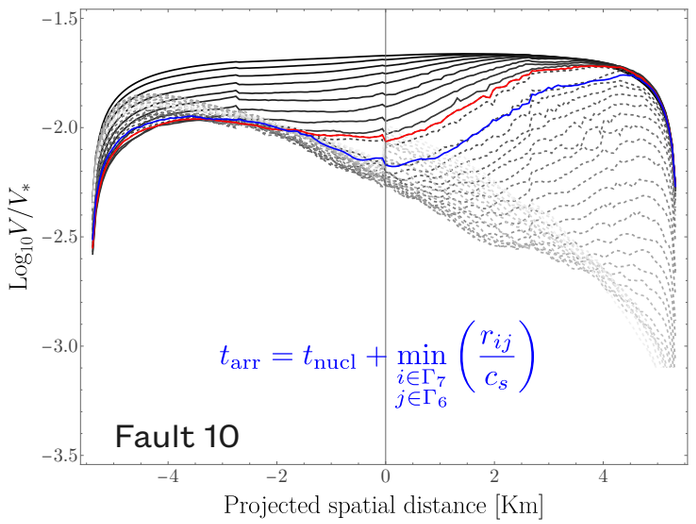


$(t - t_{\text{nuc}}) \simeq 13.14$ sec.



Rupture on Fault 11 radiates seismic waves that produce a rapid slip-rate decrease on Fault 10





3-Dimensional modeling

Missing piece before now: efficient implementation of 3D Spectral Boundary Integral Equation Method without replication

3-Dimensional non-periodic Spectral BIEM

Cochard and Rice (1997), JMPS

On-plane equilibrium:

$$\tau_j(x_1, x_3, t) = \tau_j^{\infty}(x_1, x_3, t) + \Delta\tau_j(x_1, x_3, t), \qquad j = 1, 3,$$

$$\Delta\tau_j(x_1, x_3, t) = -\frac{\mu}{2c_s}\dot{\delta}_j(x_1, x_3, t) + f_j(x_1, x_3, t)$$

Defining:

$$h_1 = \frac{\pi n_1}{\ell_1} \qquad h_3 = \frac{\pi n_3}{\ell_3}$$

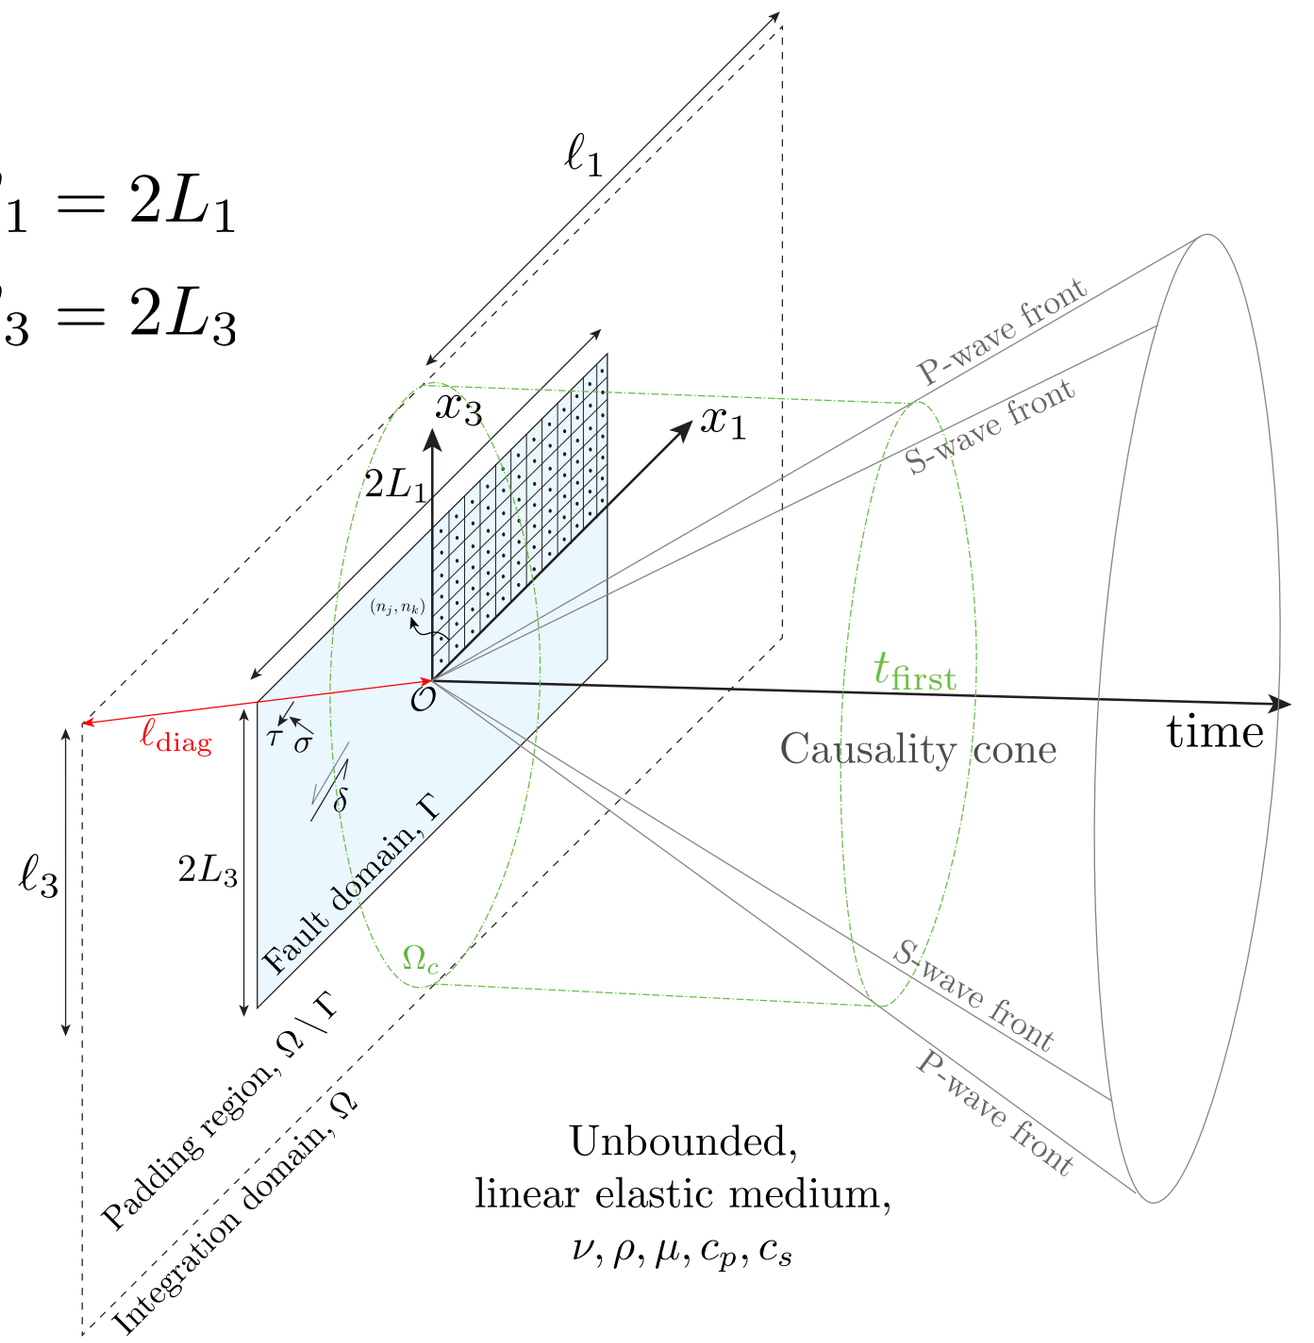
$$f_j(x_1, x_3, t) = \sum_{n_1=-\infty}^{\infty} \sum_{n_3=-\infty}^{\infty} F_j(n_1, n_3, t) e^{i(h_1 x_1 + h_3 x_3)}$$

$$F_j(n_1, n_3, t) = -\frac{\mu c_s}{2} A_{jl}^{M}(h_1, h_3) \int_{-\infty}^{t} \mathcal{K}^{M}(n_1, n_3, t-\theta) D_l(n_1, n_3, \theta)\, \mathrm{d}\theta,$$

$M = \mathrm{II}, \mathrm{III}$ (slip mode) $\qquad j = 1, 3$ (slip direction)

Since summation over repeated indices applies
-> coupling between slip along the two directions !
-> 4 time convolutions per slip direction !

$$\ell_1 = 2L_1$$
$$\ell_3 = 2L_3$$



The convolution kernels are piecewise defined. Their evaluation requires numerical integrations. We've developed robust procedures using analytical limits and Taylor expansions to treat removable singularities and numerically ill-conditioned expressions.

What is their temporal support?

SBIEM w/ & w/o replication: Comparison of convolution kernels

First transition: when the fastest wave reaches the nearest boundary of integration domain

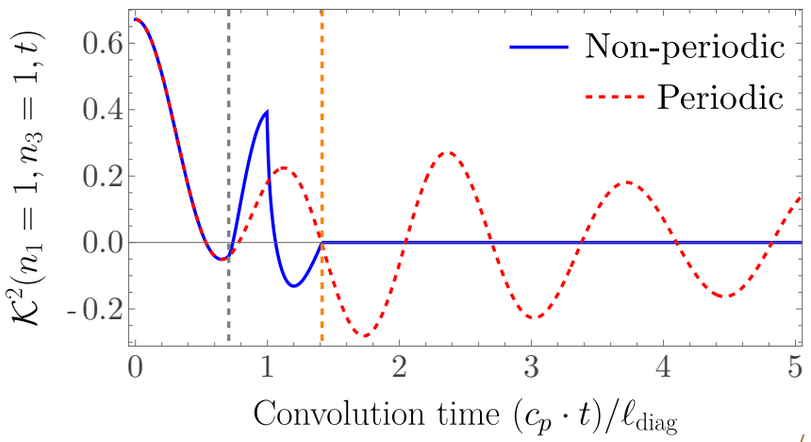
$$t_{\text{first}} = \frac{\min(\ell_1, \ell_3)}{c_{\mathcal{M}}}$$

$$c_{\mathcal{M}} = \begin{cases} c_p, & \mathcal{M} = 2, \\ c_s, & \mathcal{M} = 3. \end{cases}$$

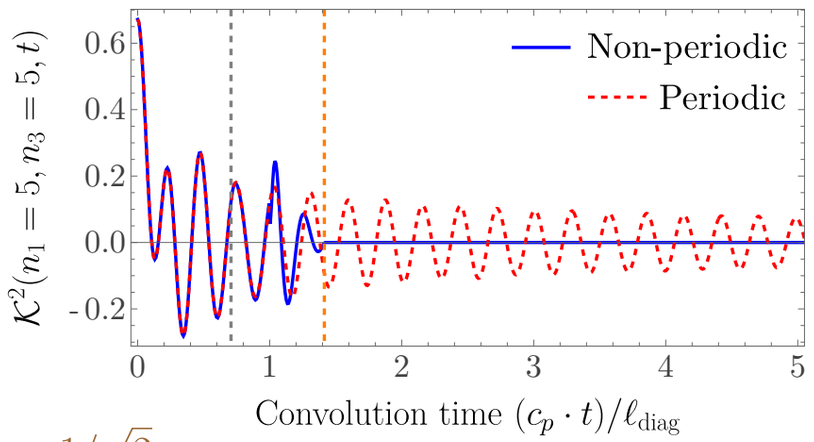
Last transition: when the slowest wave "leave" the integration domain

$$t_c = \frac{\ell_{\text{diag}}}{c_s} = \frac{\sqrt{\ell_1^2 + \ell_3^2}}{c_s}$$

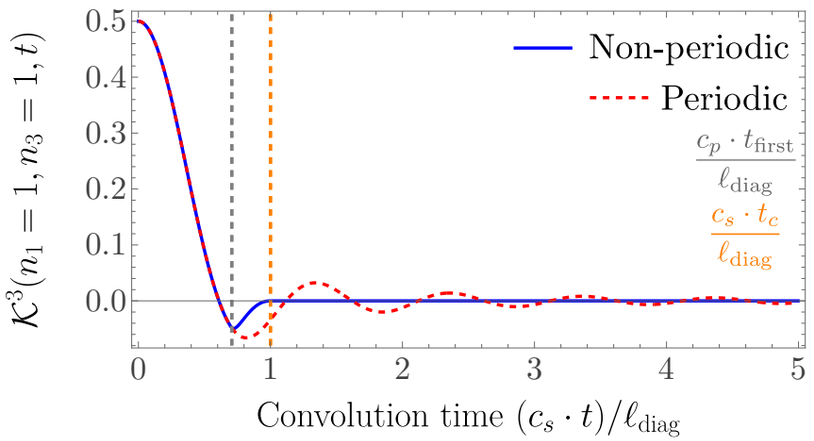
Spectral (Fourier) mode n_1, $n_3 = 1$

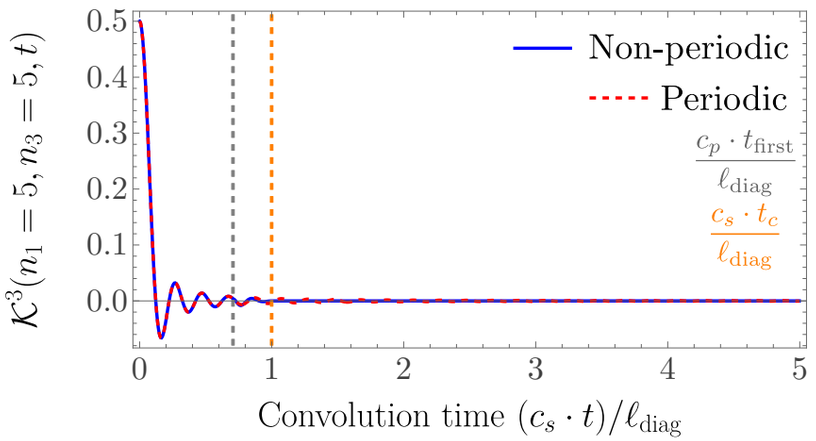


Spectral (Fourier) mode n_1, $n_3 = 5$



$c_s/c_p = 1/\sqrt{2}$





After $t_c = \ell_{\text{diag}}/c_s$, the transient convolution kernel vanishes exactly and the static response is recovered!

$$\int_{-\infty}^{t} K_M(n_1, n_3, t-\theta)\, D^l_{n_1,n_3}(\theta)\, d\theta = \int_{t-t_c}^{t} K_M(n_1, n_3, t-\theta)\, D^l_{n_1,n_3}(\theta)\, d\theta$$

The convolution truncation is exact,
unlike in the SBIEM w/ replication (see e.g, Lapusta and Liu (2009), JGR).
Furthermore, it is spectral mode-independent.

Acceleration of the temporal convolutions over their compact support

We use an explicit, second-order Runge-Kutta time integration scheme, with adaptive time step constrained to be an integer multiple of a minimum value:

$$\Delta t_{\min} = \frac{\beta_{\min}}{c_s} \min\left(\Delta x_1, \Delta x_3\right)$$

$\beta_{\min} = 0.1$ Safety factor

Lapusta et al. (2000), JGR

Lapusta & Liu (2009), JGR

1) Fast convolutions during slow deformation

- When the adaptive time step spans the complete finite elastodynamic memory window (always the case during slow deformation)

$$\begin{aligned}\mathcal{C}_l^M(n_1, n_3, t) &= \sum_{m=0}^{N_t-1} \overline{\mathcal{K}}^M(n_1, n_3, m) D_l(n_1, n_3, t - t_m) \\ &= \overline{\mathcal{K}}^M_{\text{stat}}(n_1, n_3) D_l(n_1, n_3, t), \qquad M = 2, 3\end{aligned}$$

Pre-summed modal kernels (i.e. static kernels).
They can be pre-computed before time-marching!

$\mathcal{O}(N_1 N_3 N_t) \rightarrow \mathcal{O}(N_1 N_3)$ Per Runge-Kutta stage!

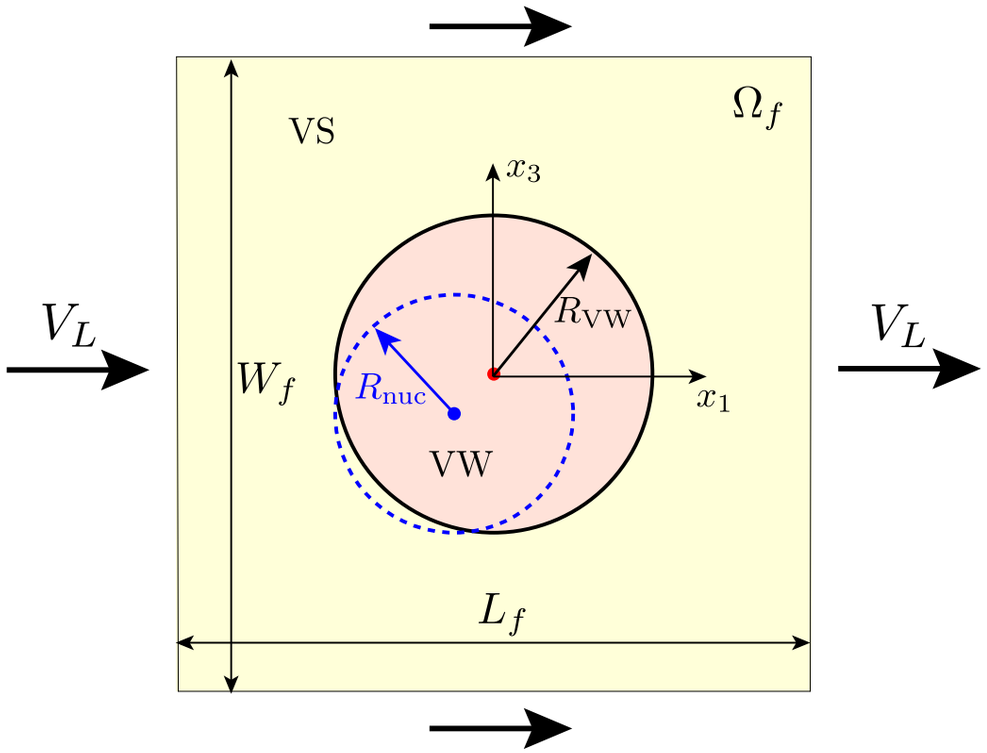
2) Fast evaluations on GPU

- Given the finite temporal support, the convolution kernels are stored as three dimensional arrays, and their pre-summed counterparts as two dimensional arrays. These are transferred once to GPU and remain in device memory throughout the simulation.

- The spectral slip histories are instead stored using fixed length circular buffers and are never physically shifted!

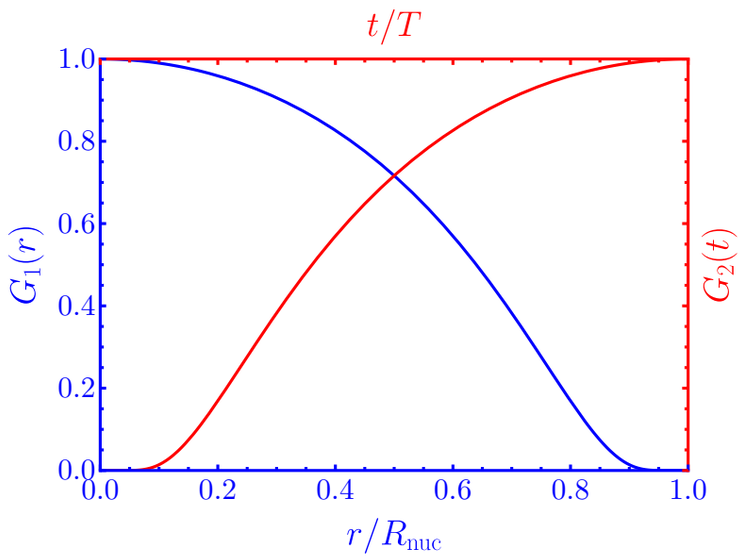
- All the convolutions are contracted on the GPU to assemble $F_j(n_1, n_3, t)$, then transferred back to CPU, and transformed back to physical space using inverse FFT.

Verification: SCEC SEAS benchmark BP7-FD-A



Parameter	Value
Density, ρ	2670 kg m^{-3}
Shear-wave speed, c_s	3.464 km s^{-1}
Shear modulus, μ	32.04 GPa
Poisson's ratio, ν	0.25
Direct-effect parameter inside the VW asperity, a_0	0.004
Direct-effect parameter in the VS region, $a_{\max}$	0.016
Evolution-effect parameter, b_0	0.01
Effective normal stress, σ'	25 MPa
Characteristic slip distance, D_c	0.50 mm
Plate rate, V_p	10^{-9} m s^{-1}
Imposed creep rate, V_L	10^{-9} m s^{-1}
Initial slip rate, V_{init}	10^{-9} m s^{-1}
Reference slip rate, V_*	10^{-6} m s^{-1}
Reference friction coefficient, f_*	0.6
Velocity-weakening radius, R_{VW}	200 m
Rate-and-state region half-length, L_f	400 m
Rate-and-state region half-width, W_f	400 m
Benchmark-recommended cell size	10 m
Final simulation time	10 yr
Nucleation perturbation amplitude, $\Delta\tau_0$	1.75 MPa
Nucleation radius, R_{nuc}	150 m
Nucleation ramp duration, T	1 s
Nucleation center, $(x_{1,\mathrm{nuc}}, x_{3,\mathrm{nuc}})$	(−50, −50) m

Time-dependent, spatially compacted perturbation for artificial nucleation



$$G_2(t) = \begin{cases} \exp\left[\dfrac{(t-T)^2}{t(t-2T)}\right], & 0 < t < T, \\ 1, & t \geq T. \end{cases}$$

$$G_1(r) = \begin{cases} \exp\left(\dfrac{r^2}{r^2 - R_{\mathrm{nuc}}^2}\right), & r < R_{\mathrm{nuc}}, \\ 0, & r \geq R_{\mathrm{nuc}}, \end{cases}$$

Rate-and-state frictional resistance, with regularized formulation:

$$\boldsymbol{\tau}_f = \sigma' f_c(V,\theta)\frac{\dot{\boldsymbol{\delta}}}{V}, \qquad V = \left\|\dot{\boldsymbol{\delta}}\right\|,$$

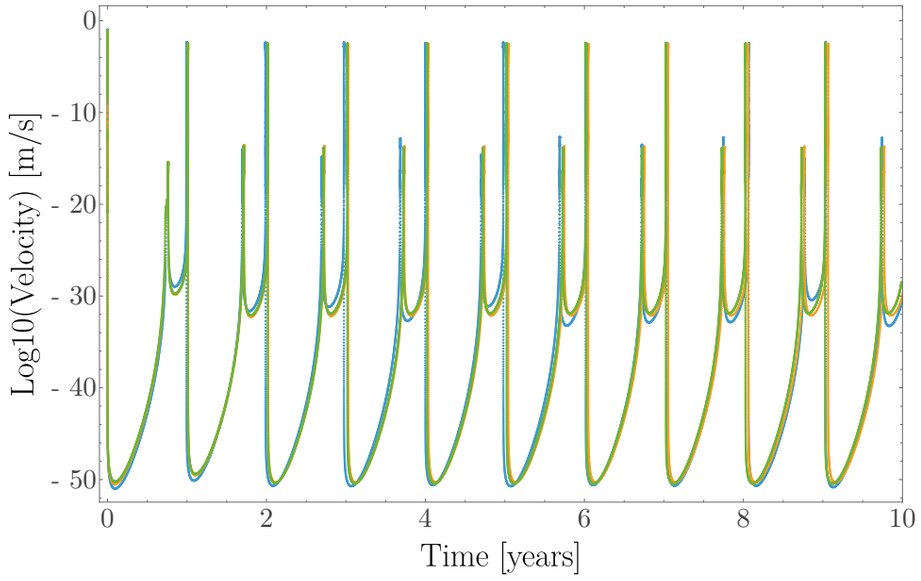
$$f_c(V,\theta) = a\,\sinh^{-1}\left[\frac{V}{2V_*}\exp\left(\frac{f_* + b\ln\left(V_*\theta/D_c\right)}{a}\right)\right]$$

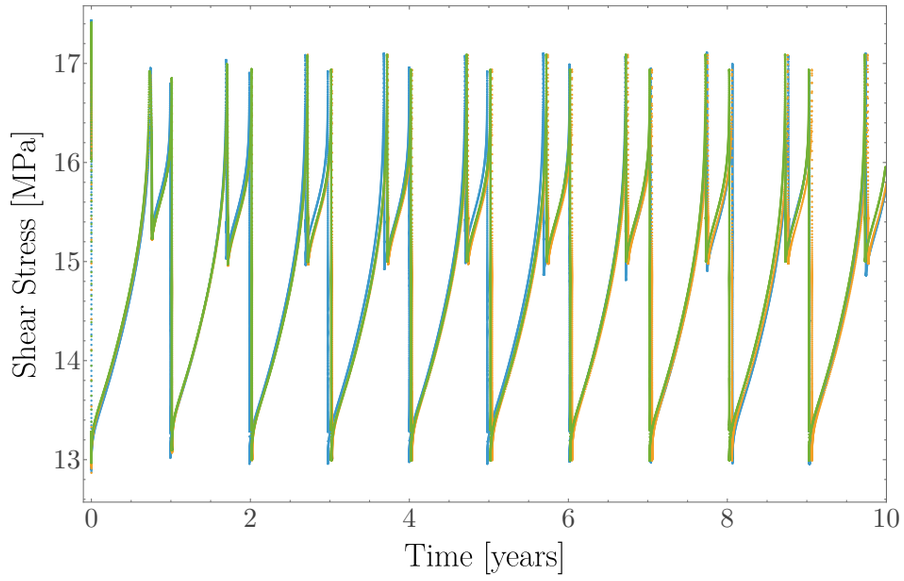
$$\frac{d\Theta}{dt} = \frac{V_0}{D_c}e^{-\Theta} - \frac{V}{D_c}$$ Aging law

Valere Lambert, Junle Jiang, and Brittany Erickson. SEAS Benchmark Problems BP7-QD/FD-A/S.
URL: https://strike.scec.org/cvws/seas/download/SEAS_BP7_May2023rev.pdf

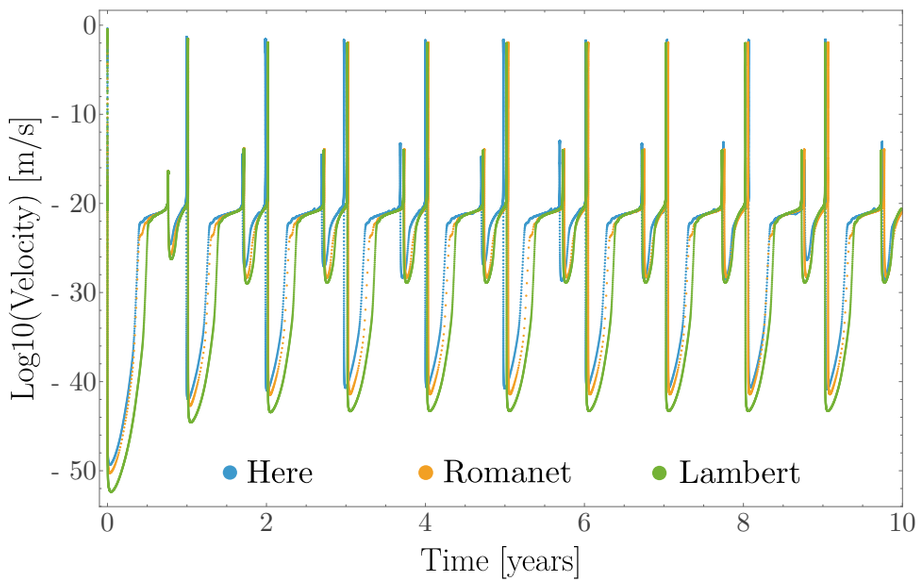
Verification: SCEC SEAS benchmark BP7-FD-A

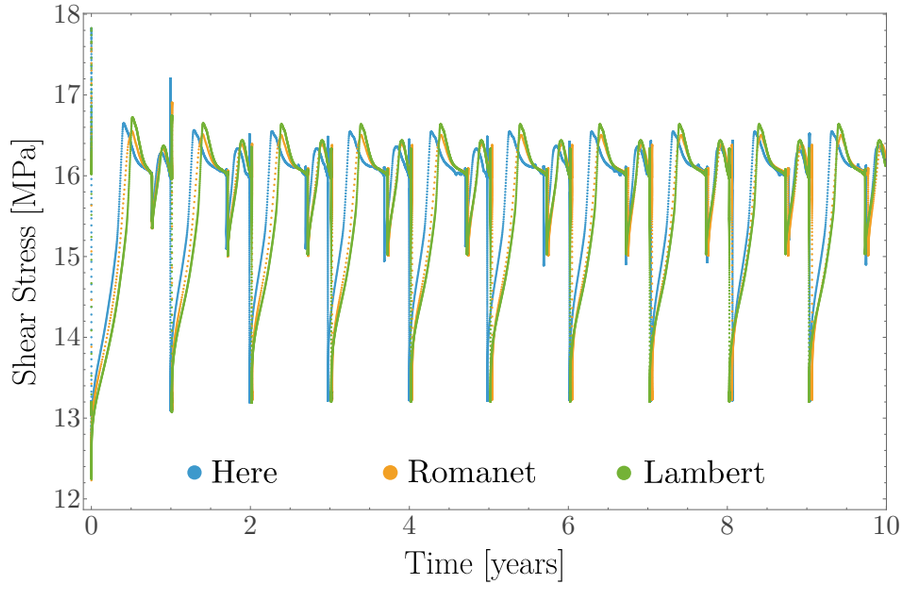
Observation point
$(x_1, x_3) = (0, 0)$ m





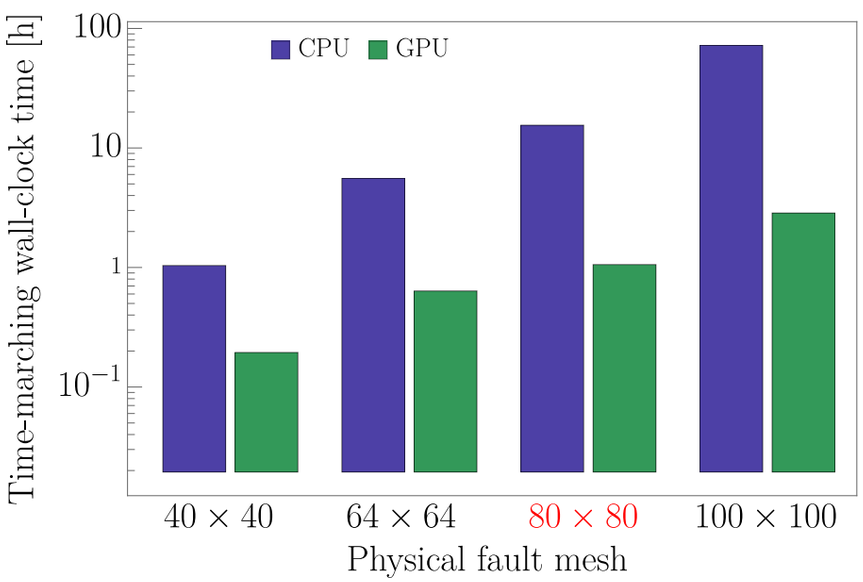
Observation point
$(x_1, x_3) = (100, 100)$ m



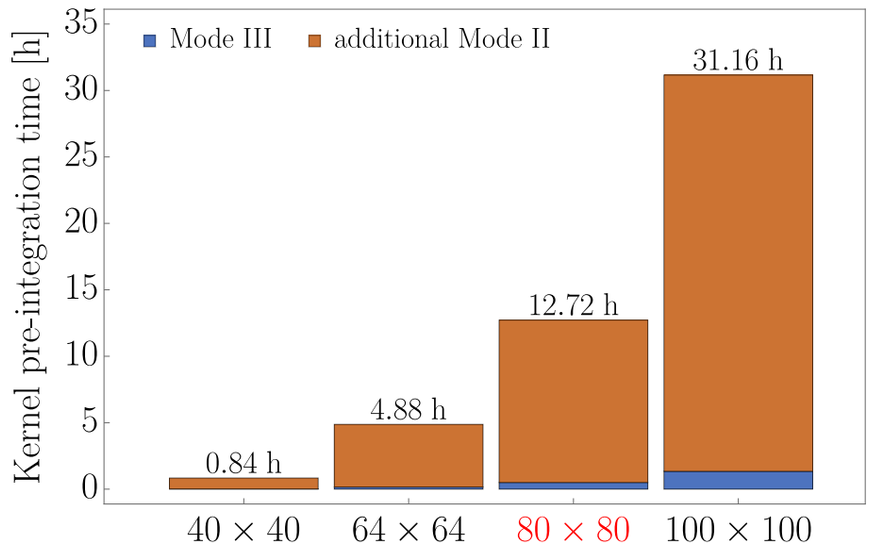


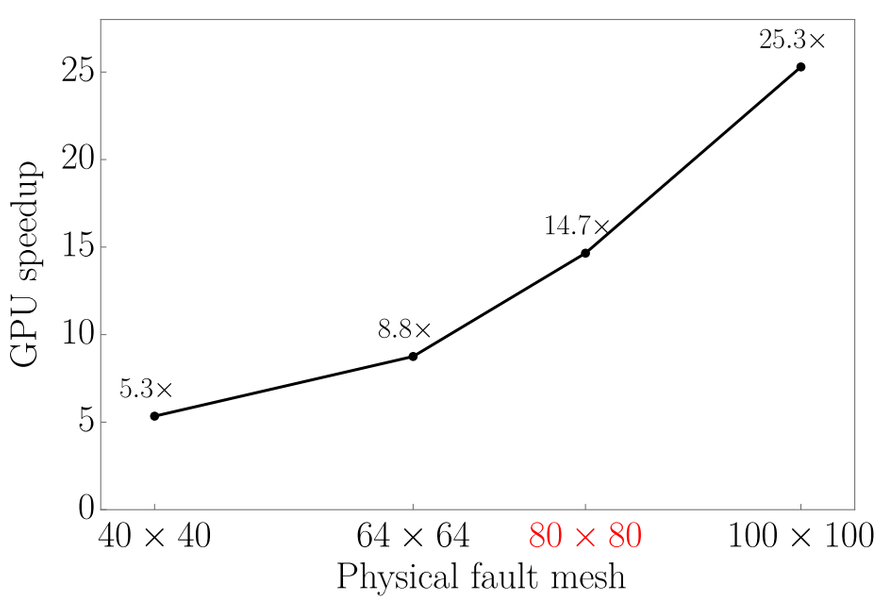
Good agreement with independently obtained numerical solutions (downloaded from SCEC SEAS code-verification repository)

Verification: SCEC SEAS benchmark BP7-FD-A

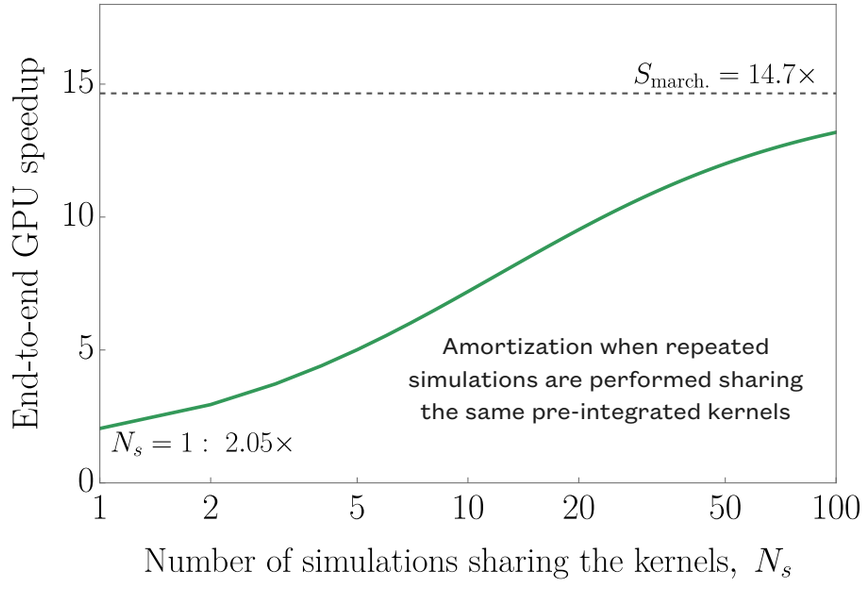


*For CPU: 16 CPU cores on Northwestern University's Quest high-performance computing cluster





**For GPU: single NVIDIA A100 PCIe GPU with 40 GB of device memory



Physical grid	Auxiliary grid	N_t	Kernels [GB]	Slip-rate histories [GB]	Spectral-slip histories [GB]	Total [GB]
40×40	80×80	566	0.058	0.058	0.116	0.232
64×64	128×128	906	0.238	0.238	0.475	0.950
80×80	160×160	1132	0.464	0.464	0.927	1.855
100×100	200×200	1415	0.906	0.906	1.811	3.622

Memory consumption very limited!

Conclusions & future steps

- **2D fault networks**: a hybrid spectral/space-time BIEM enables efficient, fully dynamic SEAS simulation on complex fault systems.

- **3D planar fault**: first implementation of non-periodic SBIEM to eliminate artificial fault replicas. It provides exact, finite-memory elastodynamic kernels.

- **GPU acceleration**: a 14.7x speedup in time marching is achieved for the SCEC BP7-FD-A benchmark. No need of supercomputer.

- **Future steps**: Extend the divide-and-conquer approach to 3D, and investigate how fault geometry and dynamic stress transfer control earthquake nucleation, propagation and interactions across complex fault networks at regional scale, over long time-scales.

Thank you for listening!

federico.ciardo@northwestern.edu