Finite-Element Methods and Exascale Algorithms for Fully Relativistic Noncollinear Pseudopotential Density Functional Theory

Nikhil Kodali | Advisor: Prof. Phani Motamarri | Computational & Data Sciences, IISc Bengaluru | SC26 Doctoral Showcase

THE HPC PROBLEM

Simulating magnetic materials with relativistic physics (DFT) means solving a **sparse complex eigenproblem**, re-solved every outer (SCF) step: up to $\sim 10^8$ **unknowns** and $\sim 10^5$ **eigenvectors**. Each electron state is a two-component complex spinor, so state-of-the-art relativistic studies typically reach only ~500–2,000 electrons.

CENTRAL IDEA

Make errors from inexact arithmetic vanish at convergence, so low precision and compressed communication become safe.

1. **FE formulation** — Phys. Rev. B 111, 195129 (2025)
2. **R-ChFSI solver** — Comput. Phys. Commun. 327, 110239 (2026)
3. **Low precision** — SC26 (accepted) arXiv:2605.30128
4. **Results** — 4a • 4b • 4c — Aurora, Frontier

My first-author contributions, all in the open-source DFT-FE code

HEADLINE RESULTS

- **~80,000** electrons, first time on 800 Aurora nodes
- **2.27×** faster full solve, same energy
- **3.5×** faster filtering vs. FP64
- **8–10×** CPU-to-CPU speedup over QE in minimum wall time

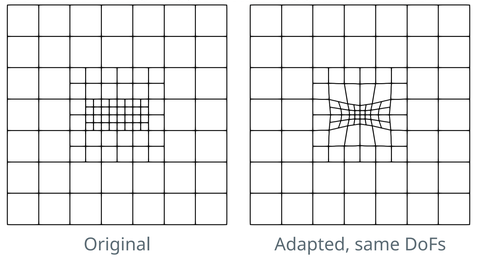
1 Finite Elements: Local and GPU-Ready

- Local basis: each cell is a small **dense batched GEMM**.
- Applying H: only **point-to-point** messages; no FFTs.
- Spinors; energies, forces, stresses in one formulation.

$$\mathbf{HX} = \mathbf{MX}\Lambda$$

Hamiltonian **H**, overlap **M**, eigenvectors **X** (electron states)

ADAPTIVE MESHES: MORE ACCURACY PER DOF



Original | Adapted, same DoFs

up to 28× lower energy error at the same number of DoFs (best case: Fe 11×, Ge 18×, Te 28×)

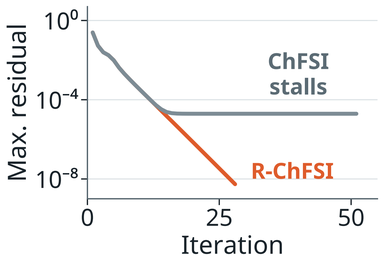
2 R-ChFSI: Robust to Inexact Math

R-ChFSI rewrites the Chebyshev recurrence around the **residual** $\mathbf{R}^{(i)} = \mathbf{HX}^{(i)} - \mathbf{MX}^{(i)}\Lambda^{(i)}$, which vanishes at convergence.

ChFSI	$\|\Delta^{(i)}\| \le \varepsilon_C$	error floor ⇒ stalls
R-ChFSI	$\|\Delta^{(i)}\| \le \varepsilon_R \|\mathbf{R}^{(i)}\|$	**error → 0**

Guarantee: once $|C_p(\lambda_n)| - |C_p(\lambda_{n+1})| > 2\|\mathbf{H}\|\,\varepsilon_R(1+\tan\theta_i)$, the angle θ_i to the wanted eigenspace shrinks at every later step.

$\Delta^{(i)}$: filter error • $\varepsilon_C, \varepsilon_R$: constants set by scaled machine precision γ_m and $\|\mathbf{D}^{-1} - \mathbf{M}^{-1}\|$ • C_p: Chebyshev filter

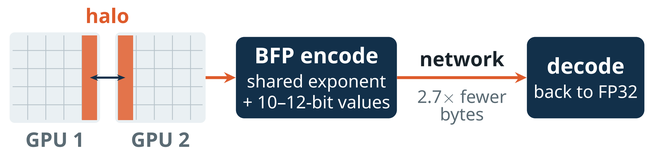


85.8M unknowns, 13,500 states, 100 Aurora nodes, approximate overlap inverse, $p = 40$.

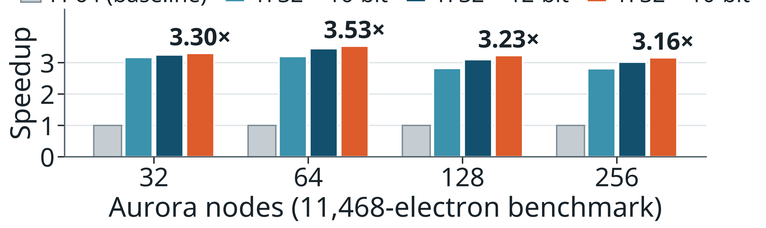
Same for 1.7M and 25.7M unknowns and all p = 20–120.

Takeaway: precision becomes a **tunable parameter**, not a risk.

3 Low Precision + Compressed Halos



Compute in FP32 on TF32 matrix engines • **Compress** point-to-point messages • **Keep** the residual in FP64.



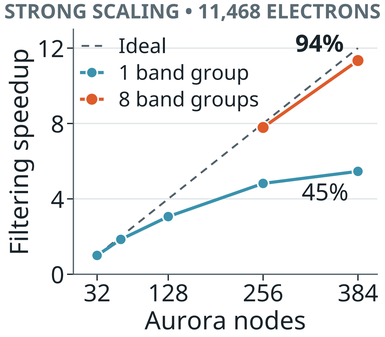
Same answer: 2.27× faster full solve, $\Delta E = 10^{-10}$ Ha/atom.

4 Results

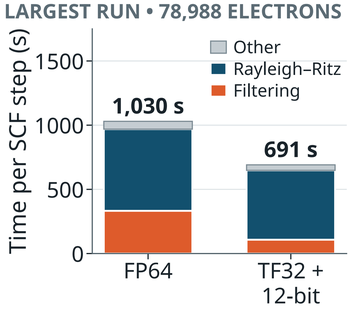
what 1 + 2 + 3 deliver together

4a Scaling on Aurora

STRONG SCALING • 11,468 ELECTRONS

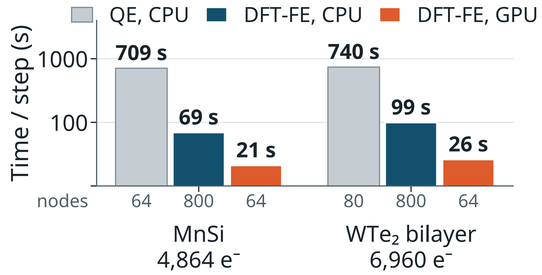


LARGEST RUN • 78,988 ELECTRONS



- **Band parallelization** restores filtering efficiency: **45% → 94%** at 384 nodes.
- On 800 nodes (9,600 GPU tiles), filtering is **3.1×** faster; **Rayleigh–Ritz** (dense subspace solve) is now the bottleneck.

4b Faster Than Plane Waves



Baseline: Quantum ESPRESSO (QE). Each code at its fastest node count (row below bars), same accuracy ~ 10^{-4} Ha/atom. CPU: Nurion (Intel KNL). GPU: Frontier (AMD).

- **Accurate** to **< 10^{-4} Ha/atom** (7 systems).
- **CPU node-hours:** QE cheaper below ~8–10k e⁻, DFT-FE above.
- **DFT-FE GPU vs. CPU: 40–70×** fewer node-hours.

4c Science: Vacancies in a 2D Magnet

Fe_3GeTe_2, a 2D magnet for spin electronics: **216-atom** cells, full relativistic spin physics.



Fe(2) preferred — the favoured vacancy site

Vacancies repel — nearby vacancies stay apart

Small canting — ferromagnetic order survives

Local response — moments heal within a few Å

Gamma-point calculations.

Take-home

1. Low precision and compression are *safe* when the solver is built for them.
2. Local FE methods suit GPUs and scale across nodes.
3. Open-source relativistic DFT at ~79k electrons.

Website | DFT-FE code

nikhilkodali@iisc.ac.in

NATIONAL SUPERCOMPUTING MISSION · PMRF · NVIDIA · Google Research · C-DAC