MMM 2026

View Abstract

CONTROL ID: 4551223

TITLE: Advanced micromagnetic simulation of several billion magnetic unknowns on multiple GPUs using JAX and AI-assisted programming

PRESENTATION TYPE: Oral

Symposium Session: Magnetic Materials Design and Discovery

AUTHORS (LAST NAME, FIRST NAME): Suess, Dieter[1]; Kraft, Robert[1]; Bruckner, Florian[1]; Abert, Claas[1]

INSTITUTIONS (ALL): 1. University of Vienna, Vienna, Austria.

STATUS:

Dieter Suess : academic

ABSTRACT BODY:

Abstract Body: We review recent developments in our group on GPU-accelerated micromagnetic simulation software built entirely in JAX, covering both finite-difference (jaxFD) and finite-element (jaxFMM [1]) formulations, explicit Dormand-Prince and implicit preconditioned BDF time integrators [2], and multi-GPU parallelization via JAX's shard_map. A key enabler of this rapid development has been agentic engineering, which dramatically accelerates the development of production-grade GPU code.

The finite-difference solver computes the demagnetization field by FFT convolution with the Newell tensor and is validated on muMAG Standard Problem 4. Multi-GPU scaling distributes both the demagnetization kernel and the FFT workspace across devices, enabling simulations with several billion finite-difference cells. On 8 NVIDIA B200-192GB GPUs, we demonstrate stray-field evaluation of a 1504^3 cube (3.4 billion cells) in 492 ms (Fig. 1a). Fig. 1(b) shows the transient domain state of a 7 x 7 x 7 um^3 Permalloy cube (one billion cells, dx = 7 nm) after 5 ns of LLG dynamics (integration time 5.6 min, setup time 6 min on 8xB200), corresponding to roughly 1 minute of simulation time per nanosecond of physical time. At these problem sizes, micromagnetic simulations bridge the gap from nanoscale devices to mesoscale structures such as thin films of millimeter lateral extent.

For unstructured meshes, jaxFMM [1] provides a fast multipole method. We compare the implicit preconditioned BDF scheme with explicit Dormand-Prince integration for both stiff and non-stiff problems on finite-difference and finite-element meshes, all running purely on GPU.

References: [1] R. Kraft et al., arXiv:2511.15269 (2025).
[2] Suess, Dieter, et al. *Journal of Magnetism and Magnetic Materials* 248.2 (2002): 298-311.

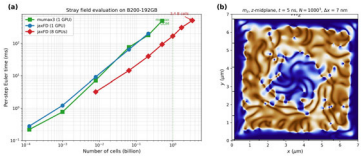

KEYWORDS: multiple GPUs, large scale micromagentics, vibe programming , AI.

CURRENT CATEGORY: 4. Magnetization Dynamics and Micromagnetics

CURRENT SUB-CATEGORY: e. Modeling of Magnetic Materials

Attendance at Conference: I acknowledge that I have read the above statement regarding the requirement that an author of this presentation is obliged to attend and present the paper at the conference.

Previous Presentation: INTERMAG 2026

Manuscript?: No

Mailing List: (none)

Early Notification for Visa: (none)

Nationality: (none)

AWARDS:

Advisor Name: (none)

Expected Graduation Date: (none)

Letter of Support: (none)

Supplementary Information: (none)
