

Nils Rottstegge

M.Sc. Student in Computer Science, ETH Zurich
High-Performance Computing | GPU Computing

nils@rottstegge.ch

linkedin.com/in/nilsrottstegge | github.com/nrottstegge

EDUCATION

ETH Zurich, Zurich, CH, February 2026 to September 2028 (expected)
Master of Science (M.Sc.) in Computer Science, major: Data Management Systems

ETH Zurich, Zurich, CH, September 2022 to March 2026
Bachelor of Science (B.Sc.) in Computer Science, final grade: 5.19/6.00 (Swiss scale)

Hochschule der Medien, Stuttgart, GER, September 2018 to August 2022
Bachelor of Engineering (B.Eng.) in Audiovisual Media, final grade: 1.3 (German scale, 1.0 best)

EXPERIENCE

Forschungszentrum Jülich, Jülich, GER, August 2026 to October 2026
Guest Student Programme, GPU Research, Jülich Supercomputing Centre (JSC)

Self-employed, Switzerland, July 2015 to Present
Video Technician / Engineer (TV)

- Operated live-broadcast video systems, resolved technical faults and led production teams.

STAN Studios, Stuttgart, GER, March 2021 to April 2022
Working Student, Process Optimisation & Business Development

- Reduced administrative tasks across all operational departments through shared responsibility for introducing Asana for project management. Coordinated projects and prepared reports for account management.

PROJECTS

Adaptive Trees on a GPU | C++20, CUDA, CUB, CUDA Graphs (JSC, JUPITER)

- Developed adaptive octrees and binary trees for the Fast Multipole Method (FMM) with work-aware refinement and traversal-free hierarchy construction. Accelerated tree updates/view construction by up to $2.4\times$ over Cornerstone (baseline) and FMM execution by up to $4.8\times$ over dense trees.

HBDIA: Hybrid Blocked-DIA | C++, CUDA, OpenMP, MPI (Bachelor's thesis, ETH SPCL)

- Designed HBDIA and a two-pass COO conversion. Implemented GPU sparse matrix-vector multiplication (SpMV) with coalesced access, CPU COO fallback and model-guided workload partitioning on GH200.
- Achieved up to $2.2\times$ speedup over cuSPARSE on stencil matrices, $1.6\times$ on a matrix from PDE-constrained optimisation, and $1.5\times$ in a 3D fluid-channel simulation.

Multi-GPU HPCG | C++, CUDA, MPI, NCCL

- Extended a DIA-based High Performance Conjugate Gradient (HPCG) solver to multiple GPUs using CUDA-aware MPI and NCCL halo exchange. Validated on **512 nodes and 2,048 GH200 GPUs** on the **ALPS supercomputer**.
- Optimised SpMV and symmetric Gauss-Seidel (SymGS) memory layouts for $1.50\times$ and $1.78\times$ speedups over the original single-GPU solver on one GH200. Ported to the AMD cluster Beverin.

Porous Convection on GPUs | Julia, ParallelStencil.jl, MPI

- Implemented 2D and 3D thermal convection solvers using finite differences, pseudo-transient relaxation and MPI domain decomposition. Validated distributed execution on **eight GPUs across two nodes**.

SKILLS

Programming: C/C++, Python (including ML), Julia, Java, SQL / PostgreSQL

Parallel & GPU computing: CUDA, CUDA Graphs, HIP/hipify, OpenACC & OpenMP, CUDA-aware MPI, NCCL

Libraries & tools: Thrust, CUB, cuBLAS, cuSPARSE, Git, CMake, Catch2/CTest, Slurm

Profiling & debugging: Nsight Systems / Compute, Compute Sanitizer, cuda-gdb, roofline analysis

HPC systems: ALPS and JUPITER supercomputers | **Languages:** German (native), English (C1)

TRAINING & AWARDS

Parallel Programming with MPI and OpenMP (27 h), JSC, August 2026

GPU course: Unified/shared memory, coalescing, streams/events, GPUDirect, communication/compute overlap.

DMV-Abiturpreis, award for outstanding Abitur mathematics performance, 2017