


ExaHyPE

Status of ChEESE-2P Flagship Codes



ChEESE




ExaHyPE (www.exahype.org)

Engine to solve hyperbolic
PDE systems using high-
order ADER-DG
discretisation on tree-
structured adaptive
Cartesian grids.

Repositories:

<https://gitlab.lrz.de/hpcsoftware/Peano>
Modified BSD license





A decorative vertical bar on the left side of the slide, featuring a dark gray background with yellow geometric patterns including circles, semi-circles, and arcs.

ExaHyPE Team:

- Michael Bader (TUM)
- Alice-Agnes Gabriel (LMU)
- **Mario Wille (TUM, ExaHyPE contact WP2)**
- Duo Li (LMU)



- CPUs: performance bottlenecks of ExaHyPE-1/Peano-3 identified in ChEERE-1P (tree traversal)
- Major rewrite Peano-4/ExaHyPE-2 in 2021/2022 – no extensive scalability tests, yet (ToDo for audit)
- GPUs: started GPU port 2022; ISC technical paper accepted on experienced OpenMP bottleneck
- Test GPU platforms: NVIDIA A100 (Alex@NHR Erlangen), AMD Instinct MI250 (JURECA@JSC Jülich)



ExaHyPE Release Information:

- Latest versions: Peano-4/ExaHyPE-2
see <https://gitlab.lrz.de/hpcsoftware/Peano>
- Documentation: “Read the Docs” (in progress)
- MPI+OpenMP (CPU; experimental for GPU)
- Mini-app: collection of small ExaHyPE benchmark applications
- Pre-built Docker images:
<https://hub.docker.com/u/peanoframework>

Dependencies:

- x86 CPUs: LIBXSMM: <https://libxsmm.readthedocs.io/>
- Python3 packages (Jinja2, SymPy, etc.)



Performance Portable ExaHyPE:

- Kernel benchmarks to evaluate performance and possible bottlenecks of numerical solvers
- Multiple GPU offloading backends available: OpenMP, ISO-C++, SYCL, HIP (planned)
- Evaluation of parallel efficiency of kernels
- Work on further improving efficiency: multiple threads offload to a GPU (cf. ISC paper)

