Scalable, Data Driven Plasma Simulations with PIConGPU

A. Huebl[1,2]

R. Widera,[1] M. Garten,[1,2] R. Pausch,[1,2] K. Steiniger,[1] S. Bastrakov,[1] F. Meyer,[1] K. Bastrakova,[1] A. Debus,[1] T. Kluge,[1] S. Ehrig,[1] M. Werner,[1,2] B. Worpitz,[3] A. Matthes,[1,2] S. Rudat,[1,2] S. Starke,[1] and M. Bussmann[1]

[1] Helmholtz-Zentrum Dresden – Rossendorf [2] Technische Universität Dresden

[3] LogMeIn, Inc.

Platform for Advanced Scientific Computing Conference (PASC19)

Zürich (Switzerland), *June 13th 2019*

 DOI: 10.14278/rodare.131

HZDR

HELMHOLTZ ZENTRUM DRESDEN ROSSENDORF

Laser-Plasma Physics

Plasma Instabilities

- Astrophysics
- Control of Laser-Plasmas

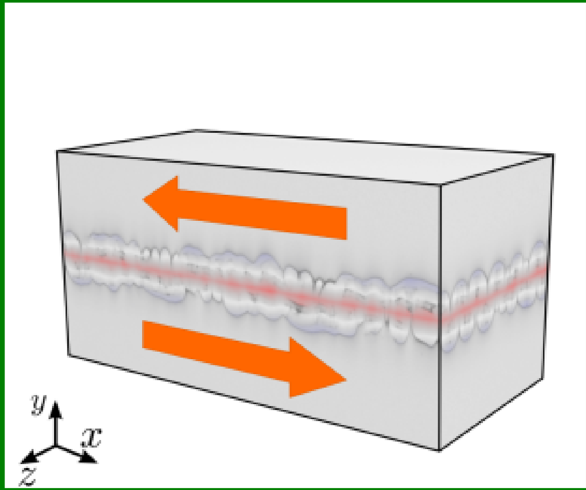

Electron Acceleration with Lasers

- Compact X-Ray sources
- Push the Energy Frontier

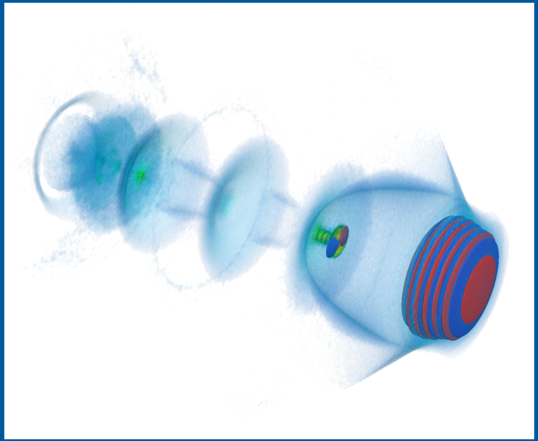

Ion Acceleration with Lasers

- Compact Ion Sources
- HED Physics

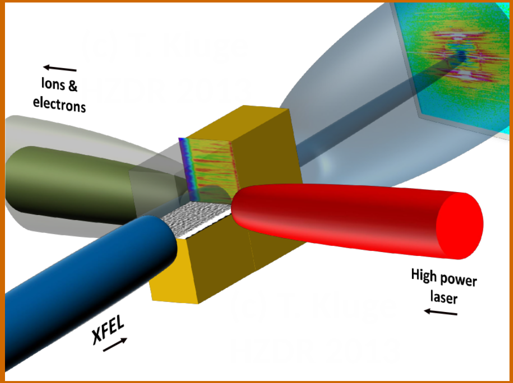

PIConGPU

Scalability

PIConGPU

- **Eulerian: electro-magnetic fields**
- **Lagrangian: particles in Vlasov-equation**

resolve:

- ω_{pe}
- ω_c

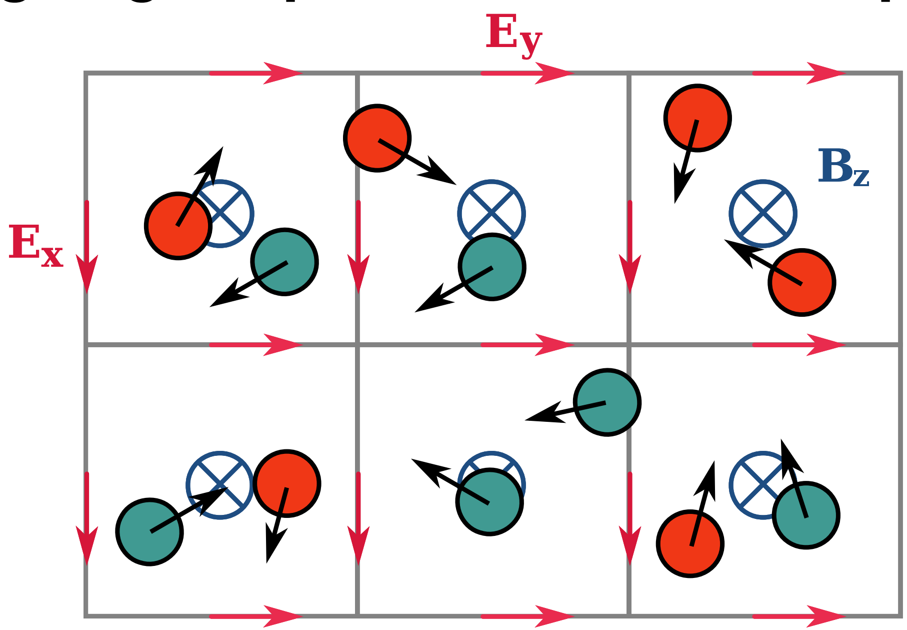

3D 3V

Weak Scaling

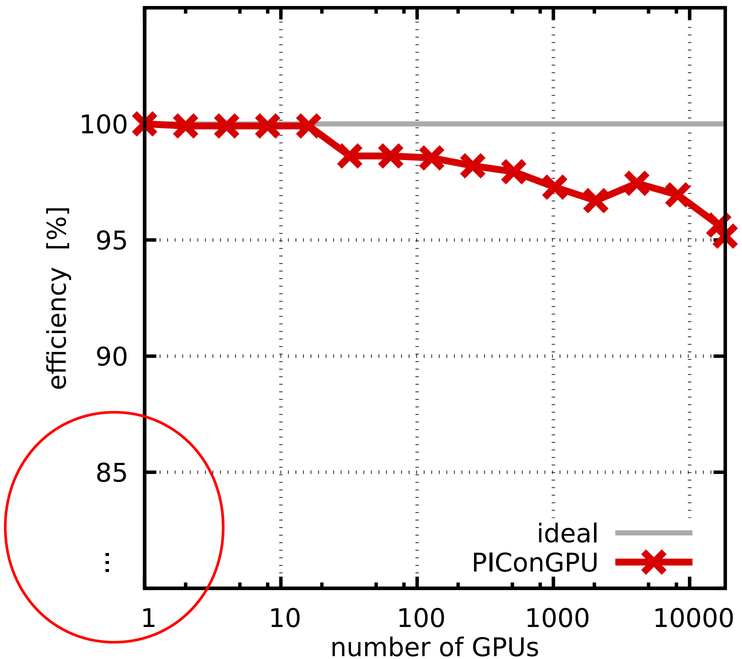

Strong Scaling

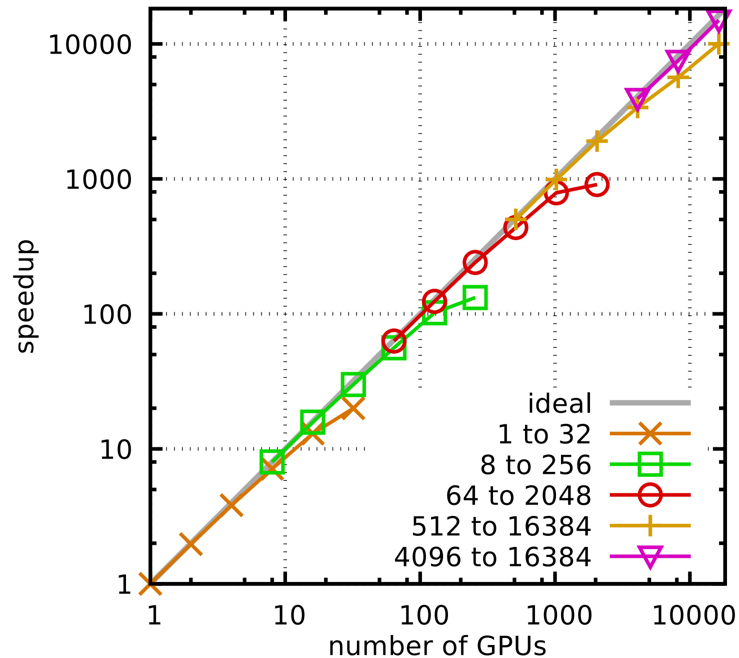

M. Bussmann, A. Huebl et al., SC'13,
DOI:10.1145/2503210.2504564

Domain Decomposition

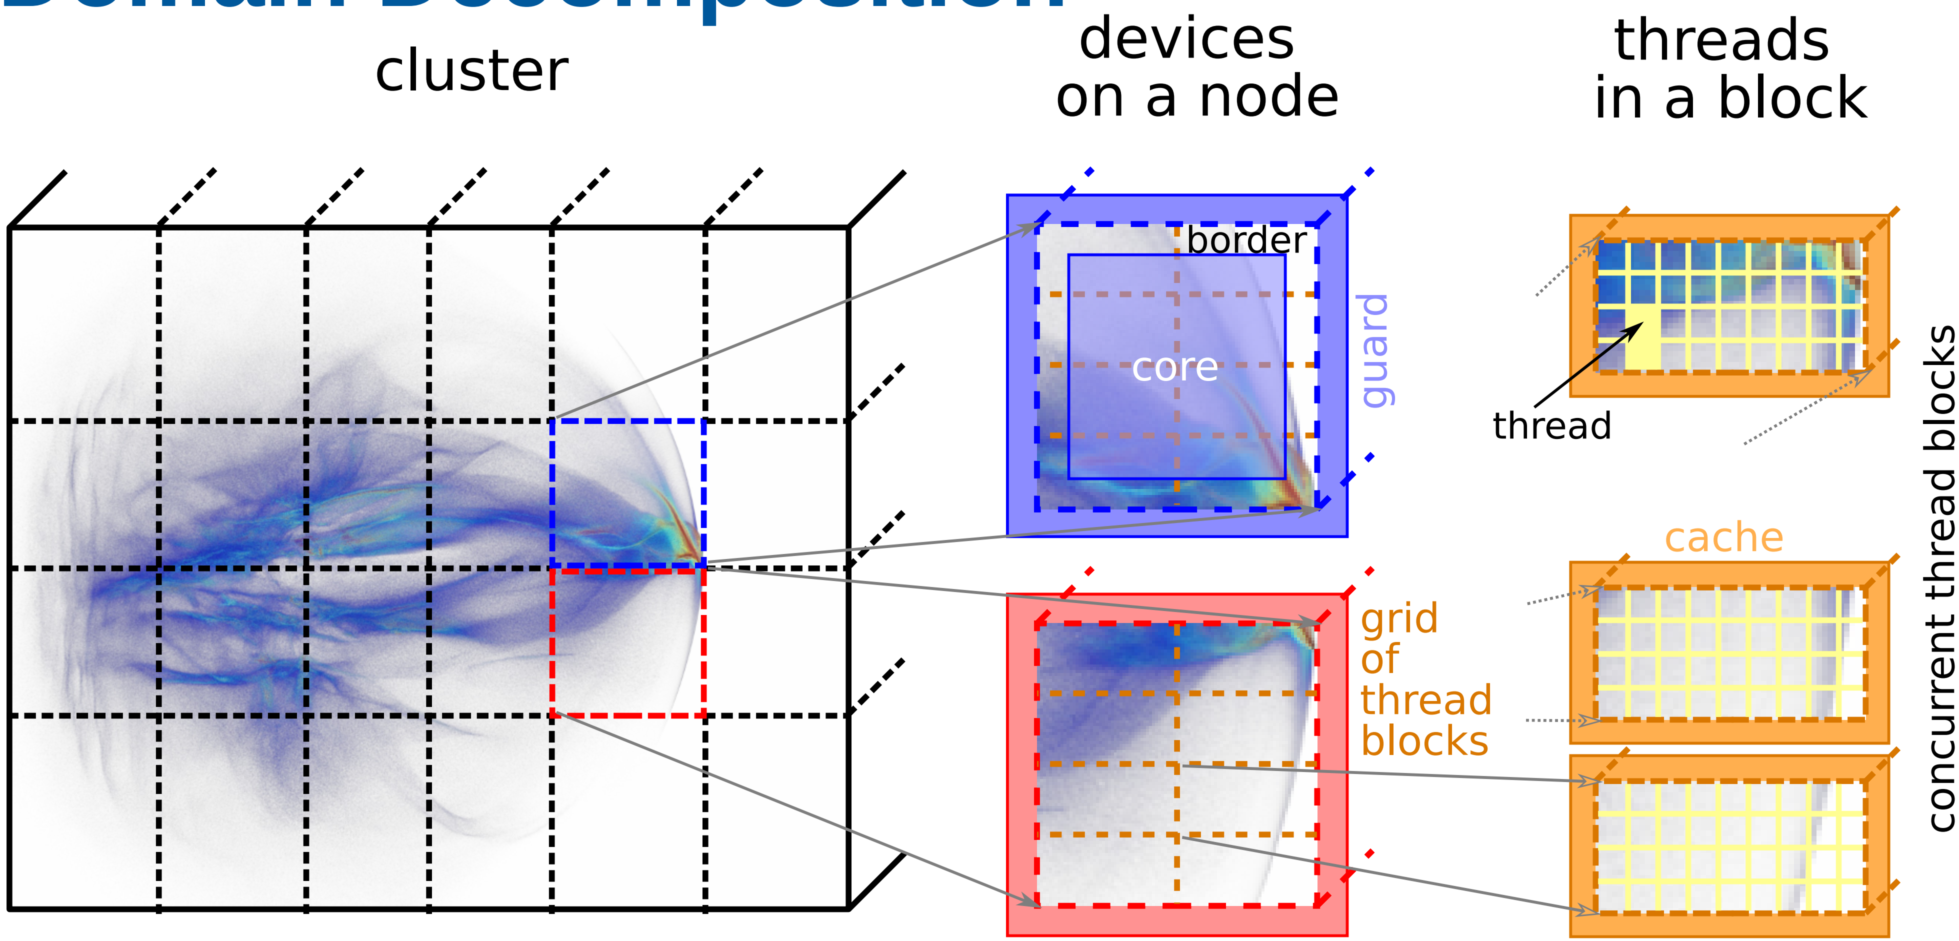

PIConGPU

Data Structures For Manycore

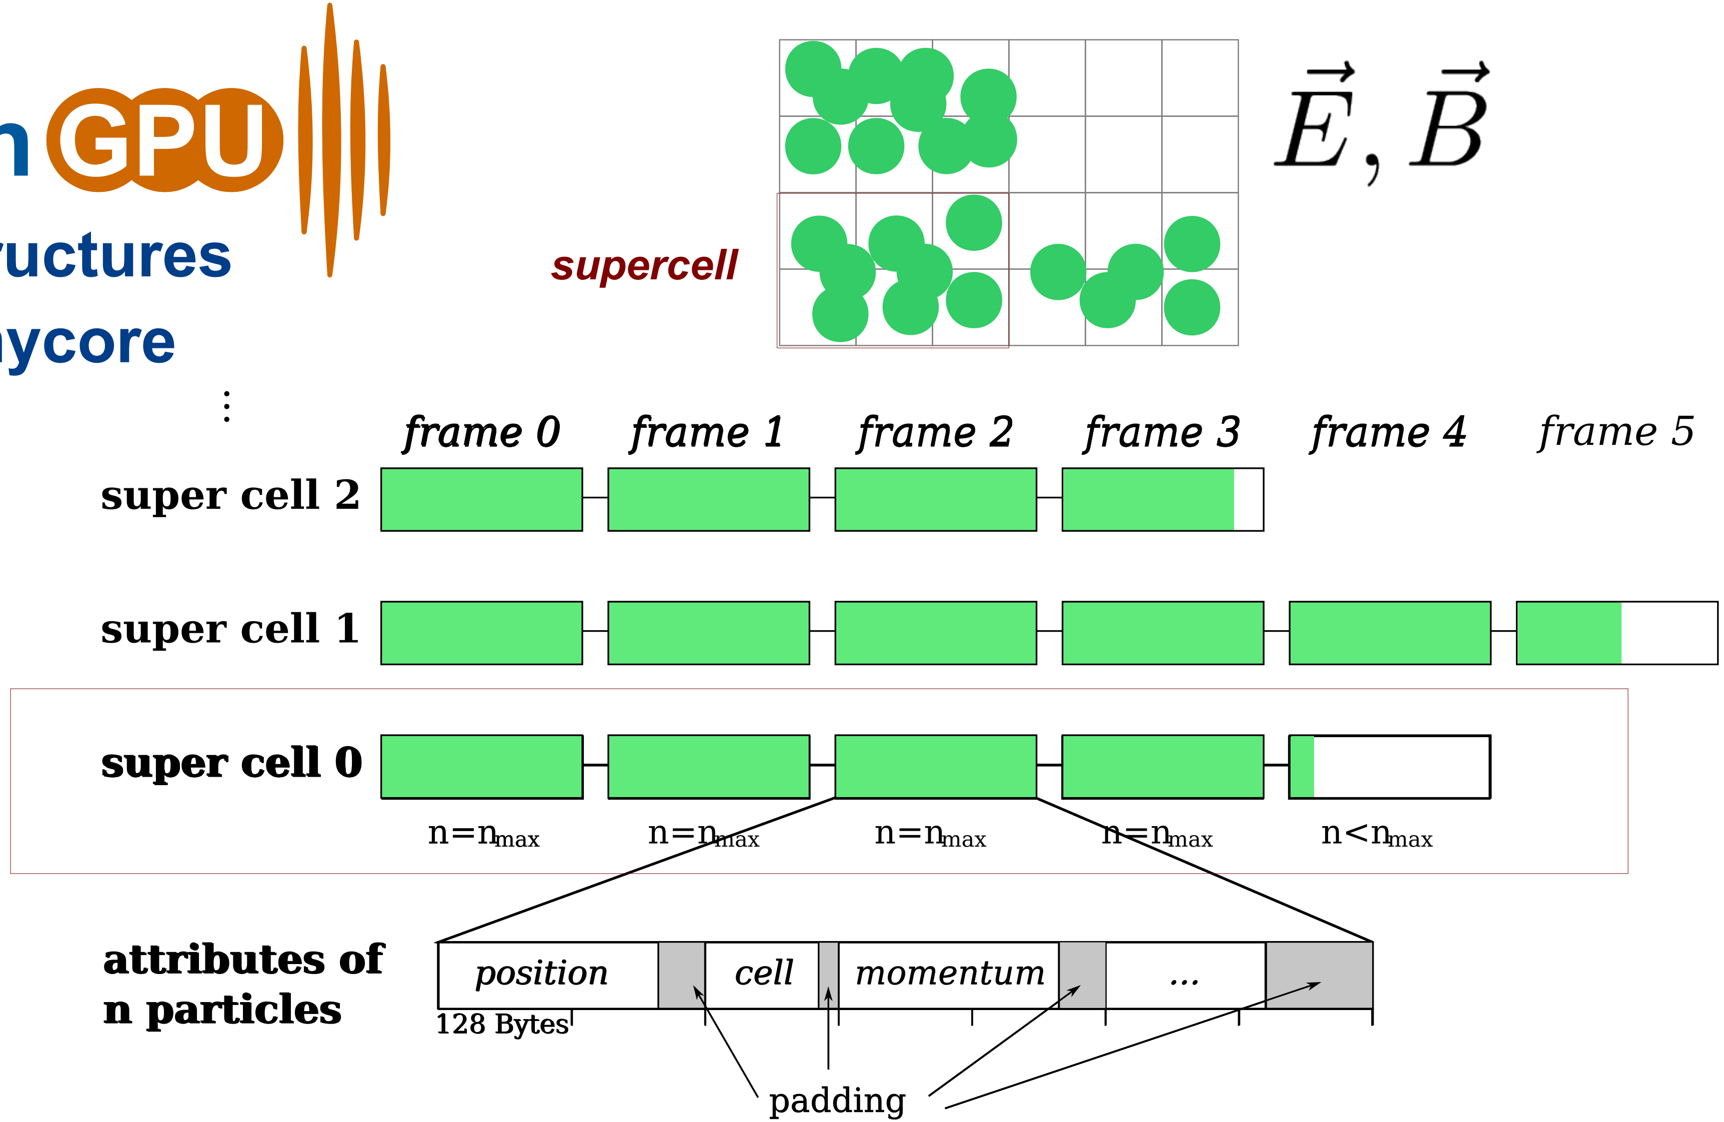

H. Burau et al., IEEE Trans. Plasma Sci. (2010), DOI:10.1109/TPS.2010.2064310
M. Bussmann et al., SC'13 (2013), DOI:10.1145/2503210.2504564
A. Huebl, Diplomarbeit (2014), DOI:10.5281/zenodo.15924

HPC Application Software Stack

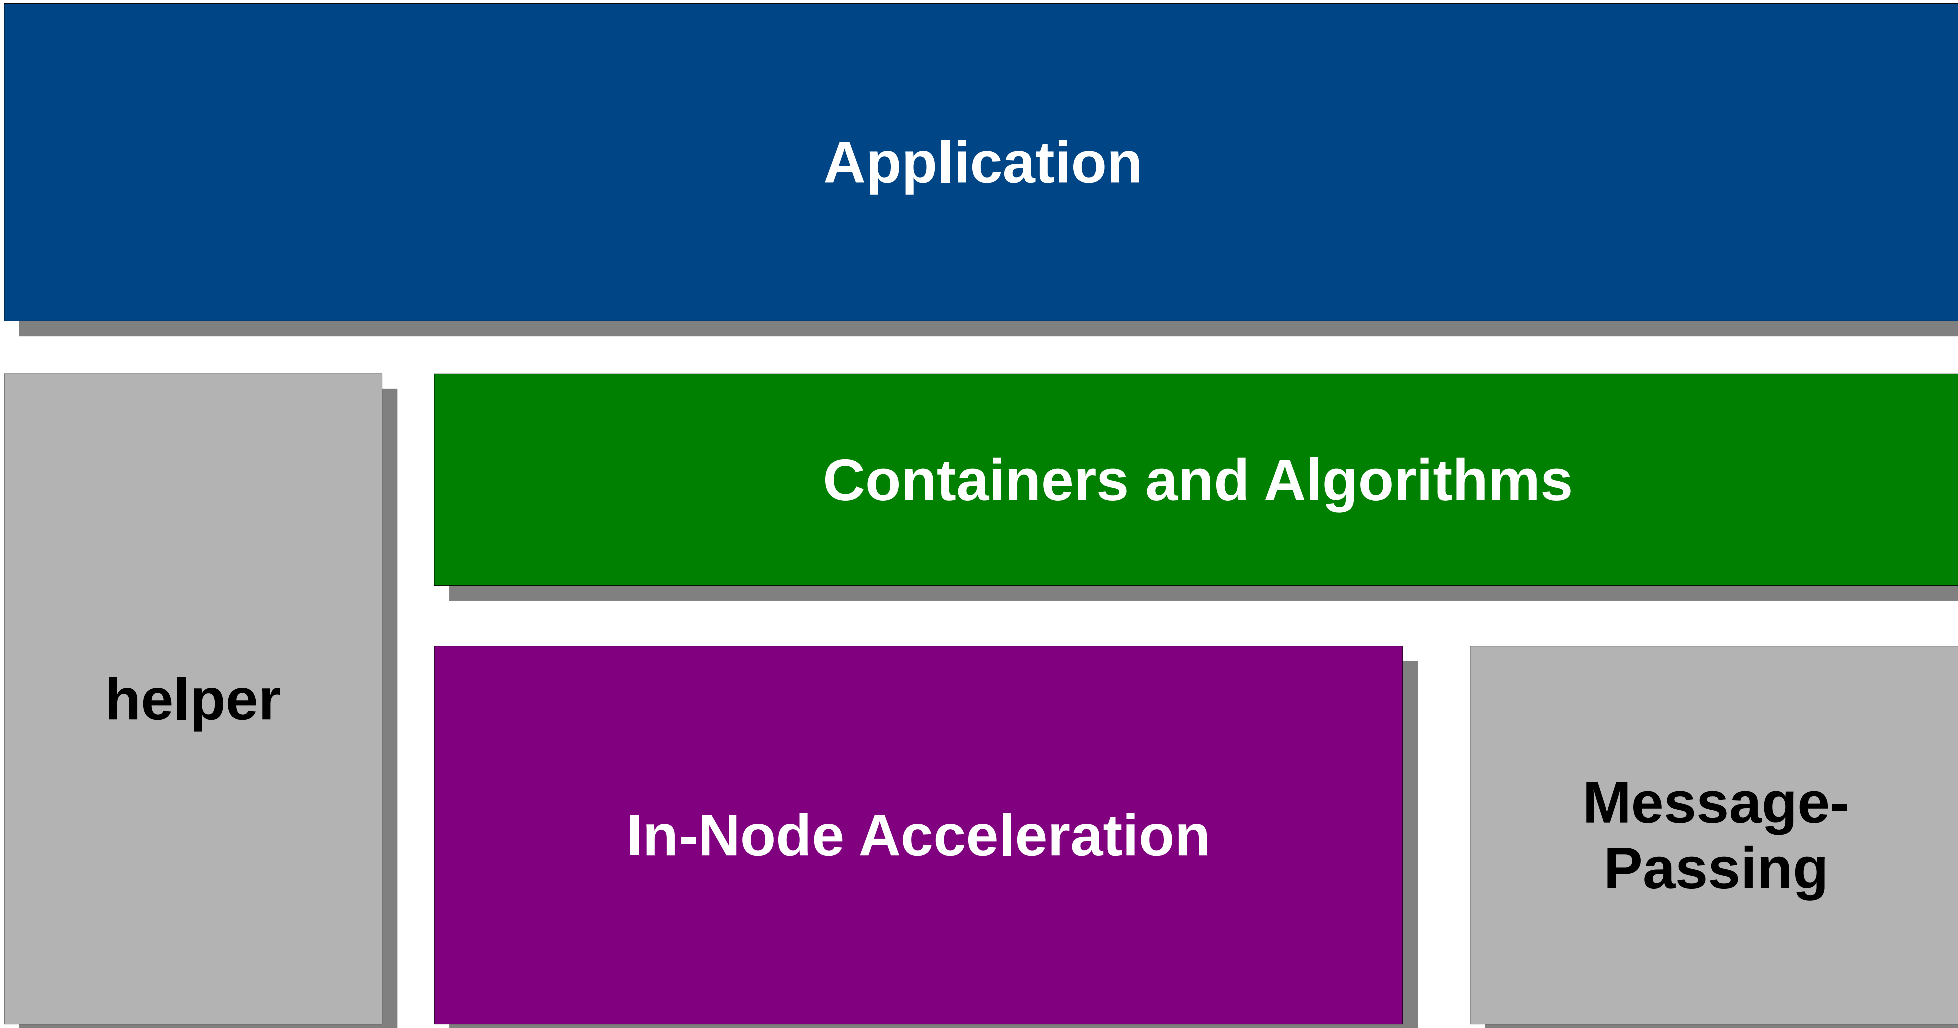

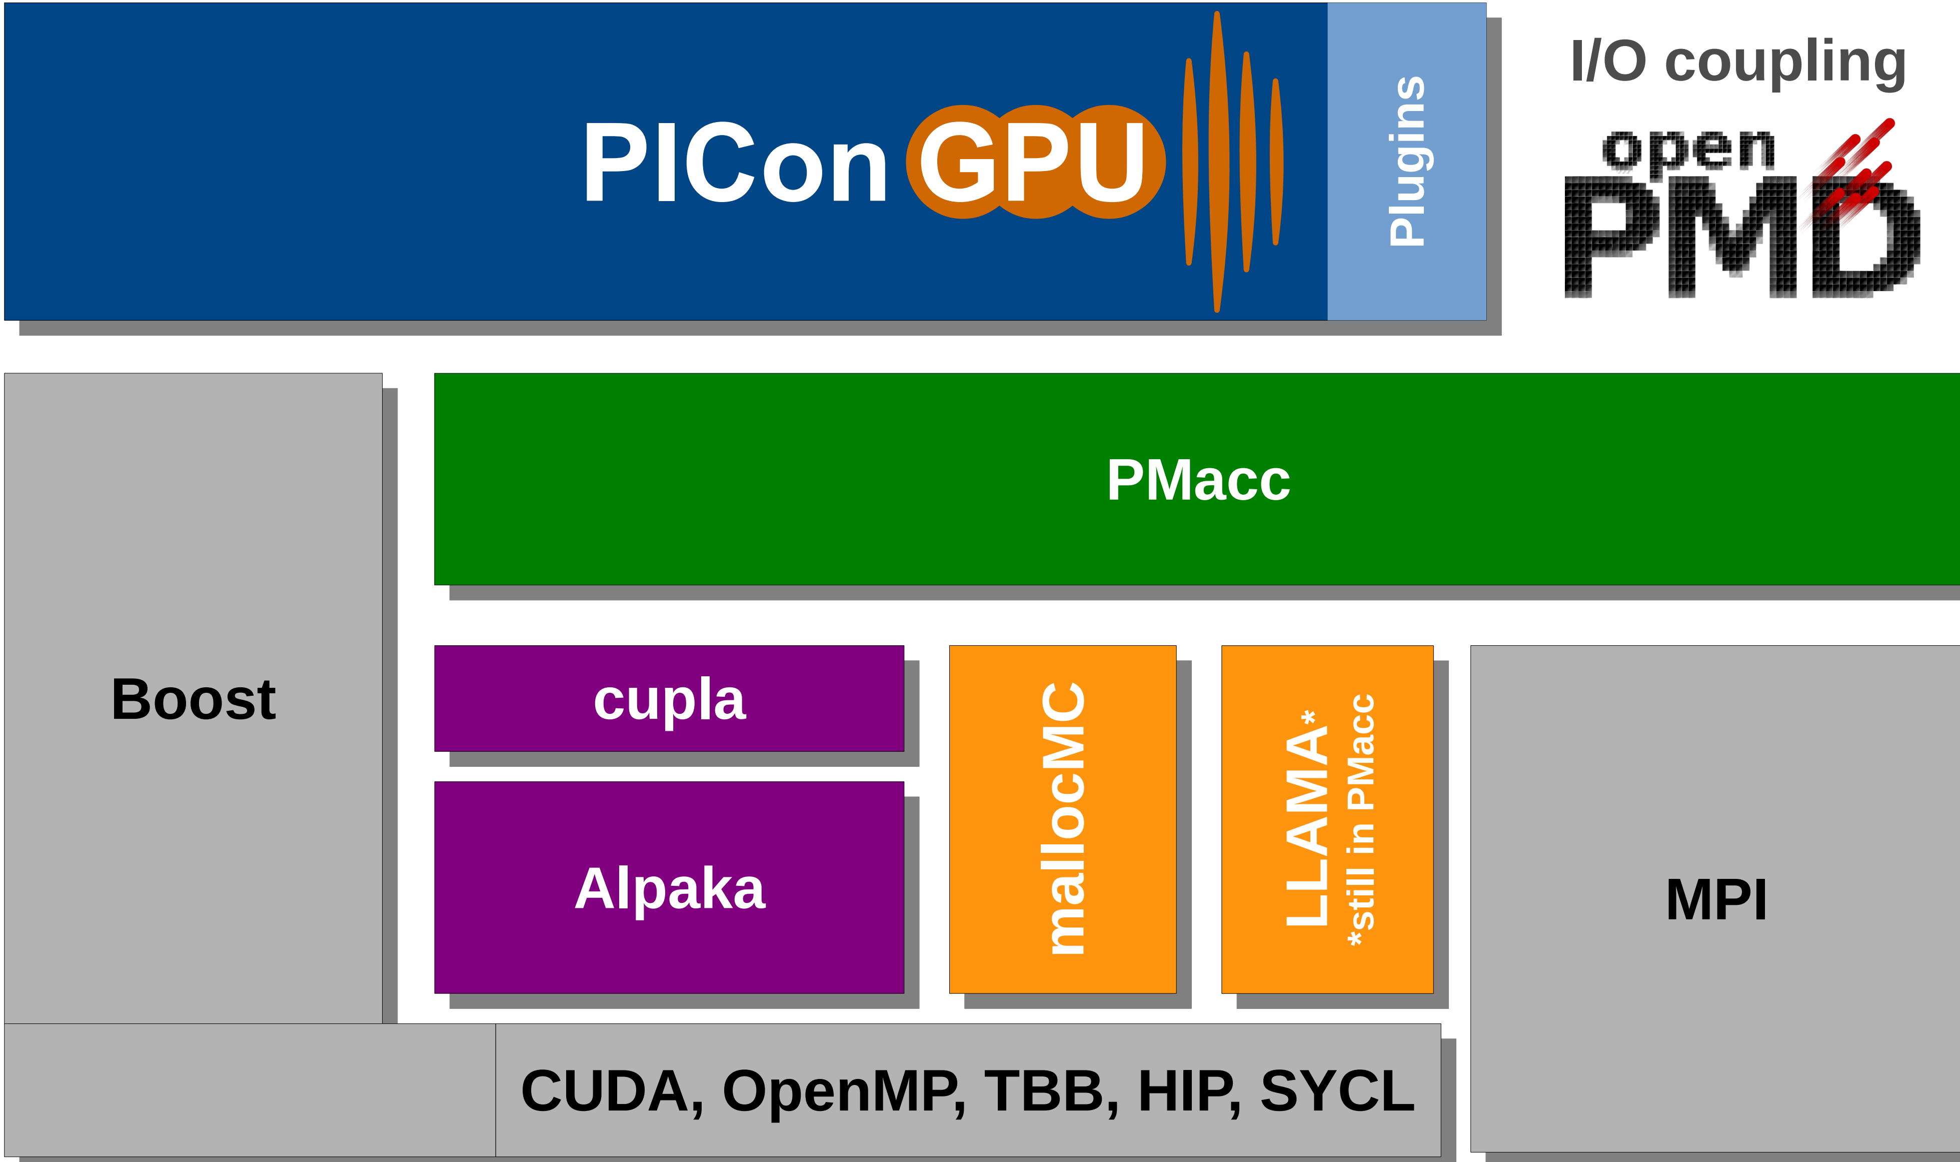
PIConGPU
Plugins
I/O coupling
openPMD
PMacc
Boost
cupla
Alpaka
mallocMC
LLAMA*
*still in PMacc
MPI
CUDA, OpenMP, TBB, HIP, SYCL

Performance Portability

Parallel Programming Models

- **re-implementation (Nx)**
  - not affordable for domain-science teams

```
#ifdef CUDA_ENABLE
    // CUDA Kernel implementation
    // …

#elif OPENMP_ENABLE
    // OpenMP implementation
    // …

#else
    // Sequential CPU implementation
    // …

#endif
```

- **single-source manycore**
  - CUDA: CUDAx86 (PGI), GPUOcelot (GaTech); today: HIP (AMD)
  - *Abstract & <u>configure</u> the task parallelism!*

Alpaka's Task Parallelism

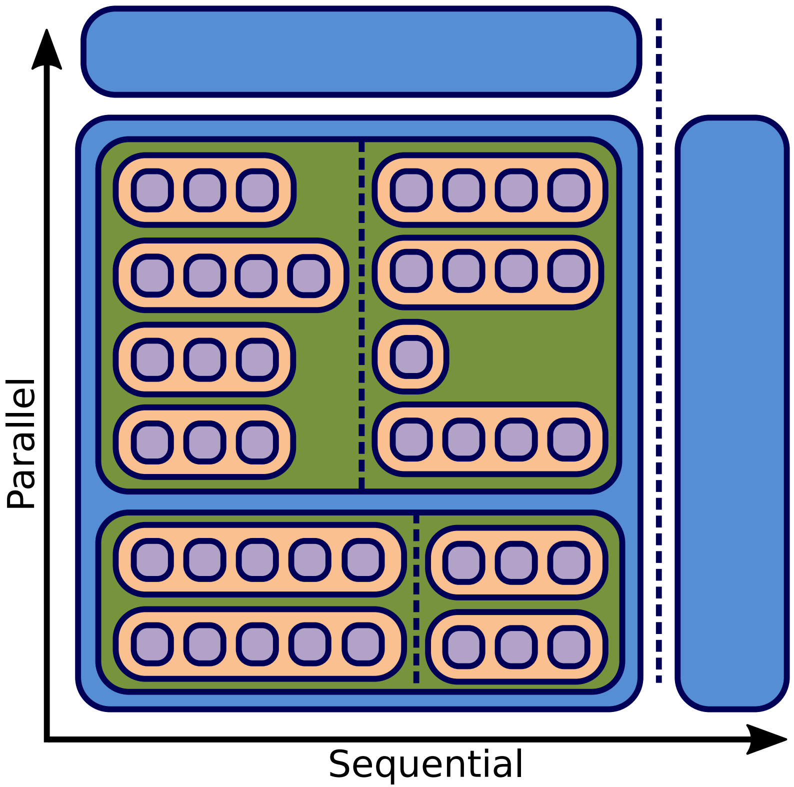

Grid	whole parallel collection
Block	fully independent part of the grid
Thread	executed concurrently
Element	sub-thread, sequential lock-step

---- Synchronize

Alpaka
Hardware Mapping

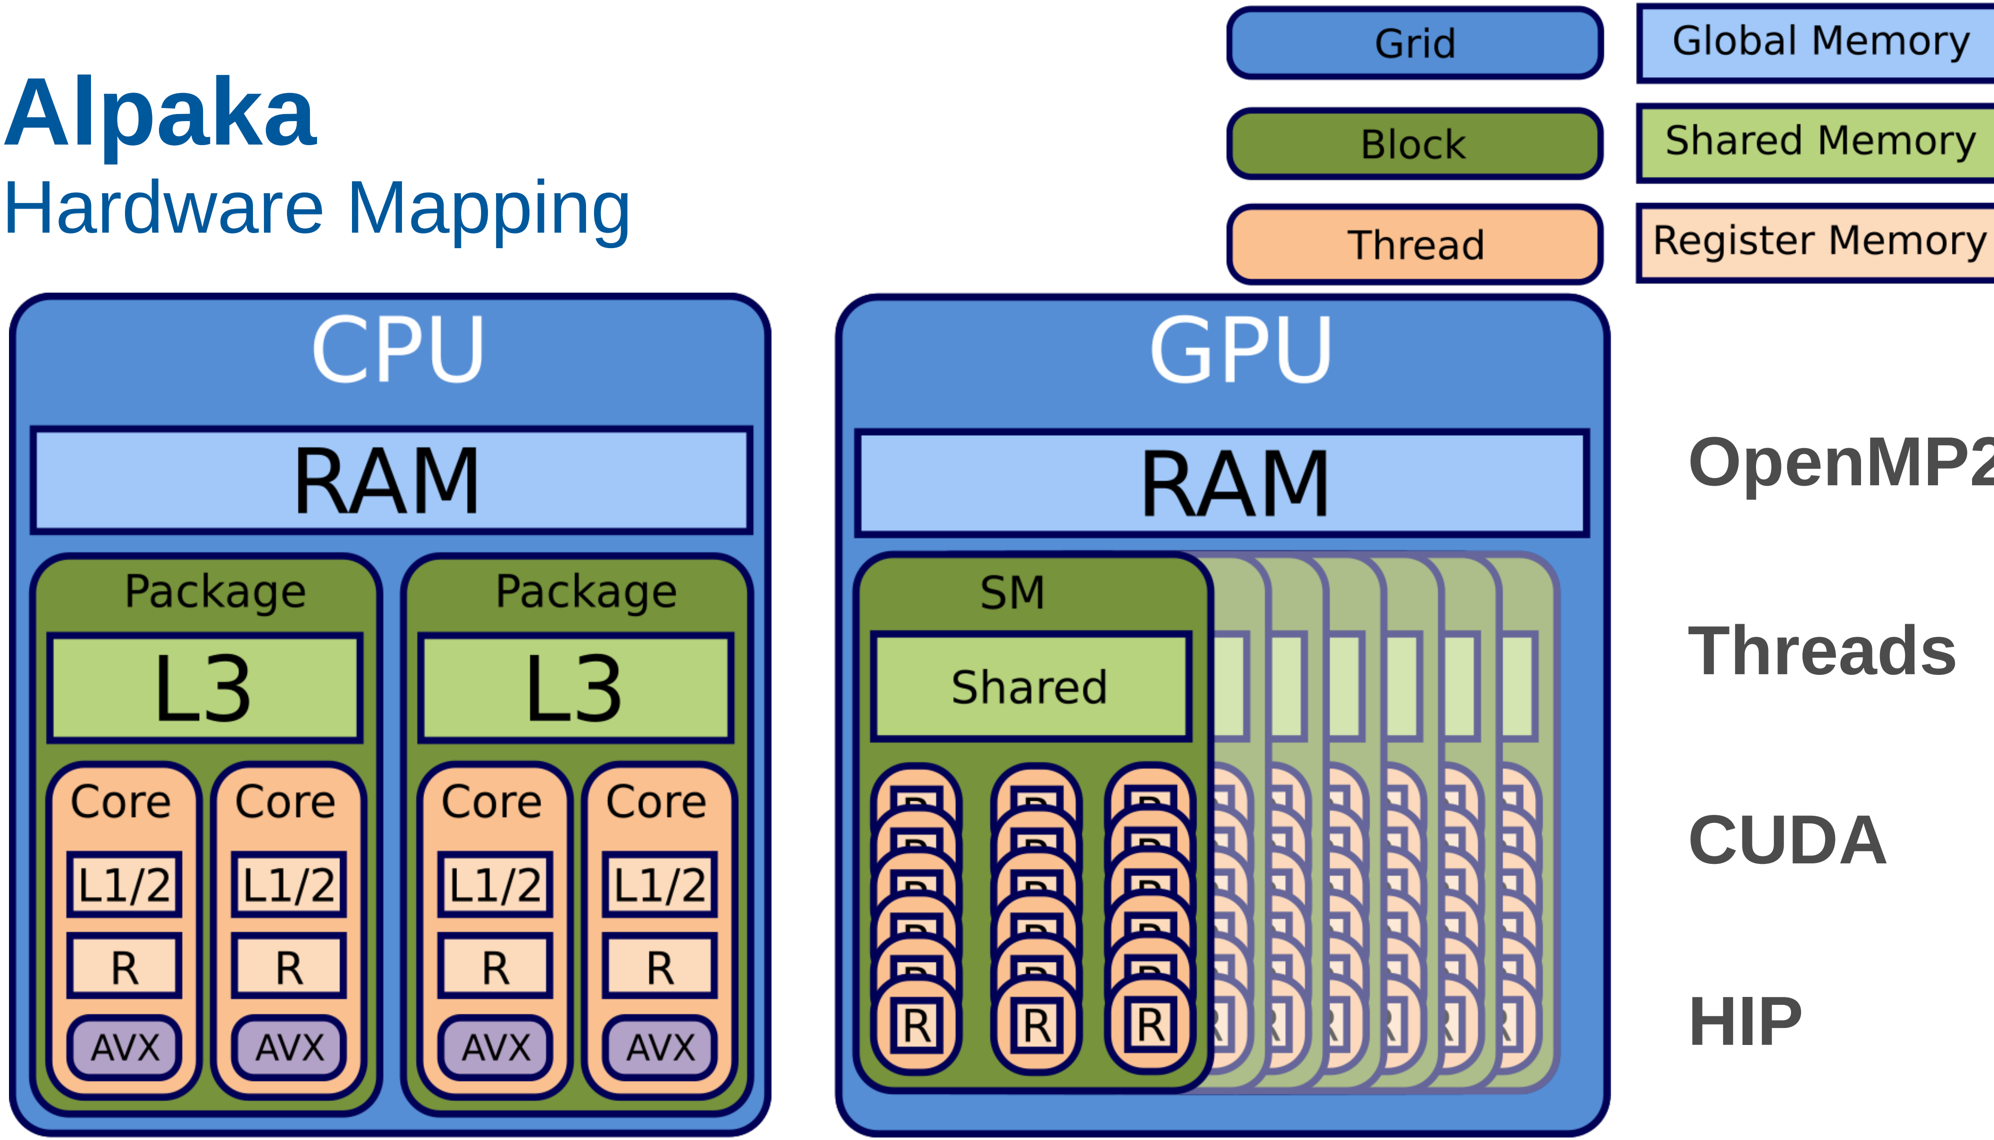

OpenMP2

Threads

CUDA

HIP

Alpaka

... e.g. MIC

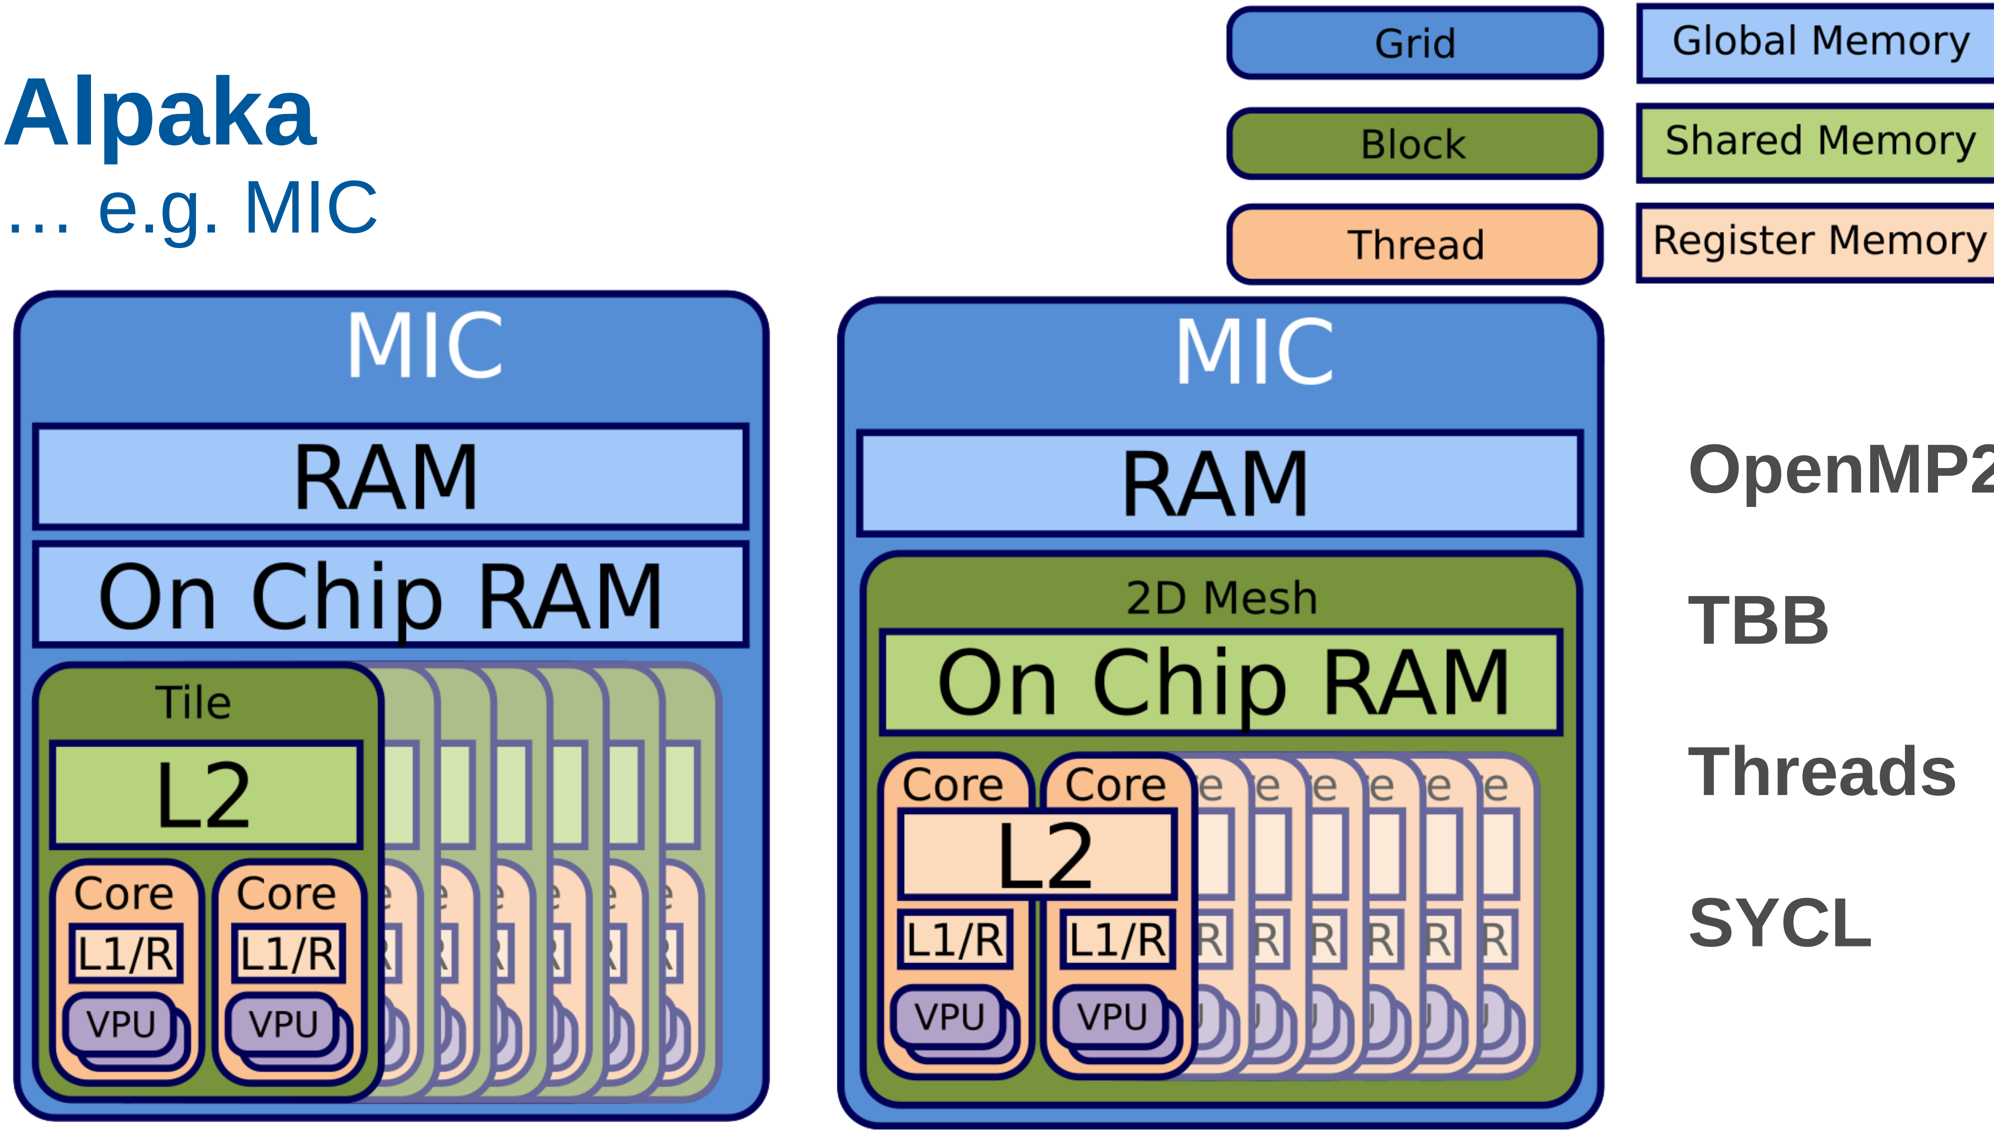

OpenMP2

TBB

Threads

SYCL

PIConGPU on Alpaka

Performance after only 3 (!) weeks

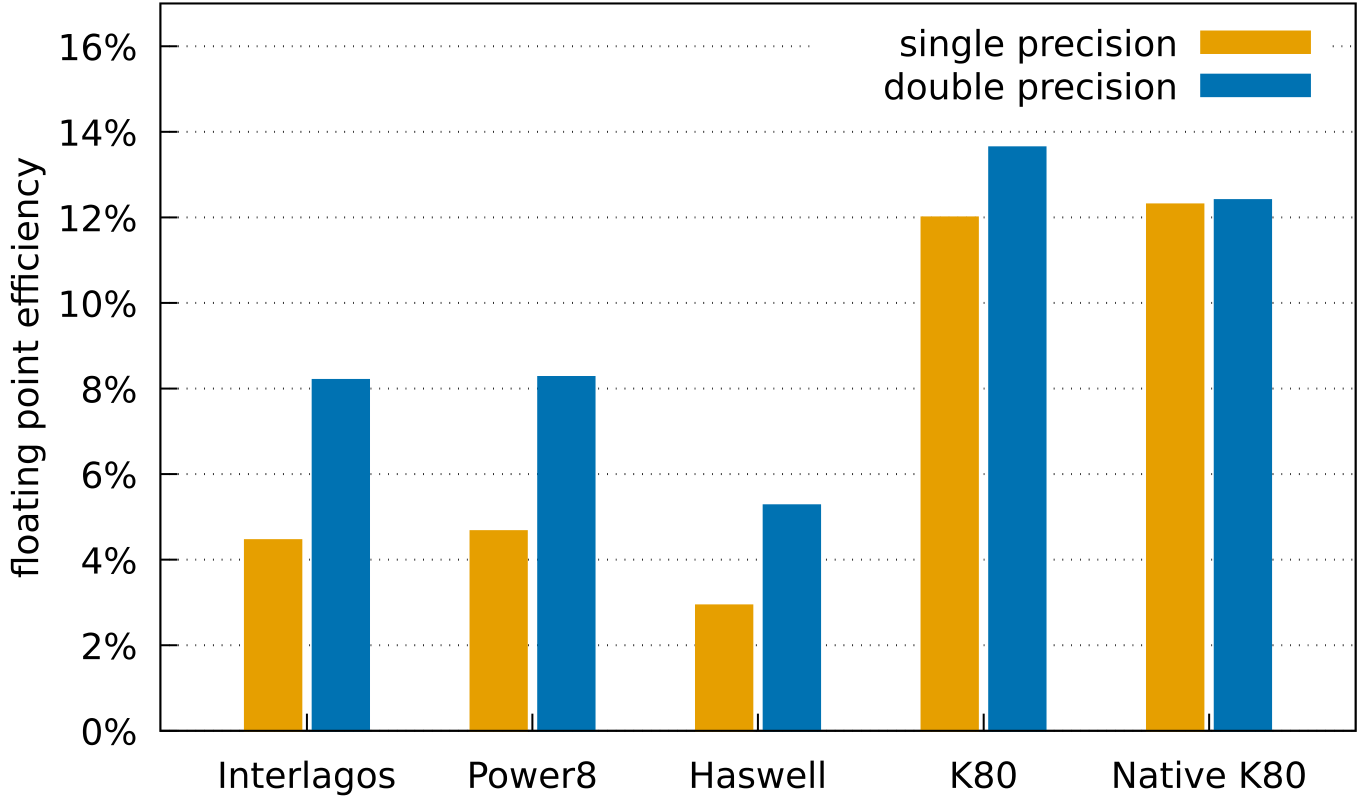

Maintainability
PIConGPU + PMacc code lines

Before: 80k LOC (20k in kernels)
After: 50k LOC (1 year)

E. Zenker et al., ISC (2016), DOI:10.1007/978-3-319-46079-6_21
A. Matthes et al., ISC (2017), DOI:10.1007/978-3-319-67630-2_36

DRESDEN concept HZDR

Scientific Data Workflows

The actual HPC challenge

The I/O Challenge Continues!

www.olcf.ornl.gov/frontier/

SYSTEM SPECS	TITAN	SUMMIT	FRONTIER
Peak Performance	27 PF	200 PF	> 1.5 EF
Storage	32 PB, 1 TB/s, Lustre Filesystem	250 PB, 2.5 TB/s, GPFS™	2-4x performance and capacity of Summit's I/O subsystem. Frontier will have near node storage like Summit.

1/3x 1/3x

In situ approaches: **tightly** versus **loosely** coupled workflows

Virtual Detectors: Tightly Coupled Plugins

Binning of a **spectrogram**
Creation of a **phase space image**

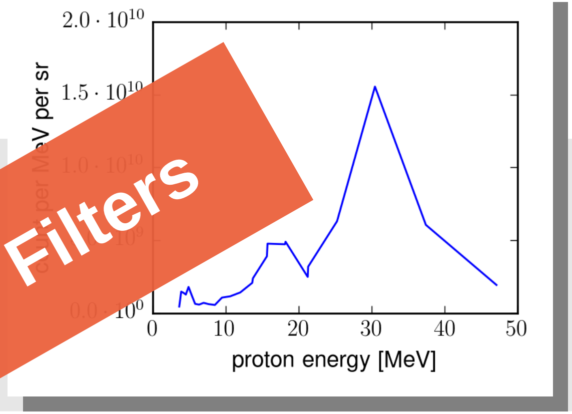

In situ **radiation diagnostics**

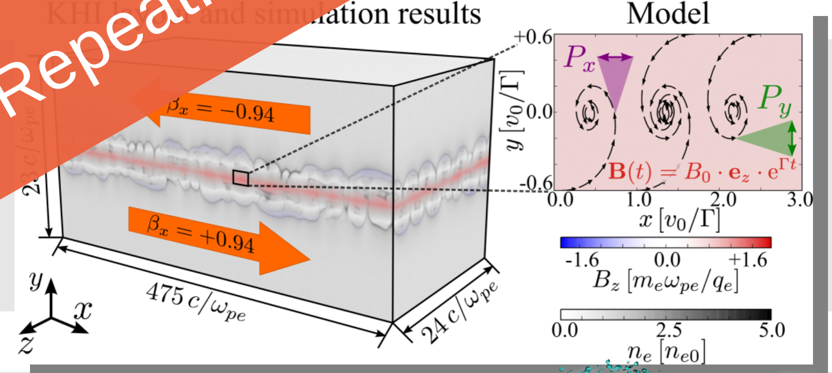

Observe, Correlate & Repeat: Particle Filters

Ray-cast or photo-realistic ray-trace
Lossy data **compression**

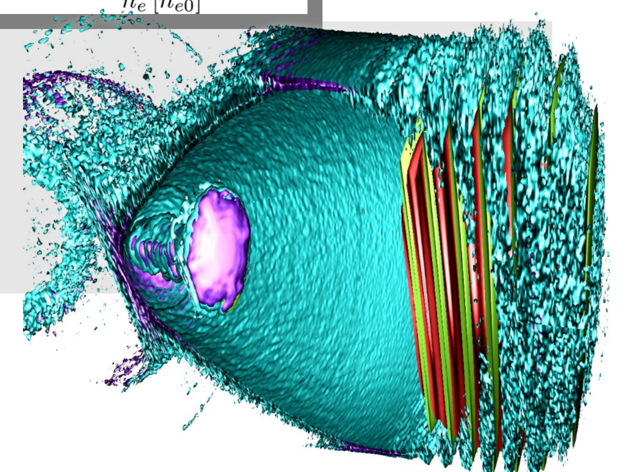

A. Huebl et al. (2014), DOI:10.1109/TPS.2014.2327392
R. Pausch et al. (2017), DOI:10.1103/PhysRevE.96.013316
A. Matthes, A. Huebl et al., ISC'16 (2016), DOI:10.14529/jsfi160403
A. Huebl et al., ISC'17 (2017), DOI:10.1007/978-3-319-67630-2_2

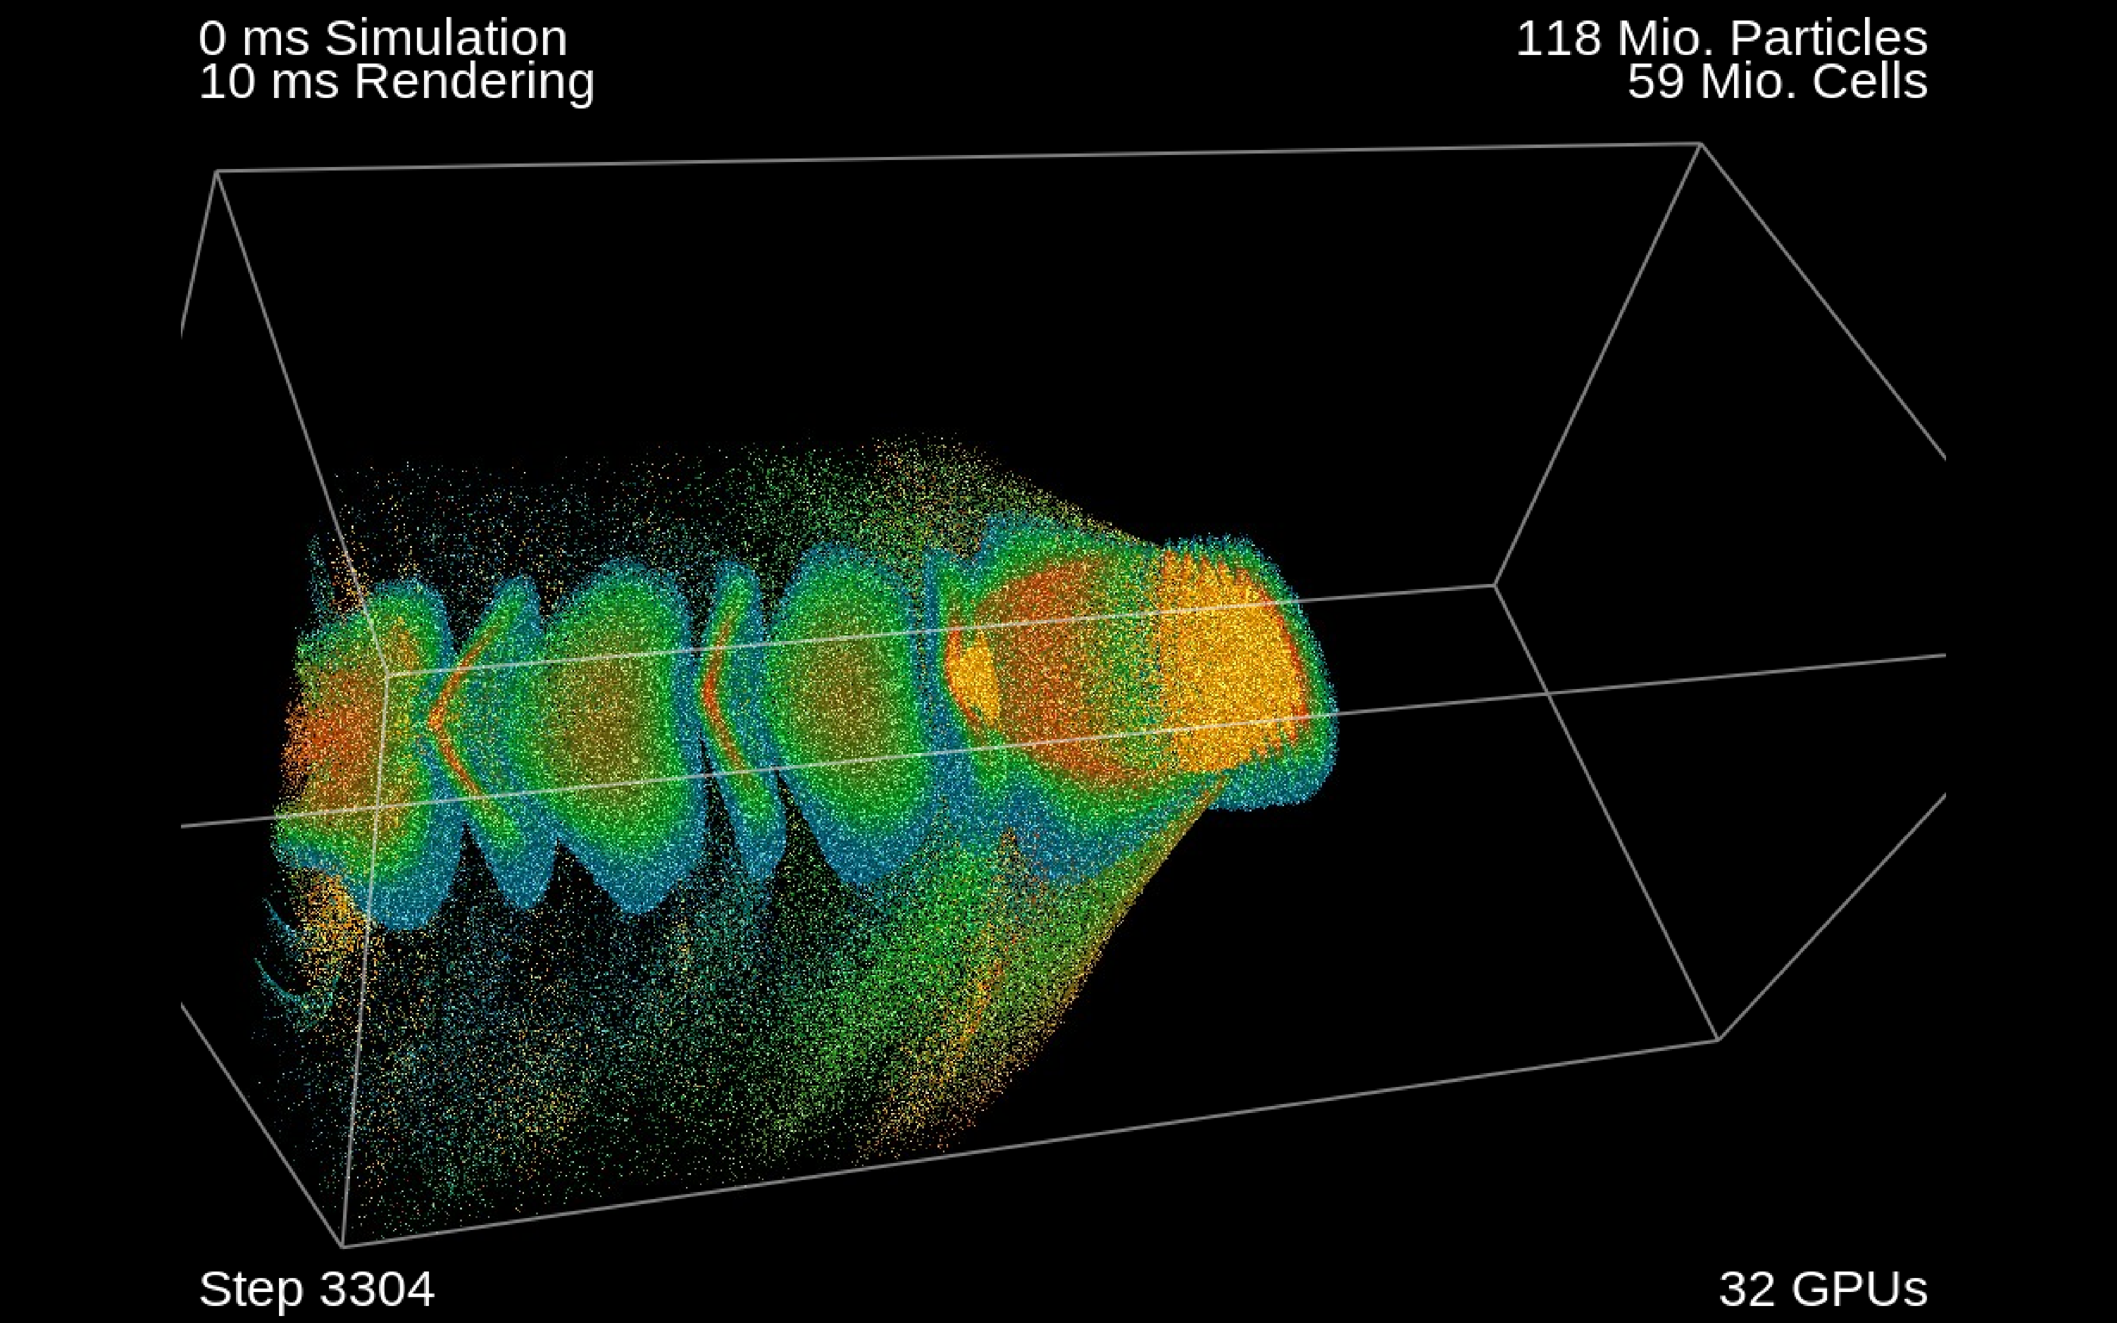
0 ms Simulation
10 ms Rendering
118 Mio. Particles
59 Mio. Cells
Step 3304
32 GPUs

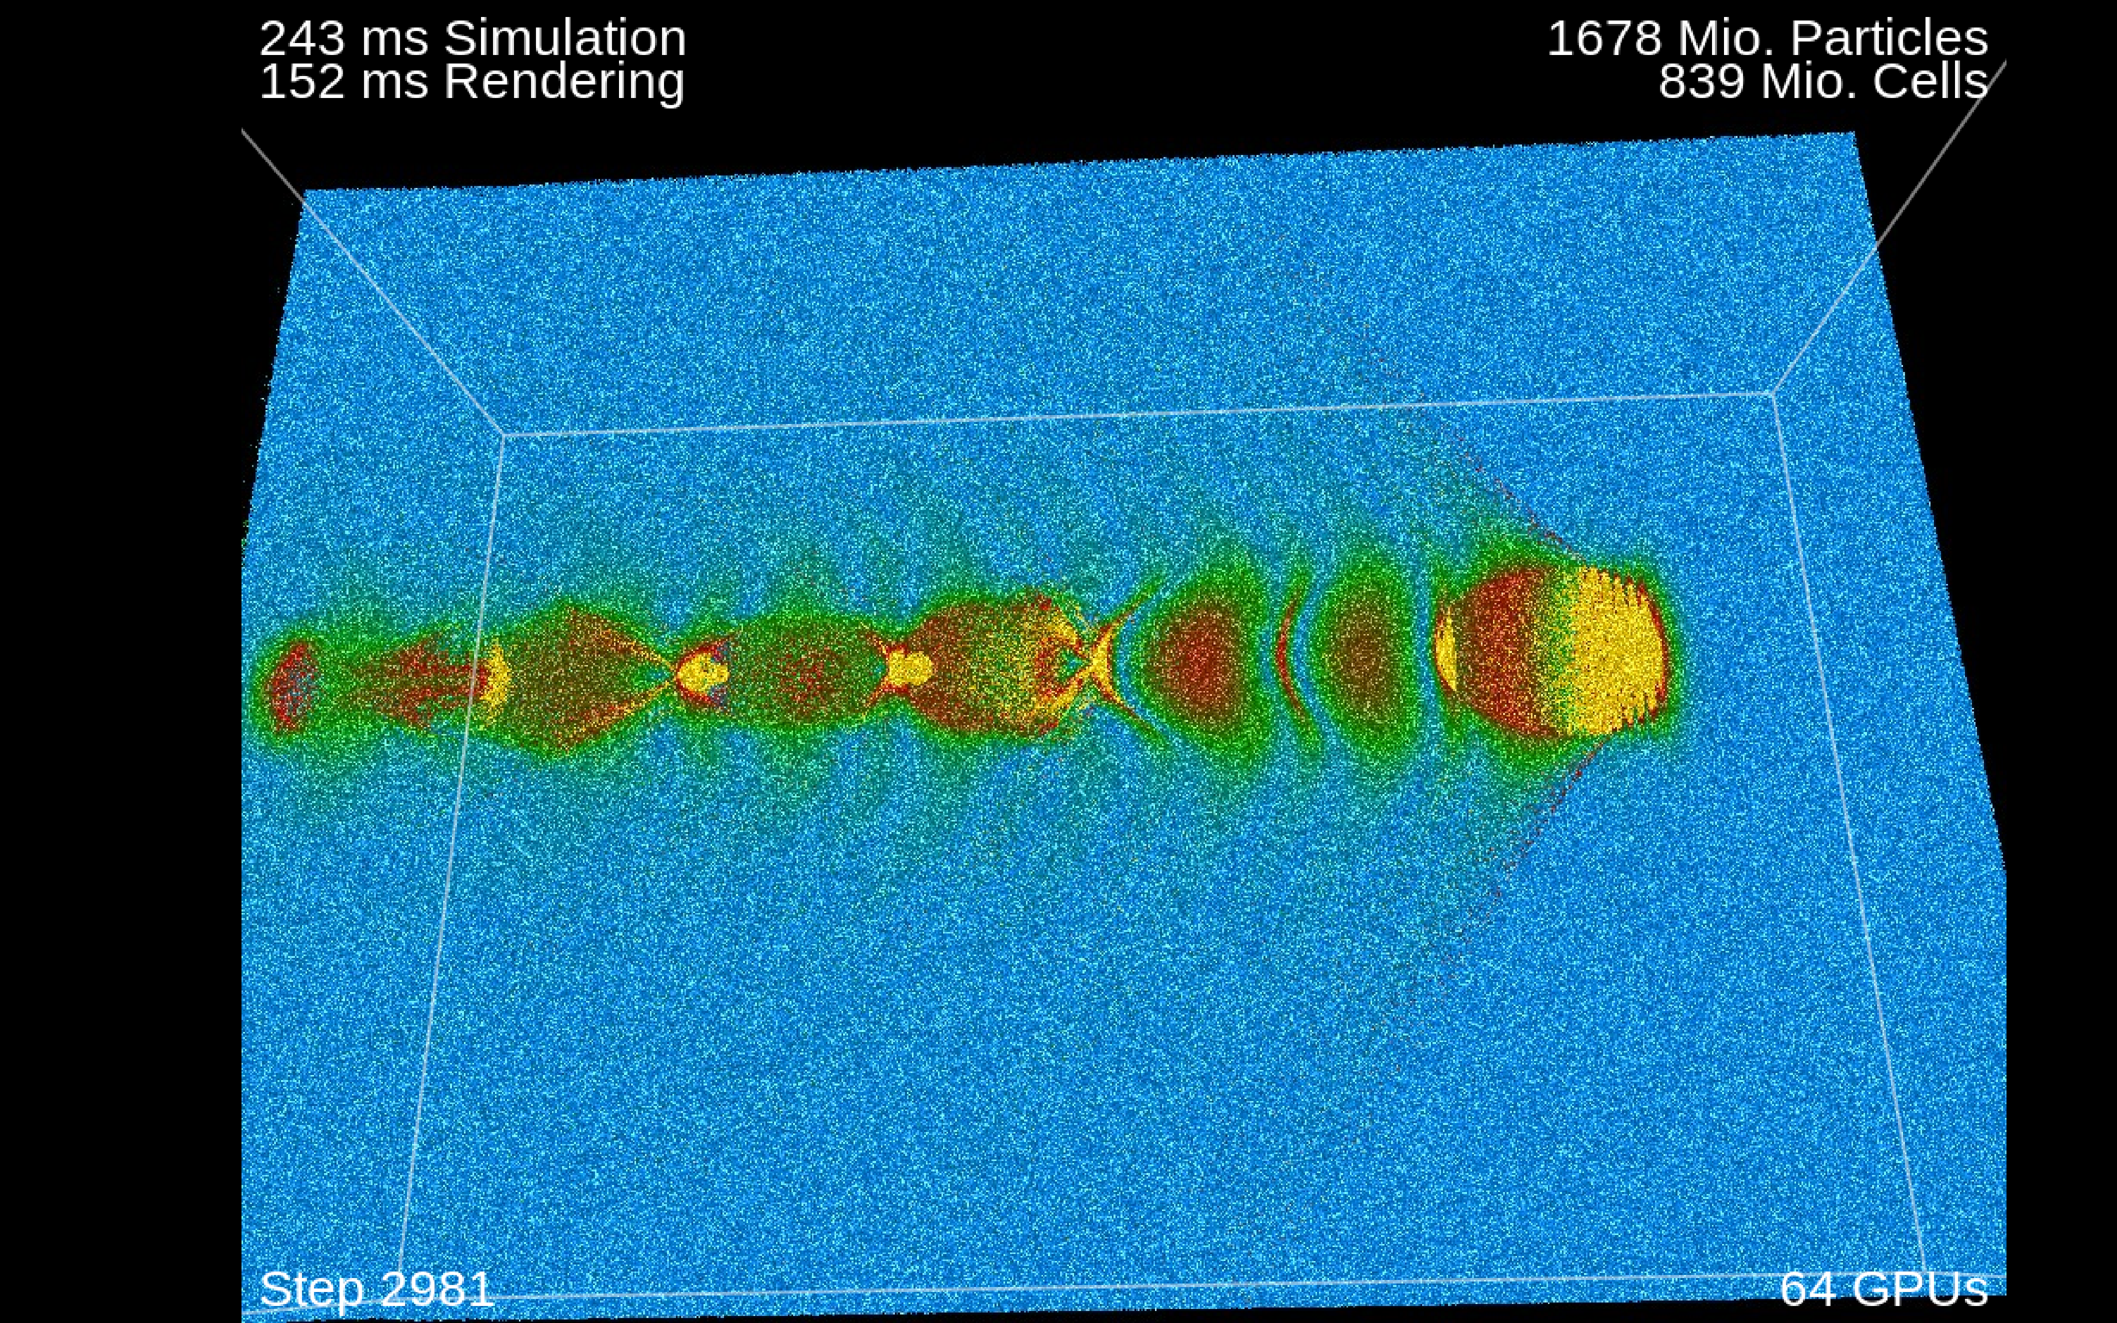
243 ms Simulation
152 ms Rendering
1678 Mio. Particles
839 Mio. Cells
Step 2981
64 GPUs

Cling CUDA

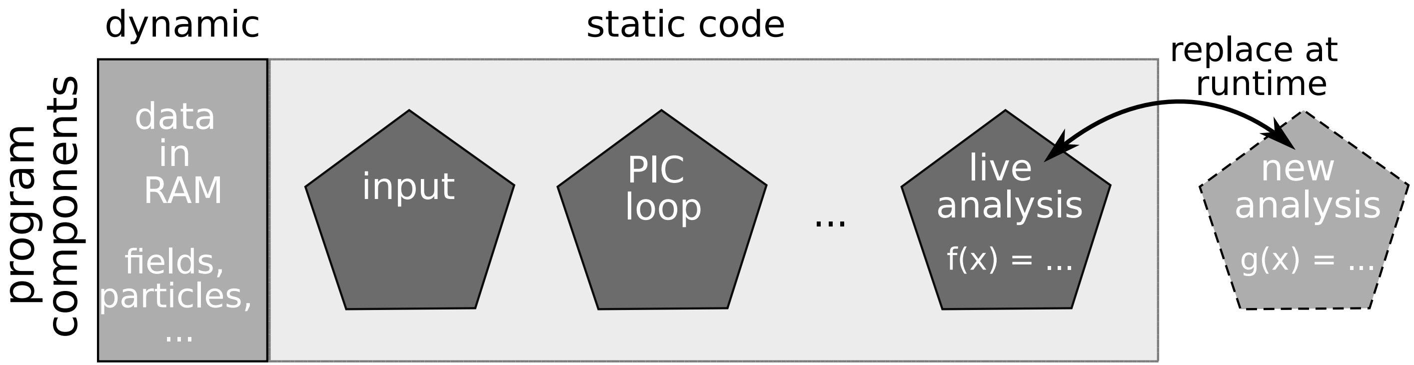

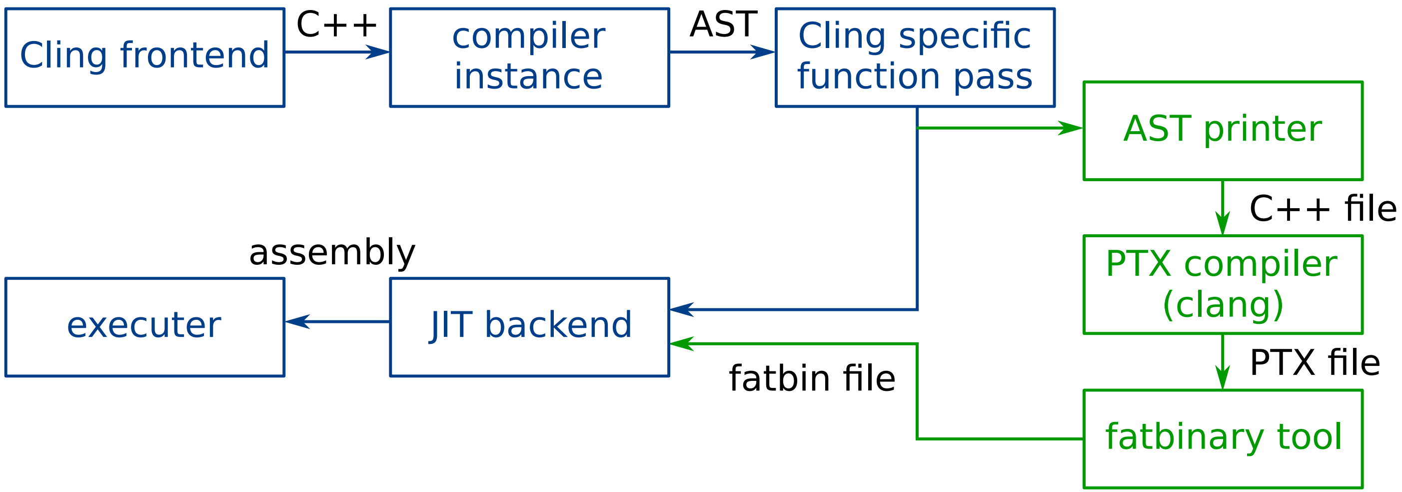

S. Ehrig (2018), HZDR/TU Dresden, Diploma Thesis

Cling CUDA: We implemented a GPU JIT!

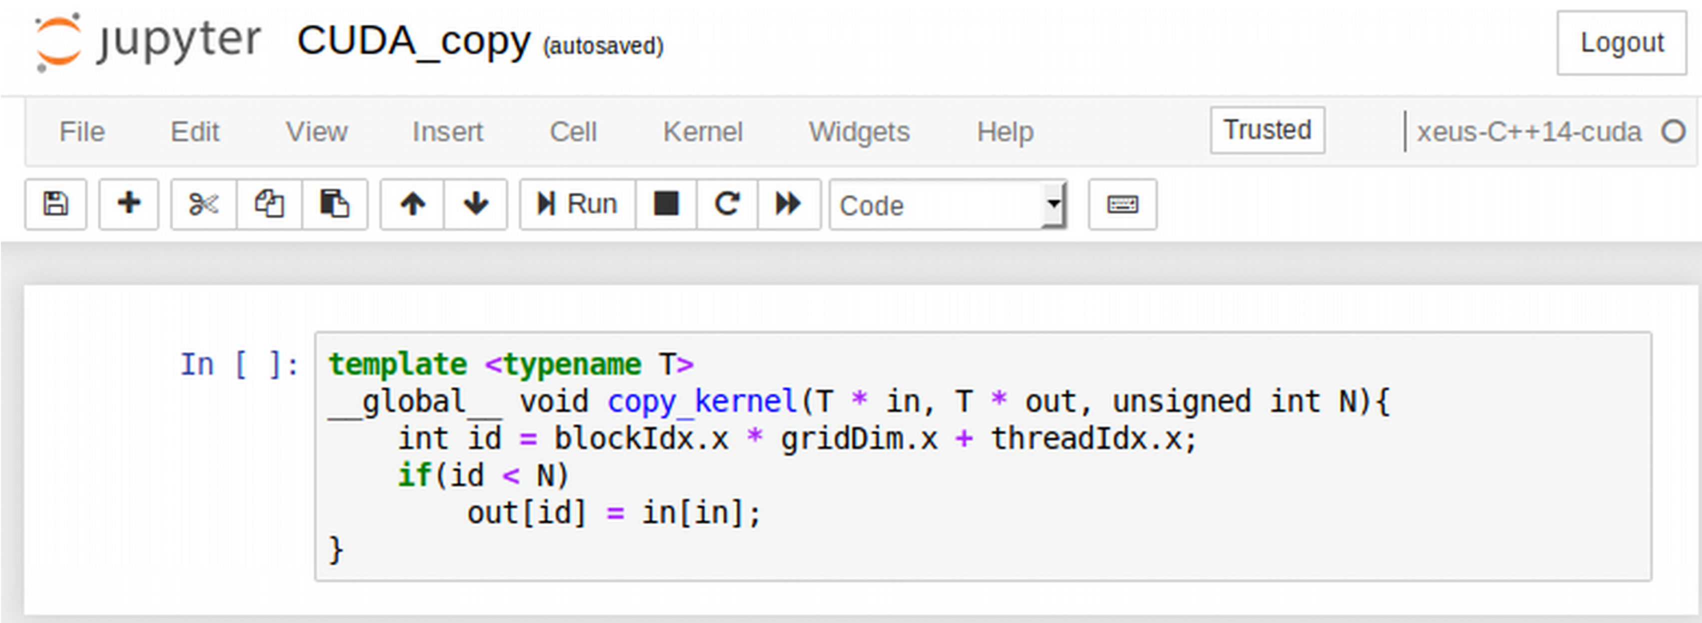

S. Ehrig (2018), HZDR/TU Dresden, Diploma Thesis

Exascale Challenge: I/O Scalability

Titan I/O Weak Scaling with PIConGPU

$$T_{\text{eff}} \equiv \frac{N \times S}{t_{\text{I/O}}}$$

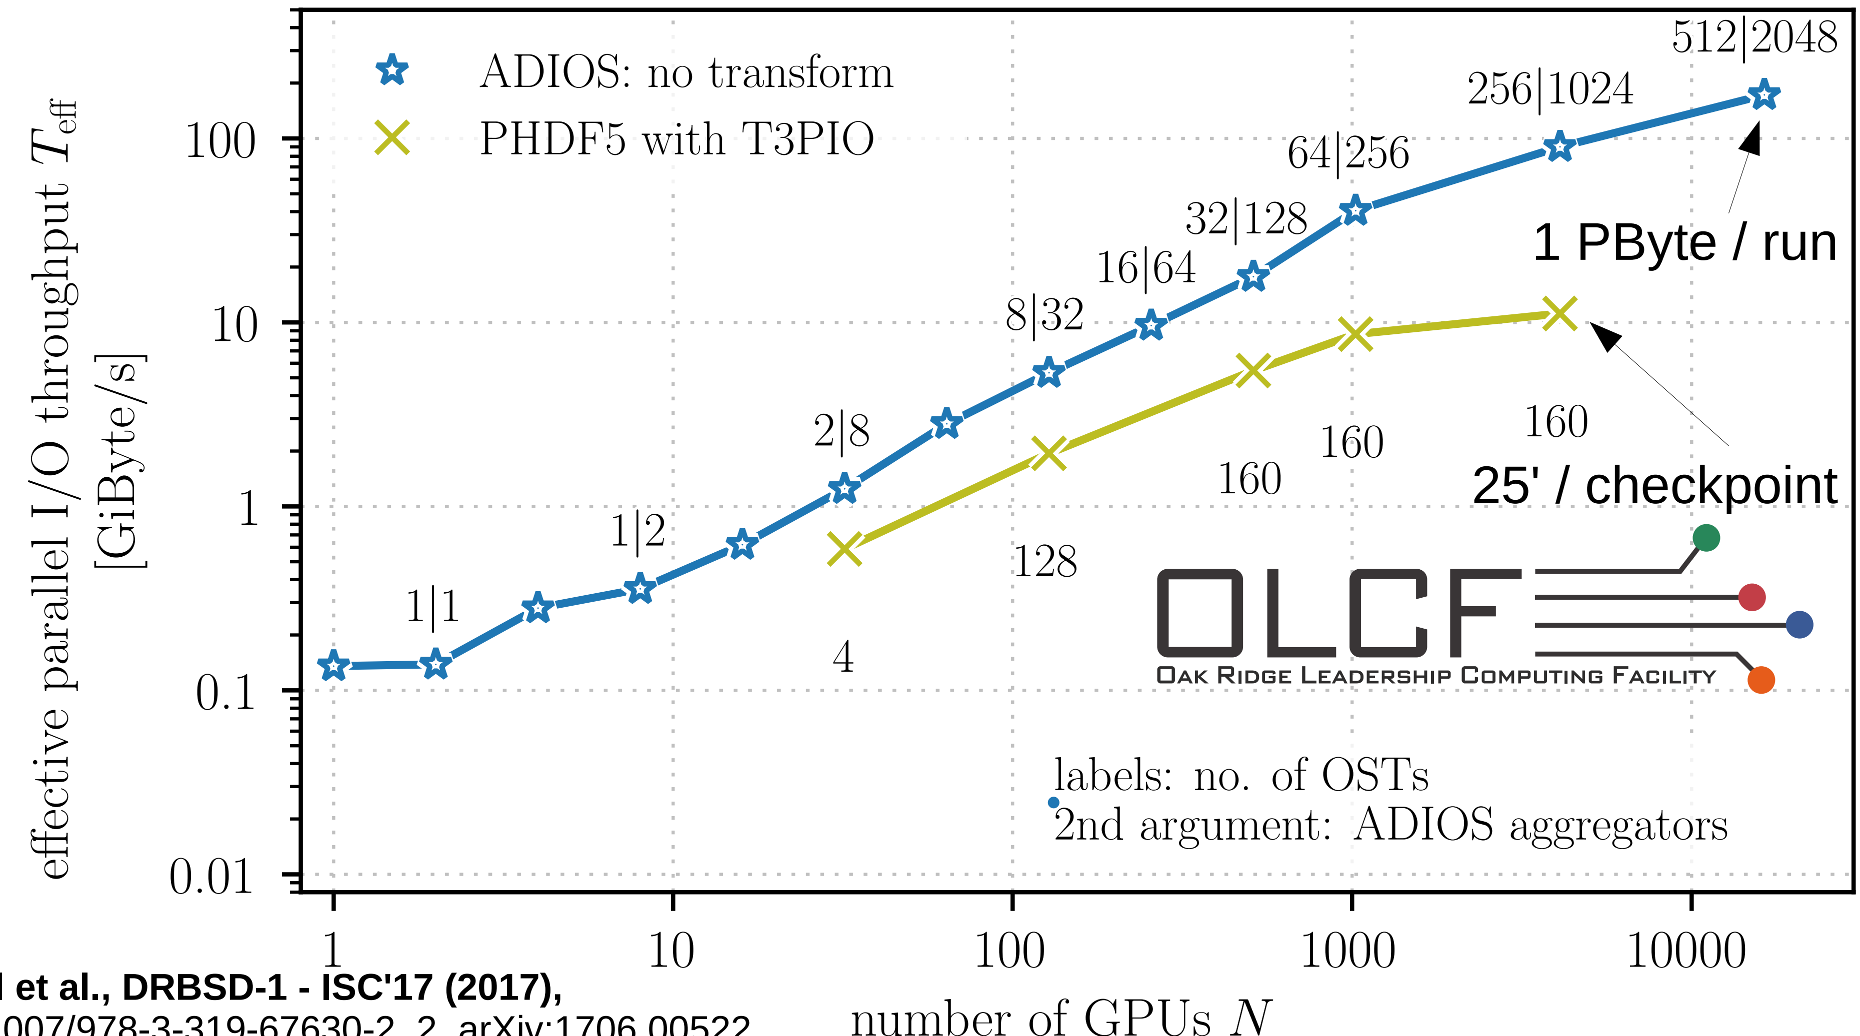

A. Huebl et al., DRBSD-1 - ISC'17 (2017),
DOI:10.1007/978-3-319-67630-2_2, arXiv:1706.00522

github.com/openPMD

www.openPMD.org

- markup / schema for arbitrary hierarchical data formats
- truly, *scientifically* self-describing
- basis for open data workflows

Batteries included: **openPMD-api**

ADIOS

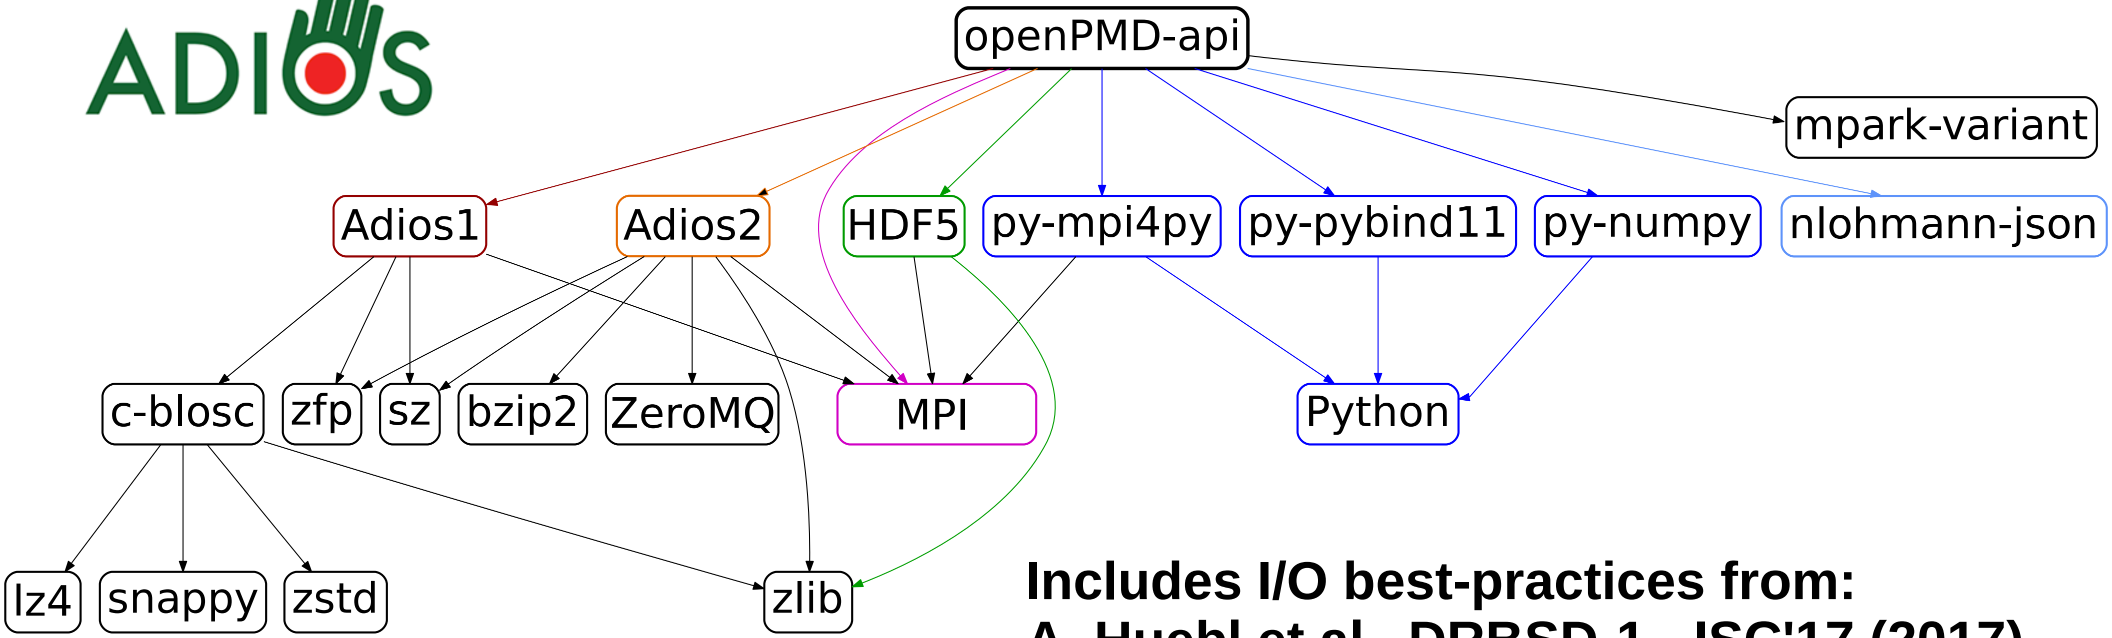

Includes I/O best-practices from:
A. Huebl et al., DRBSD-1 - ISC'17 (2017),
DOI:10.1007/978-3-319-67630-2_2, arXiv:1706.00522

Available via: Spack CONDA-FORGE PyPI

Productivity

PIC as a Service

PIConGPU as a Service

- full **focus on a science case**: expressive, reproducible
  - ***reduce*** options to a ***virtual experiment***
  - implement a **sharable data workflow** from best-practices
  - *transparently hide the HPC workflow*

- **collaborative service**: usability & productivity
  - accessible for **experimentalists** and **theoreticians**
  - **hackable** for power-users
  - **auto-documenting** and **auto-archiving**
    - **… ML, digital twins, etc.**

RODARE
ROSSENDORF DATA REPOSITORY

Authors: S. Rudat, S. Starke, J. Kelling, A. Huebl

A. Huebl, G. Juckeland et al., SIAM'18 (2018)

DRESDEN concept HZDR

PIConGPU as a Service Demo

Combine Qualitative and Quantitative Analysis

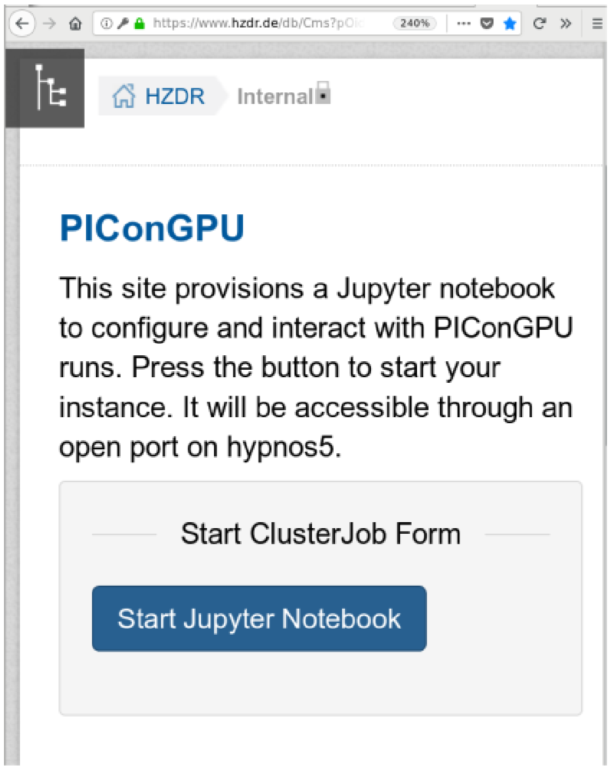

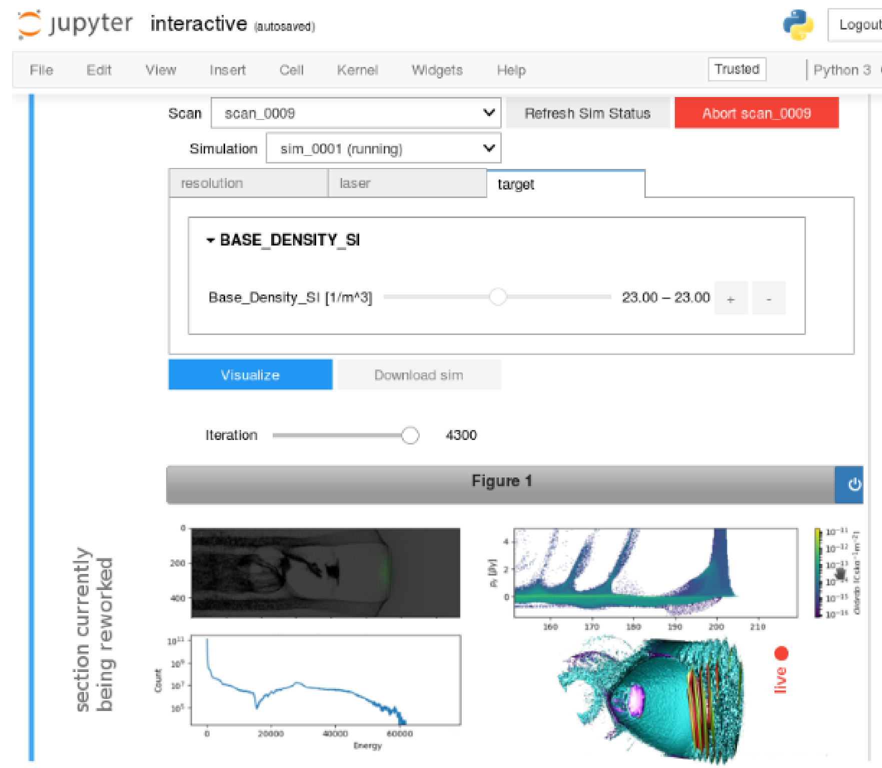

Summary

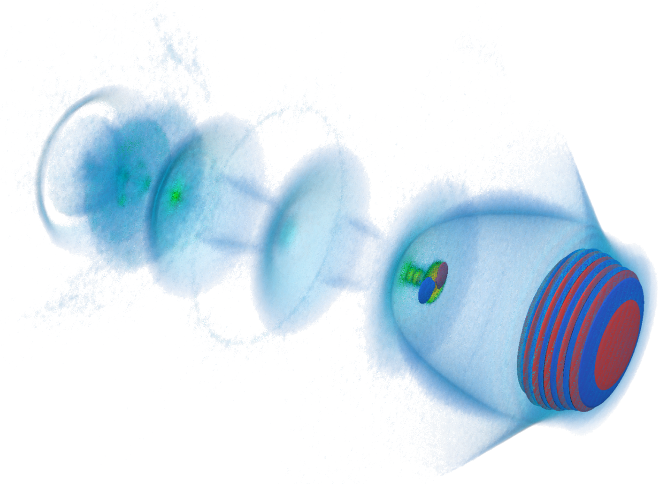

- **Scalability of Algorithms first**
  - Use abstract programming models
  - Use and build community libraries
- **In situ data reduction for all production jobs unavoidable**
  - Start loosely coupled, implement hot-spots tightly
  - Join the openPMD community for data description :)
- **Build intuitive, reproducible workflows**
  - Modular and transparent for power-users

talk by **A. Huebl** (HZDR) et al., **a.huebl@hzdr.de**

picongpu.hzdr.de
github.com/ComputationalRadiationPhysics

This project has been enabled by many people in open-source and open-science communities. Great thanks to the communities and developers of: PIConGPU, Alpaka, ROOT/Cling, Jupyter, the SciPy ecosystem, ADIOS, HDF5, Boost, CMake, openPMD, Spack, ...

This research used resources of the Oak Ridge Leadership Computing Facility located in the Oak Ridge National Laboratory, which is supported by the Office of Science of the Department of Energy under Contract DE-AC05-00OR22725.

This project received funding within the MEPHISTO project (BMBF-Förderkennzeichen 01IH16006C).
This project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No 654220.