SPH-EXA: A Framework for Smoothed Particle Hydrodynamics and Gravity at Exascale

Swiss SKA Days 2022
October 3rd - 4th

Presented by Sebastian Keller and Osman Seckin Simsek

https://github.com/unibas-dmi-hpc/SPH-EXA

Universität Basel

Florina Ciorba (PI) *florina.ciorba@unibas.ch*
Ruben Cabezon (Co-PI) *ruben.cabezon@unibas.ch*
Osman Seckin Simsek *osman.simsek@unibas.ch*
Lukas Schmidt *luke.schmidt@unibas.ch*

Universität Zürich

Lucio Mayer (Co-PI) *lmayer@physik.uzh.ch*
Noah Kubli *noah.kubli@uzh.ch*
Darren Reed *darren.reed@uzh.ch*

CSCS

Sebastian Keller *sebastian.keller@cscs.ch*
Jean-Guillaume Piccinali *jgp@cscs.ch*
Jean Favre *jean.favre@cscs.ch*
Benjamin Cumming *benjamin.cumming@cscs.ch*
John Biddiscombe *john.biddiscombe@cscs.ch*

José Escartin (ESO)
Axel Sanz (UPC)
Joseph Touzet (Paris-Saclay)

SPH-EXA: Smoothed Particle Hydrodynamics at Exascale

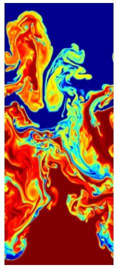

Hydrodynamics

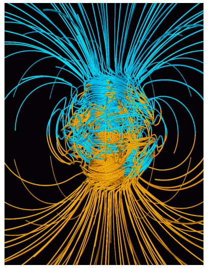

Magnetic fields

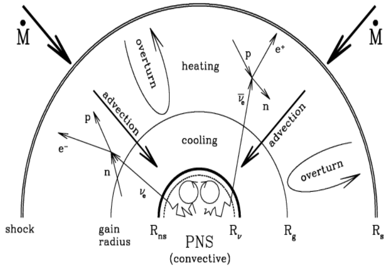

Neutrino transport

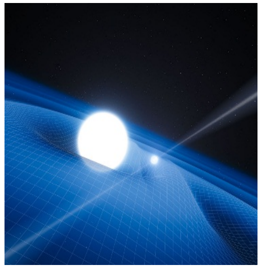

General Relativity

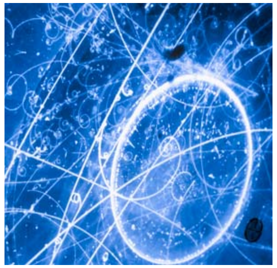

Nuclear Physics

Vision

First **trillion particle simulation** of galaxy formation with SPH, gravity, and radiation running at Exascale.

Philosophy

Interdisciplinary **codesign** between computer scientists, astrophysicists, cosmologists, and visualization specialists to implement the SPH method (and additional physics) **for Exascale**, instead of optimizing already existing codes.

SPH-EXA Strategy

- Implements state-of-the-art SPH method, leveraging characteristics of existing codes
- Composable framework written in C++20
- Communication with MPI and parallelization with OpenMP|OpenACC|CUDA|HIP
- Support and optimization for all various types of hardware architectures, in-situ visualization
- No external dependencies other than core compiler/language components and MPI
- Dynamic schedluling, adaptive load balancing, and fault-tolerance
- Designed for easy extensibility with additional physical effects and observables
- Extensive test coverage with unit and integration tests

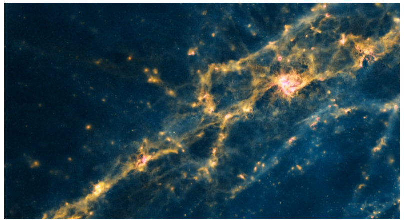

Formation of a galaxy in the GigaERIS simulation (with ChaNGa, Mayer et al.)

SPH-EXA project

(1) PASC SPH-EXA1 (2017-2021)
(2) PASC SPH-EXA2 (2021-2024)
(3) SKACH (2021-2024)

SPH-EXA1: https://hpc.dmi.unibas.ch/en/research/sph-exa/
SPH-EXA2: https://hpc.dmi.unibas.ch/en/research/pasc-sph-exa2/

SPH-EXA SKACH: https://hpc.dmi.unibas.ch/en/research/skach/
SPH-EXA download: https://github.com/unibas-dmi-hpc/SPH-EXA

~~Classical~~ Scientific Software Development Cycle

New technologies
- New performance concerns
- New scalability issues

Performance Optimization
- Analysis
- Parallelization and scaling
- Porting and acceleration

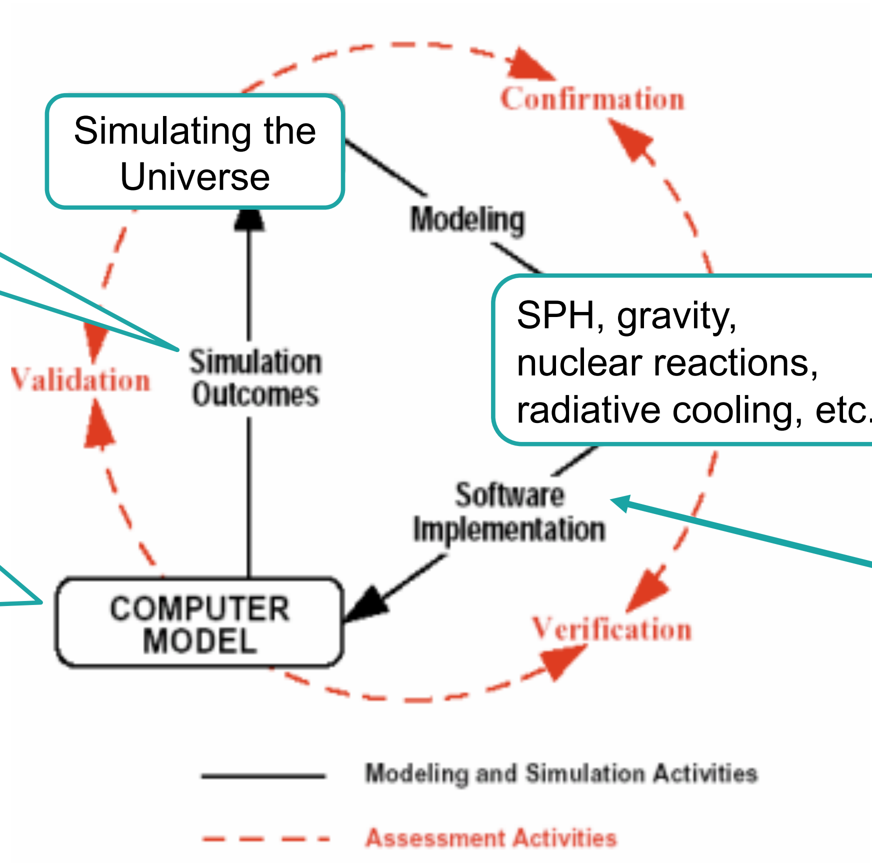

[Adapted from: Schlesinger, S., “Terminology for Model Credibility,” Simulation, Vol. 32, No. 3, 1979.]

<u>OPTIMIZATION</u>
- Performance analysis
  - Visualization
- Parallelization and scaling
  - Domain decomposition
  - Multithreading
  - Communication
  - Scheduling
  - Load balancing
- Portability and acceleration
  - Asynchronous execution
  - Mixed precision
  - Vectorization

SPH-EXA Software Development: Interdisciplinary Codesign

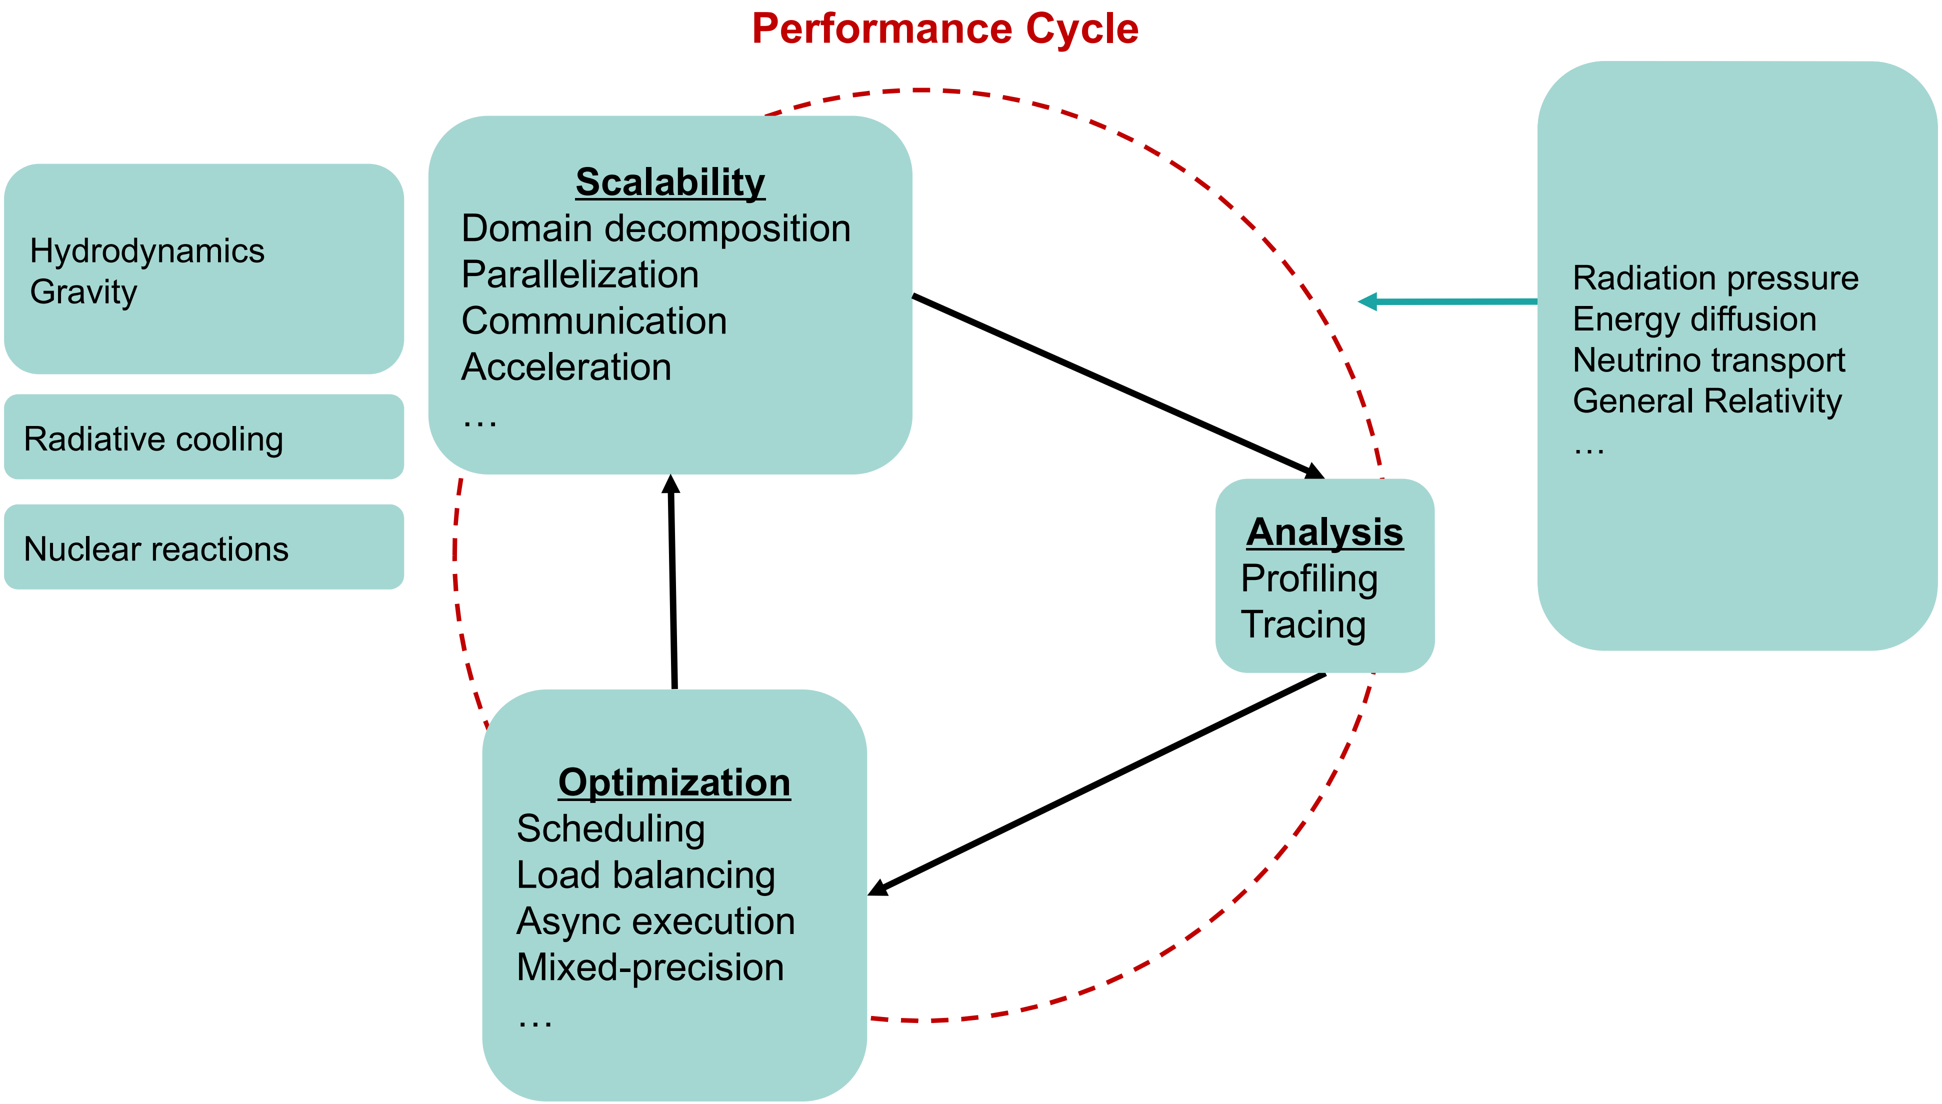

- Highly-optimized and scalable components for hydrodynamics, gravity, and other physics
- High performance simulations

SPH-EXA Performance Optimization Strategy

Scalability

weak, strong

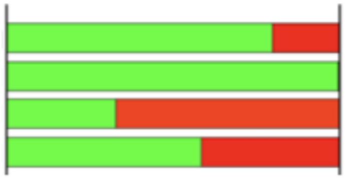

Scheduling & Load-balancing

dynamic, adaptive,

asynchronous execution

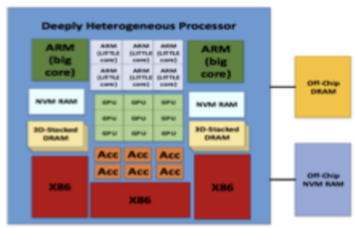

Heterogeneity

portability on CPUs and GPUs of various architectures

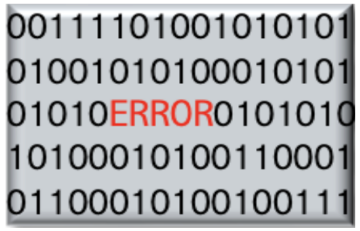

Fault-tolerance

detection and recovery of silent errors

SPH-EXA Framework Components

SPH-EXA application front-end

Parallel I/O and test case setup

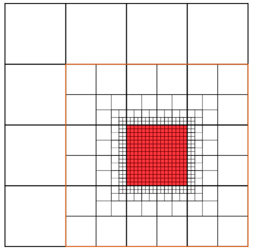

Cornerstone

Octree and domain decomposition framework

SPH

Hydrodynamics solver

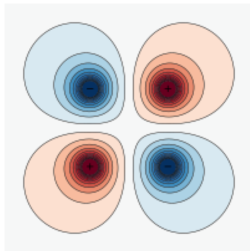

Ryoanji

N-body gravity solver

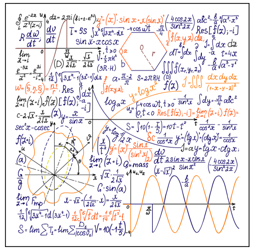

Physics modules

Radiative cooling (GRACKLE)
Nuclear reactions
Star formation
Stellar feedback
(in development)

SPH-EXA Framework Components Overview

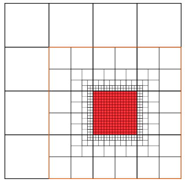

Cornerstone octree

Domain Decomposition

- Space-filling curves and octrees
- Global and locally essential octrees
- Octree-based domain decomposition
- 18’500 lines of code

SPH Solver

Modern SPH implementation with key features (astro.physik.unibas.ch/en/people/ruben-cabezon/sphynx/):

- Generalized volume elements
- Integral approach to derivatives
- Artificial viscosity with switches
- 4’000 lines of code

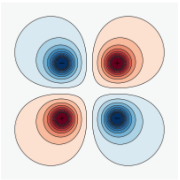

Ryoanji N-body solver

Gravity-solver on GPUs with:

- Cornerstone octrees
- Breadth-first traversal inspired by Bonsai (https://github.com/treecode/Bonsai)
- EXA-FMM multipole kernels (https://github.com/exafmm)
- 3’000 lines of code

SPH-EXA

SPH-EXA application front-end

- Handling of initial conditions, checkpointing and I/O
- Flexible combination and addition of additional physics for domain scientists
- 5’500 lines of code

Interface for Domain Decomposition

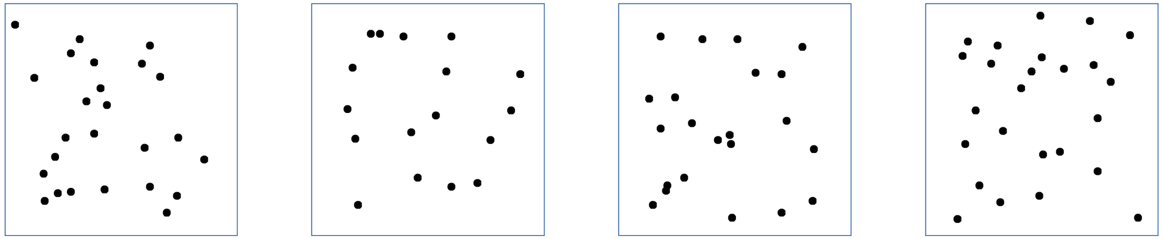

```
domain.sync(x, y, z, h_smoothingLengths, masses, charges, … );
```

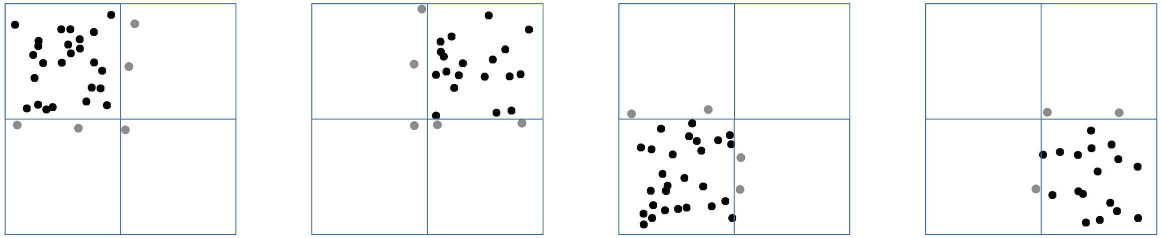

Features:

- Decompose global domain into compact subdomains
- Distribute arbitrary number of particle properties along with coordinates
- Identify and transfer halo particles
- Full GPU support: accept input data on GPU, perform octree builds and communication out of GPU
- Provide read-only access to underlying octree, for use by the N-body solver

Interface for Domain Decomposition

unordered particle data buffers (initial, or moved after time-step) on each compute node (can be on GPU)

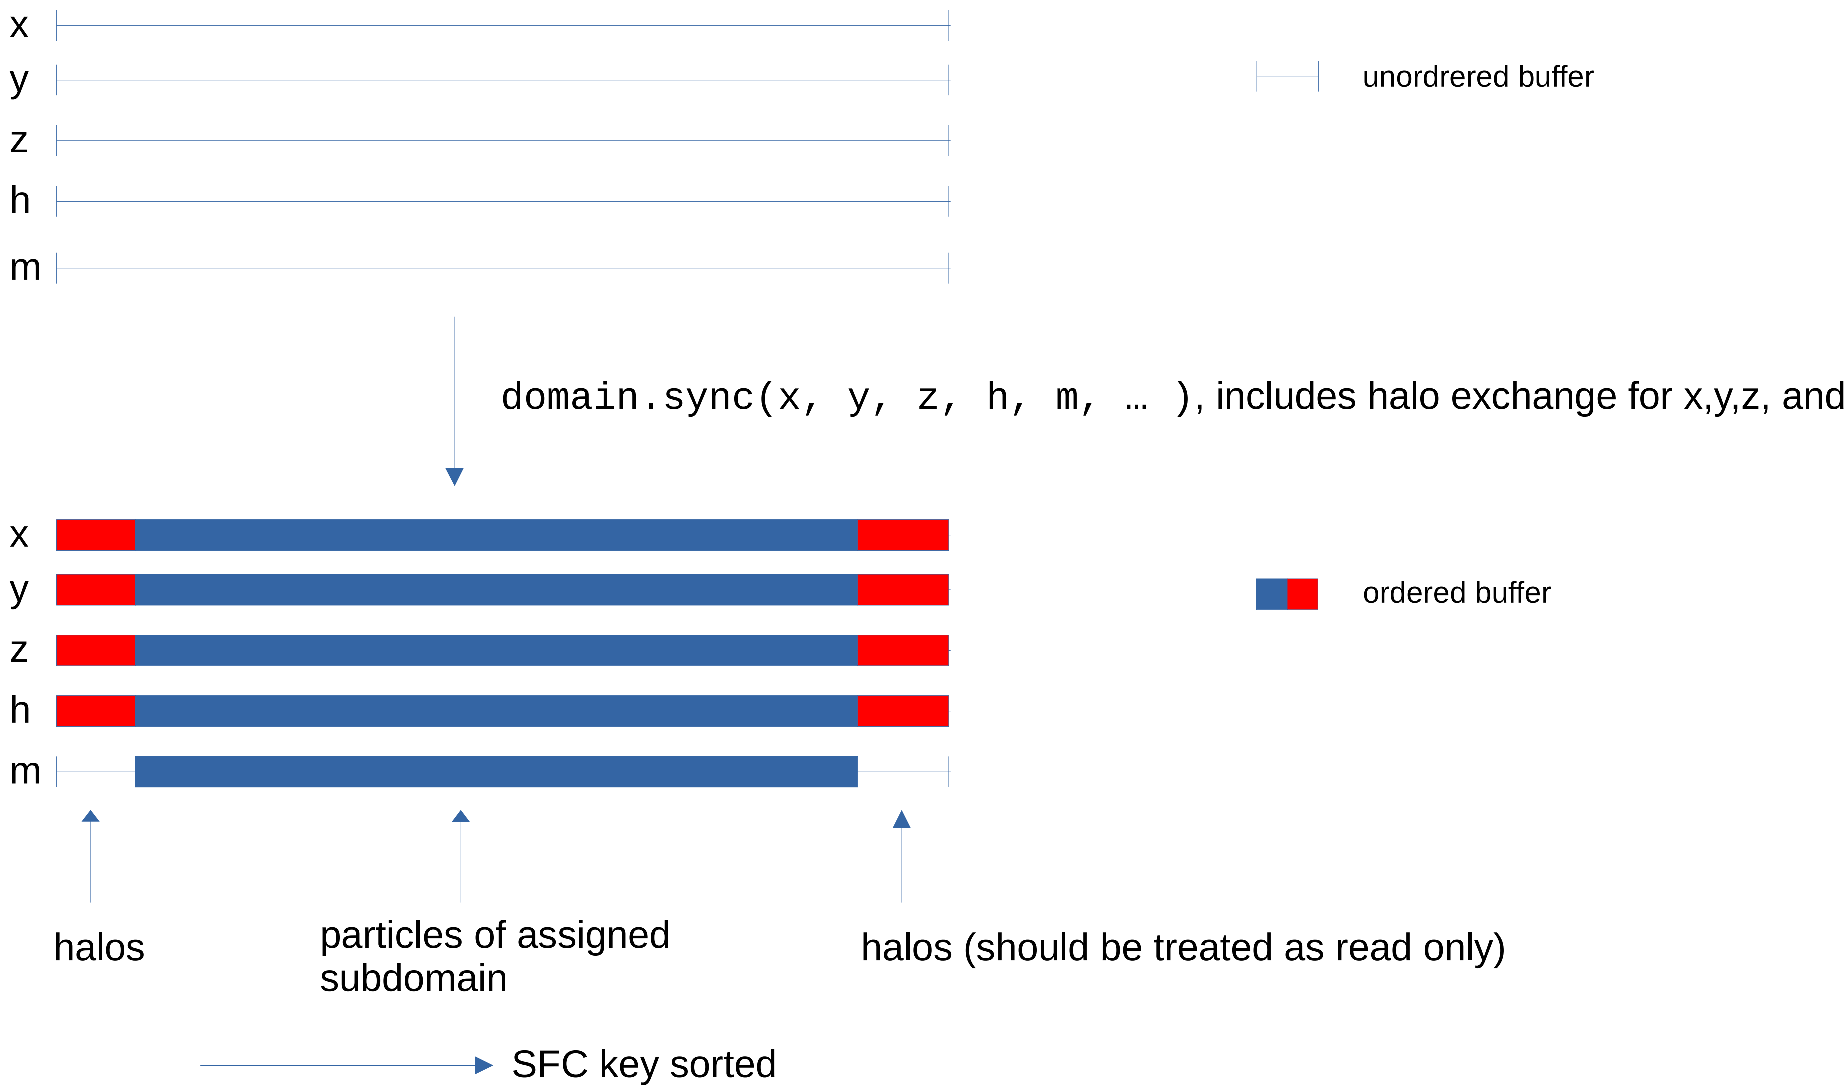

Halo Exchange

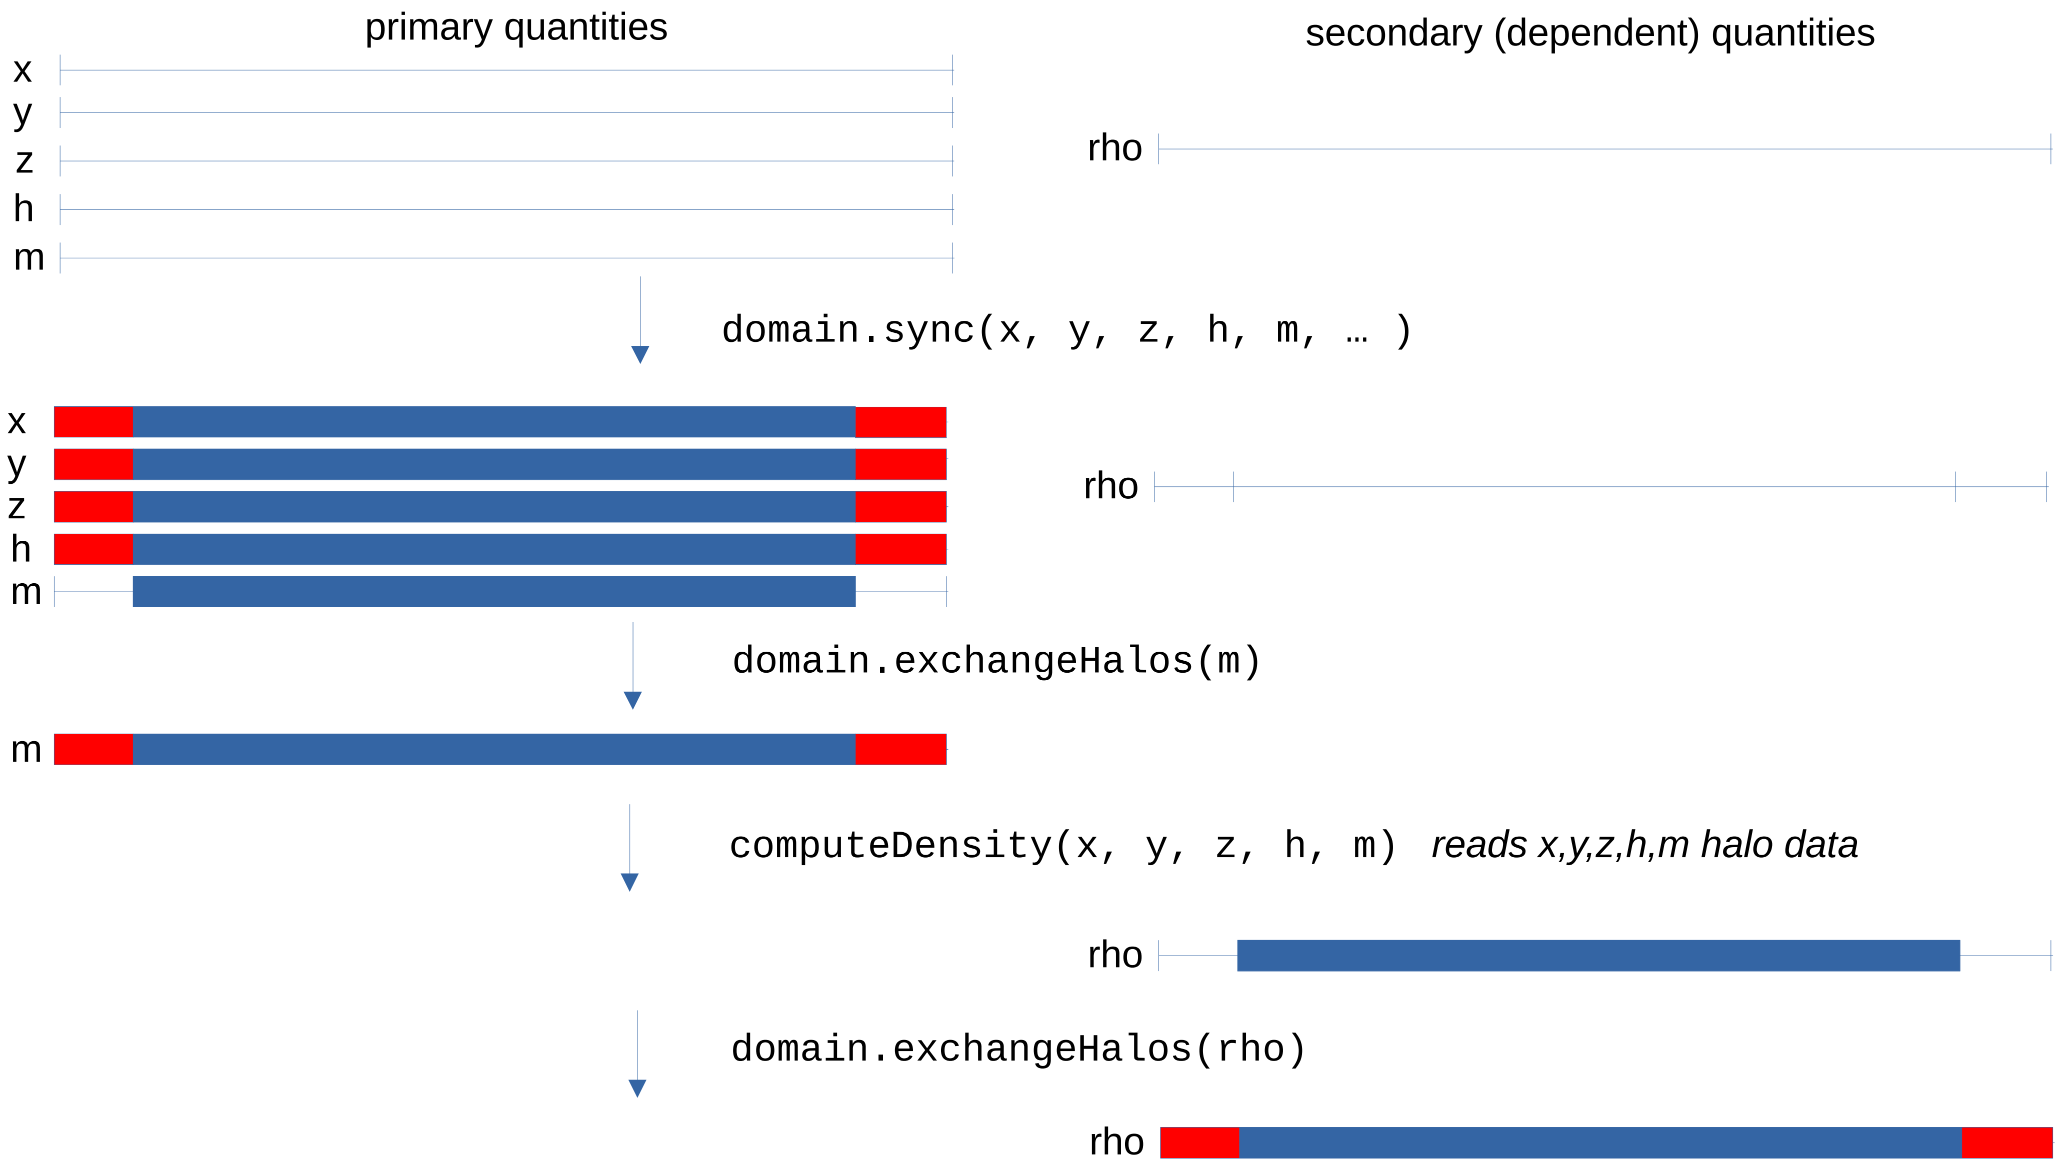

Locally Essential Octrees

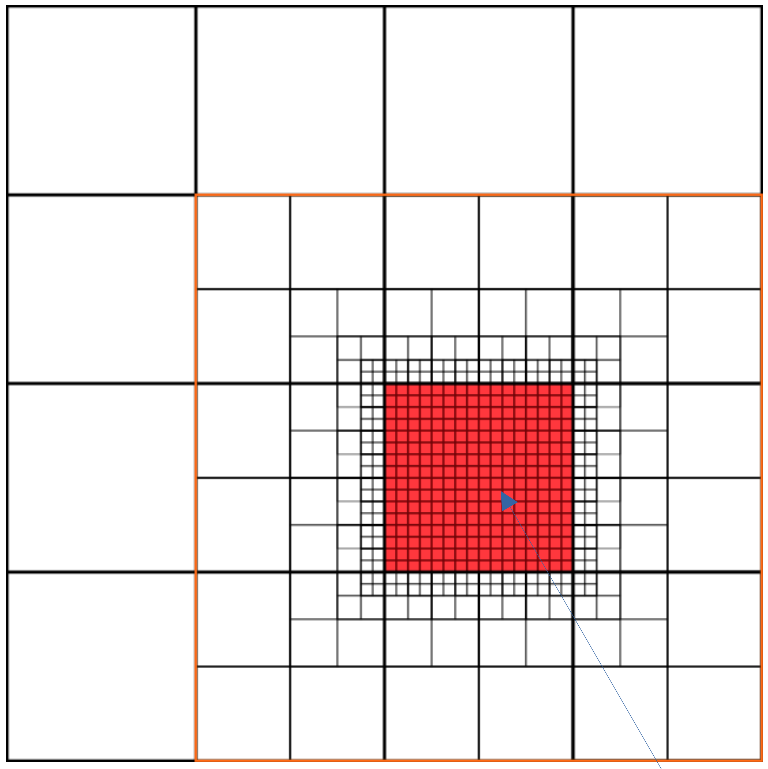

Octree requirements for calculating forces on process A:

- Must cover the global domain
- High resolution (N_{crit}) only inside the subdomain
- Angle-based (MAC) decay of resolution outside

Particles in far-away non-halo cells are not present on A

→ remote (non-halo) cells must not fail the multipole acceptance criterion (MAC) used for long-range forces computed with Barnes-Hut or FMM

→ locally essential tree is refined based on cell-counts (N_{crit}) inside the subdomain and based on MAC outside

Simplified drawing for homogeneous particle distribution:

- subdivision pattern may be irregular
- subdomains may have fractal shape

subdomain of process A in red area

Impact of the Locally Essential Octree on the Global Octree

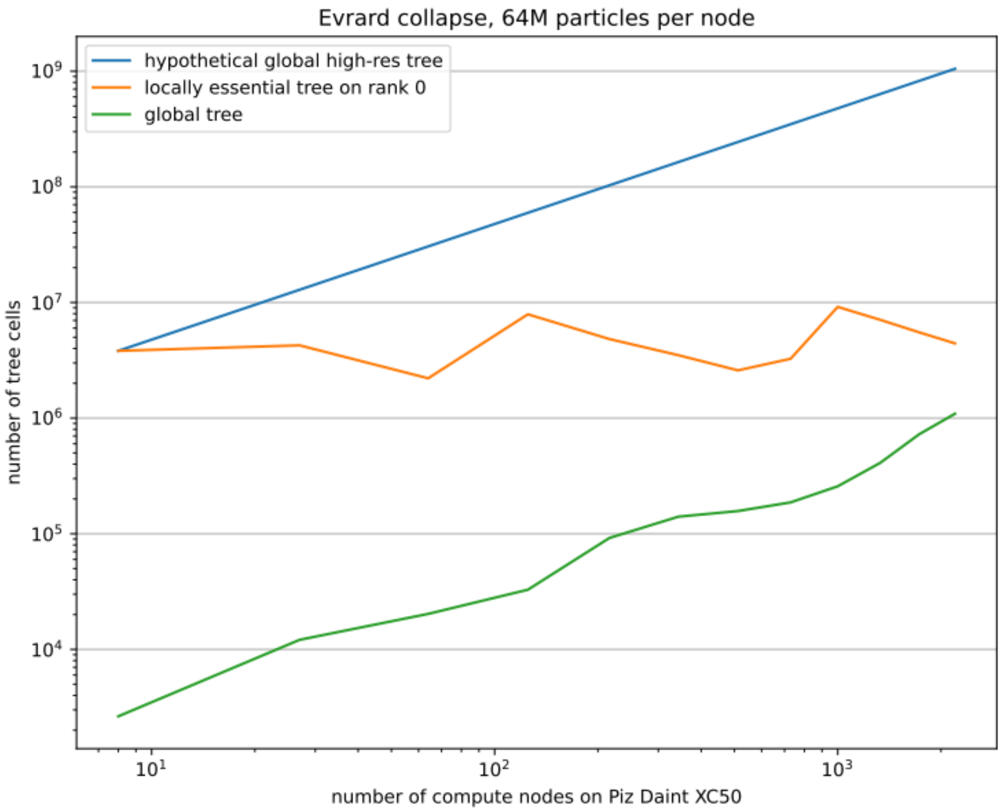

- Global octree still grows linearly with system size, but is 1’000x smaller thanks to the locally essential tree
- Locally essential octree only grows logarithmically, $N_{crit} = 64$

Support for Various GPU Architectures

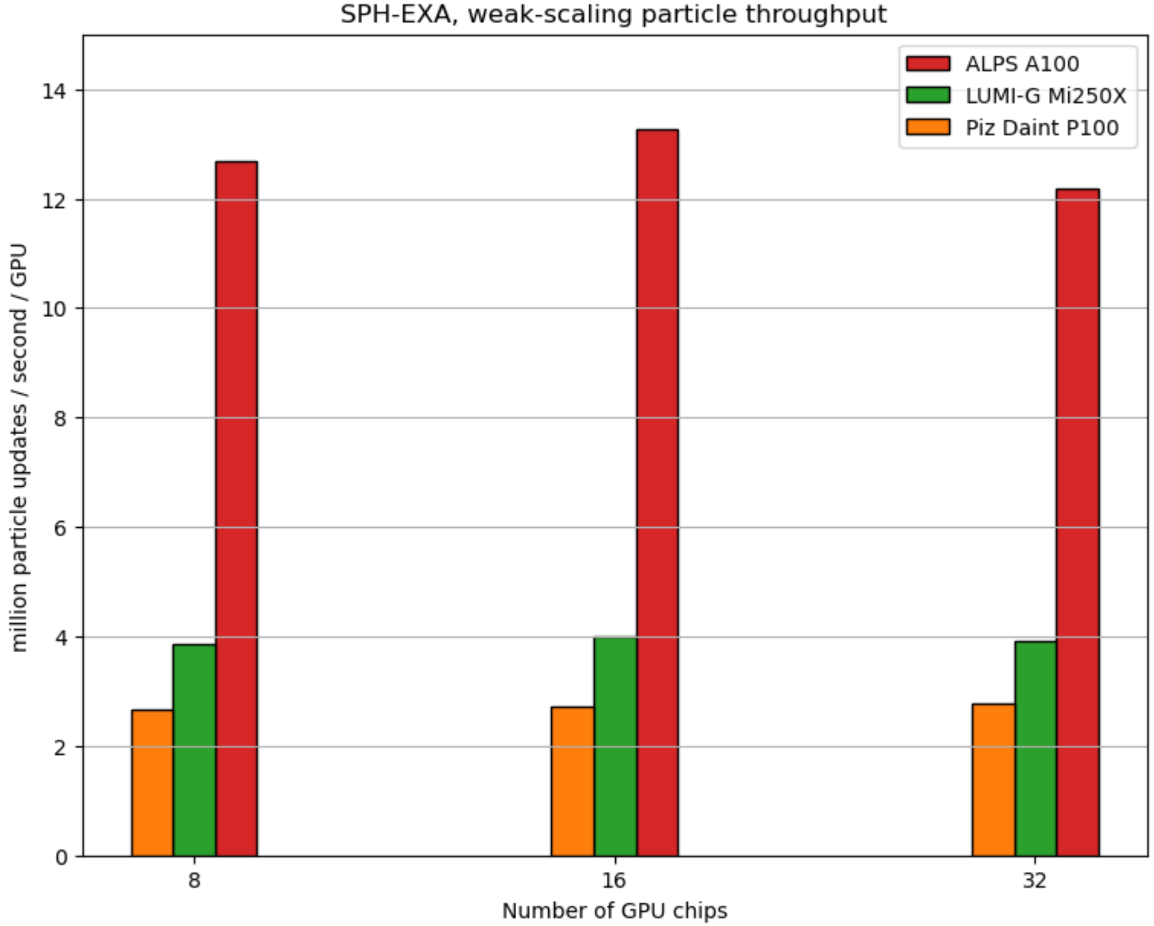

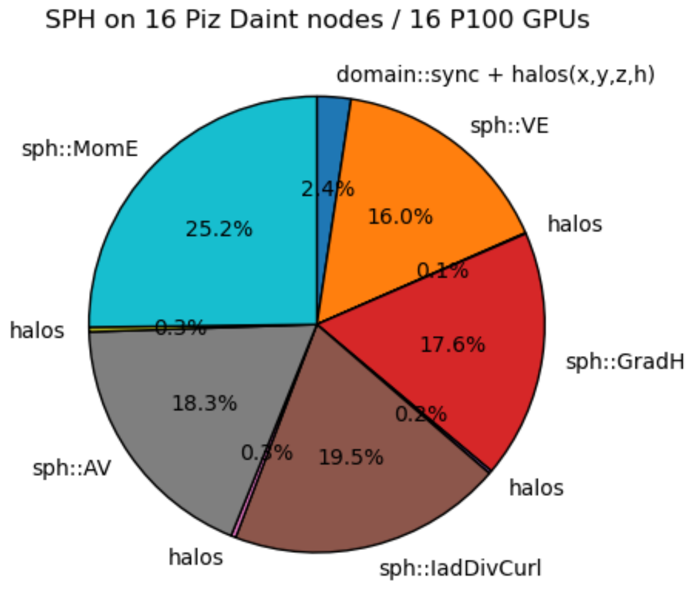

GPU performance comparison in an SPH simulation (Sedov blast)

SPH-EXA Scalability for Hydrodynamics and Gravity

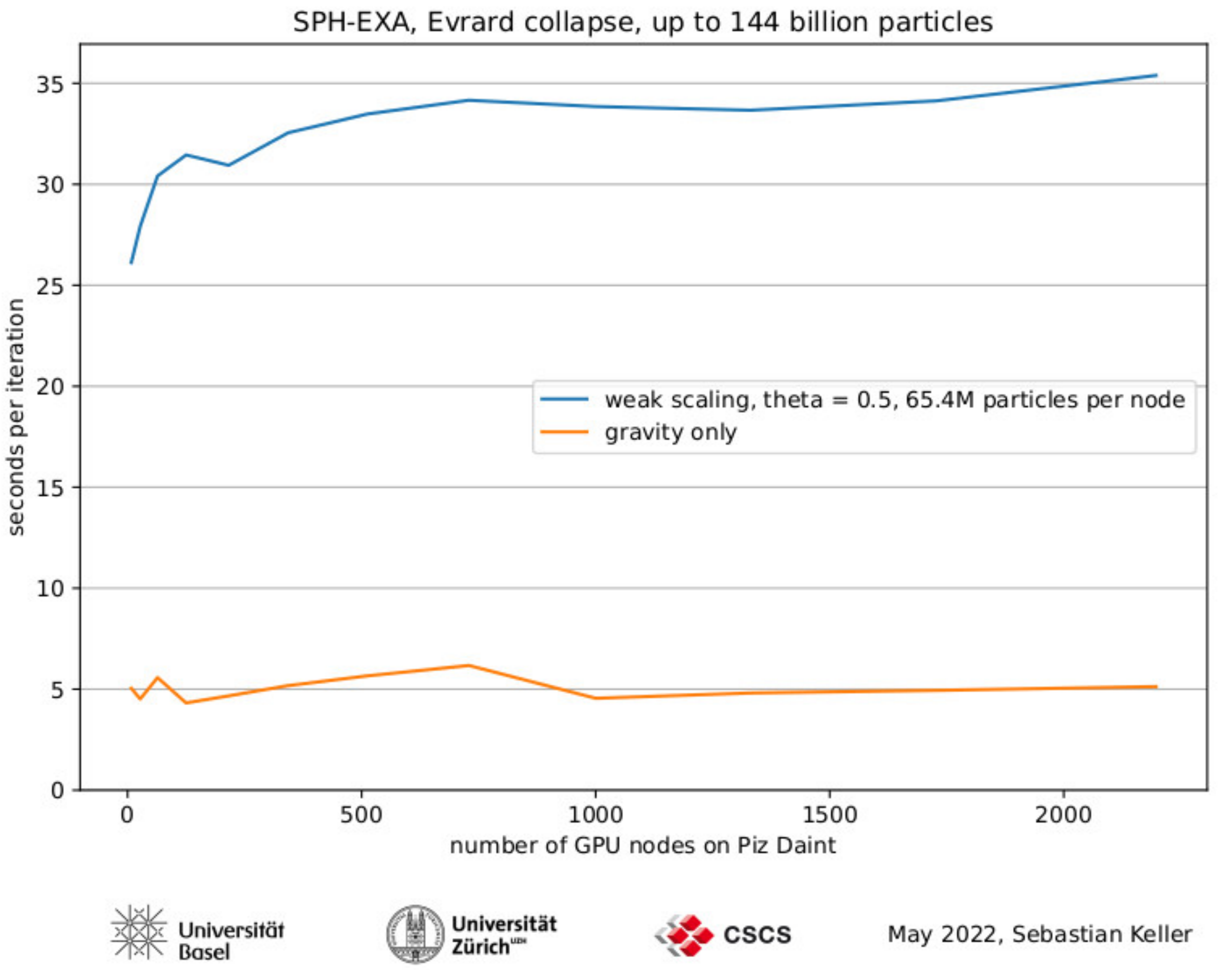

- Despite the global nature of gravitational forces, SPH-EXA scales to the full Piz Daint machine

SPH-EXA: Simulation of Subsonic Turbulence

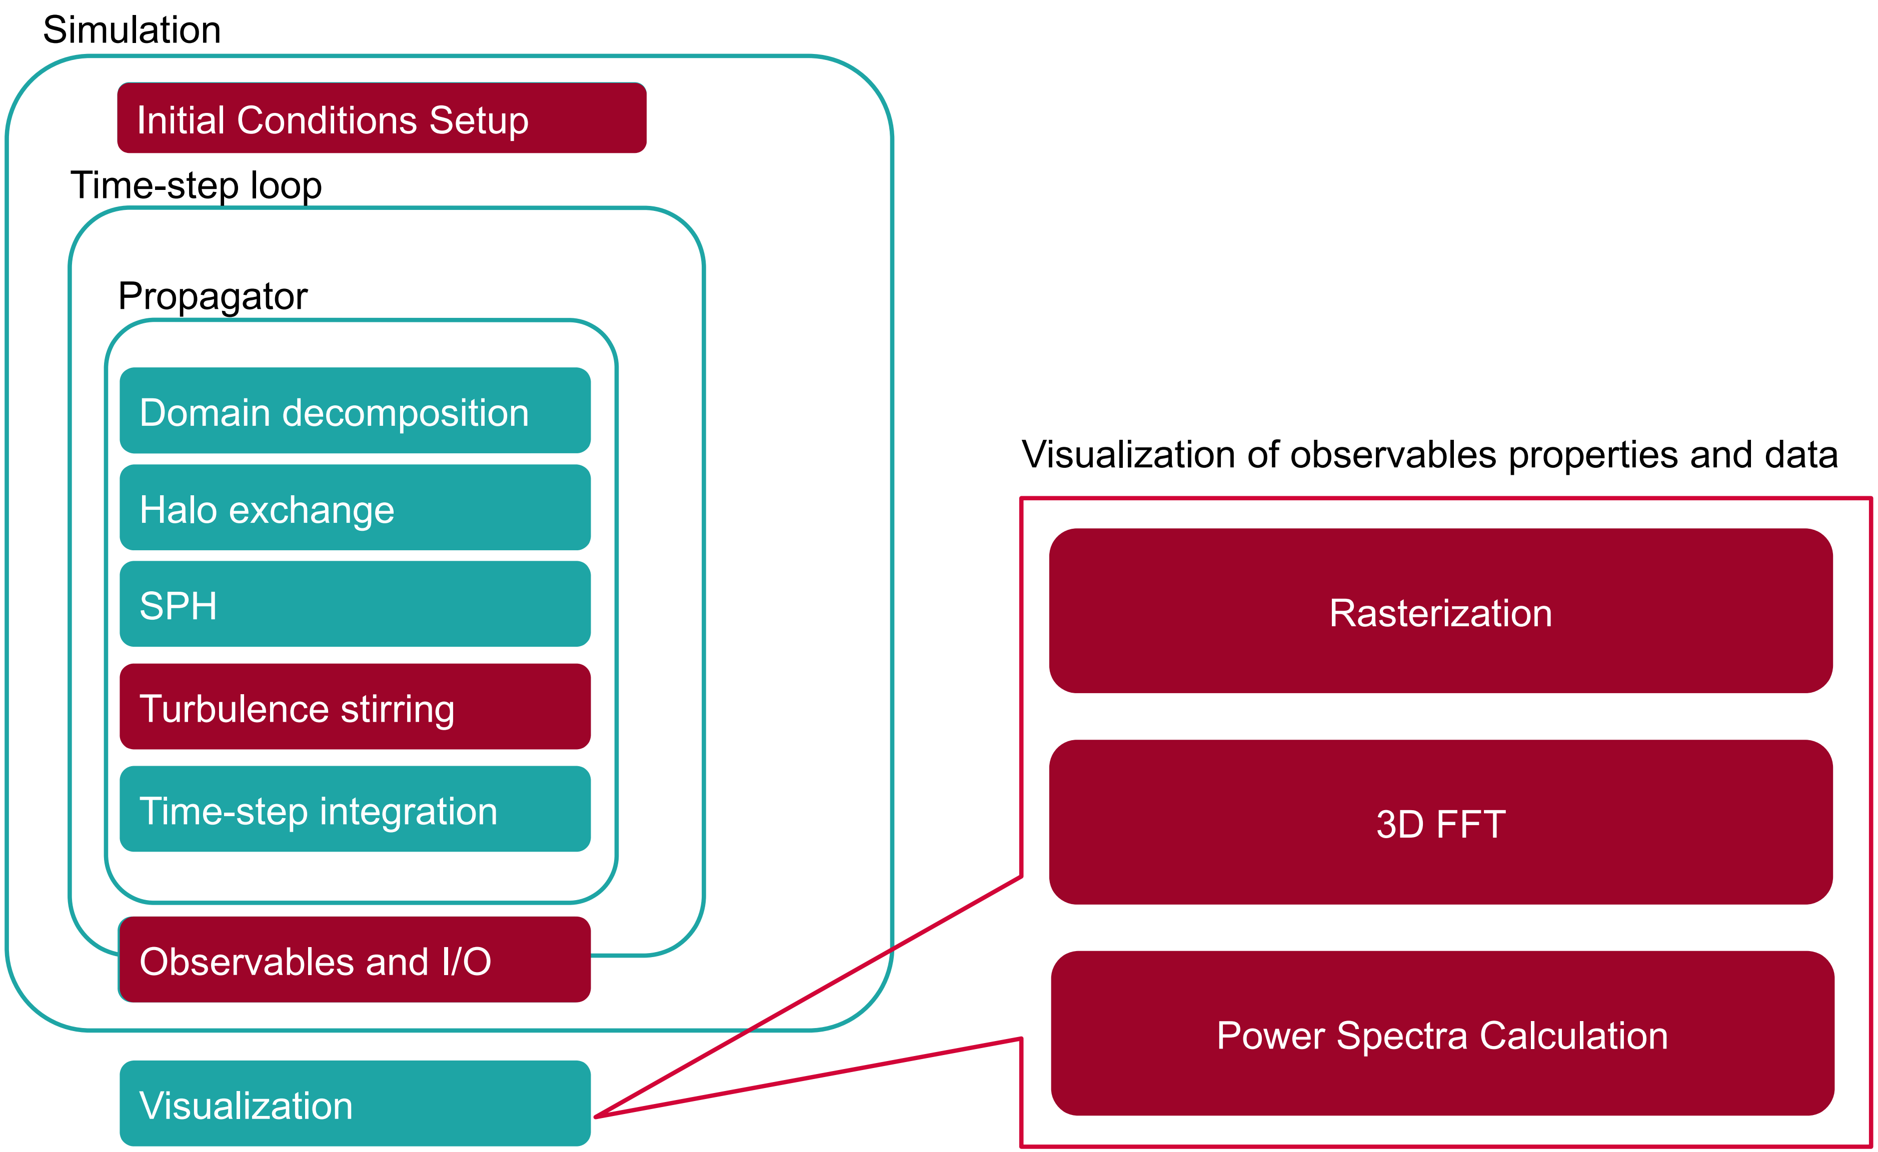

SPH-EXA: Power Spectra Calculation for Turbulence Simulations

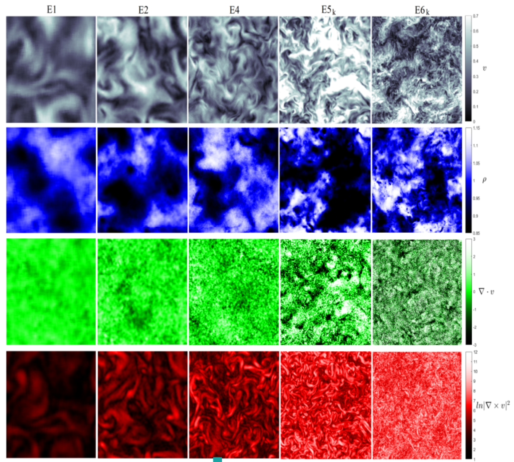

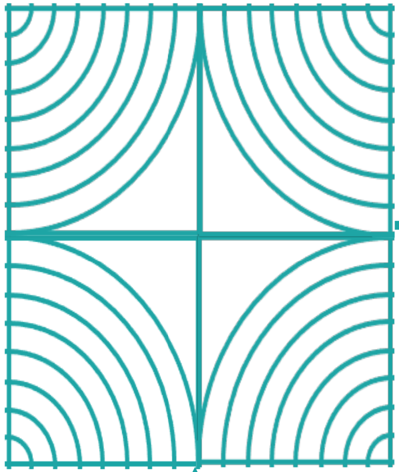

Shift/Transform

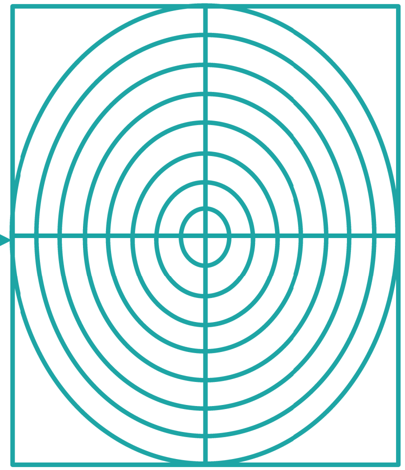

Rasterization

Parallelized Discrete 3D FFT

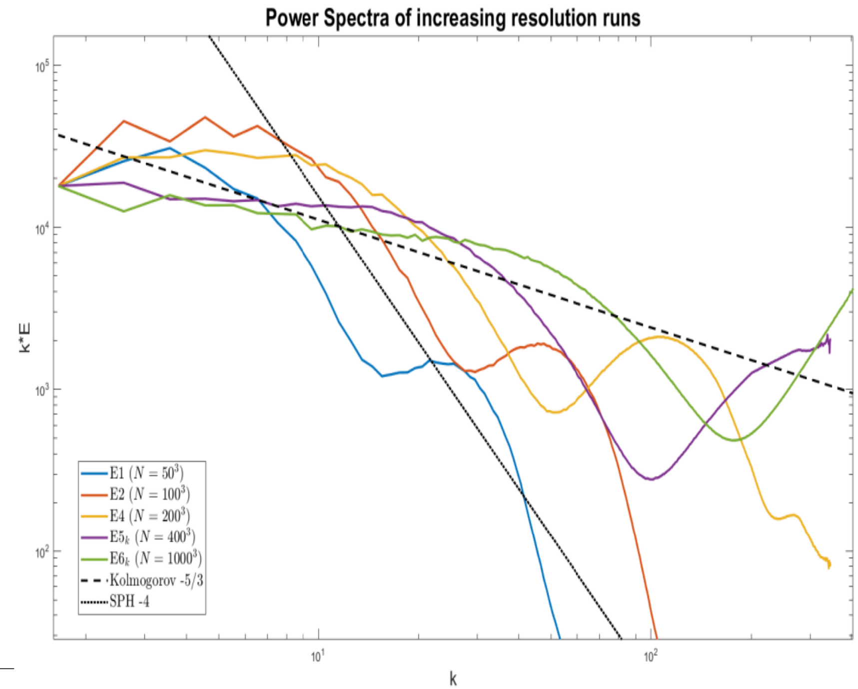

SPH-EXA: Next Steps

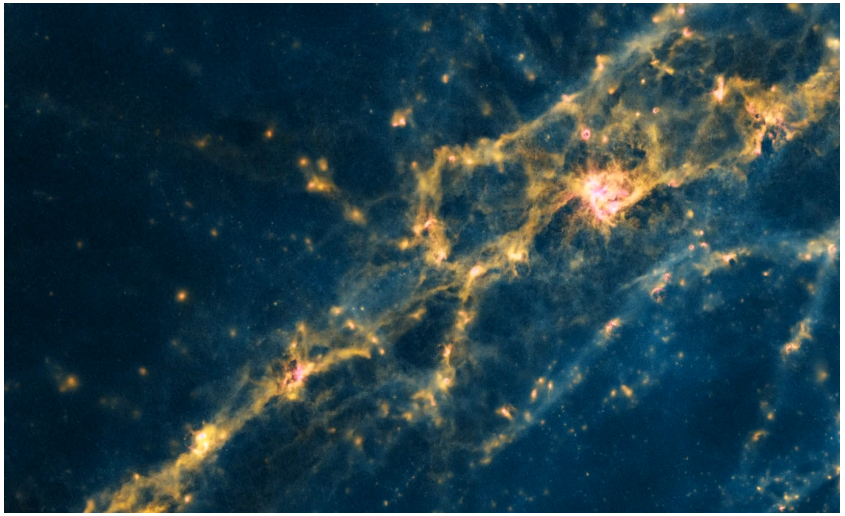

Galaxy formation simulations (PHOEBOS)
- Radiative cooling
- Stellar feedback
- Star formation

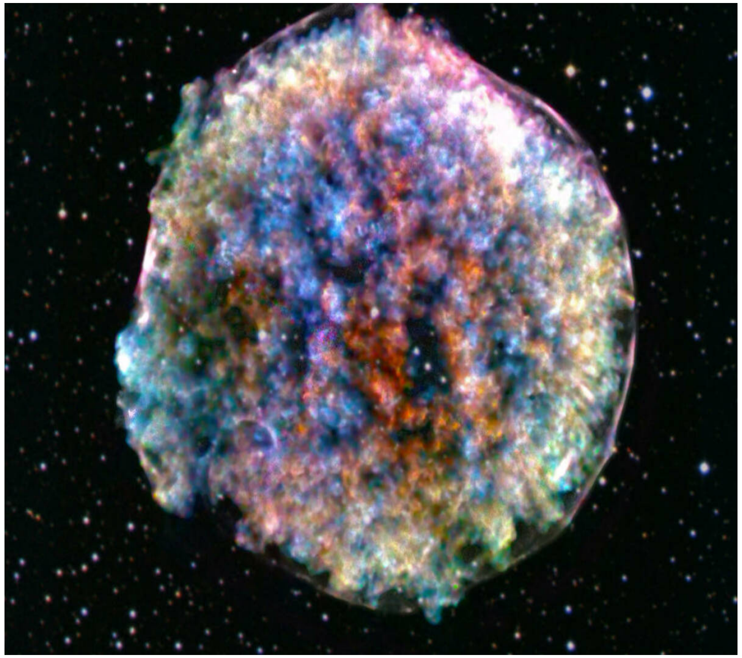

Type IA supernova simulations
- Nuclear reactions
- Nuclear flame tracking