Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

GPU porting of FFT and collapse time calculations within PINOCCHIO

Marius D. Lepinzan, P. Monaco, T. Castro and L. Tornatore

Spoke 3 III Technical Workshop, Perugia 26-29 Maggio, 2025

Scientific Rationale

PINOCCHIO is a code, based on Lagrangian Perturbation Theory (LPT), for simulating Dark Matter halos in cosmological boxes and past light cones (*Monaco et al. 2002, 2013; Munari et al. 2017)*

Comparison with full N-body simulations:

- ~1000 faster
- 5 – 10% accuracy in reproducing 2-point halo statistics, halo mass function and halo bias
- 5 – 10% accuracy in reproducing cosmic void statistics (*Lepinzan et al. in prep*)

GADGET

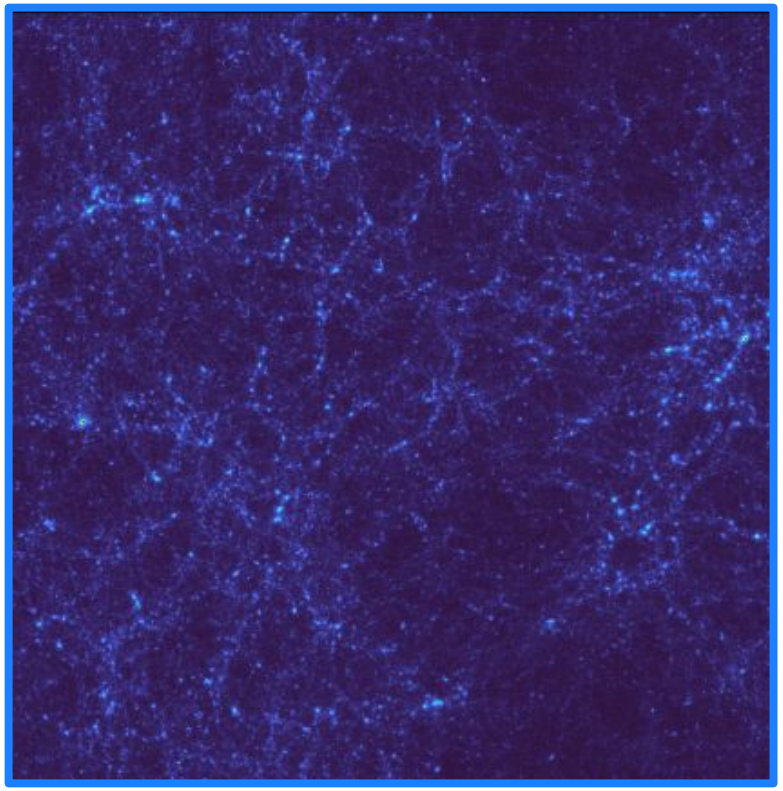

PINOCCHIO

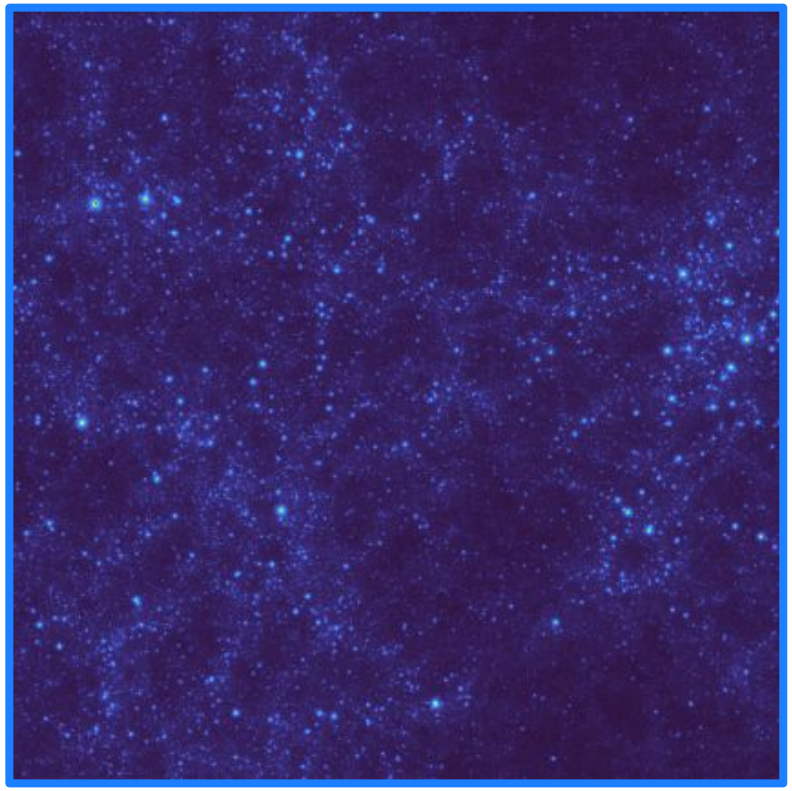

Code workflow

1. Generation of a **linear density field** on a **regular grid**
2. Computation of **collapse time** using an ellipsoidal model based on **LPT**
3. **Fragmentation** of the collapsed medium for **halo identification**
4. Creation of **halo catalogs** (box and light cones)

→ NGenIC code

→ FFTs + custom solver

→ Custom algorithm

Technical Objectives, Methodologies and Solutions

- **Optimize the code to fully leverage modern HPC infrastructure, including GPUs:**

  - **Improve code performance**: *suitable threadization? main bottlenecks ?*

  - **Identify off-loadable regions**: *what can be ported to GPUs?*

  - **Improve scientific output**: *Adopting new algorithm?*

- **Adopted solutions:**

  - **Improve** the **MPI** framework: **OpenMP**

  - **Porting** of **FFT** and **collapse times** solver to **GPU**: **heFFTe** + **OpenMP**

  - **Optimize** and **investigate** a new fragmentation algorithm: **profiling** + **ADP**

  - **Testing, bug fixing, testing, bug fixing… !!**

Main Results

Offloading of collapse times calculation on GPU with OpenMP

- GSL does not have support for GPU offloading
- Need of custom cubic spline and bilinear splines interpolation FULL GPU
- GPU offloading tested out of PINOCCHIO and compared with GSL

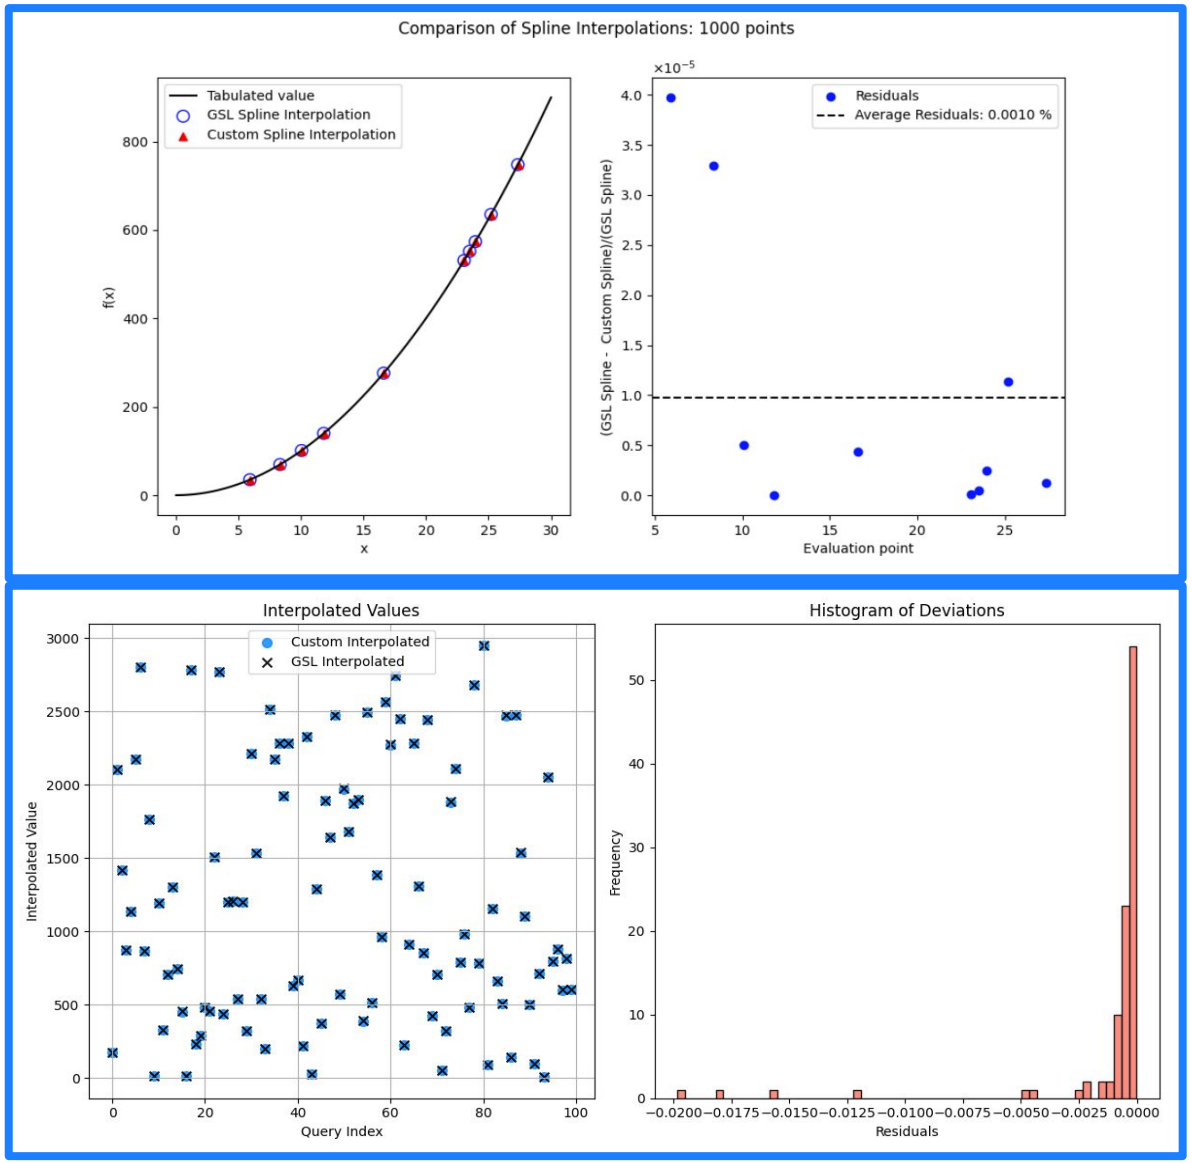

Main Results

Offloading of collapse times calculation on GPU with OpenMP

- Test on NVIDIA and AMD platform. Offloading and performance portability achieved: ~10x speedup (*Lepinzan et al. in prep*)
- Euclid Box (box ~ 3.5 Gpc, 6144^3 particles) number of halos variations: 10^{-3}
- Euclid Box gain for single run: ~ 110 s
- 1000 of Euclid Box gain: ~ 30 hours (wall-time)

(Collaboration with Giovanni Lacopo and David Goz)

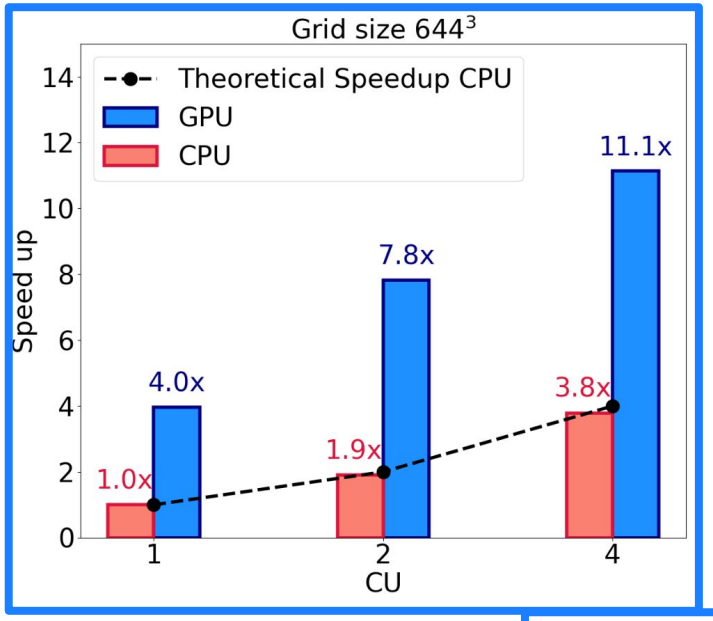

NVIDIA *(Leonardo)*

AMD *(Setonix)*

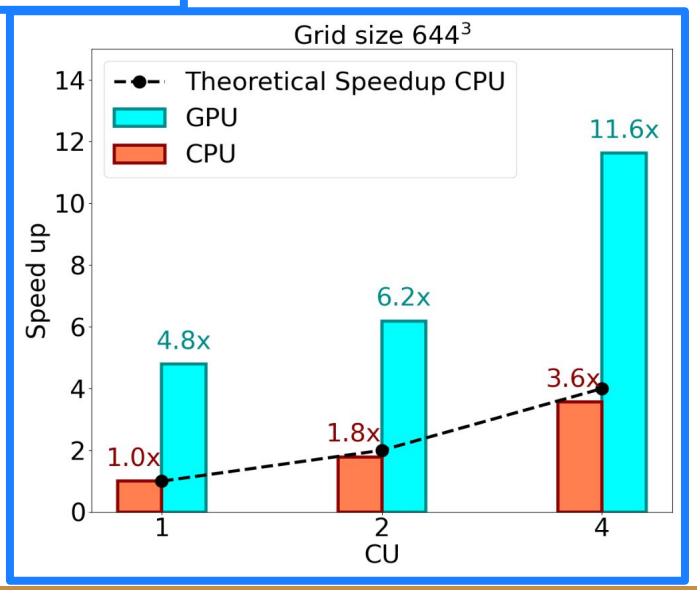

Main Results

Offloading of FFT calculation on GPU with heFFTe into a PINOCCHIO mini-app

- Test on NVIDIA platform. Offloading achieved: ~28x speedup (*Lacopo, Lepinzan et al. in prep*)
- Euclid Box (box ~ 3.5 Gpc, 6144^3 particles) Fmax distribution: binary equal
- Euclid Box gain for single run: ~ 370 s
- 1000 of Euclid Box gain: ~ 100 hours (wall-time)

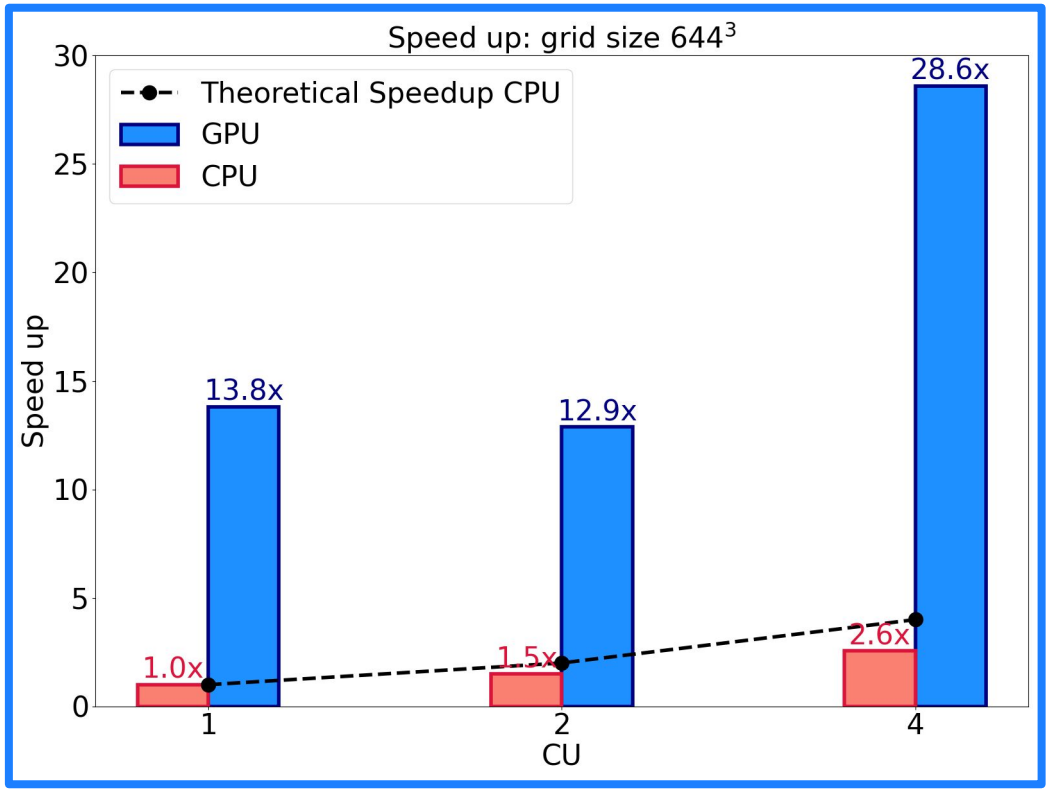

NVIDIA
(Leonardo)

(Collaboration with Giovanni Lacopo and David Goz)

Final Steps (short timescale)

Optimize the **legacy fragmentation** algorithm

- **Identified** the main bottleneck using **Score-P** profiling: **binary-search**
- **Parallelize** the **CPU** version using **multi-threading**
- **Replace** the **binary-search CPU** version with a **GPU-based** implementation

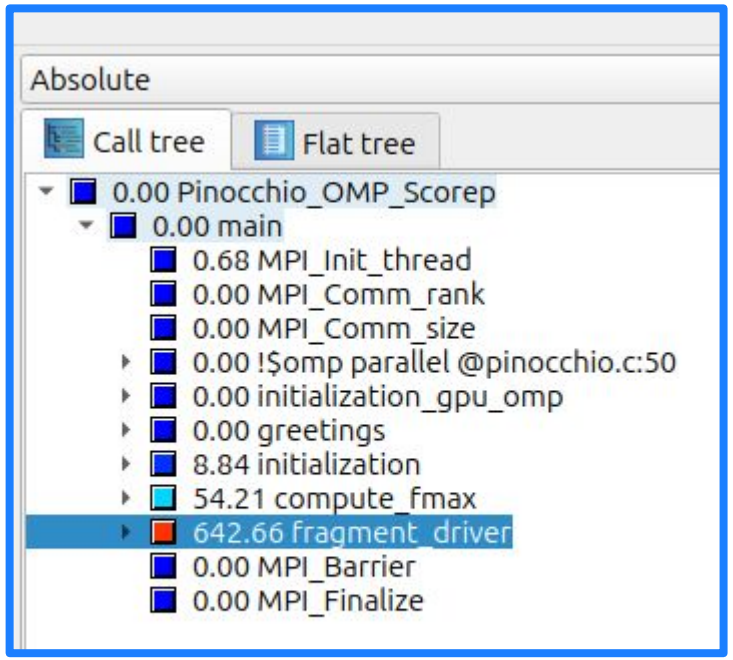

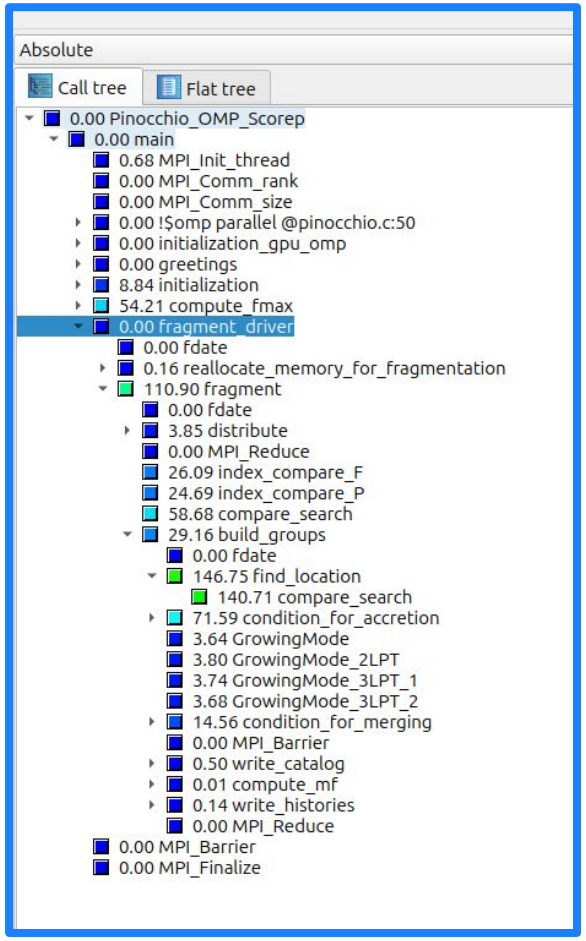

Final Steps (longer timescale)

New methodology for the fragmentation

- **Clustering algorithm** (Advance Density Peak) for a **domain decomposition**: identify **Eulerian patches** that will **end up in halos** according to **PINOCCHIO**
- Apply the **current algorithm** for **fragmentation** on every **independent domain**

Eulerian patches

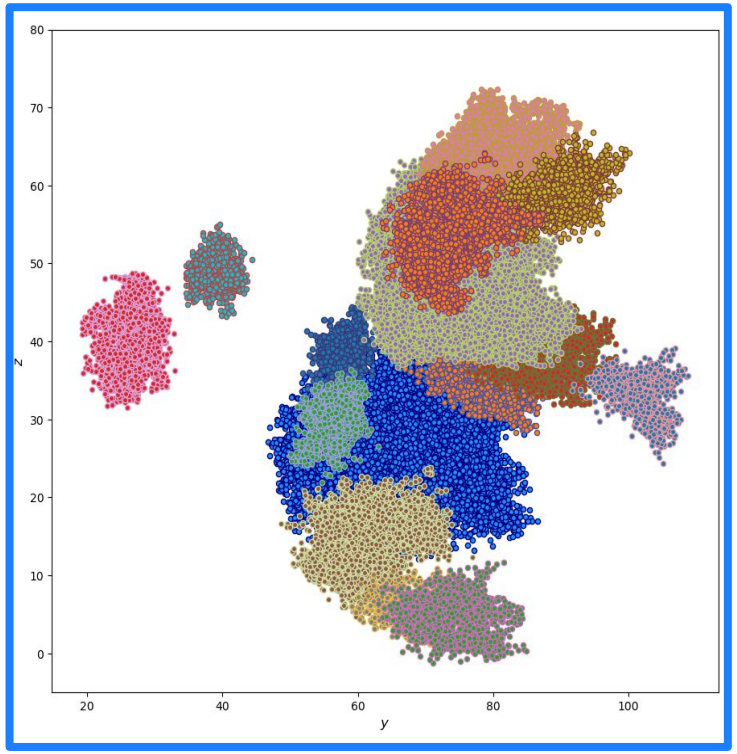

Lagrangian patches

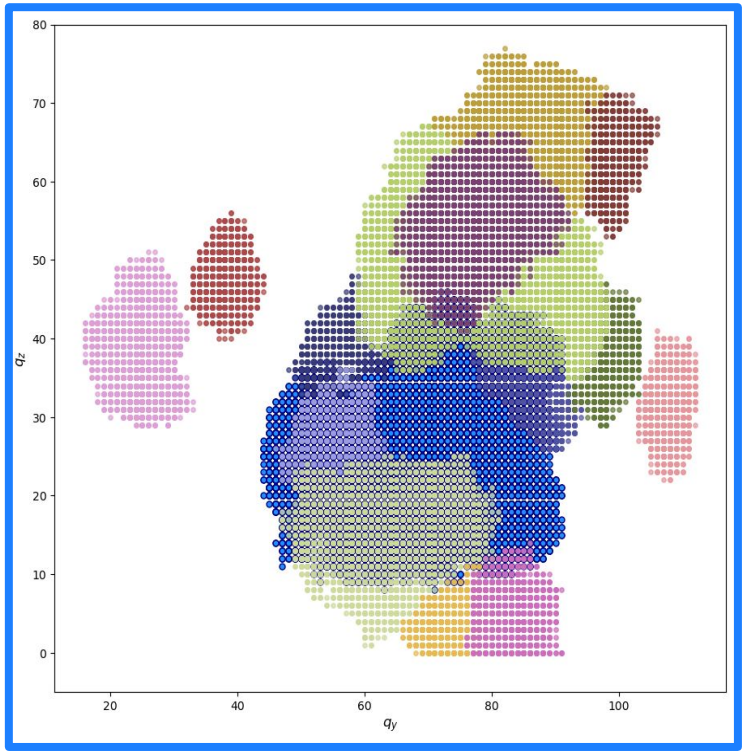

Final Steps (longer timescale)

New methodology for the fragmentation

- **Bypass** the **Eulerian space** and apply the **clustering algorithm** directly to a regular **3D grid** of points, using the **collapse time (FMAX)** for each particle provided by PINOCCHIO

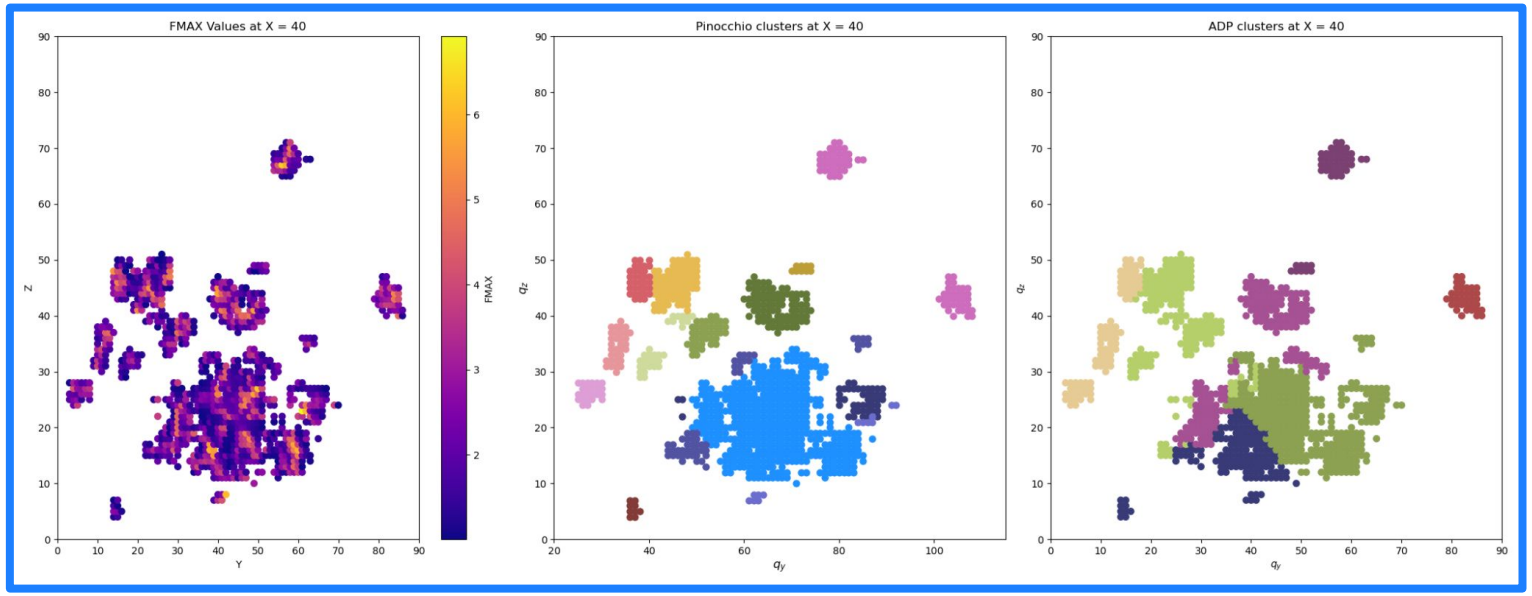

Timescale, Milestones and KPIs

04/2023 | 09/2023 | 02/2024 | 06/2024 | 10/2024 | 08/2025

MS 6

1. Familiarization with the code
2. Bug fixing
3. Identification off-loadable regions and bottleneck

MS 7

1. Threadization of collapse time with OpenMP
2. GPU offloading of collapse times with OpenMP: delayed because of GSL incompatibility

MS 8

1. GPU offloading of collapse times with OpenMP: done
2. Machine learning approach for fragmentation: ongoing
3. New stable release of the code (v5.1): normal, neutrinos and modified gravity (MG) test done

MS 9

1. Collapse times GPU: Multi-MPI, Multi-GPU version tested on NVIDIA and AMD platforms
2. Machine learning approach: ADP vs HDBSCAN on going
3. Stable 5.1 version: ready without MG

MS 10

1. New fragmentation
2. FFTs on GPUs
3. 100 simulations up 1000^3 particles
4. Technical paper on the code

KSP *EuMocks: Mocking the Universe for Euclid++*

Pierluigi Monaco, UniTS (+INAF, INFN, IFPU, ICSC),
M. Lepinzan, T. Castro, L. Tornatore, G. Taffoni, C. Carbone

coming soon: (Euclid) paper presenting the largest collection of halos on the past light cone ever produced

Euclid preparation. Simulating thousands of Euclid spectroscopic skies

Euclid Collaboration: P. Monaco[1,2,3,4]*, G. Parimbelli, M. Y. Elkhashab, J. Salvalaggio, T. Castro, M.D. Lepinzan, E. Sefusatti, L. Stanco, L. Tornatore, G. Addison, S. Bruton, C. Carbone, F. Castander, J. Carrettero, S. de la Torre, P. Fosalba, G. Lavaux, S. Lee, K. Markovich, K. McCarthy, F. Passalacqua, W. Percival, I. Risso, C. Scarlata, P. Tallada-Crespí, Y. Wang, et al.

(Affiliations can be found after the references)

May 26, 2025

we plan a press release on this release
these are being used to prepare the analysis of Euclid's DR1 spectroscopic sample
ongoing work with WP4 to make the simulations available to the community through IA2
ongoing work with WP5 to make online processing of simulations possible in the datalake
preparation for a demonstrator set of massive simulations for Euclid's DR2 is in progress