HPC & Big Data @ UniTS

Stefano Borgani
Dept. of Physics - University of Trieste
INAF – Astronomical Observatory of Trieste
ICSC - National Center for HPC, Big Data and Quantum Computing

1. HPC & large simulations of cosmic structures @ DF-UniTS
   - Collaboration with INAF-OATs
   - ICSC Spoke-3 and beyond

2. HPC @ Dept. of Physics : activities in Spoke-2 and Spoke-7

3. HPC @ UniTS at large: activities in Spoke-7

Talk @ INAF USC-C Assembly - Trieste March 9th - 13th 2026

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

Franklin Aldas (INAF-TS)
Paramita Barai (INAF-TS)
Veronica Biffi (INAF-TS)
Stefano Borgani (UniTS, INAF-TS)
Tiago Castro (INAF-TS, Univ. São Paulo)
Alice Damiano (INAF-TS)
Michela Esposito (SISSA)
Alessandra Fumagalli (INAF-TS)
David Goz (INAF-TS)
Gianluigi Granato (INAF-TS)
Giovanni Lacopo (INAF-TS)
Marius Lepinzan (INAF-TS)
Pierluigi Monaco (UniTS, INAF-TS)
Giuseppe Murante (INAF-TS)
Antonio Ragagnin (INAF-TS)
Cinthia Ragone-Figueroa (U.Cordoba, INAF-TS)
Elena Rasia (U.Michigan, INAF-TS)
Alex Saro (UniTS, INAF-TS)
Tommaso Tarchi (INAF-TS)
Giuliano Taffoni (INAF-TS)
Luca Tornatore (INAF-TS)
Milena Valentini (UniTS, INAF-TS)

LPT approximate simulation: Pinocchio (see talk by P. Monaco) – Euclid
N-body simulations: OpenGadget3 (see talk by L. Tornatore) – Euclid
Hydrodynamical simulations: OpenGadget3 – Euclid, galaxy formation

Computational Resources: CINECA (Leonardo Booster & DCGP), Pleiadi
O(10^7) core-h/year *(ISCRA, EuroHPC, ICSC)*
Storage: DRES-CINECA, IA2: ~***400 TB***

What is a cosmological simulation?

ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

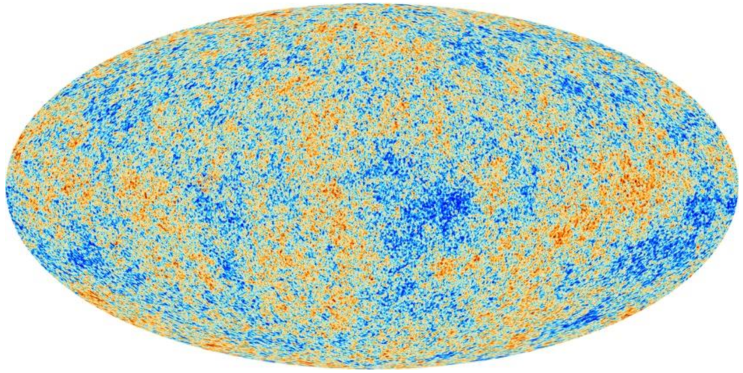

380.000 years after the Big Bang

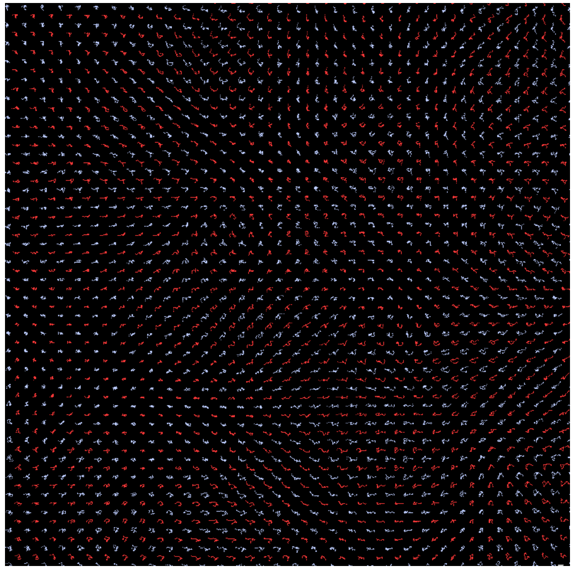

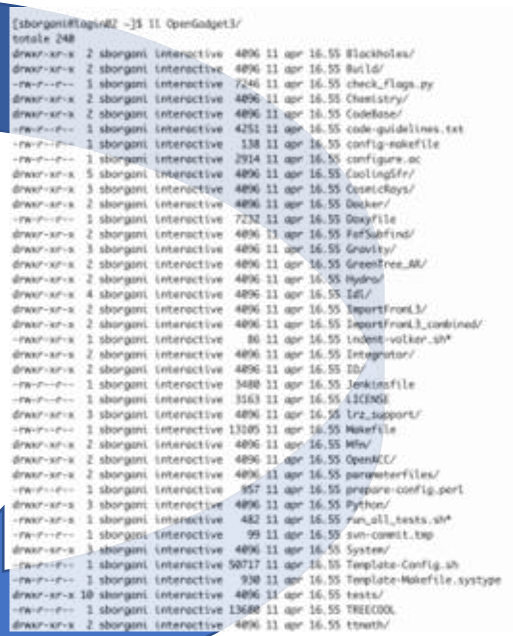

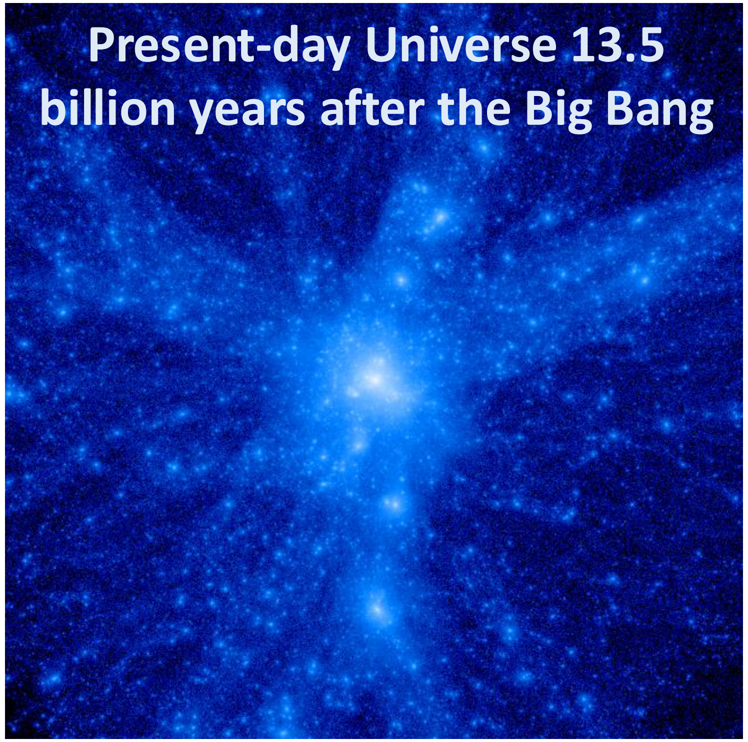

What is a N-body + hydro code ?

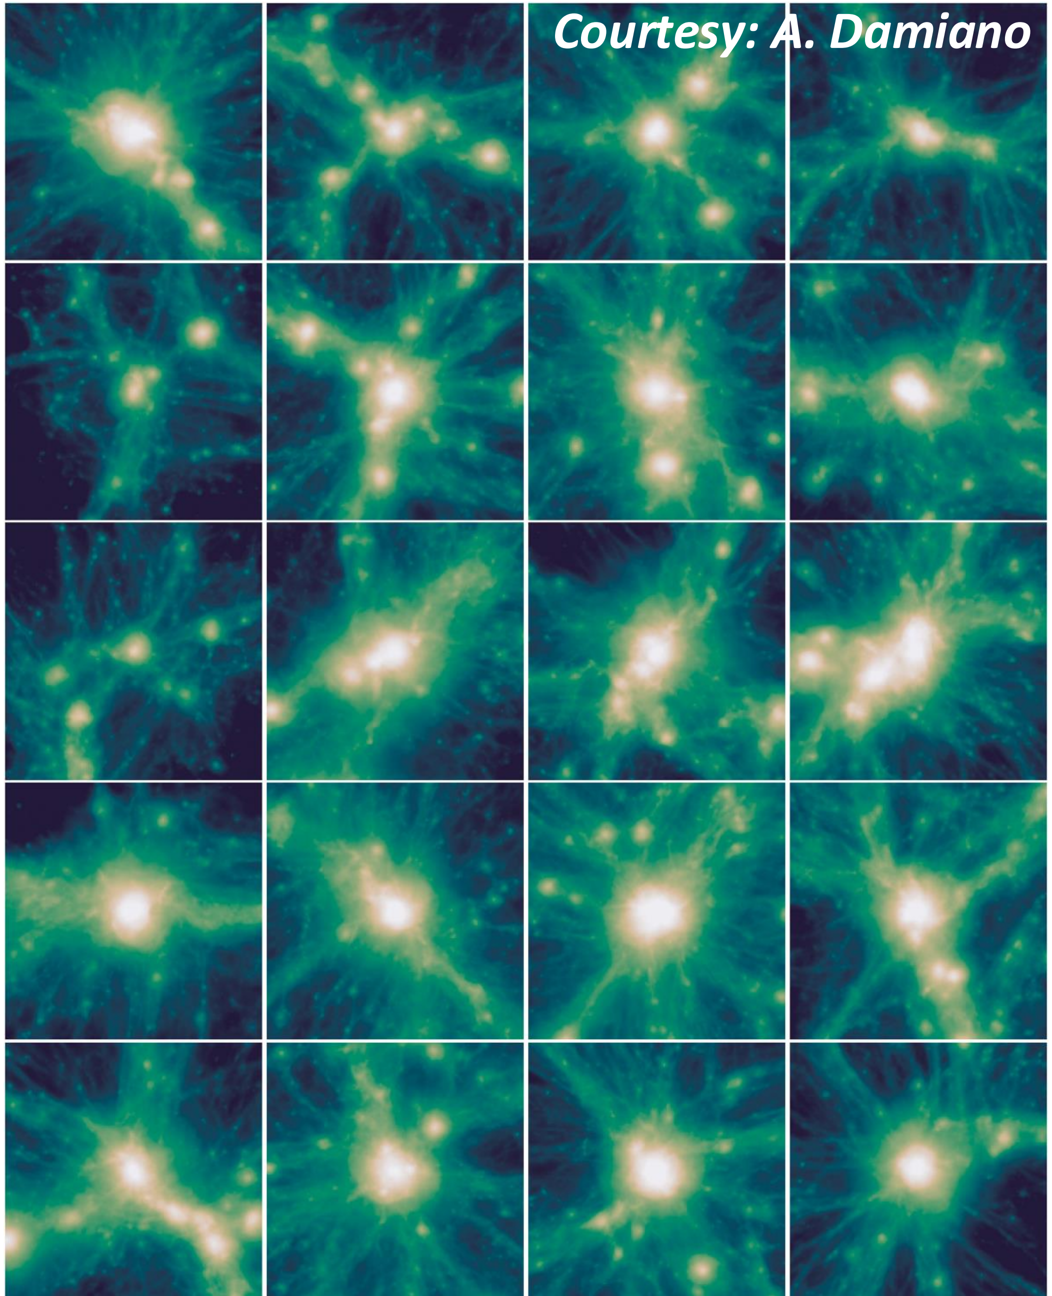

Example of OpenGADGET3 code (Dolag+26; Tornatore+26)
Collab. With USM/LMU - See talk by L. Tornatore

Gravity solver: TreeCode (short) + ParticleMesh (long)

Hydro solver:
to integrate fluid Euler Equations

1. Smoothed Particle Hydrodynamics
2. Meshless Finite Mass (MFM)

\+ Higher-order kernels, "Wake-up" for time-step of gas particles, Time-dependent artificial viscosity, Artificial thermal diffusion, Plasma Effects

Additional (Astro)physics:
Radiative Cooling, Star Formation, SN feedback, Chemical enrichment, dust formation and evolution, SMBH evolution (spin, dynamical friction), AGN feedback (thermal and kinetic), non-equilibrium chemistry, MHD, cosmic rays + ….

Parallelism: Hybrid MPI/OpenMP/OpenACC

The first “big splash”

40 k CPU-h on IBM-Power4 @ CINECA ; First INAF-CINECA agreement for supercomputing resources

2 x 480^3 DM + gas particles in a 192 h^{-1} Mpc ; full physics

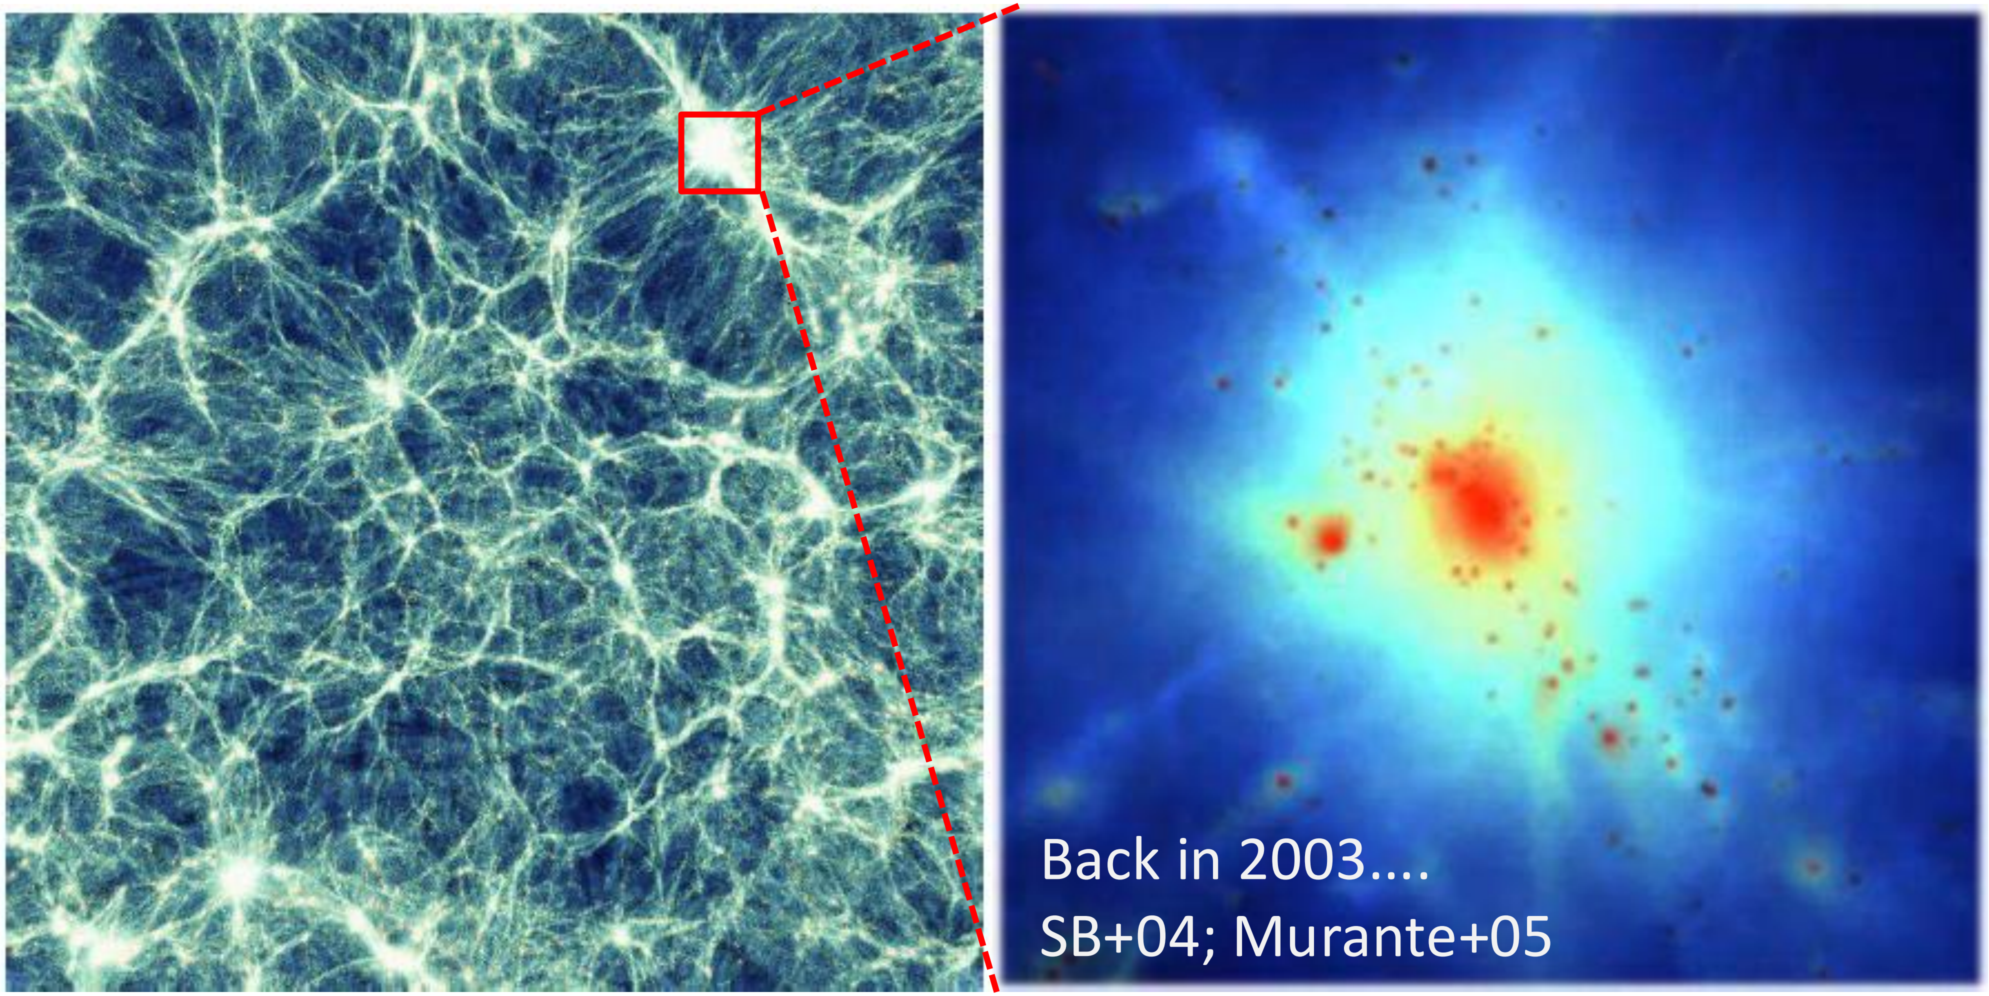

High-res simulations of galaxy clusters

ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

***A KSP supported by ICSC** (Esposito+25; SB+26, Damiano+26, Rasia+26, …)*

Model	ϵ_f	(α_h, α_c)	Eddington	Evaporation	Regions
M1	0.10	(10,100)	yes	no	29
M2	**0.05**	(10,100)	yes	no	29
M3	**0.05**	**(3,30)**	yes	no	2
M4	**0.05**	(10,100)	yes	**yes**	4
M5	0.1	(10,100)	yes	**yes**	3
M6	**0.2**	(10,100)	yes	**yes**	3
M7	**0.05**	(10,100)	**no**	no	8

- Changing the BH feedback efficiency ε_f: fraction of BH radiated energy which is thermally coupled to the surrounding gas
- Changing the ''boost factor'' by which Bondi accretion is multiplied for the cold and the hot accretions modes
- Removing the Eddington limit
- Allow BH feedback to evaporate the cold clouds in the multi-phase particles (*see talk by A. Damiano*)

→ Total of 336 halos with $M_{200}>1.5\times10^{13}\ M_\odot$ in a complete set

→ Total of 78 regions simulated, with an average cost of 200 kcore/h per region

Region	M_{200} [$10^{14} M_\odot$]	R_{200} [Mpc]	N_{13}	Models
D1	18.61	2.49	6	M1,M2,M7
D2	4.87	1.59	6	M1,M2,M3,M4,M5,M6,M7
D3	7.36	1.83	4	M1,M2,M7
D4	5.70	1.68	2	M1,M2,M4,M5,M6,M7
D5	1.79	1.14	3	M1,M2,M7
D6	15.19	2.33	12	M1,M2
D7	16.58	2.40	15	M1,M2
D8	18.04	2.46	14	M1,M2,M4
D9	1.44	1.06	3	M1,M2,M3,M4,M5,M6,M7
D10	16.51	2.39	10	M1,M2,M7
D11	13.17	2.22	14	M1,M2
D13	17.89	2.46	15	M1,M2
D14	20.07	2.55	7	M1,M2,M4
D15	19.65	2.54	11	M1,M2
D16	38.79	3.18	21	M1,M2
D17	17.90	2.46	4	M1,M2
D18	12.03	2.15	12	M1,M2,M7
D19	16.19	2.38	12	M1,M2
D20	19.83	2.54	12	M1,M2
D21	17.21	2.43	5	M1,M2
D22	21.92	2.63	14	M1,M2
D23	15.39	2.34	8	M1,M2
D24	14.93	2.31	4	M1,M2
D25	11.39	2.11	5	M1,M2
D26	15.97	2.37	9	M1,M2
D27	18.11	2.47	4	M1,M2
D28	22.39	2.65	20	M1,M2
D29	17.17	2.42	16	M1,M2
Box	1.76	1.13	25	M1,M2
Box	1.69	1.12	–	–

Dianoga Simulations

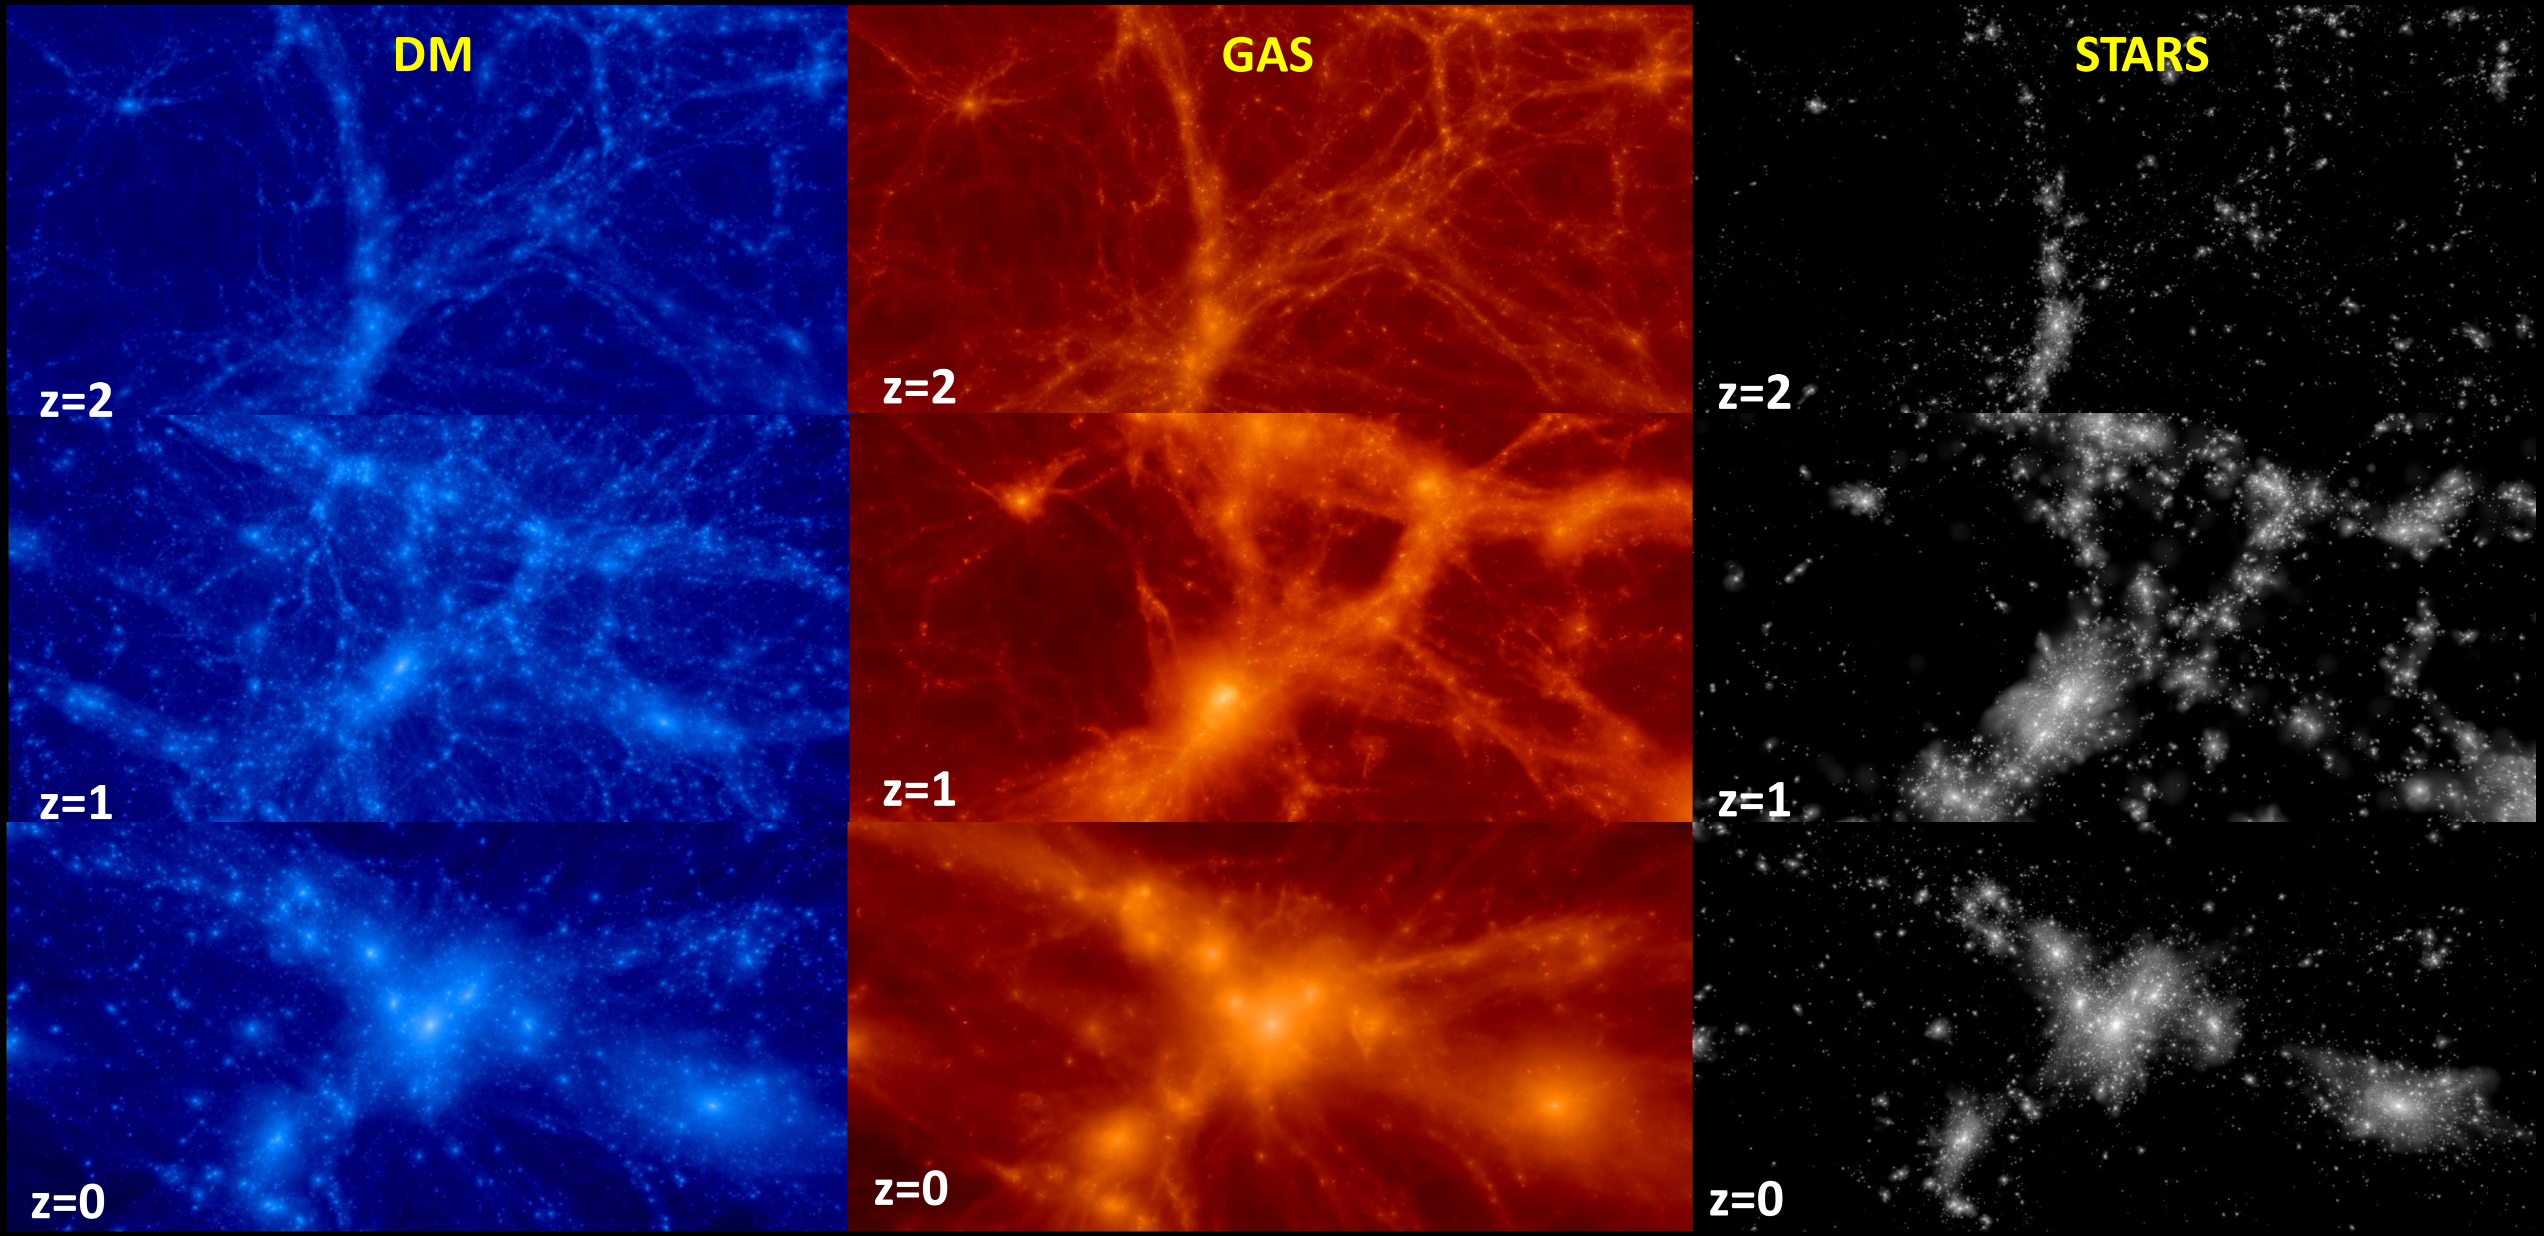

Star formation in proto-cluster regions

(Esposito et al. 2025) ***– INAF PhD Fellowship***

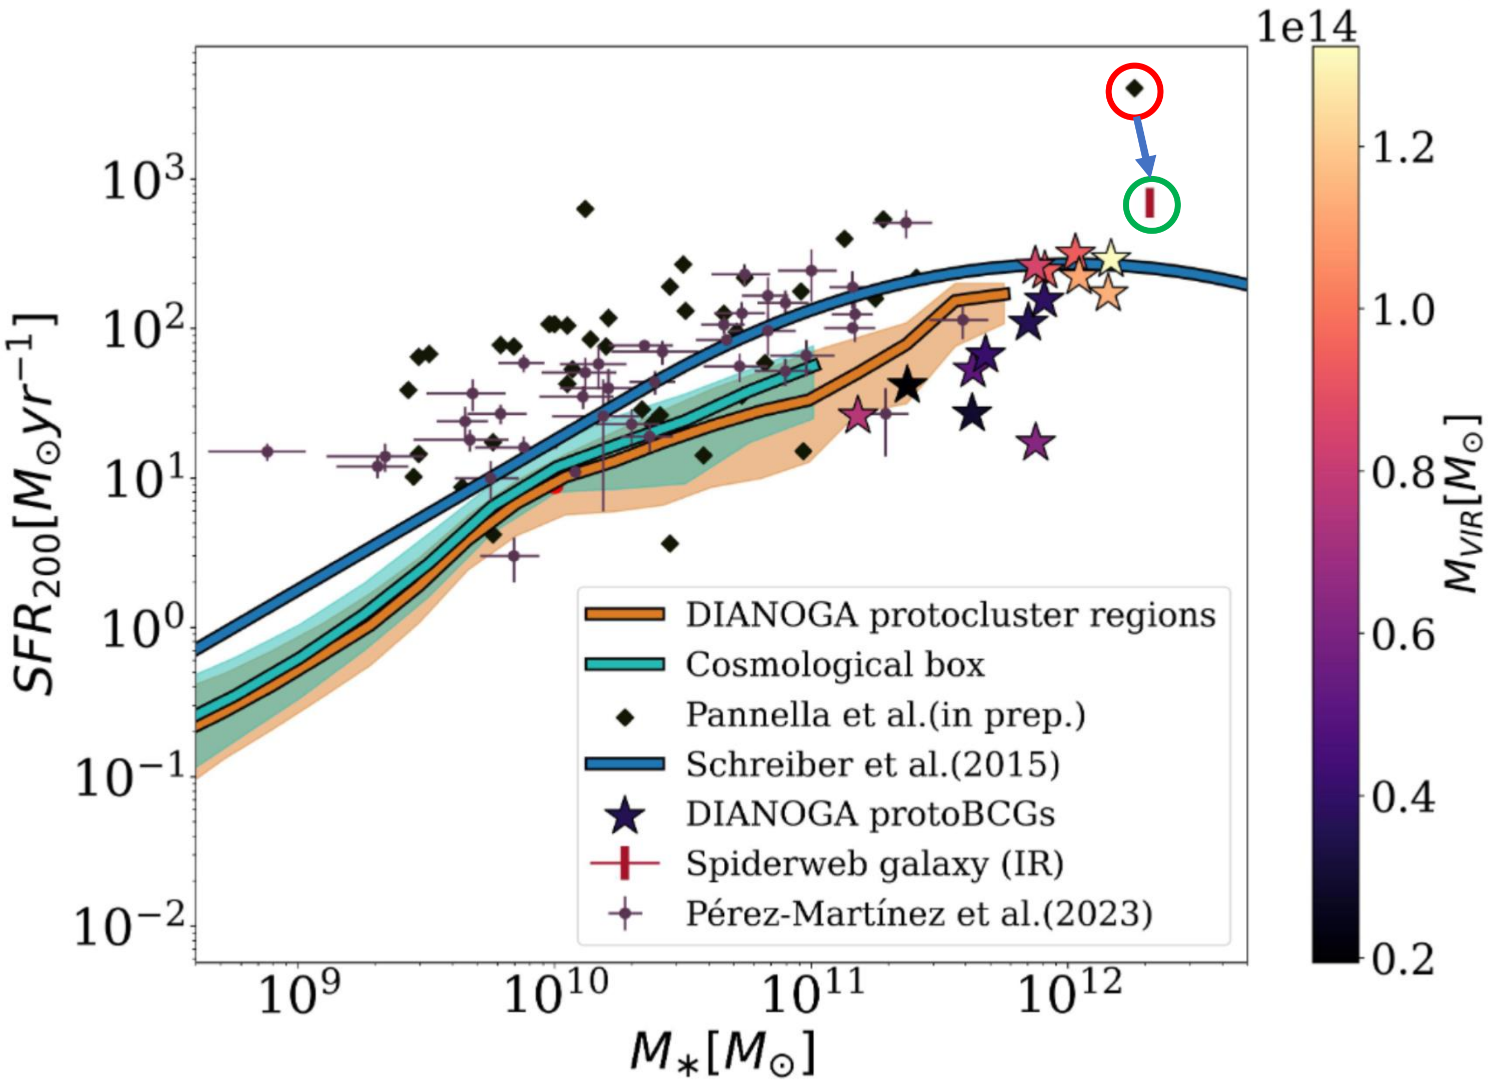

→ Model-prediction of the main sequence at z=2.2 below the observed one, both in the field and in protocluster

→ Result almost independent of the adopted model of SF

→ SFR of the Spiderweb much reduced when including FIR data (Seymour+2012; Drouart+2014), besides UV dust-corrected fluxes *(Pannella+ 2026, in prep)*

→ ''Only'' a factor 2-3 above simulation predictions

Unresolved dynamical friction & BH orbits

Alternative ad-hoc prescription:

Large dynamical mass: enhance by hand the BH dynamical mass at seeding to amplify the resolved DF
➔ *Significant change in the local potential*

Continuous repositioning: at every time-step pin the BH on the local minumum of the potential
➔ *Merging time-scales completely wrong*

➔ Significant implications for the MBH-MBH merger events (as GW sources)

Damiano+2024, 2025a,b ***- ICSC PhD Fellowship***

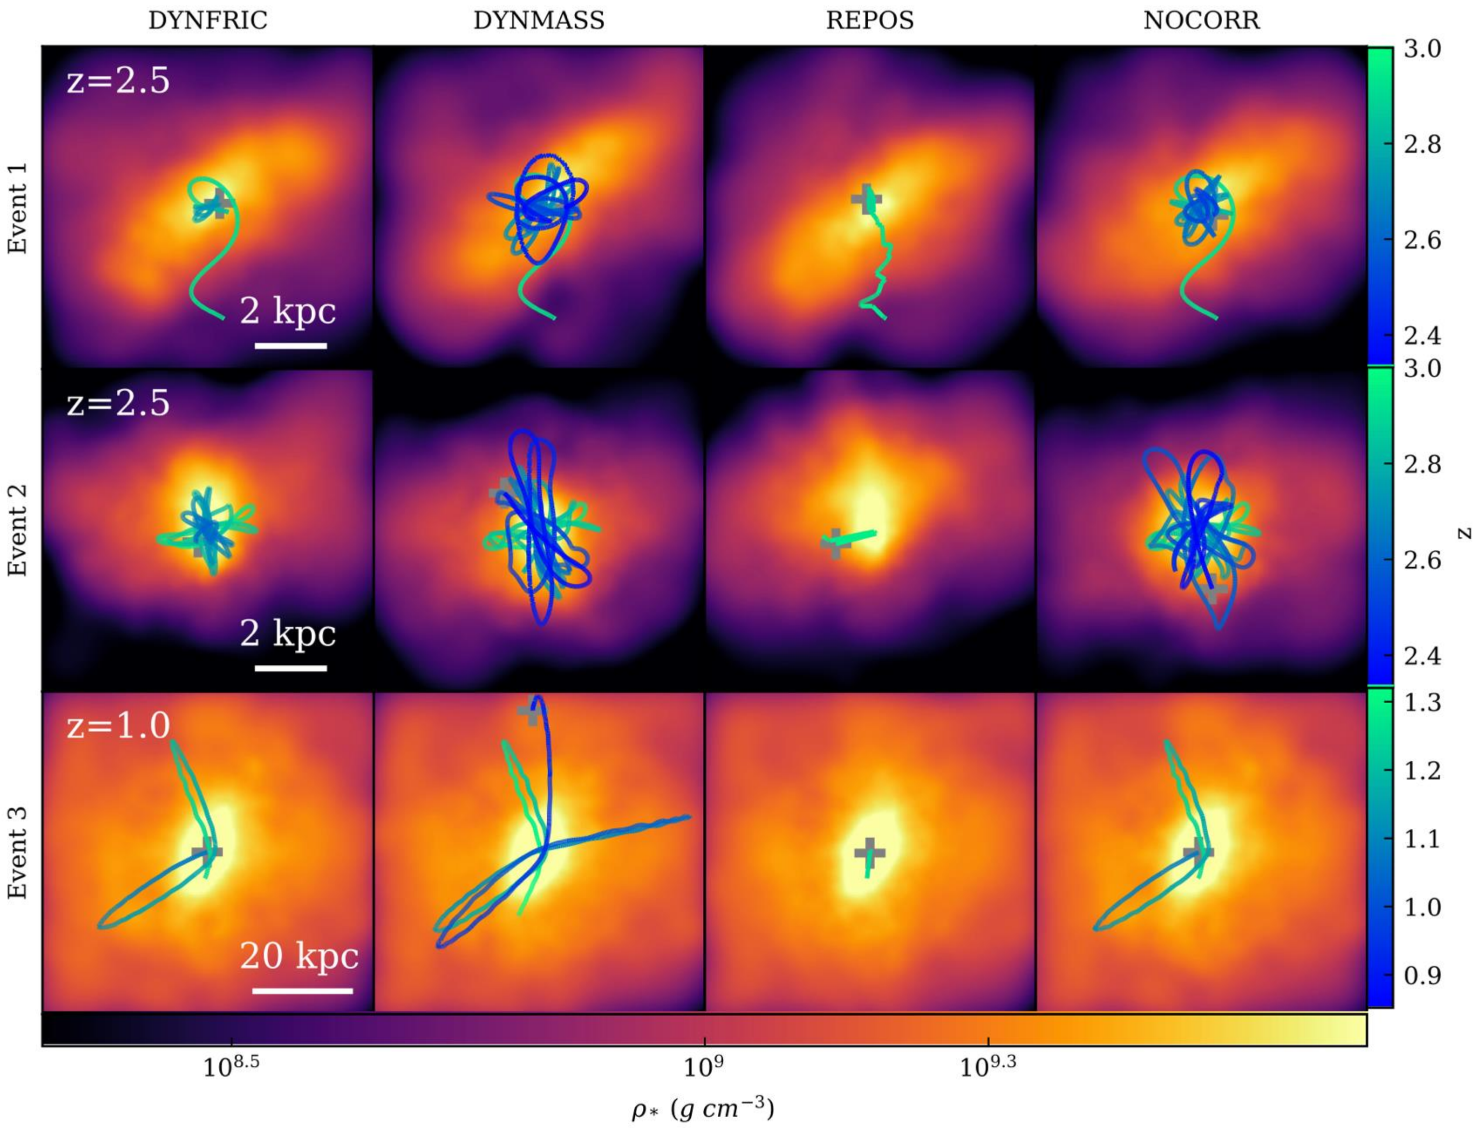

N-BODY - Euclid: A 1% precise HMF

(Euclid: Castro+2023; **PRIN-PNRR TD at INAF – PI: UniTS**)

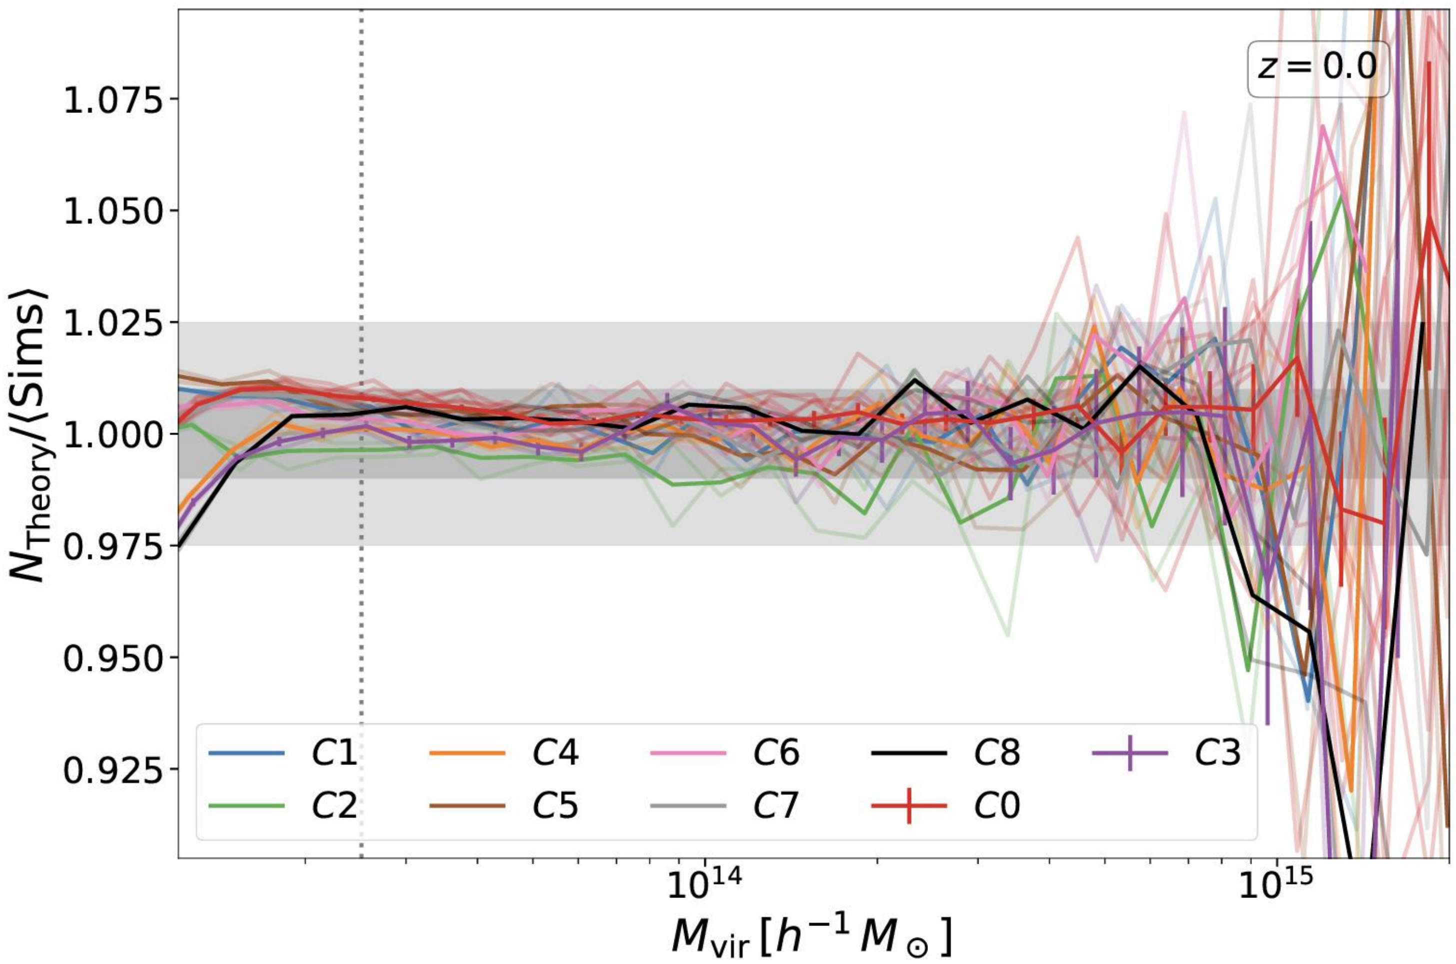

PICCOLO - PathfInder Cluster COsmoLOgy:

Code: OpenGadget 3/monofonic

Name	Om0	h	Ob0	ns	sigma_8
C0	0.3158	0.6732	0.0494	0.9661	0.8102
C1	0.1986	0.7267	0.0389	0.9775	0.8590
C2	0.1665	0.7066	0.0417	0.9461	0.8341
C3	0.3750	0.6177	0.0625	0.9778	0.7136
C4	0.3673	0.6353	0.0519	0.9998	0.7121
C5	0.1908	0.6507	0.0527	0.9908	0.8971
C6	0.2401	0.8087	0.0357	0.9475	0.8036
C7	0.3020	0.5514	0.0674	0.9545	0.8163
C8	0.4093	0.7080	0.0446	0.9791	0.7253

- L=2 h^{-1}Gpc ; $N_p=2048^3$
- 10 realizations for C0, 2 for the other models

HYDRO - Euclid: Impact of baryons on the HMF

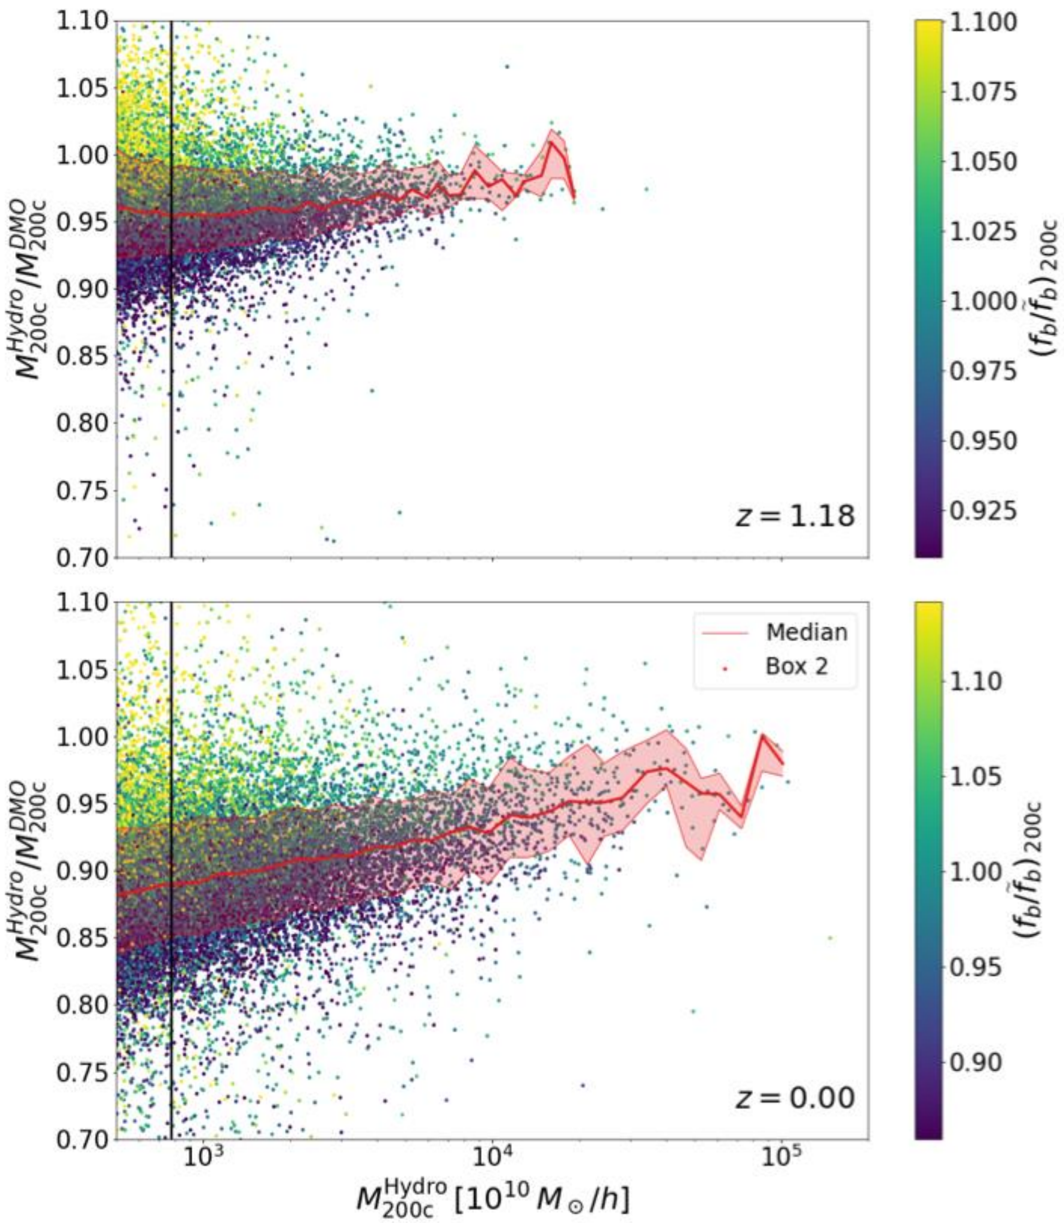

Q: How do we know the baryonic correction to be applied to observed cluster masses?

Castro+2021; Castro+2024

Based on the *Magneticum Simulations* (Dolag+2025)

→ 5-10% reduction of halo masses induced by the effect of baryons (see also Cui+2015; Velliscig+2016; Bouquet+2016; ….)

→ Effect slightly larger at lower redshift

→ Mass reduction inversely correlated with the baryon fraction

Efficient GW transient detections (LVK Coll. - Spoke-2 at DF)

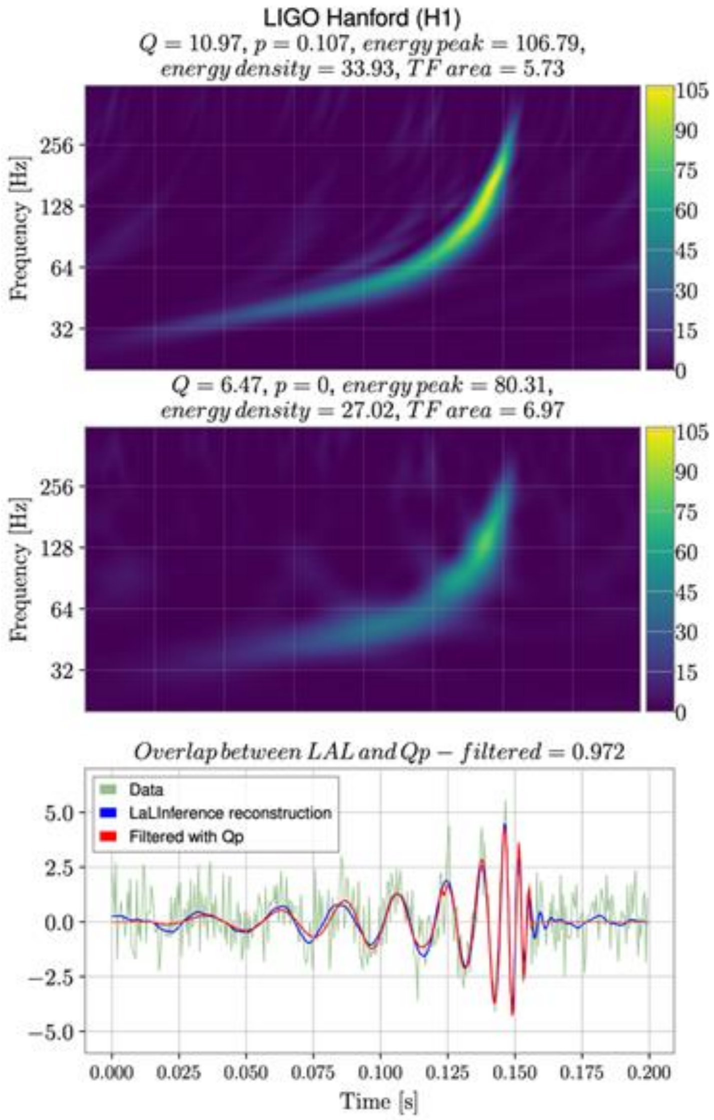

Performance of the Qp transform and of the standard Q transform on GW150914 data, plus comparison between a PE reconstruction and the Qp reconstruction.

Unmodeled methods like coherent WaveBurst (cWB) are used to detect **generic gravitational-wave transients** (GWT), but they can also be applied to well-modeled GW signals such as those originating from compact binary coalescences (CBC). Overall, our project aimed at improving the efficiency of cWB.

What we achieved:

- In-depth study of the wavelet-based part of cWB. We were able to greatly improve the performance of other methods of continuous wavelet inversion (see this paper and the related code). In the future, we plan to use this method to replace the current WDM-wavelet-based code in cWB.
- We contributed to the development of the PycWB code, a Python+C++ version of the cWB C++ code which is modular and easily extensible. While the new code already achieves a better performance with respect to the older code, we are still working on GPU extensions of the code. The PycWB code is publicly available on this github repository.
- Exploratory results on the use of symbolic regression applied to cWB data. In this part of the project we aim at performing symbolic regression on cWB waveform reconstructions of CBCs to obtain the relevant CBC parameters.

Courtesy: E. Milotti

Materials simulations (Spoke 7 at DF)

using ⇨ **ab-initio quantum mechanical calculations**
(very accurate, for systems of limited size)

⇩ **molecular dynamics with classical empirical potentials**
(large scale, but not very accurate or transferable)

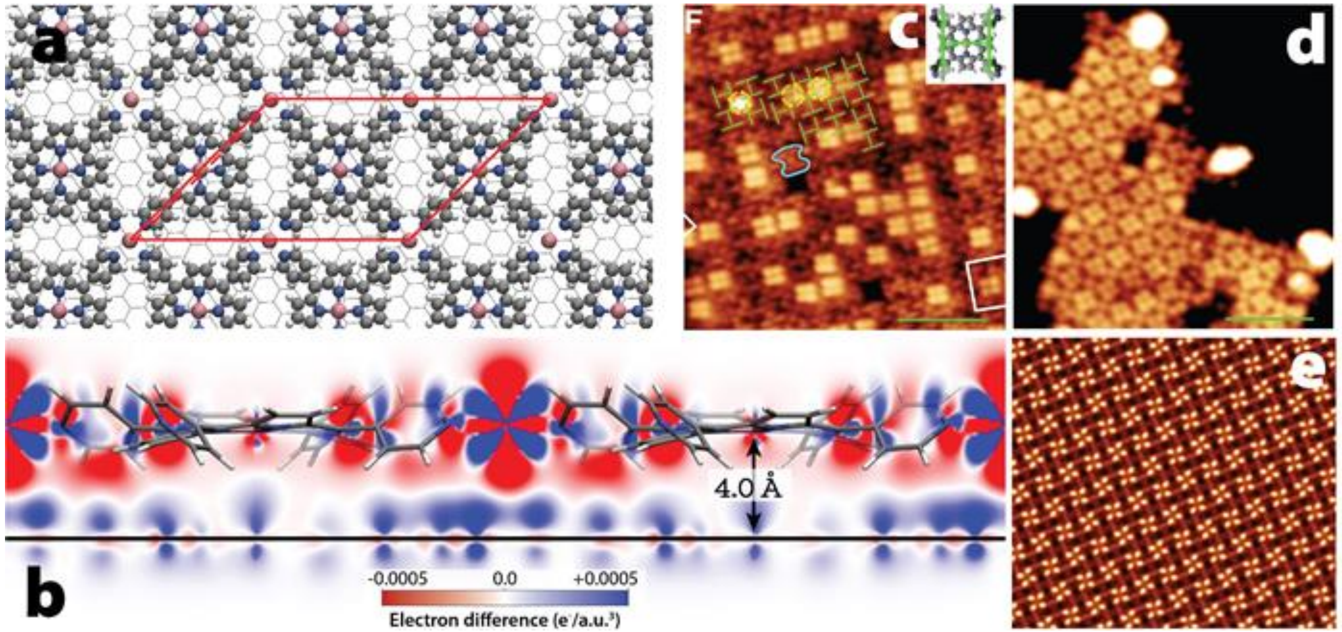

Example: Nitrogen molecule impinging on (pristine or defected) graphene-coated alumina.

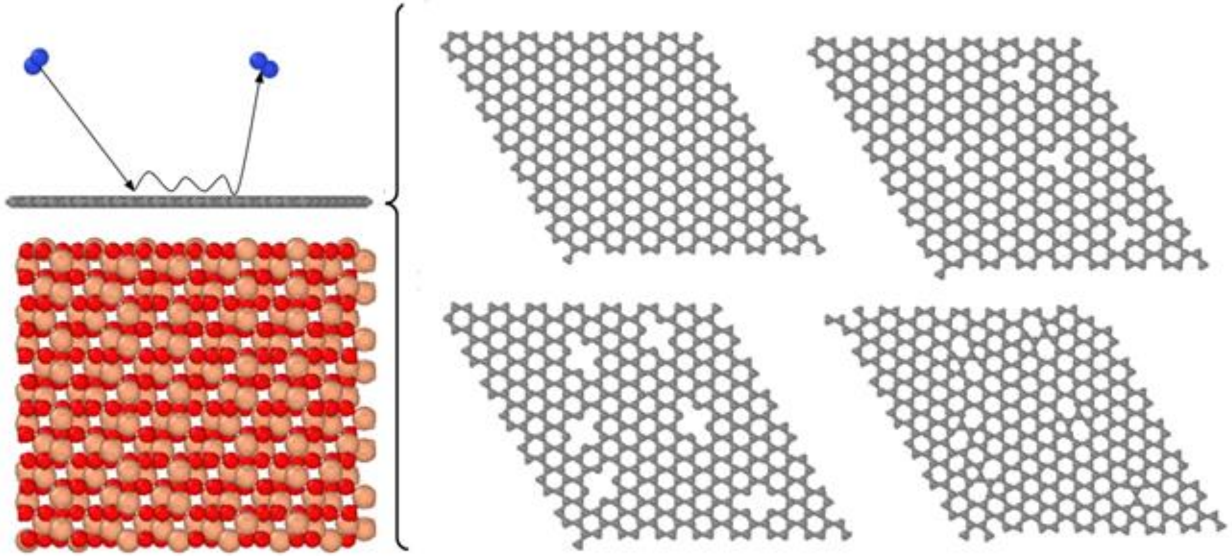

Example: Molecular layer of cobalt tetra-pyridyl porphyrins on graphene: a) optimized model, b) electron density rearrangement due to graphene-molecular layer interaction c,d) experimental and e) simulated scanning tunnelling microscopy images.

simulations performed on the CINECA Leonardo supercomputer

Courtesy: M. Peressi

Theoretical & Computational Chemistry (Spoke-7 at DSCF)

Mauro Stener, Giovanna Fronzoni, Daniele Toffoli, Emanuele Coccia, Lucia Gardossi, Emanuele Carosati

- Simulation of properties (including optical) of molecular, nanoscale, and condensed-phase materials.
- Development of chemical computational codes (DFT, TDDFT, ab-initio, and QMC).
- New algorithms for parallel architectures.
- Ultrafast electronic processes in molecules and complex systems.
- Molecular nanoplasmonics.
- Developers of ADF-AMS (https://www.scm.com/) and WaveT (https://github.com/stefano-corni/WaveT_TDPlas) codes

WaveT: strong scaling up to 16K CPU

Time-resolved dynamics

Nanoplasmonics

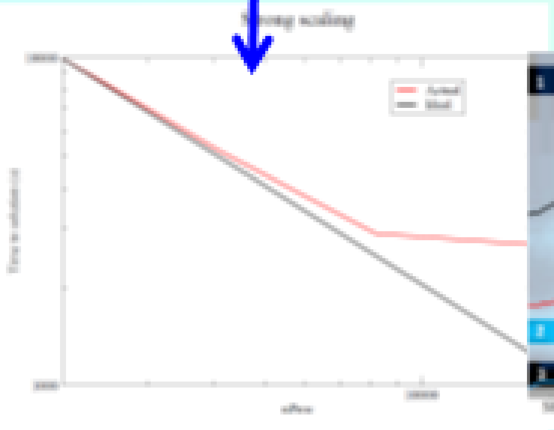

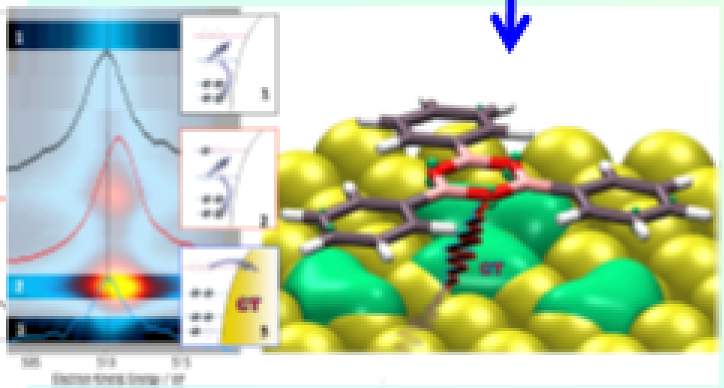

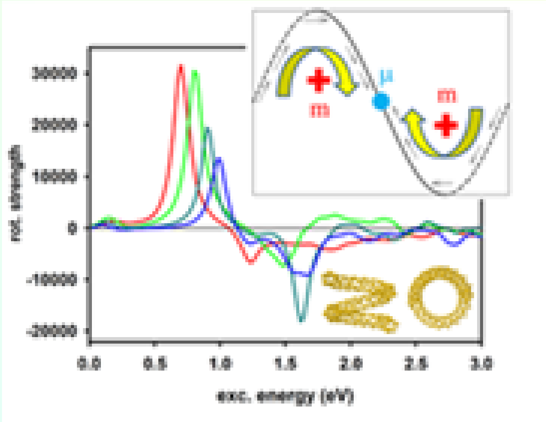

Courtesy: M. Stener

Molecular Biology and Nanotechnology Lab (Spoke-7 at DIA)

Sabrina Pricl, Erik Laurini

- Multiscale modelling for managing the complexity of process
- Transport and thermodynamic property
- Carbon capture and sequestration, hydrogen production, decarbonization of steel industry

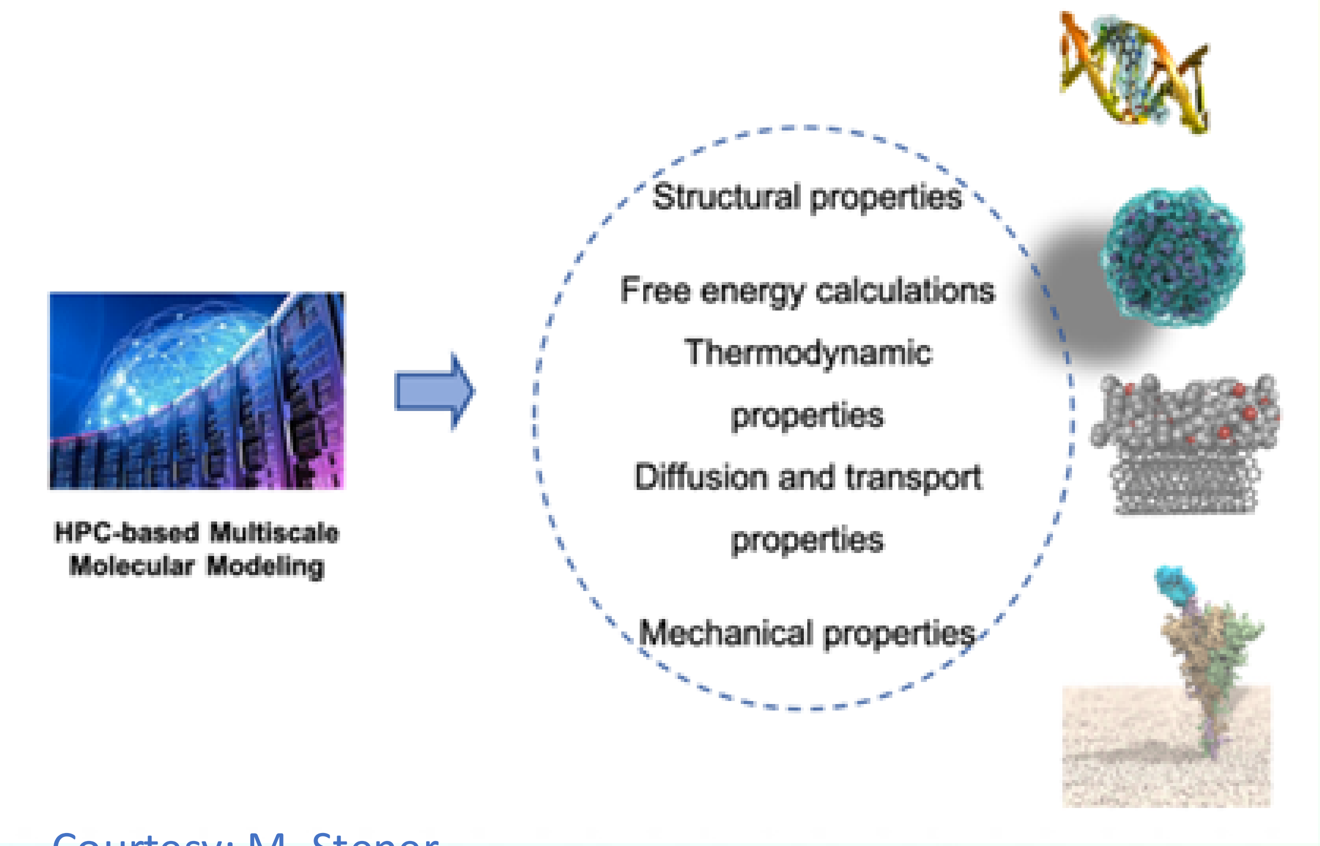

Courtesy: M. Stener

Bora HPC cluster @ DF-UniTS

Courtesy: D. Coslovich

Number crunching @DF since 2024...

Contacts:

D. Coslovich, M. Costanzi, G. Principe

Bora is the **HPC cluster** of the **Department of Physics** (DF) at the University of Trieste

- Funded by the **Department of Excellence** grant of the MUR (2023-2027)
- Hybrid and extensible computing infrastructure:
  - **CPU**, **FAT** and **GPU** nodes
  - 216 cores, 2.3TB RAM, 100 Tb HD, 2 x A30 + upcoming extension
- Suitable for
  - **development/test** of software before large-scale deployment
  - **familiarizing students** with an HPC environment
  - **small-scale projects** by students and young researchers @DF

AI and HPC education @ UniTS

Courtesy: L. Bortolussi

BACHELOR PROGRAM

ARTIFICIAL INTELLIGENCE & DATA ANALYTICS

~100 freshers/yr

MASTER PROGRAM

DATA SCIENCE & ARTIFICIAL INTELLIGENCE

~50 freshers/yr

PHD PROGRAM

APPLIED DATA SCIENCE & ARTIFICIAL INTELLIGENCE

~15 freshers/yr

AI and data science: foundations, generative AI, scientific applications.
HPC (Area Science Park and INAF lecturers): CPU, GPU, cloud, data management, FAIR data
Applications in science and industry.

Academic Partners

ICTP · AREA SCIENCE PARK · INAF ISTITUTO NAZIONALE DI ASTROFISICA

Programmatics & take-home messages

DF-UniTS & INAF-OATs: an established tradition of collaboration in *Computational Astrophysics*

- Solid scientific collaborations
- Coordinated participation to funding projects: EU (e.g. Euro-HPC) and national (ICSC, PRIN)
- PhD-level training: co-supervision and co-funding of PhD fellowships
- Training at undergrad level: co-supervision of student, teaching, mentoring, tutoring

UniTS & INAF: A virtuous example of mutual benefit *(not only in HPC & Big Data)*

- Access to computing infrastructures and data-storage facilities
- Training of a next generation scientists and technologists expert in HPC & Big Data
  → not optional for the future of astrophysics and for a high-profile role of INAF in one of the most strategic directions of societal and economic development
- Possibility to collaborate on the definition of strategies and priorities for education and training