Unveiling the sonic scale with Smoothed Particle Hydrodynamics

EUROHPC USER DAY 2023 Brussels 11.12.23

EuroHPC Joint Undertaking

Project: “*TGSF: The Role of Turbulence and Gravity in Star Formation*”

EuroHPC used: LUMI-G

Speaker: *Rubén CABEZÓN (sciCORE – Univ.*

- What do we want to do in TGSF?
- How are we going do it?
- What is our current state?
- Which are the expected outcomes?

TGSF: first large-scale simulation of turbulence + gravity

Study the role of turbulence **and** gravity in the process of forming pre-stellar cores.

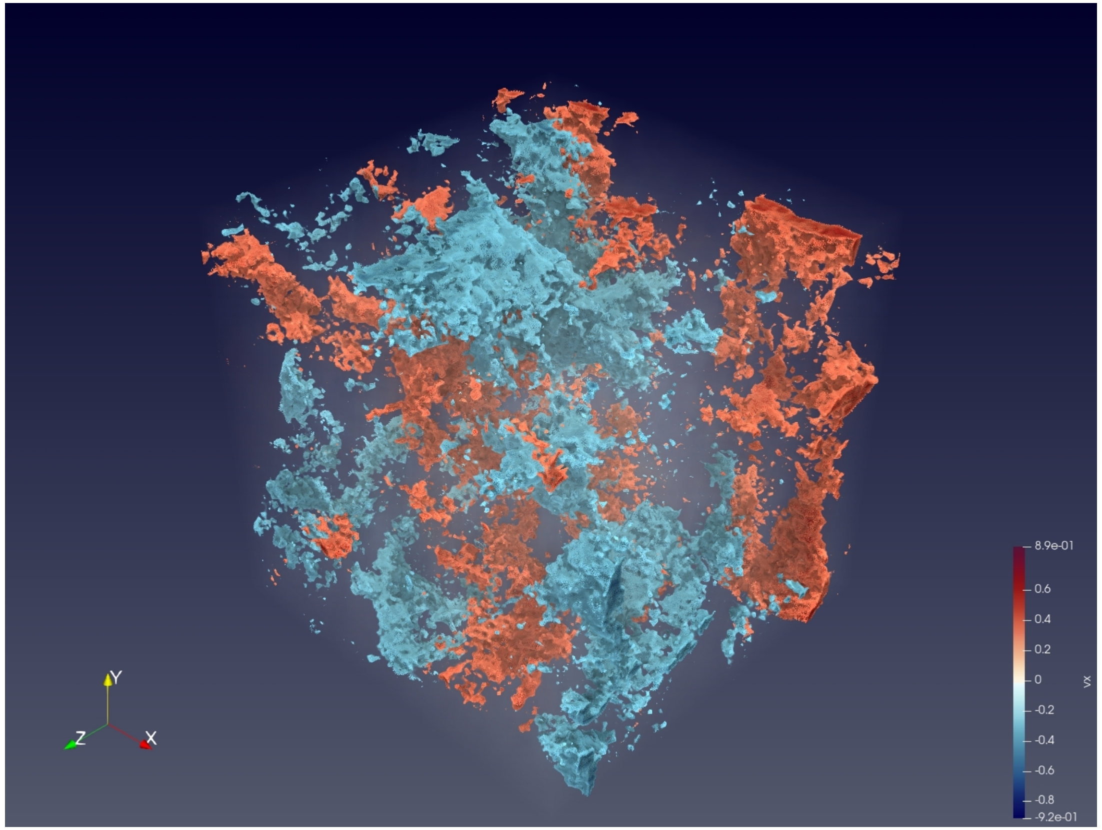

Subsonic turbulence. 1000^3 volumetric rendering of the X component of the velocity field. (SPH-EXA team)

+

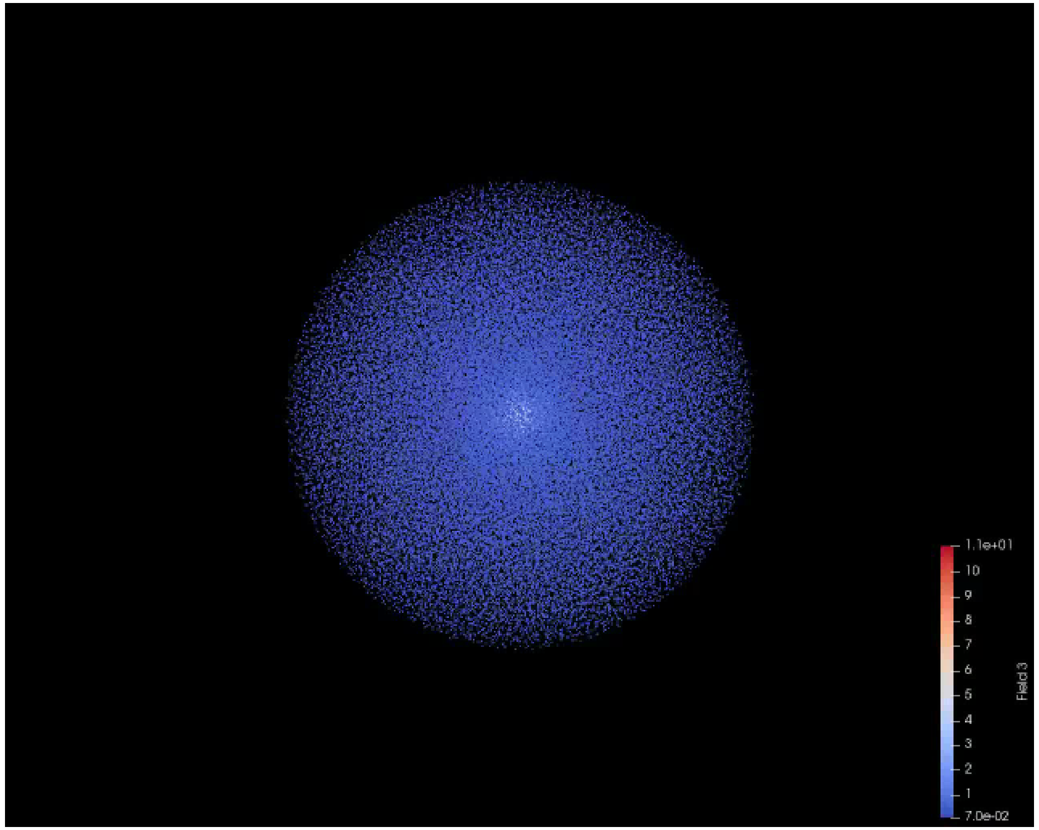

Gravitational collapse of an isothermal cloud. Evrard test. (SPH-EXA team)

Definition

Sonic scale (l_s) : is the scale at which the transition from supersonic to subsonic turbulence occurs.

$$l_s = \phi_s L(\mathcal{M})^{-2}$$

ϕ_s encompases our lack of knowledge about the exact position of the sonic scale. Usually taken as $\phi_s = 1$

Nevertheless, a large-scale simulation (Federrath et al. 2021) has directly measured ϕ_s to be x2.4 smaller.

This pushes the collapse scale to smaller scales than previously considered and it has a critical relevance in the predictive power of star formation theories.

In order to test this, **self-gravity** must be included in such simulation.

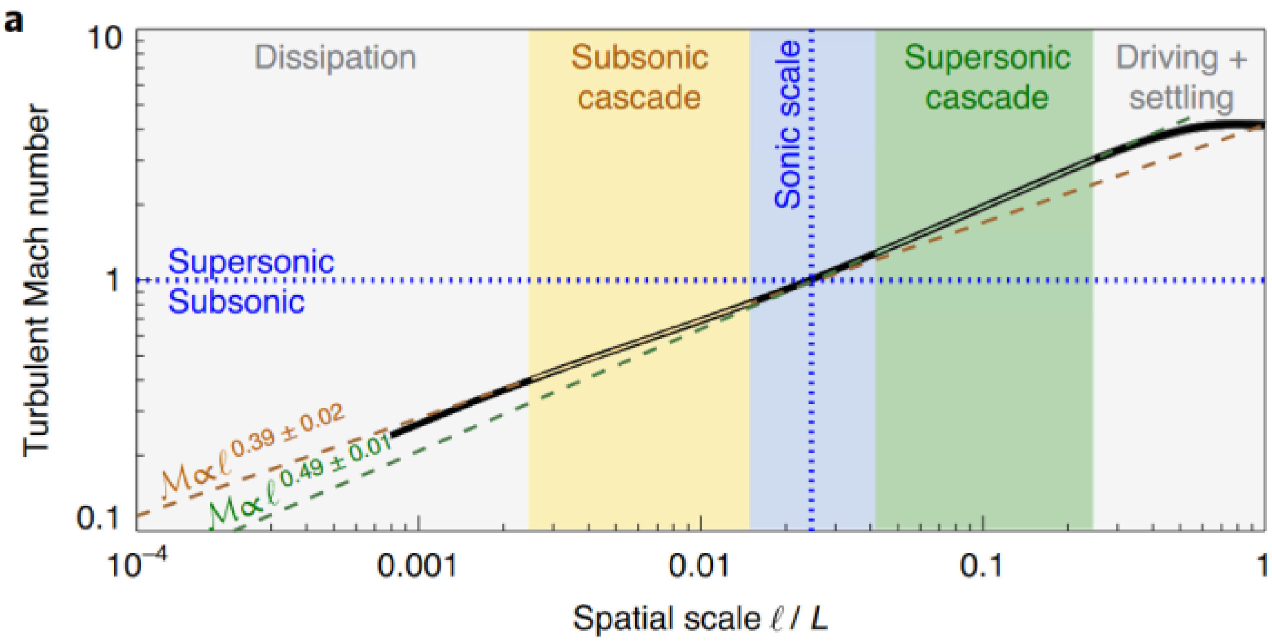

(Federrath et al., Nature Astronomy, 5, 2021)

TGSF: first large-scale simulation of turbulence + gravity

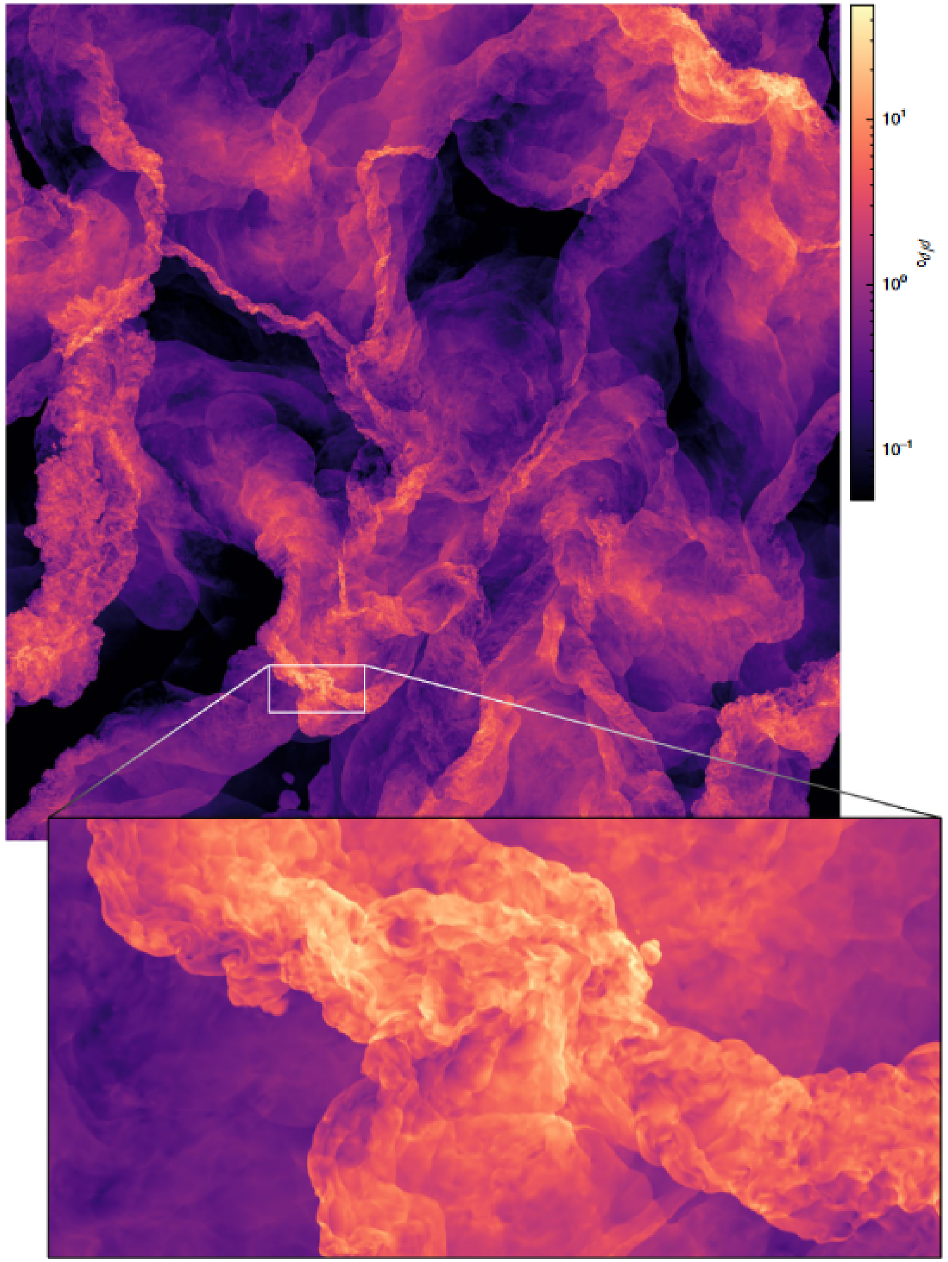

FLASH
Eulerian code
$10{,}048^3$ grid cells
Hydrodynamics only (no gravity)
CPU only (65,536 cores)

Gas density contrast distribution of ISM turbulence. (Federrath et al., Nature Astronomy, 5, 2021)

SPH-EXA
Lagrangian code
$10{,}079^3$ SPH particles
Hydrodynamics + gravity
CPU + GPU (16,416 GPUs)

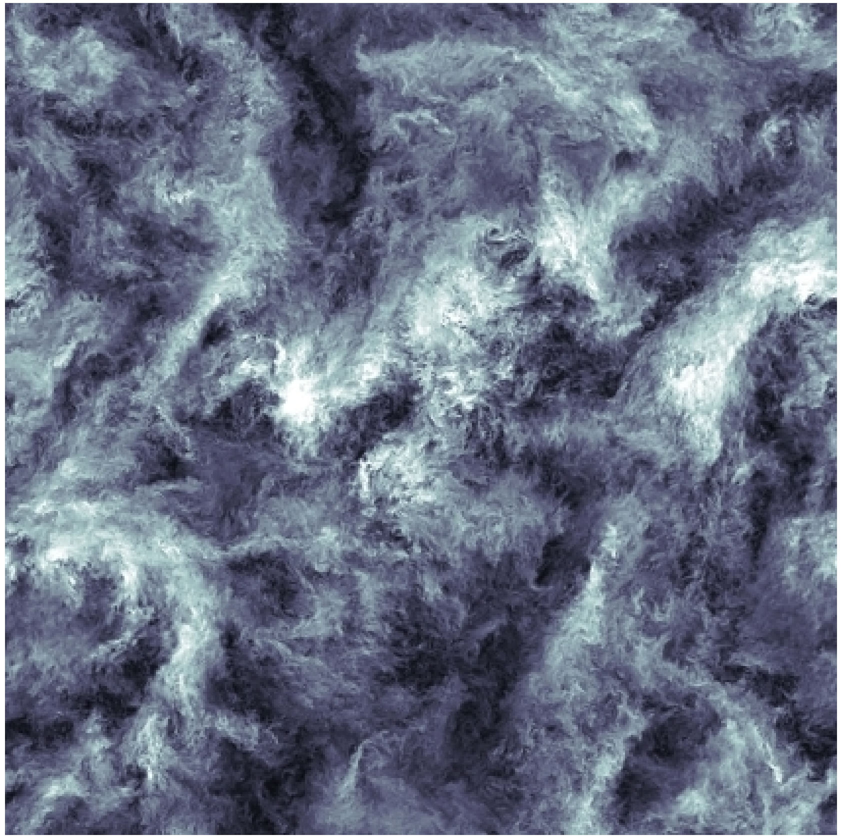

Velocity field distribution of subsonic turbulence. 3000^3 particles (SPH-EXA team, 2023)

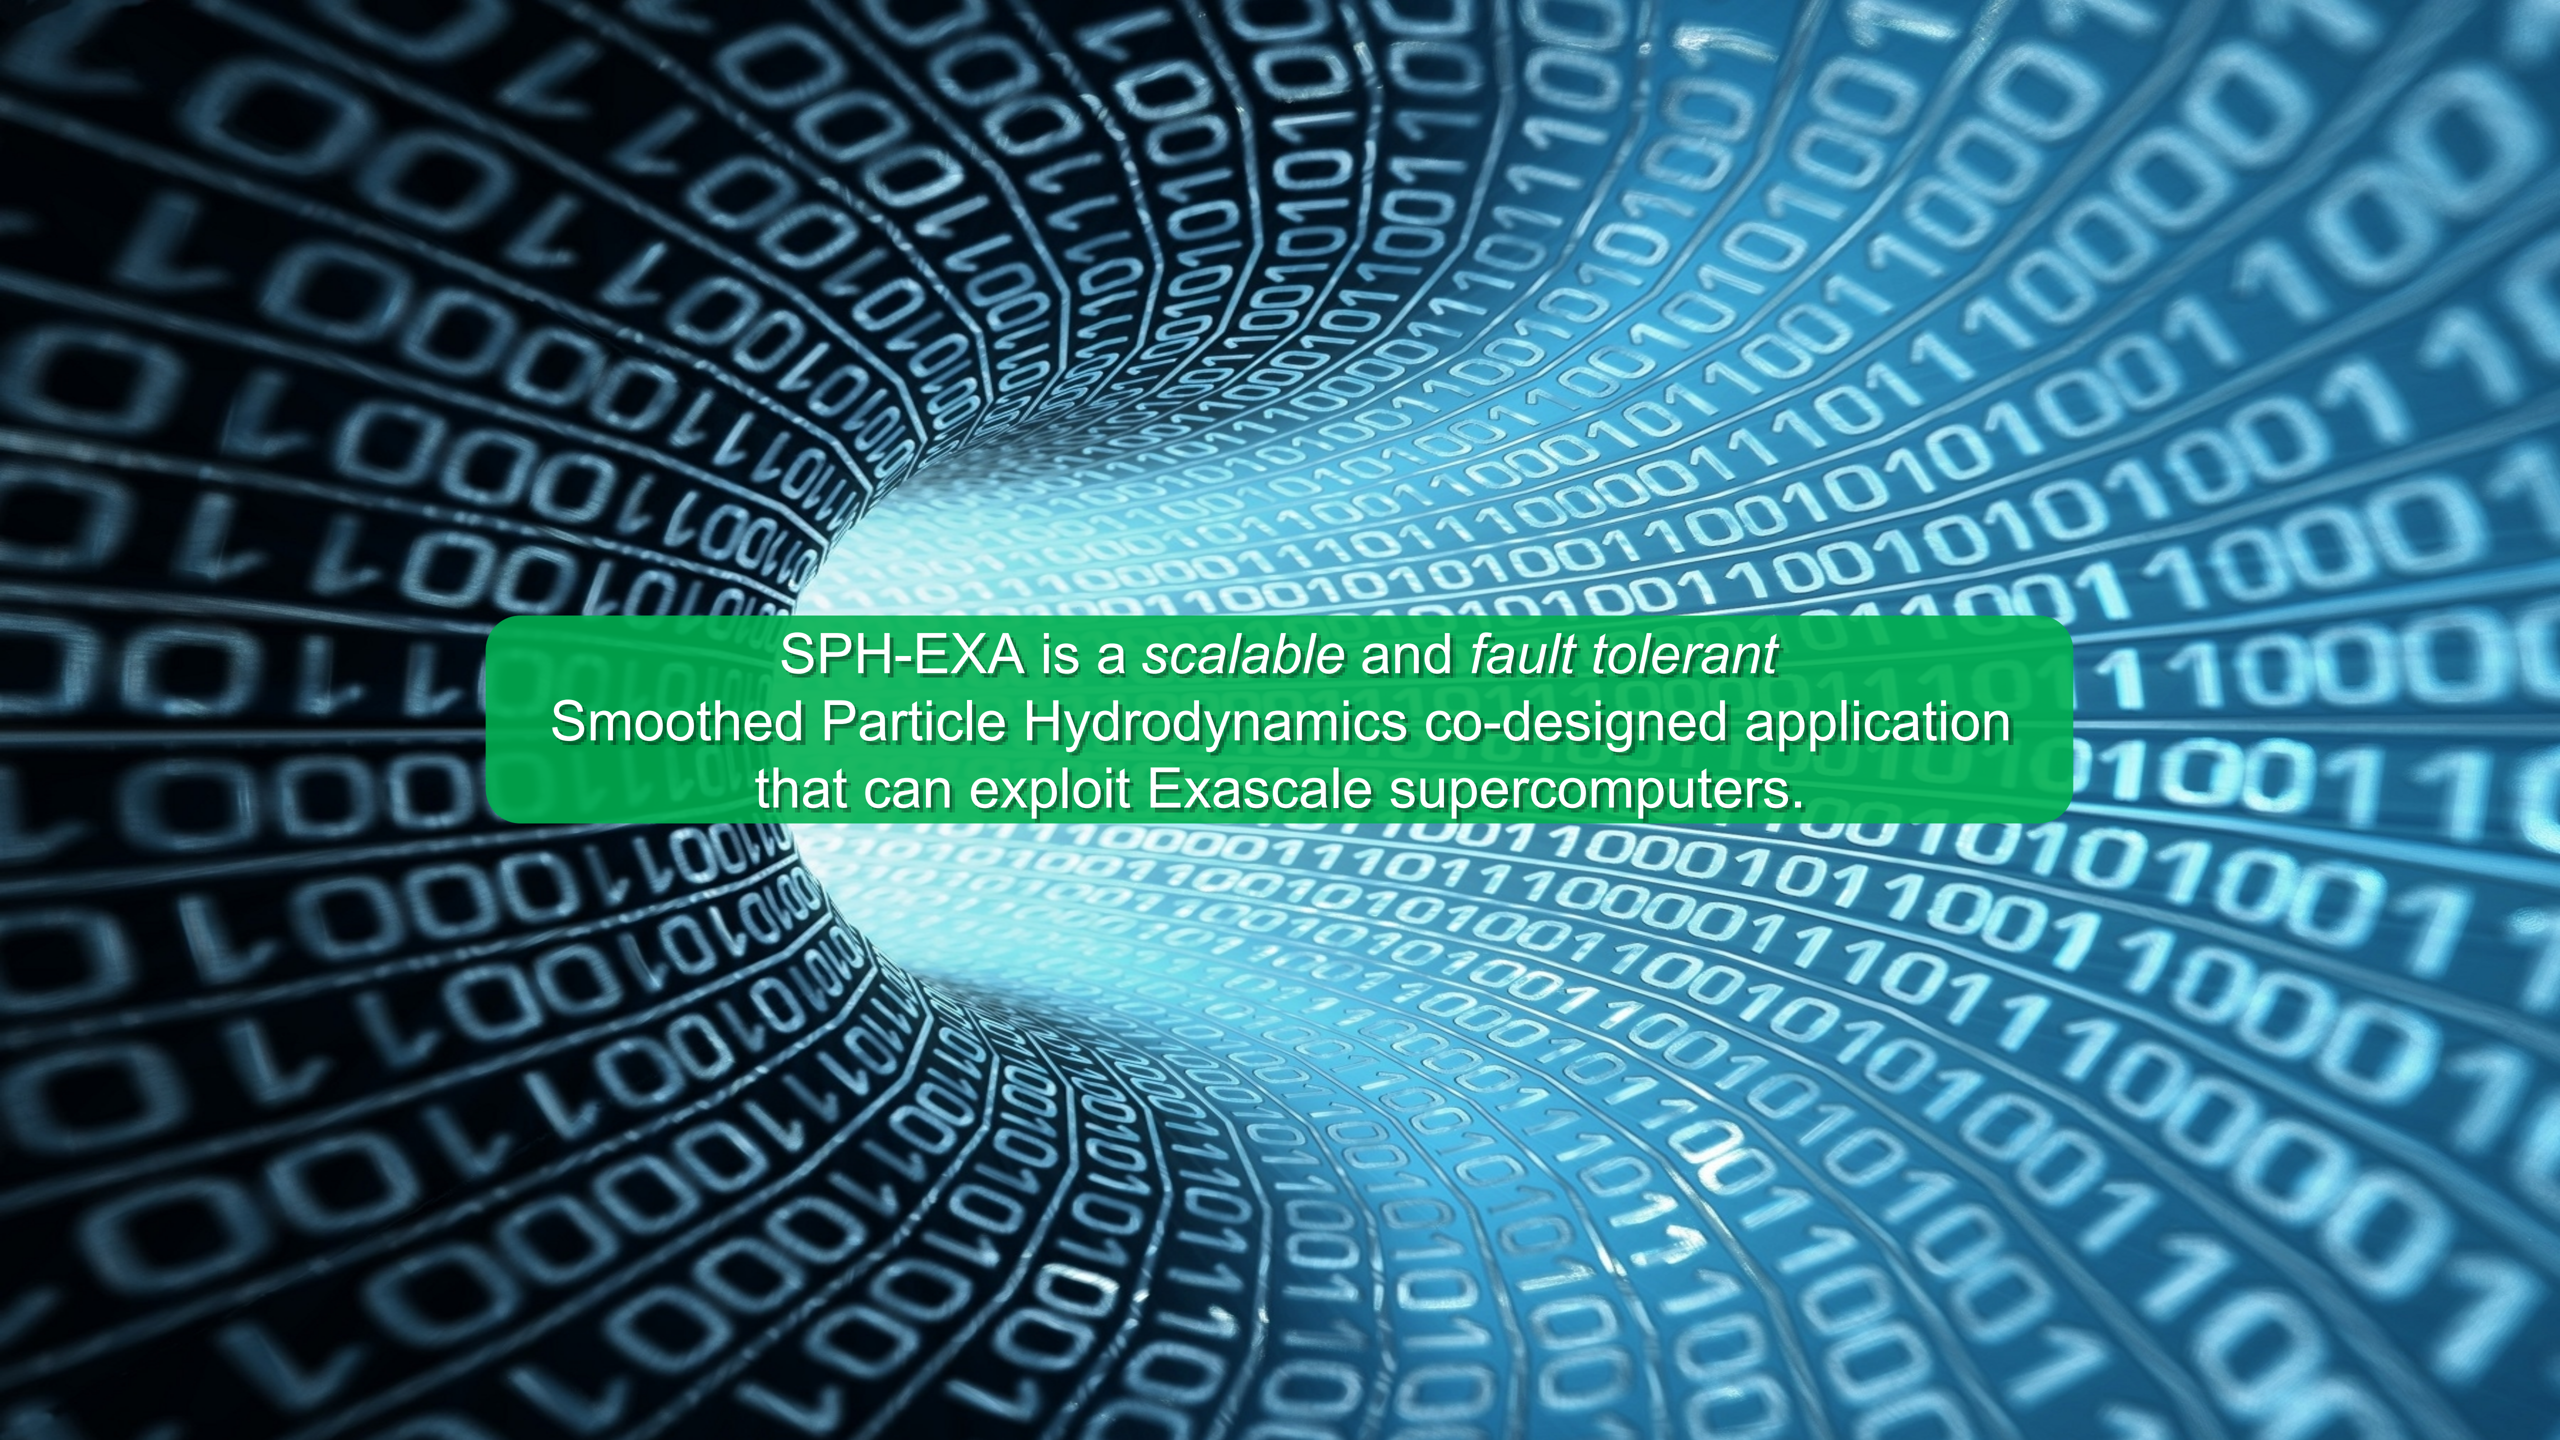

SPH-EXA is a *scalable* and *fault tolerant*
Smoothed Particle Hydrodynamics co-designed application
that can exploit Exascale supercomputers.

SPH-EXA framework component overwiew

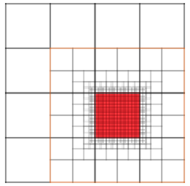

Cornerstone octree

Domain Decomposition

- Space-filling curves and octrees
- Global and locally essential octrees
- Octree-based domain decomposition
- 21’600 lines of code

ChaNGa

SPH Solver
Physics modules

Modern SPH and physics implementation with key features (astro.physik.unibas.ch/sphynx, github.com/N-BodyShop/changa):

- Generalized volume elements
- Integral approach to derivatives
- Artificial viscosity with switches
- Sub-grid physics
- 3’800 lines of code

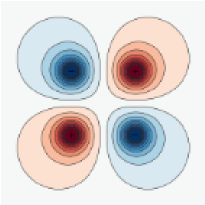

Ryoanji N-body solver

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)
- 4’100 lines of code

SPH-EXA

SPH-EXA application front-end

- Initial conditions generation, checkpointing, parallel I/O, compression
- Flexible combination and addition of additional physics for domain scientists
- Performance data and energy consumption measurements
- In-situ visualization
- 7’200 lines of code

Using SPH-EXA

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

C++ 20 (GCC 11+)

Cmake 3.22+

CUDA 11.2+

HDF5 1.10+

OpenMP

MPI

AMD ROCm

```
$> git clone https://github.com/unibas-dmi-hpc/SPH-EXA.git
$> cd SPH-EXA
$SPH-EXA> mkdir build
$SPH-EXA> cd build
$SPH-EXA/build> cmake ..
.
. Output
.
$SPH-EXA/build> make -j
.
. Output
.
$SPH-EXA/build> cd main/src/sphexa
$SPH-EXA/build/main/src/sphexa> ls
sphexa
sphexa-cuda
$SPH-EXA/build/main/src/sphexa>
```

OPTIONAL [Ascent ParaView Catalyst ADIOS2]

SPH-EXA: a next-generation SPH code for astrophysics, cosmology, and fluid dynamics

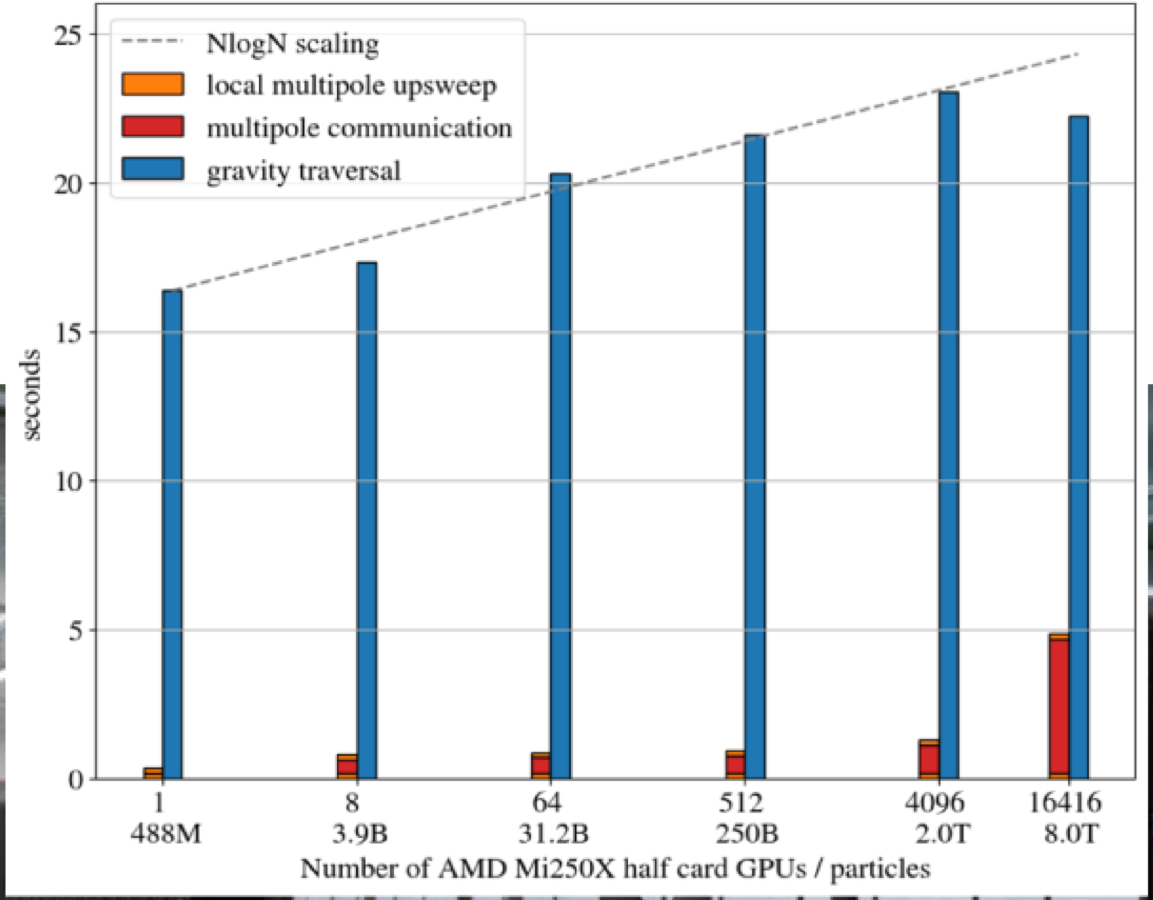

Close to logarithmic (Nlog(N)) weak scalability of the gravity solver in SPH-EXA on LUMI -G.
(Keller et al. PASC'23 Proceedings, 18, 2023)

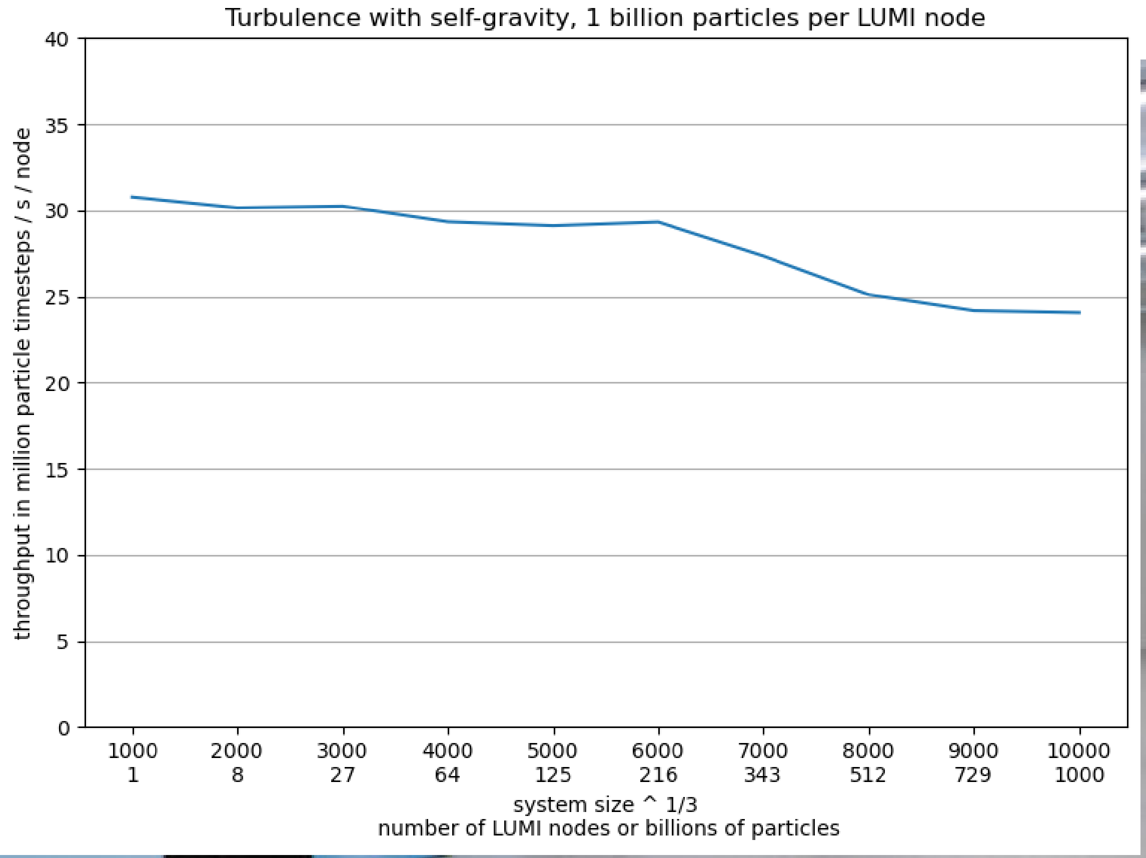

Weak scaling of SPH-EXA on LUMI-G, running 8 MPI ranks/node (1 MPI rank/ GPU half-card) and 1 billion particles/node (SPH-EXA team, 2023)

TGSF: the plan

Ornstein-Uhlenbeck stirring

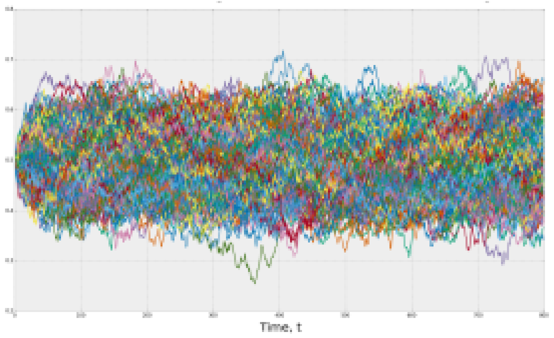

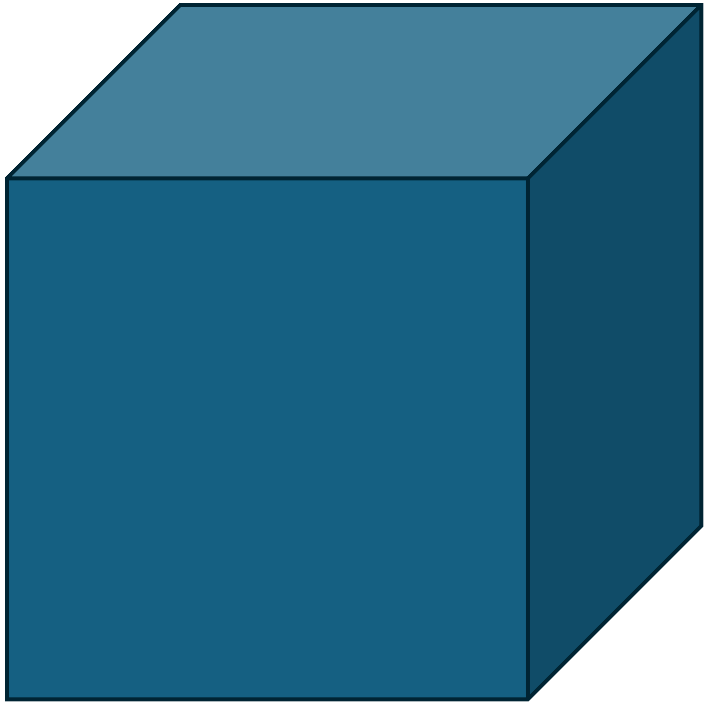

Homogeneous box with PBC

Relaxation to achieve stable Mach number
T ~ 10 sound-crossing times

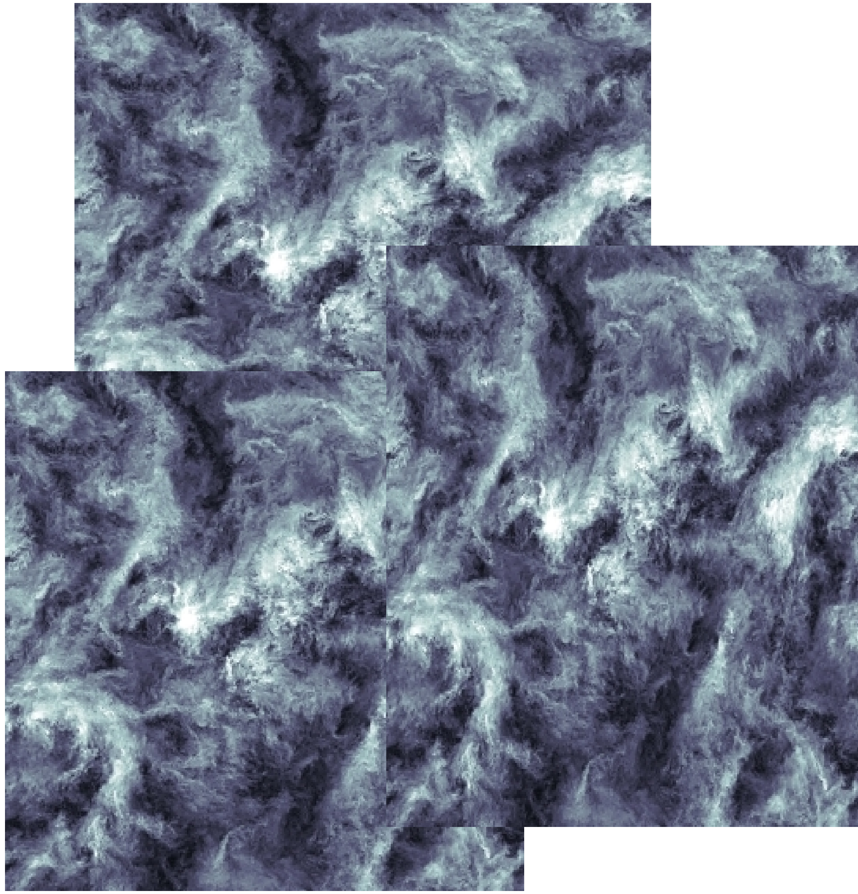

Turbulent box with PBC

For $10{,}000^3$ particles, relaxing the ICs already consumes 1.4 million node hours in LUMI-G

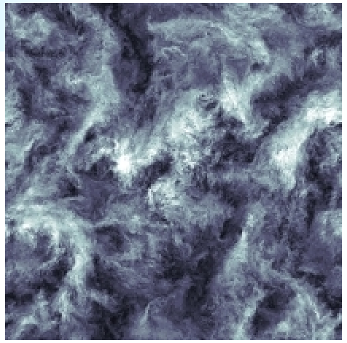

Simulations set

$1{,}000^3$
$1{,}260^3$
$1{,}587^3$
$2{,}000^3$
$2{,}520^3$ → LR → + self-gravity
$3{,}175^3$
$4{,}000^3$
$5{,}040^3$ → MR → + self-gravity
$6{,}350^3$
$8{,}000^3$
$10{,}079^3$ → HR → + self-gravity

$1{,}000^3$

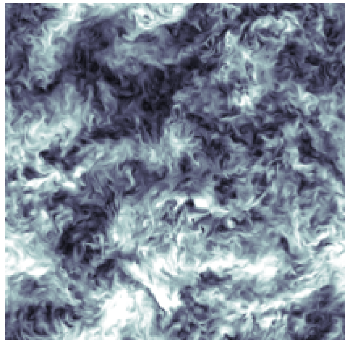

Relaxation
$t \sim 10t_{sc}$

Particle splitting
(x2 particles)

$1{,}260^3$

Relaxation
$t \sim 2t_{sc}$

Particle splitting
(x2 particles)

$1{,}587^3$

…

Particle splitting
(x2 particles)

$8{,}000^3$

Relaxation
$t \sim 2t_{sc}$

Particle splitting
(x2 particles)

$10{,}079^3$

This process to generate the ICs consumes **x4 fewer NHs** than the direct approach.
It also enables **simultaneous simulations**.

Extreme Scale Access

Topic: The role of turbulence and gravity in star formation
Duration: 1 yr
Allocation: 22,000,000 GPUh in LUMI-G

Study the formation of stellar cores and their initial mass function at unprecedented resolution

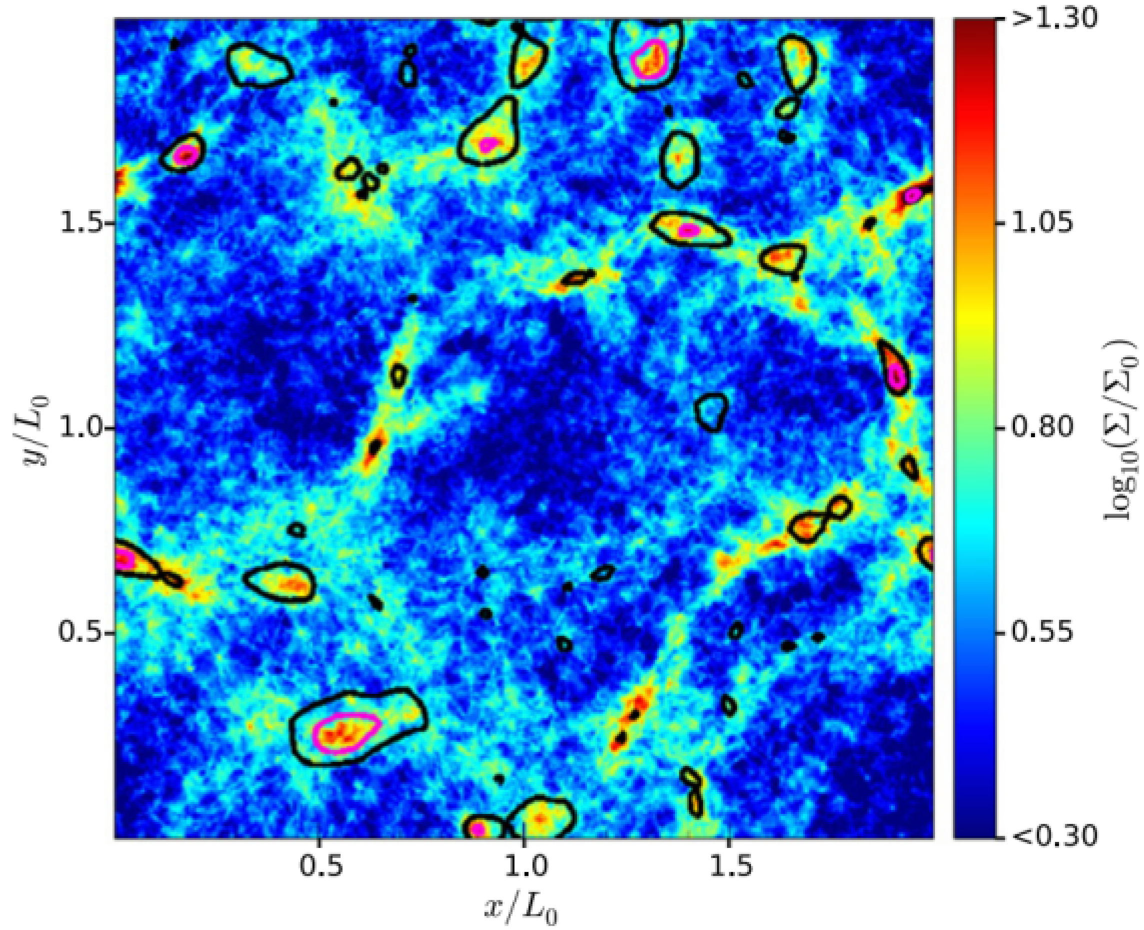

Collapsing cores in a supersonic turbulent medium, with the ATHENA code. (Gong and Ostriker, ApJ, 806, 2015)

TGSF: the objectives

EuroHPC Joint Undertaking

Extreme Scale Access

Topic: The role of turbulence and gravity in star formation
Duration: 1 yr
Allocation: 22,000,000 GPUh in LUMI-G

Study the formation of stellar cores and their initial mass function at unprecedented resolution

Study turbulent transport and mixing

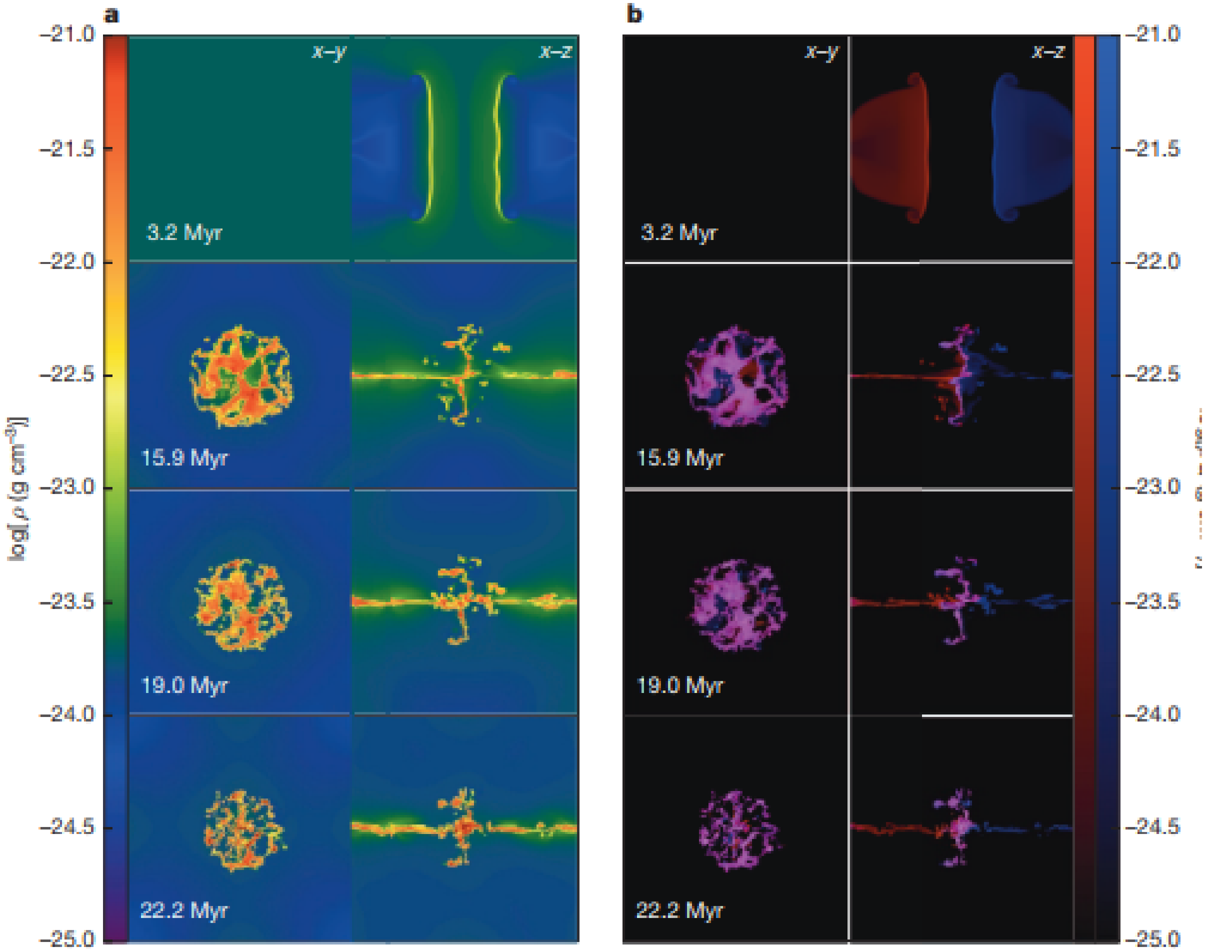

Gas density (left) and tracers density (right). Mixing raising from turbulence in colliding flows. (Feng and Krumholz, Nature, 513, 2014)

TGSF: the objectives

Extreme Scale Access

Topic: The role of turbulence and gravity in star formation
Duration: 1 yr
Allocation: 22,000,000 GPUh in LUMI-G

Study the formation of stellar cores and their initial mass function at unprecedented resolution

Study turbulent transport and mixing

Contribute to the general theory of turbulence

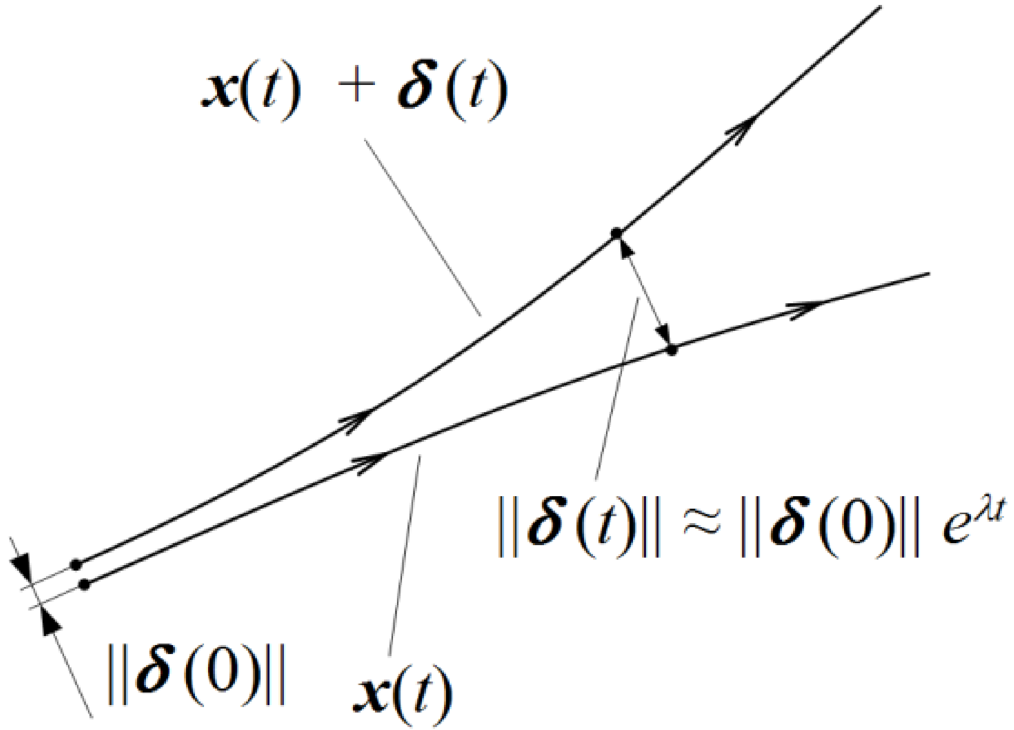

Maximal Lyapunov exponent (λ)

TGSF: the objectives

EuroHPC Joint Undertaking

Extreme Scale Access

Topic: The role of turbulence and gravity in star formation
Duration: 1 yr
Allocation: 22,000,000 GPUh in LUMI-G

- Study the formation of stellar cores and their initial mass function at unprecedented resolution
- Study turbulent transport and mixing
- Contribute to the general theory of turbulence
- Study the load imbalance, performance, and energy consumption at unprecedented scales

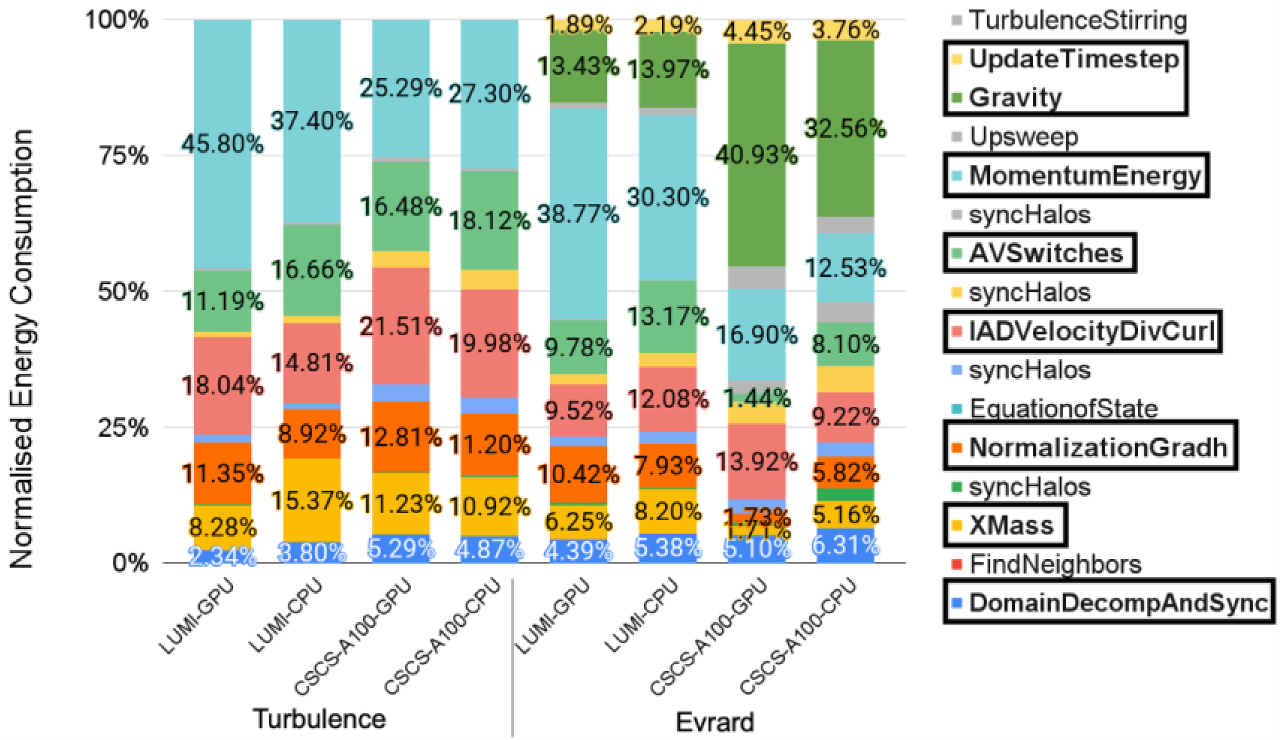

Accurate measurement of application-level energy consumption for energy-aware large-scale simulations with SPH-EXA.
(Simsek et al., SC-W 2023)

TGSF: the objectives

EuroHPC Joint Undertaking

Extreme Scale Access

Topic: The role of turbulence and gravity in star formation
Duration: 1 yr
Allocation: 22,000,000 GPUh in LUMI-G

- Study the formation of stellar cores and their initial mass function at unprecedented resolution
- Study turbulent transport and mixing
- Contribute to the general theory of turbulence
- Study the load imbalance, performance, and energy consumption at unprecedented scales
- Study compression techniques for large-scale checkpointing, compression, and visualization

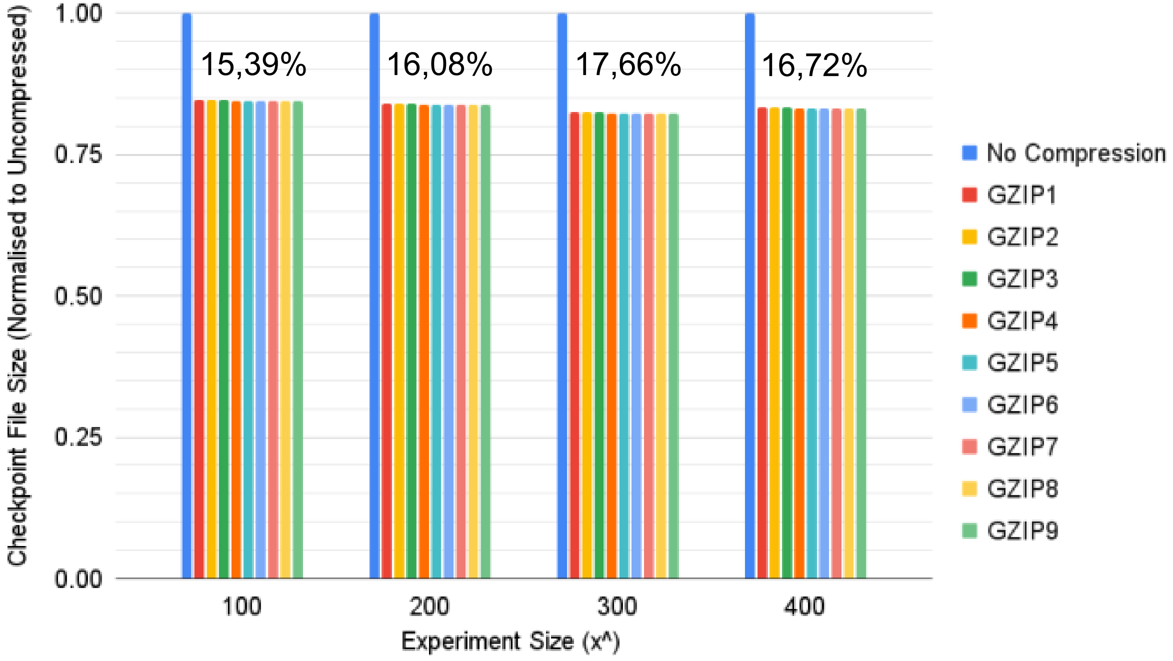

Checkpoint compression using HDF5 built-in lossless compression.
Subsonic turbulence tests in Piz Daint (CSCS).
(SPH-EXA team, 2023)

Acknowledgements

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

FUNDING

Platform for Advanced Scientific Computing

SKACH

Universität Basel

Florina Ciorba
florina.ciorba@unibas.ch
Rubén Cabezón
ruben.cabezon@unibas.ch
Osman Seckin
Yiqing Zhu
Lukas Schmidt

Universität Zürich[UZH]

Lucio Mayer
lmayer@physik.uzh.ch
Noah Kubli
Darren Reed

CSCS

Sebastian Keller
Jean-Guillaume Piccinali
Jonathan Coles
Jean Favre
John Biddiscombe

UNIVERSITAT POLITÈCNICA DE CATALUNYA BARCELONATECH

Domingo García-Senz
Axel Sanz

ESO

José Antonio Escartín

université PARIS-SACLAY

Joseph Touzet

MISSISSIPPI STATE UNIVERSITY

Ioanna Banicescu

UNIVERSITY of WASHINGTON

Tom Quinn

UNIVERSITÄT HEIDELBERG

Ralf Klessen

Performance and scaling of PION for modelling Colliding-Wind Binary systems

EUROHPC USER DAY 2023 Brussels 11.12.23

Project: "*Colliding winds in massive binary systems*"

EuroHPC used: Karolina (IT4I, Czech Republic)

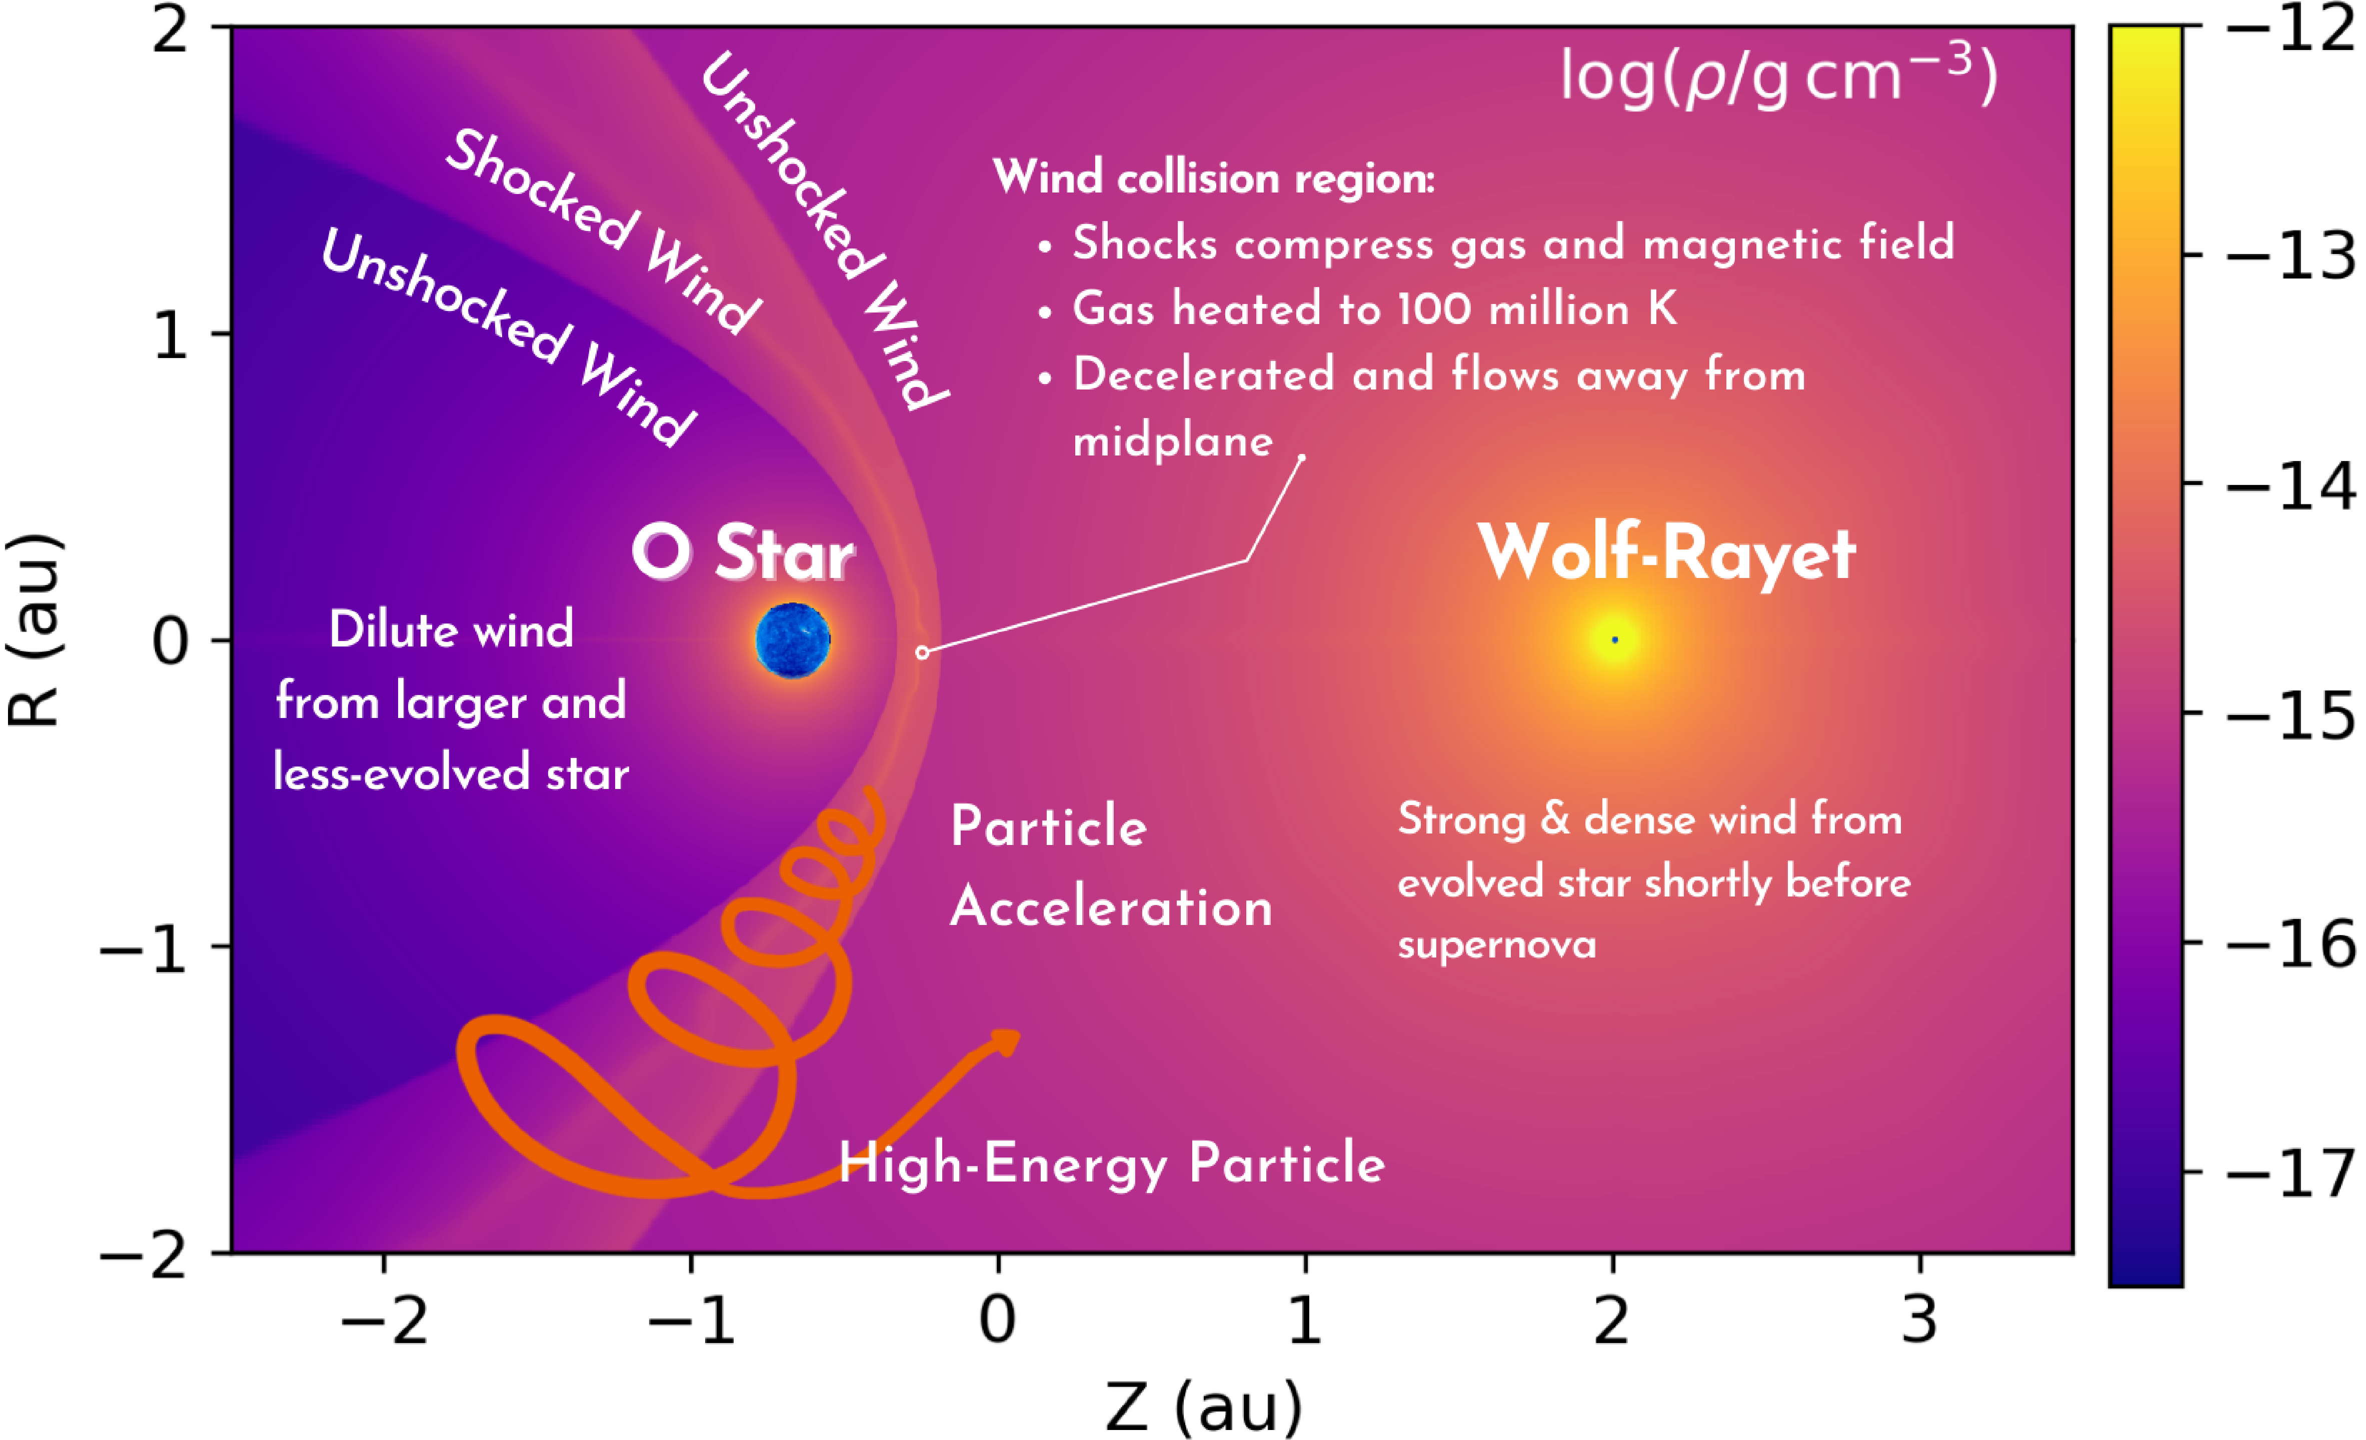

$\log(\rho/\mathrm{g\,cm^{-3}})$
Unshocked Wind
Shocked Wind
Unshocked Wind
Wind collision region:
• Shocks compress gas and magnetic field
• Gas heated to 100 million K
• Decelerated and flows away from midplane
O Star
Wolf-Rayet
Dilute wind from larger and less-evolved star
Particle Acceleration
Strong & dense wind from evolved star shortly before supernova
High-Energy Particle
R (au)
Z (au)

Colliding-wind binaries

- CWBs probe winds of both stars:
  - Eccentric orbits → predictable time-varying shock conditions
  - Constrain winds of massive stars
    - Wind clumping, acceleration, mass-loss rates
    - Final pre-supernova evolutionary stage
  - Intense dust production
  - Particle acceleration in shocks
  - Progenitor systems for NS/NS, NS/BH, BH/BH mergers.

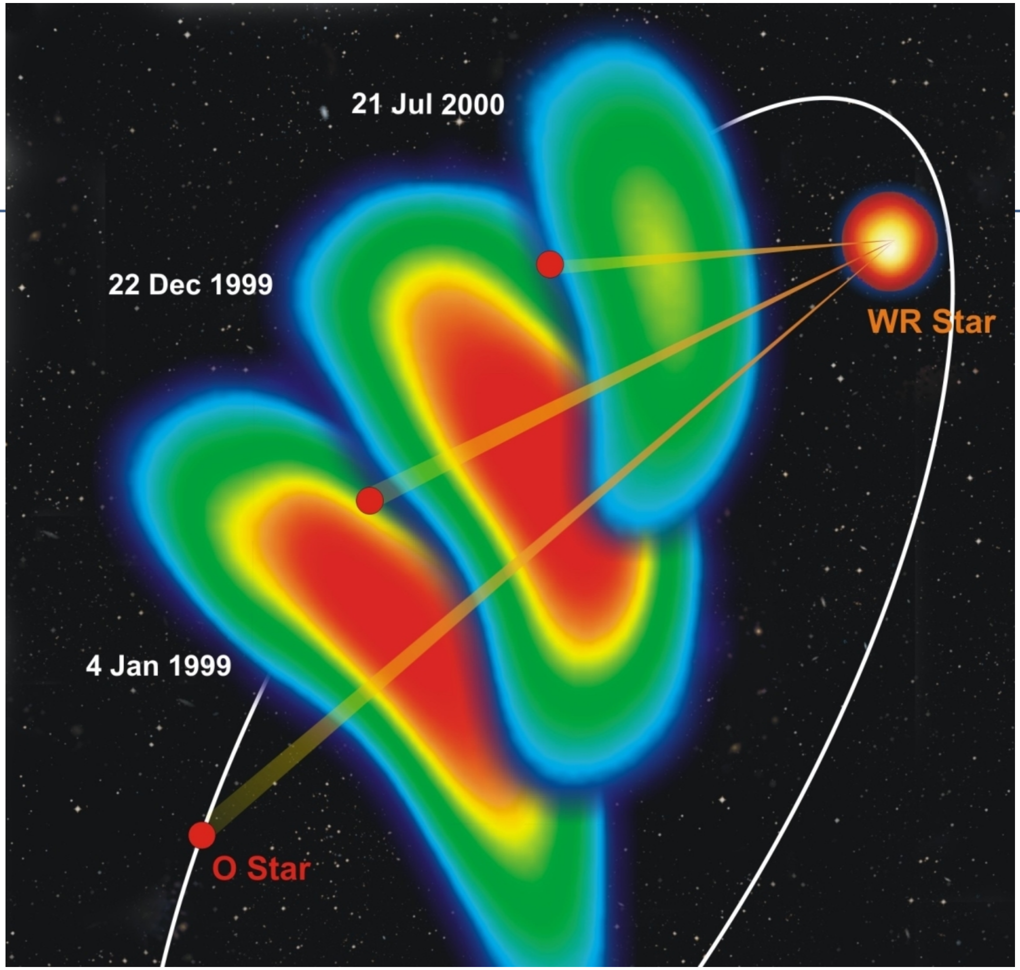

Synchrotron radio observations of WR140
Dougherty et al. (2006)

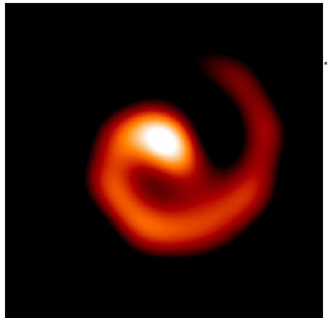

Inner dust nebula of WR104
Credit: Peter Tuthill.
http://www.physics.usyd.edu.au/~gekko/wr104.html

DIAS
Institiúid Ard-Léinn Bhaile Átha Cliath | Dublin Institute for Advanced Studies

Colliding-wind binaries

- CWBs probe winds of both stars:
  - Eccentric orbits → predictable time-varying shock conditions
  - Constrain winds of massive stars
    - Wind clumping, acceleration, mass-loss rates
    - Final pre-supernova evolutionary stage
  - Intense dust production
  - Particle acceleration in shocks
  - Progenitor systems for NS/NS, NS/BH, BH/BH mergers.

Concentric dust rings in WR 140 from successive periastron passages (JWST, Lau+2022)

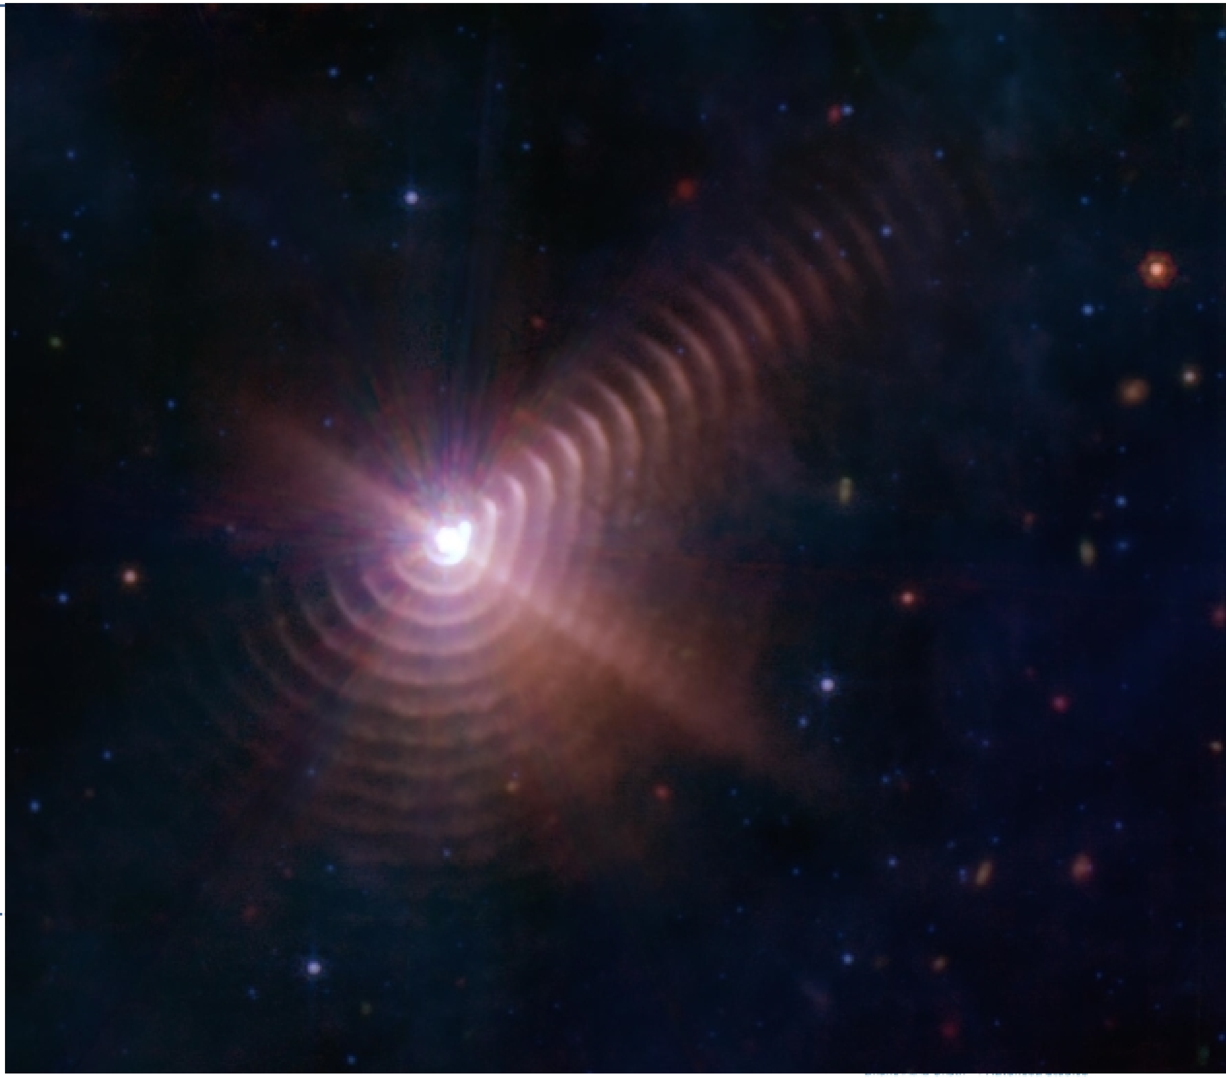

X-ray lightcurve of WR140 (Pollock+,2021)

https://doi.org/10.3847/1538-4357/ac2430

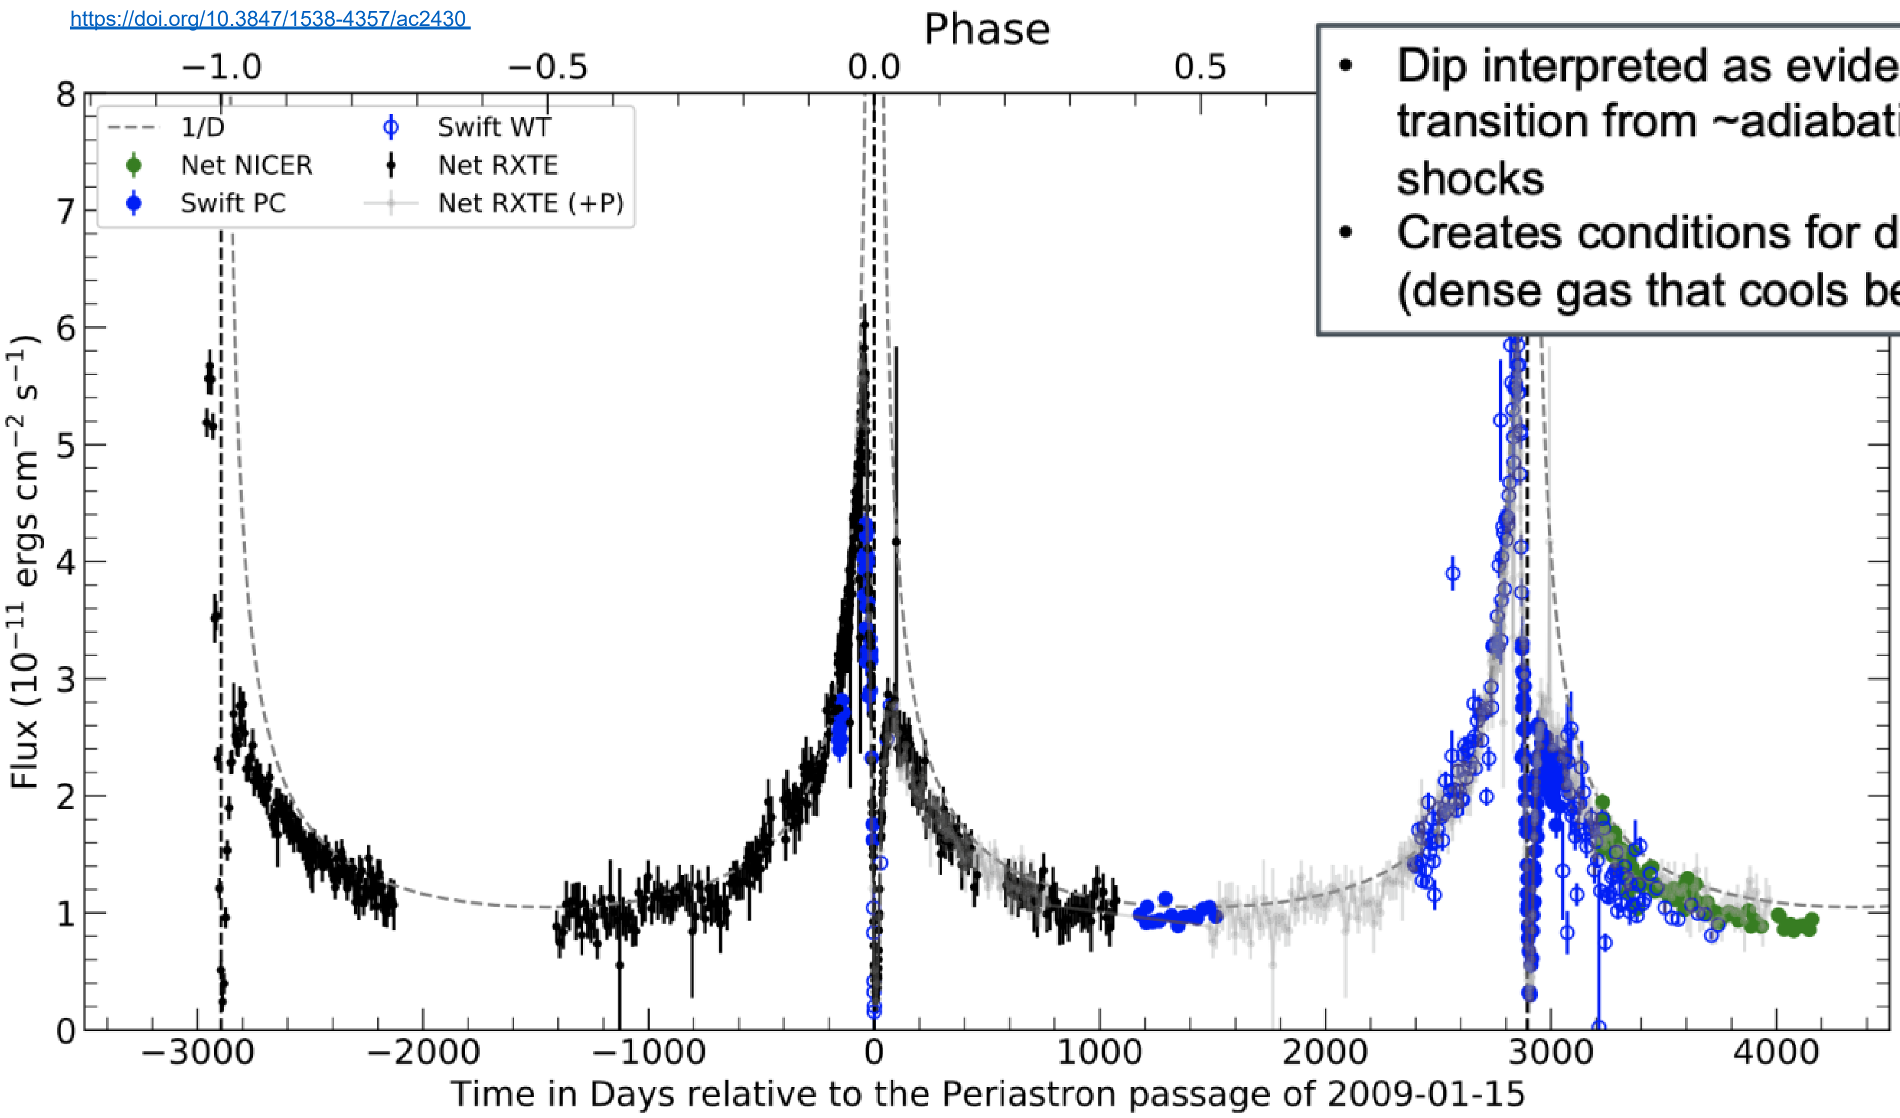

- Dip interpreted as evidence for transition from ~adiabatic to radiative shocks
- Creates conditions for dust formation (dense gas that cools below 10^4 K)

Figure 2. The RXTE, Swift, and NICER light curve of WR 140, 2000–2020. The time zero-point corresponds to the periastron passage JD 2,454,846.727 = 2009 January 15 05:26. The gray points are the RXTE fluxes advanced by one period. The black dashed vertical lines are the times of periastron passage, while the gray dashed curves show the expected $1/D$ variation in flux for an unobscured adiabatic system of colliding winds.

The need for MHD simulations

- Interpretation of radio and gamma-ray observations requires:
  - Where are the shocks and what are their properties?
  - Conditions in the thermal plasma
  - Magnetic field configuration (for acceleration and transport)
  - System timescales (dynamical, thermal, acceleration)
- Phenomenological or one-zone models often inconclusive or inconsistent
- 3D simulations with particle acceleration+transport very challenging
  - Galactic-scale algorithms don’t translate to these small scales
  - PIC simulations don’t get to system scale
- 1st step: MHD simulations with postprocessing for CRs (test-particle limit)
  - Assume acceleration is inefficient → no back-reaction on MHD flow

PION MHD code

- Developed to model MHD of photoionized and wind-driven nebulae around massive stars.
- Finite-volume method
- 1D spherical / 2D cylindrical / 3D Cartesian
- Euler or ideal-MHD equations
- Statically refined, nested grid, adaptive timesteps
- MPI+OpenMP parallelized
- Radiative transfer from point sources with energy deposition and ionization
- Methods paper and public code release (Mackey *et al.* 2021, MNRAS)
- pyPION module: read snapshots to numpy arrays

https://www.pion.ie
https://git.dias.ie/massive-stars-software/pion
https://git.dias.ie/massive-stars-software/pypion

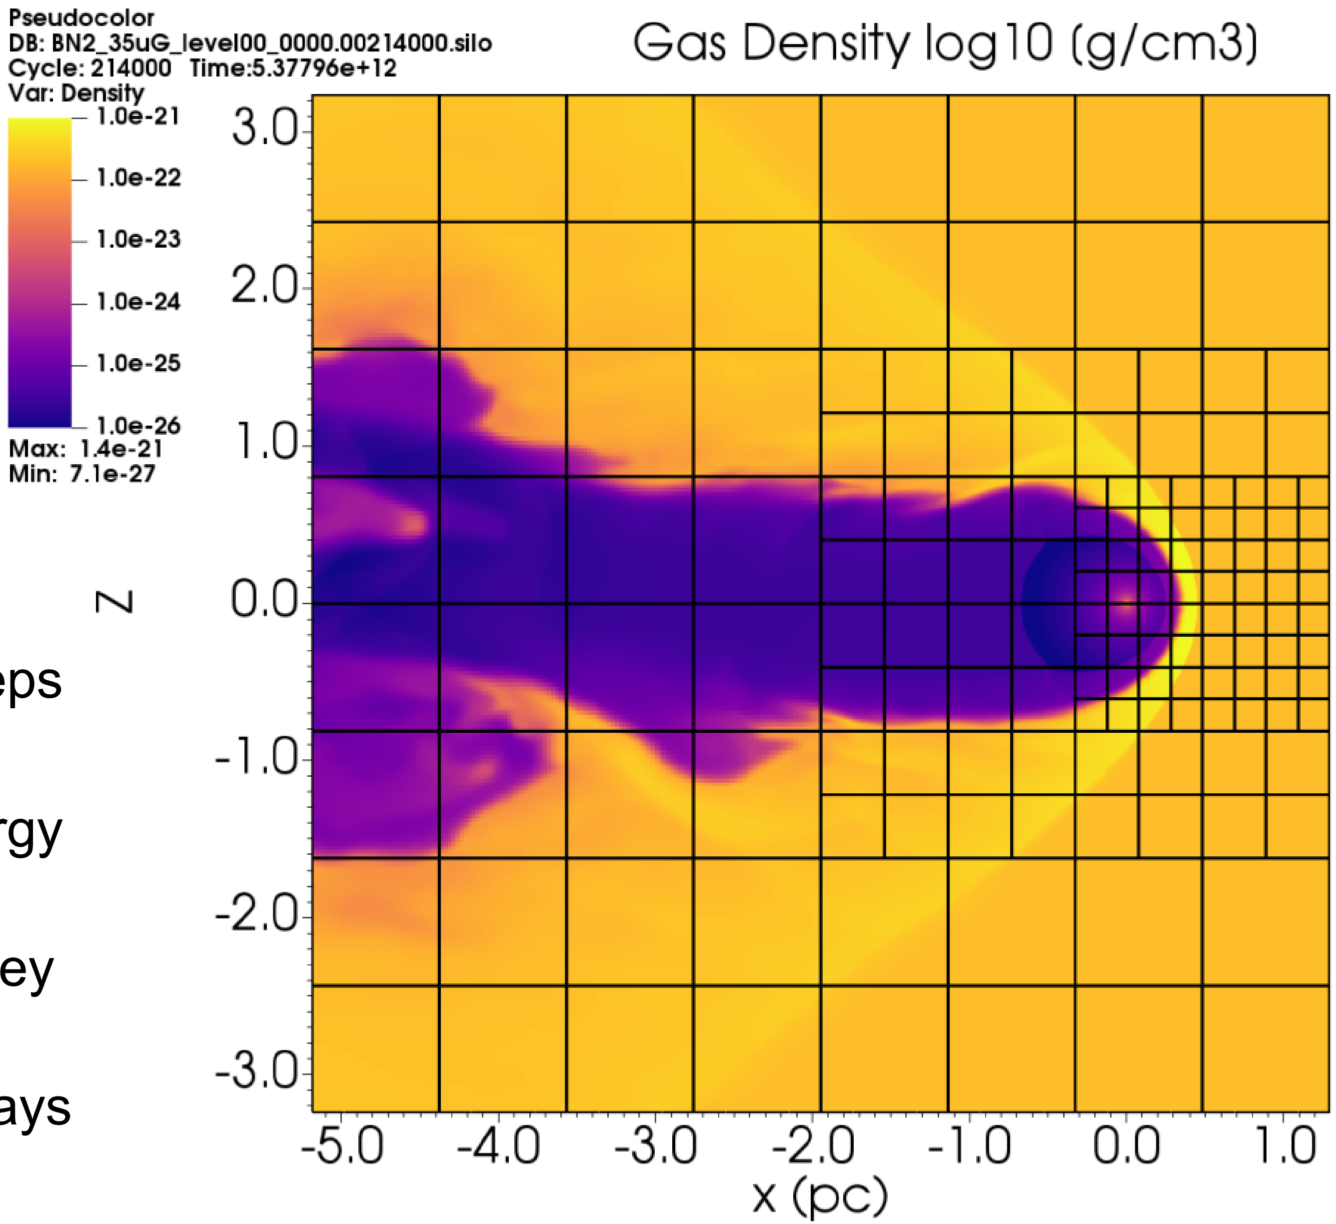

3D Simulation by S. Green of a bow shock from a runaway O star. 256^3 with 3 levels of refinement.

New Capabilities of PION

- Development support from ICHEC through EuroHPC *"Academic Flagship Programme"*
- Binary systems:
  - Moving stars (simple N-body integrator)
  - Wind acceleration (β-law)
  - Inverse-Compton cooling
- Hybrid MPI-OpenMP parallelization
  - Better strong scaling
- New release v3.0 in preparation

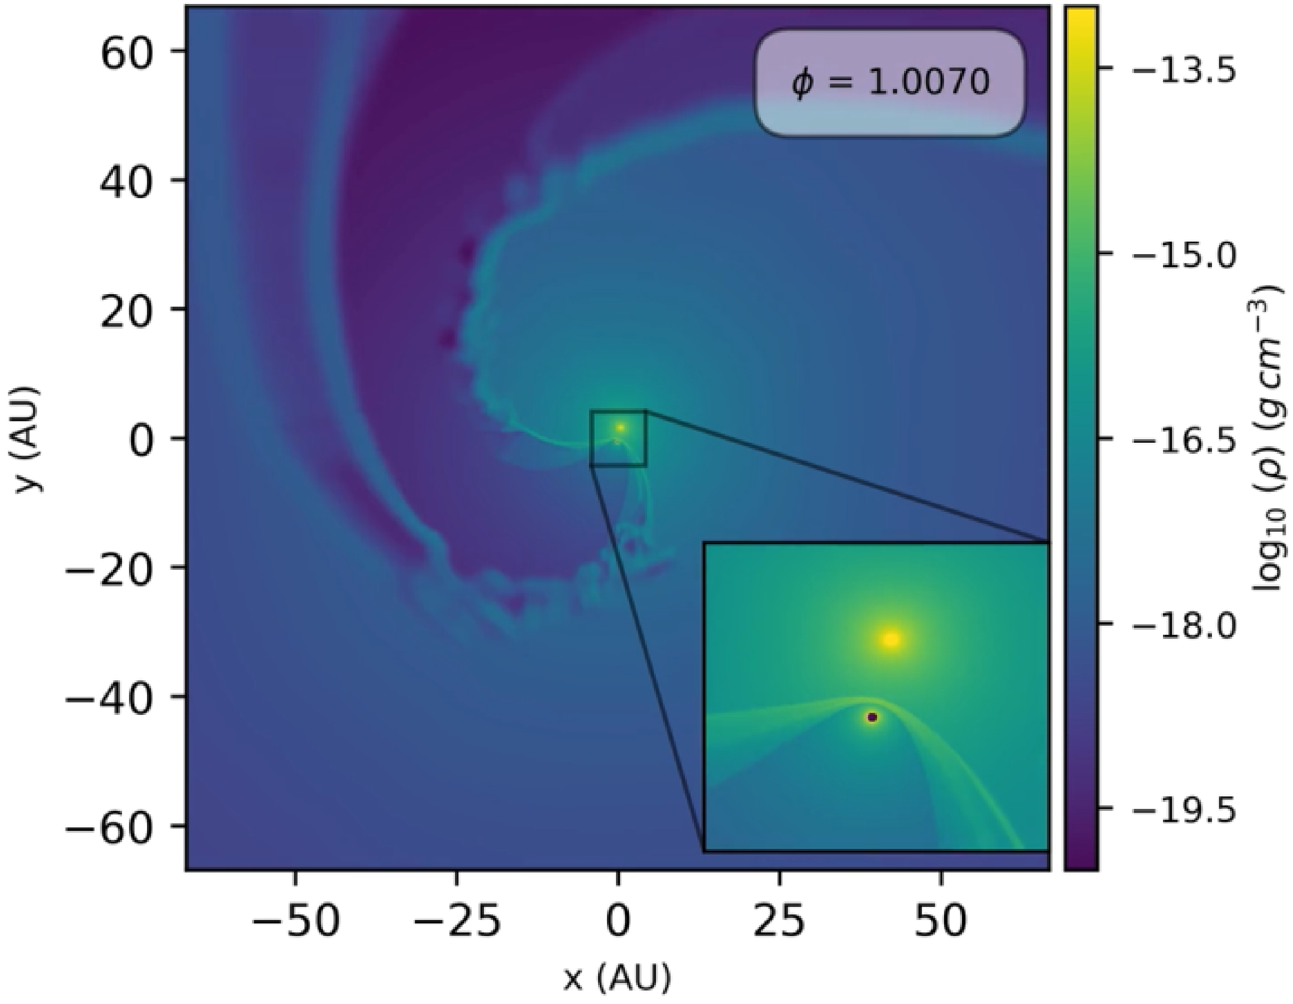

Mackey et al. (2023, MNRAS, **526**, 3099)
https://arxiv.org/abs/2301.13716

3D simulations of WR140 periastron passage

- 3D MHD simulations with static mesh-refinement, 256×256×64 cells per level
- 7 grid levels centred on centre of mass
- Both stars on finest level at periastron
- *First 3D MHD simulations of CWBs with orbital motion, stellar rotation, IC cooling in the literature*
- Start 140 days before periastron and finish similar time after periastron

Mackey et al. (2023, MNRAS, **526**, 3099)
https://arxiv.org/abs/2301.13716

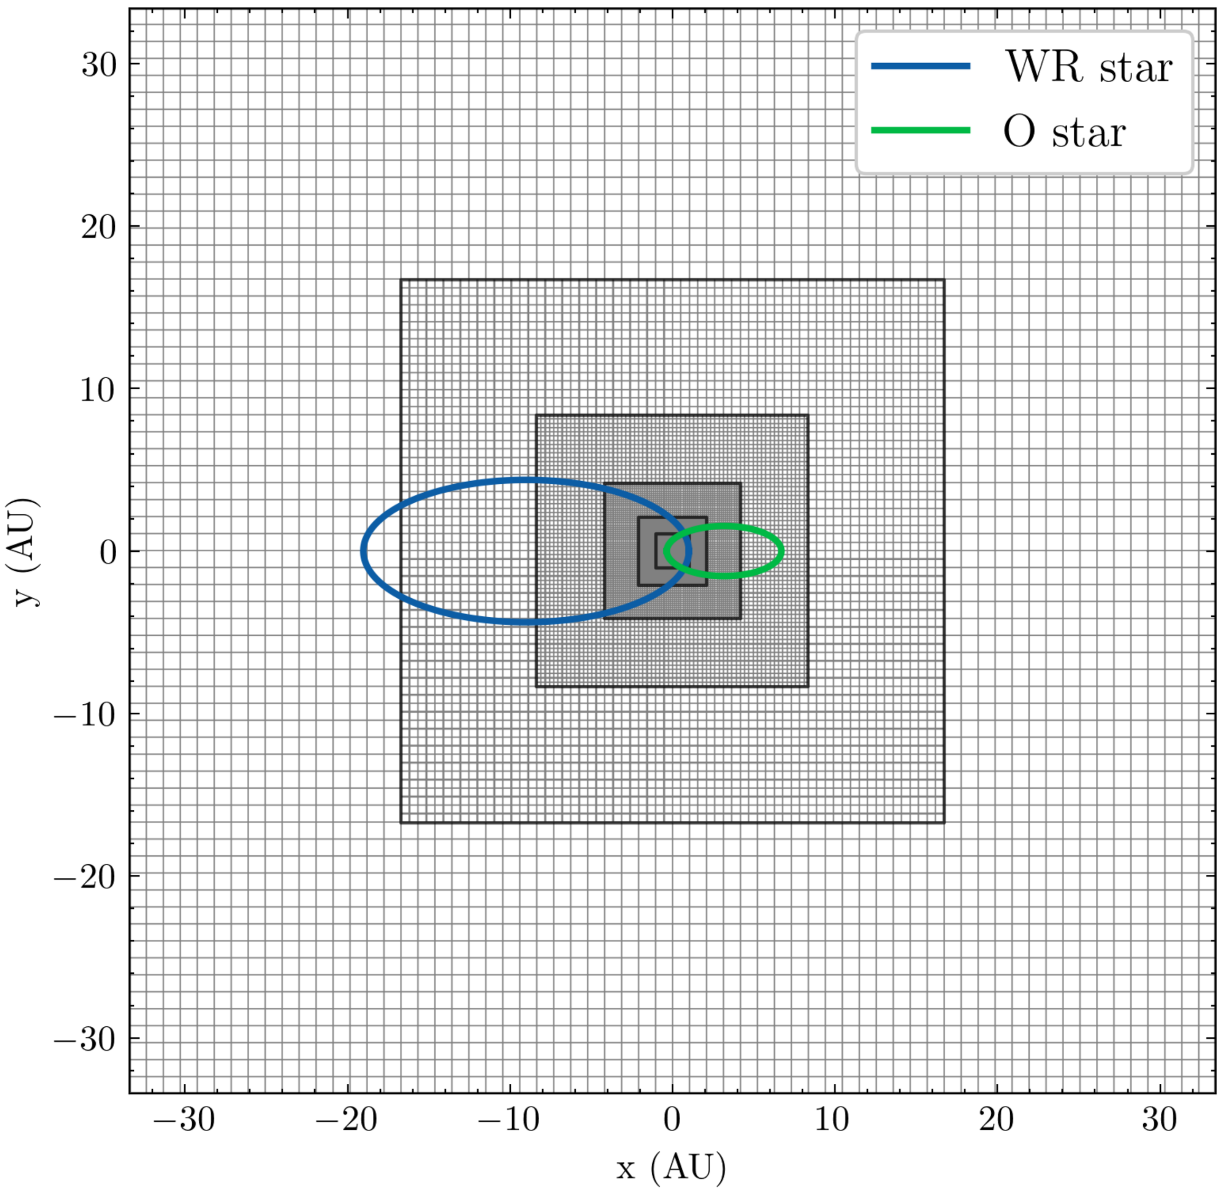

https://www.it4i.cz/en/infrastructure/karolina

VSB TECHNICAL UNIVERSITY OF OSTRAVA | IT4INNOVATIONS NATIONAL SUPERCOMPUTING CENTER

EUROCC E-INFRA CZ EXTRANET | CZ

ABOUT INFRASTRUCTURE RESEARCH INDUSTRY COOPERATION FOR USERS EDUCATION EVENTS

KAROLINA

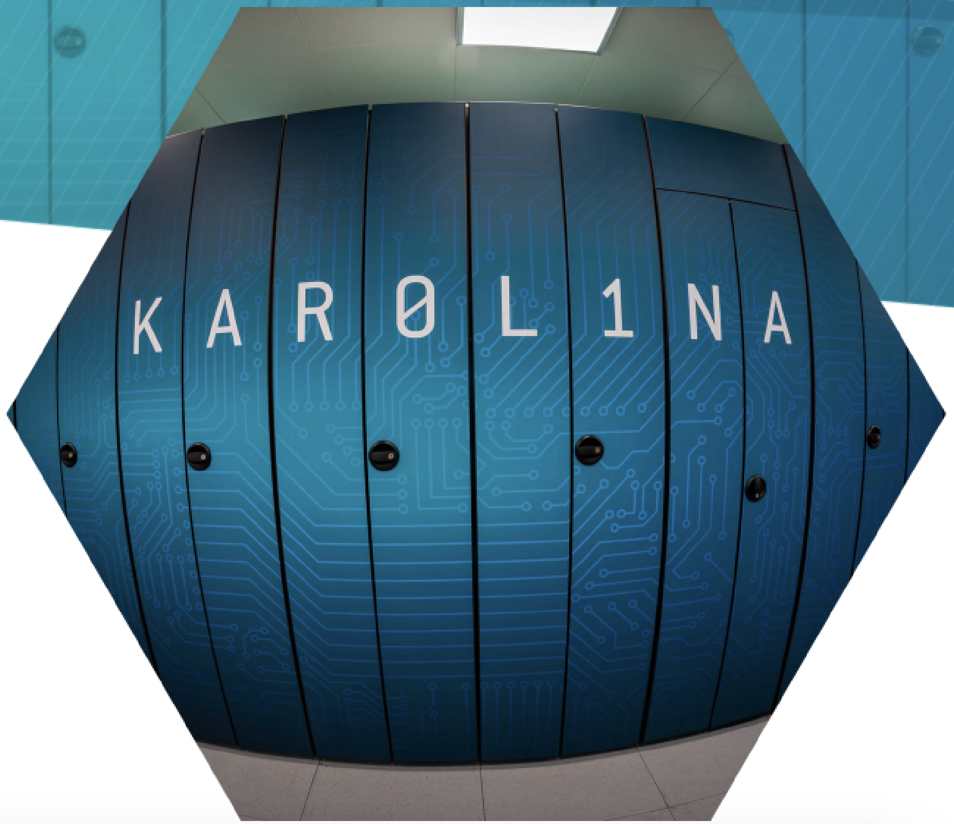

The petascale system Karolina, acquired as part of the EuroHPC Joint Undertaking, was installed in 2021. In the TOP500 list, which evaluates supercomputers in terms of their performance, it ranked 69th worldwide, 19th in Europe, and in the Green500 list of the most energy-efficient supercomputers, it even ranked 8th in 2021. The HPC system is designed to respond coherently to the needs of its user communities, addressing complex scientific and industrial challenges, including standard numerical simulations, demanding data analysis, and artificial intelligence applications.

DIAS

Institiúid Ard-Léinn Bhaile Átha Cliath | Dublin Institute for Advanced Studies

WR 21a binary system

- Very massive CWB system, a partially eclipsing binary (Barba+,2022)
    - 93 $M_\odot$ primary star, WNh type
    - 53 $M_\odot$ secondary, O3 type
    - 31.7 day orbit, $e = 0.695$
- X-ray luminosity among the highest of all known systems, $L_X \approx 5\times10^{34}$ erg/s (Gosset+,2016)
- No published hydro/MHD simulations

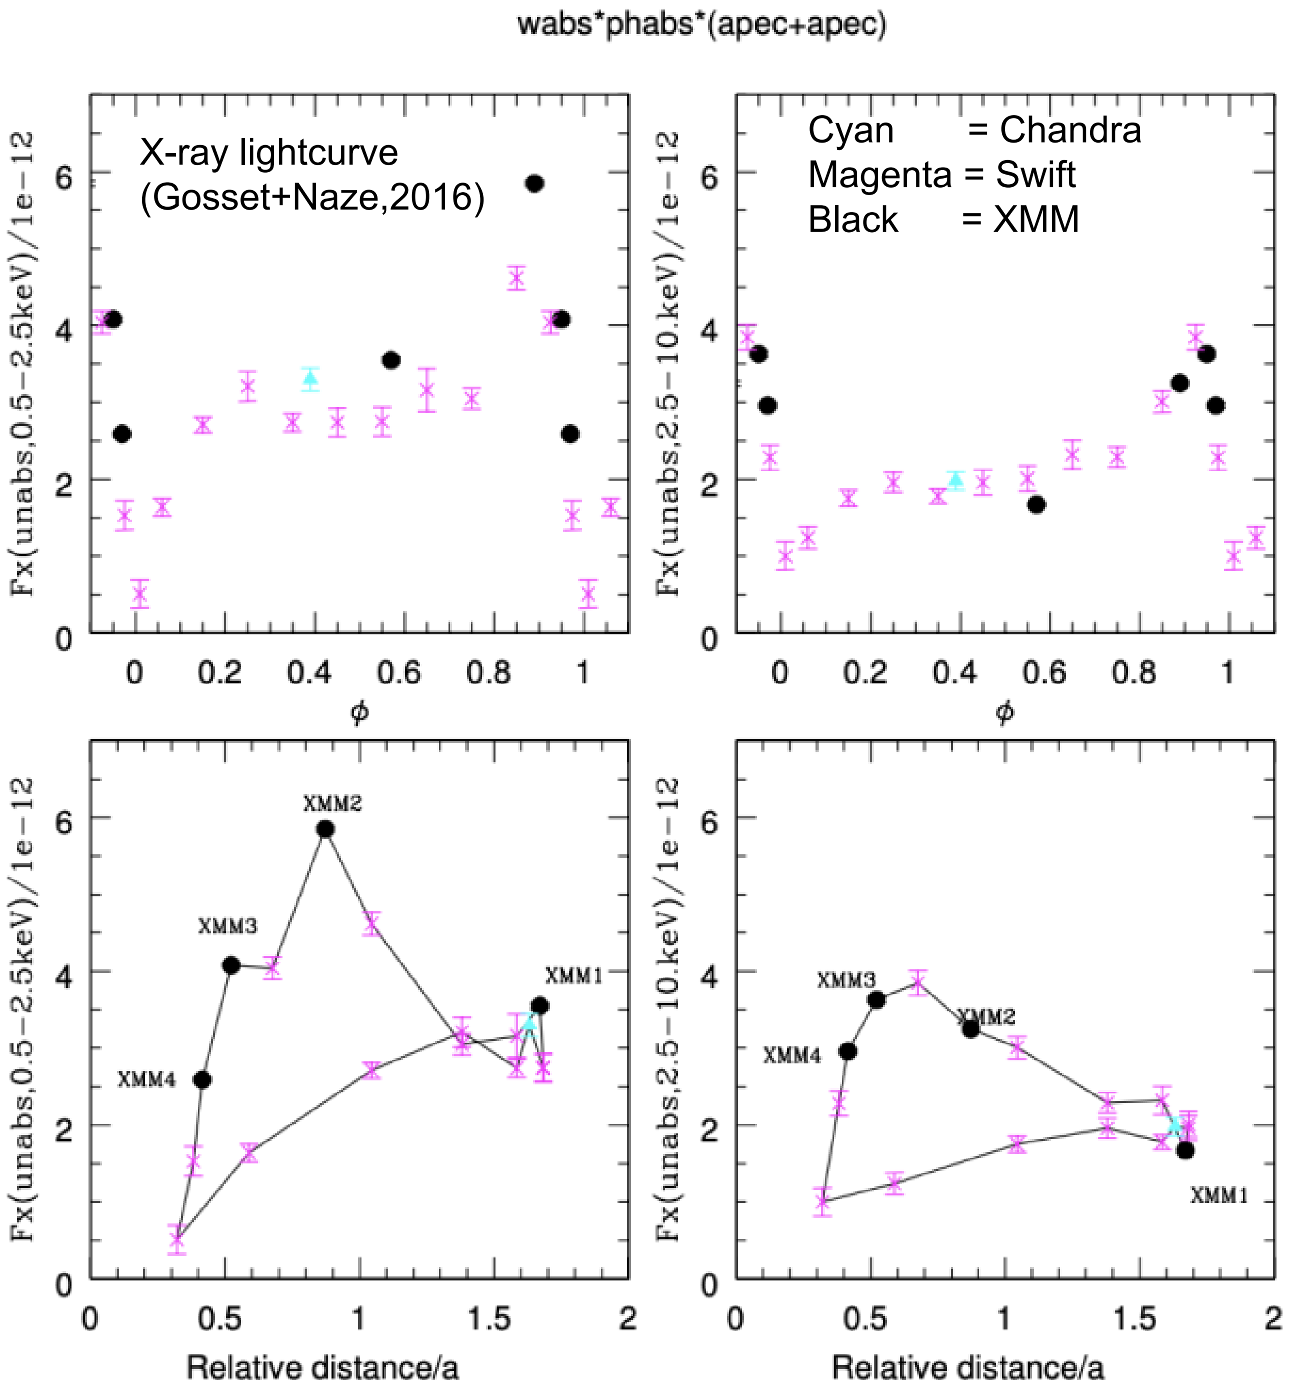

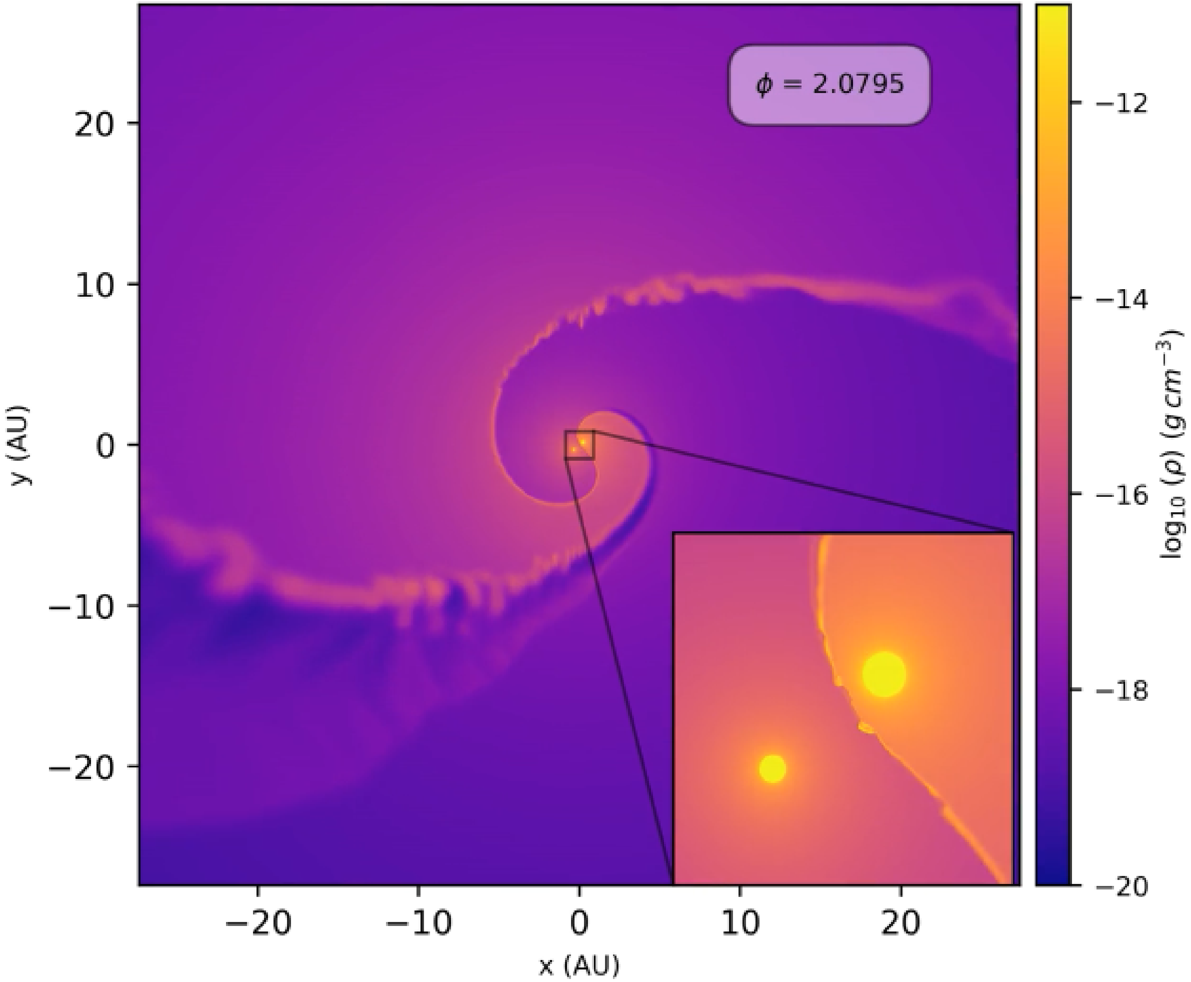

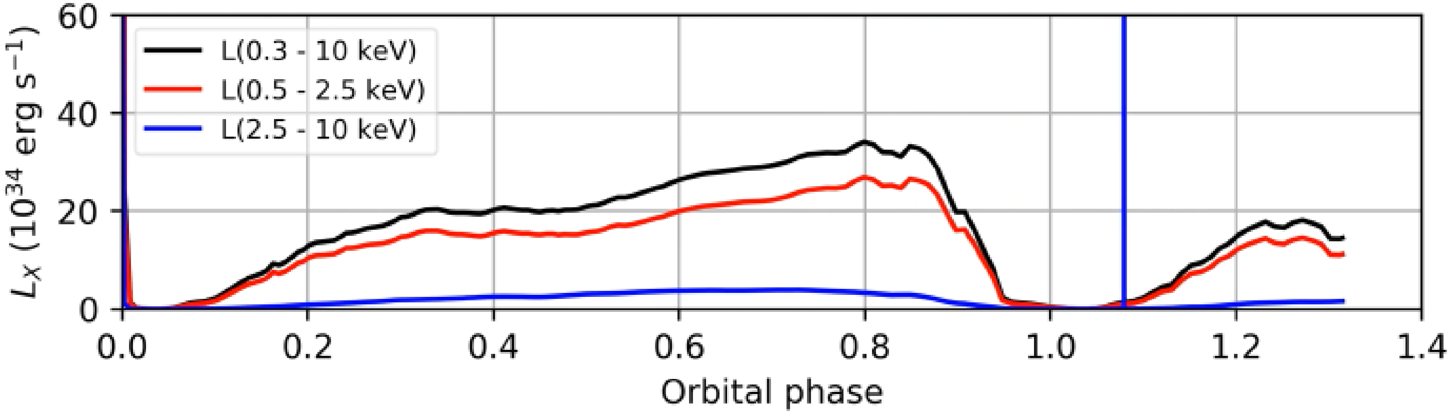

Problem Setup

- 3D hydrodynamics simulation with static mesh-refinement
- Assume radiation field of both stars accelerate/decelerate both winds (perfect coupling)
- $320^2 \times 160$ cells per level
- 8 refinement levels (2x each time)
- Gas density plotted on log scale
- Slice through z=0, orbital plane
- Yellow circles == stellar radius

EuroHPC
Joint Undertaking

Strong scaling for 3D hydrodynamics

- $320^2 \times 160$ cells per level, 8 levels
- 199 coarse-level steps; 25,472 on finest-level
- Same starting and finishing snapshot
- 1, 2, 4, 8, 16, 32, 64 nodes (128 core/node)
- 1, 2, 4, 8 OpenMP threads / MPI process
- Strong scaling plot, normalized to 1 node, 1 OpenMP thread/proc.

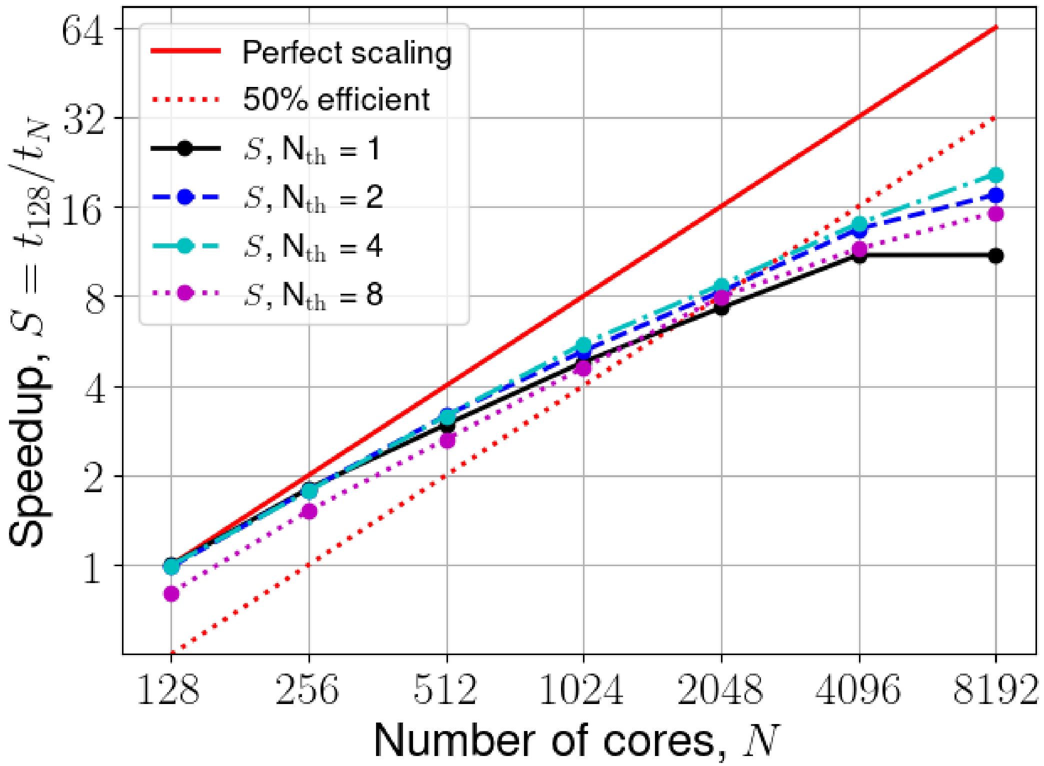

EuroHPC Joint Undertaking

DIAS
Institiúid Ard-Léinn Bhaile Átha Cliath | Dublin Institute for Advanced Studies

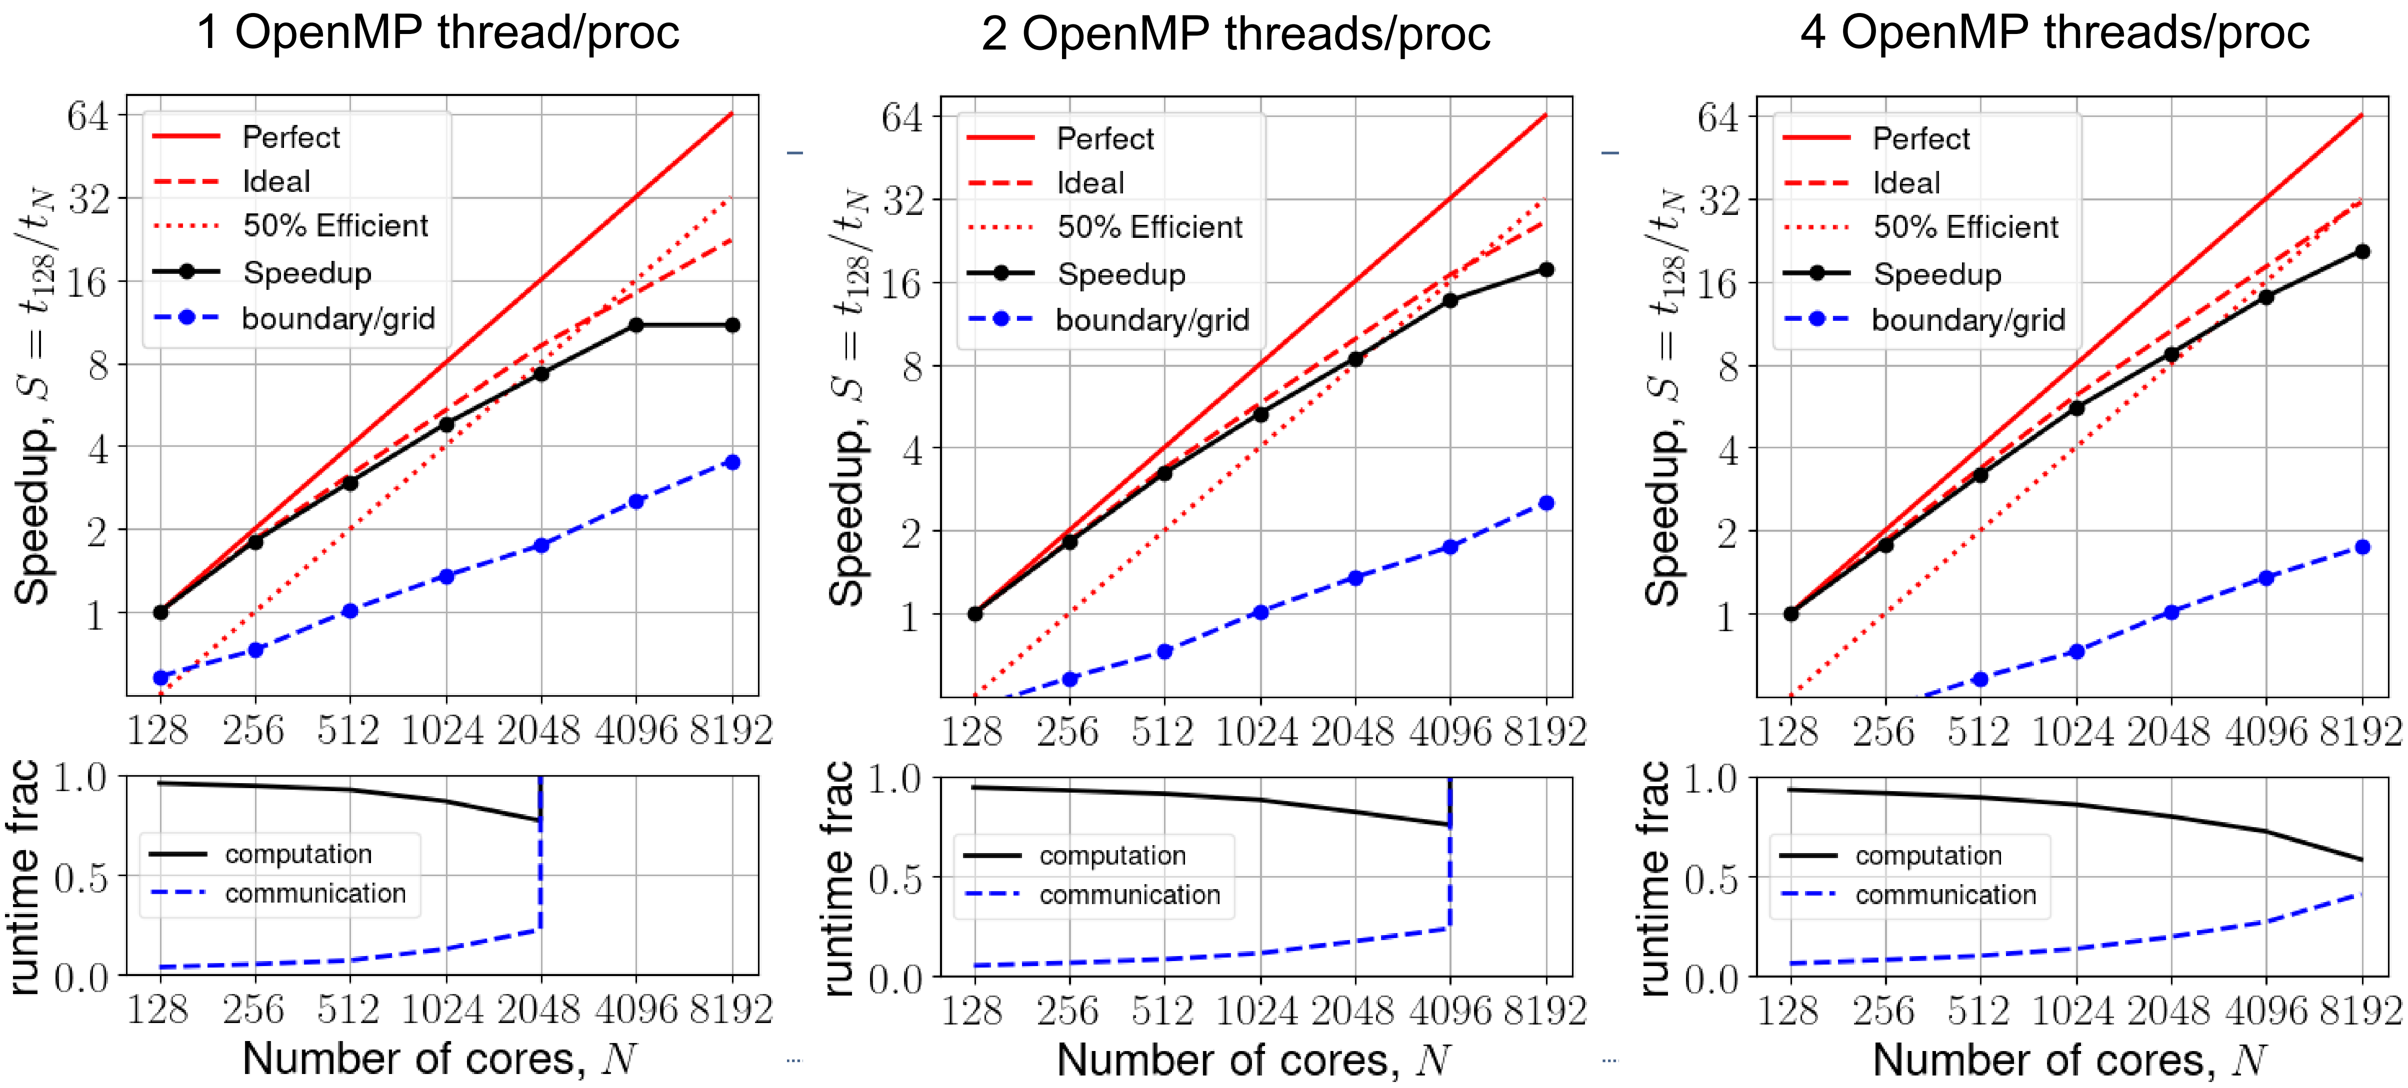

- Boundary/grid line is the ratio of boundary cells to grid cells in each sub-domain
- ”Ideal” is the best speedup we can expect given the extra cells to calculate from the boundary data

Weak scaling for 3D hydrodynamics

- Weak scaling for W21a hydro simulation
- 8 OpenMP threads per MPI process
- Performance good on 128-2048 cores
  - 1 node: $128^2 \times 64$ cells per level
  - 2 nodes: $160^2 \times 80$ cells per level
  - 4 nodes: $192^2 \times 96$ cells per level
  - 8 nodes: $256^2 \times 128$ cells per level
  - 16 nodes: $320^2 \times 160$ cells per level

EuroHPC
Joint Undertaking

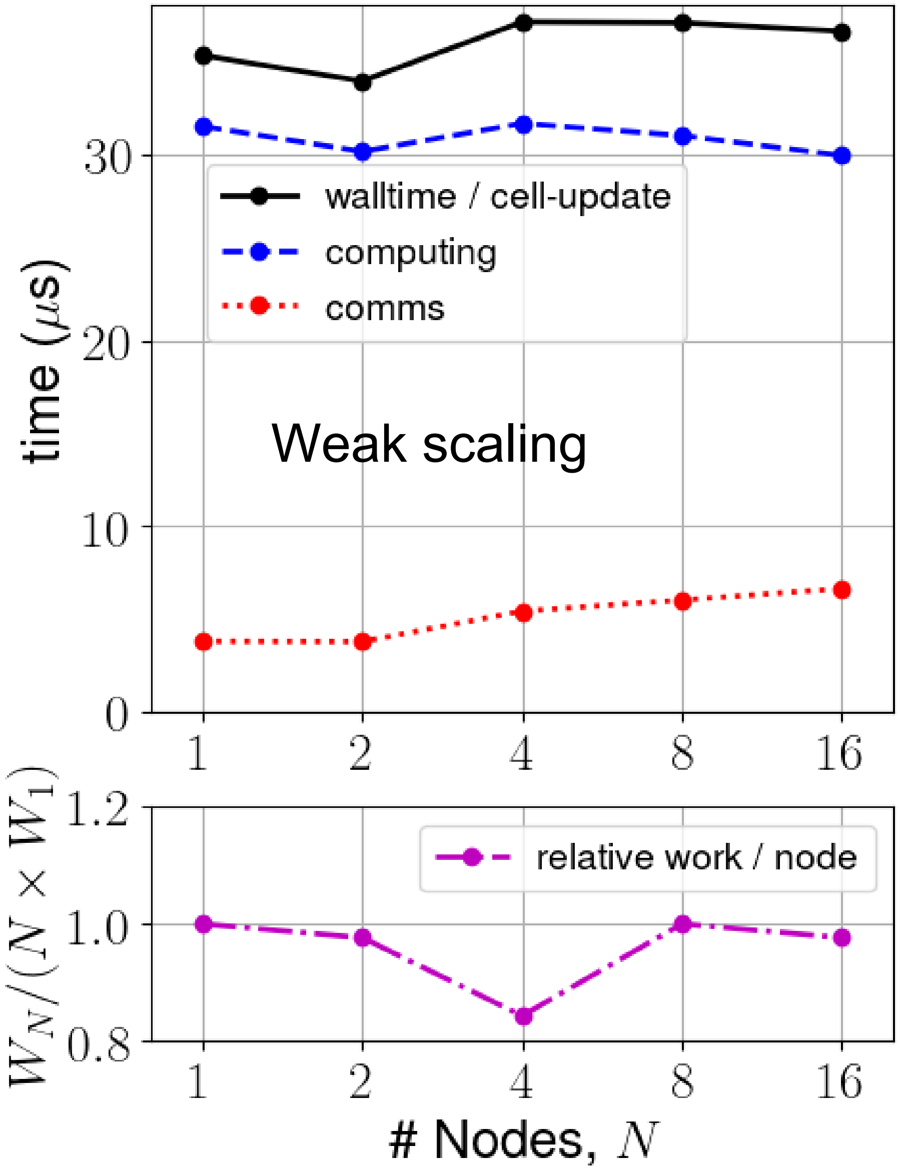

Summary / Outlook

- Simulation code prepared to make first **3D MHD models** of colliding winds in binary systems with orbital motion
- Weak scaling excellent up to 2048 cores on Karolina
- Strong scaling >50% efficient on 16 nodes vs. 1 node using 4 OpenMP threads / MPI proc
- Room for improvement of multithreading with 8-16 OpenMP threads (level communication)

Outlook: Higher-resolution simulations and multi-wavelength synthetic observations:

- Predictive power for X-ray lightcurves, spectra;
- Test models for wind acceleration, wind clumping;
- Radiative heating and cooling of shocked plasma;
- Condensation model to study dust formation;
- Post-processing to study particle acceleration and radiation in test-particle limit.

Thanks!

Collaborators:
Robert Brose,
Arun Mathew,
Thomas Jones,
Luca Grassitelli (Uni. Bonn),
Brian Reville (MPIK, Heidelberg)

We acknowledge EuroHPC JU for awarding this project access to Karolina hosted by IT4Innovations"

Studying planet formation on HPC systems with GPU computing

EUROHPC USER DAY 2023 Brussels 11.12.23

Project: *“2022R03-047”*

EuroHPC used: Vega

Speaker: *Ramon BRASSER (CSFK)*

The Solar System's planets

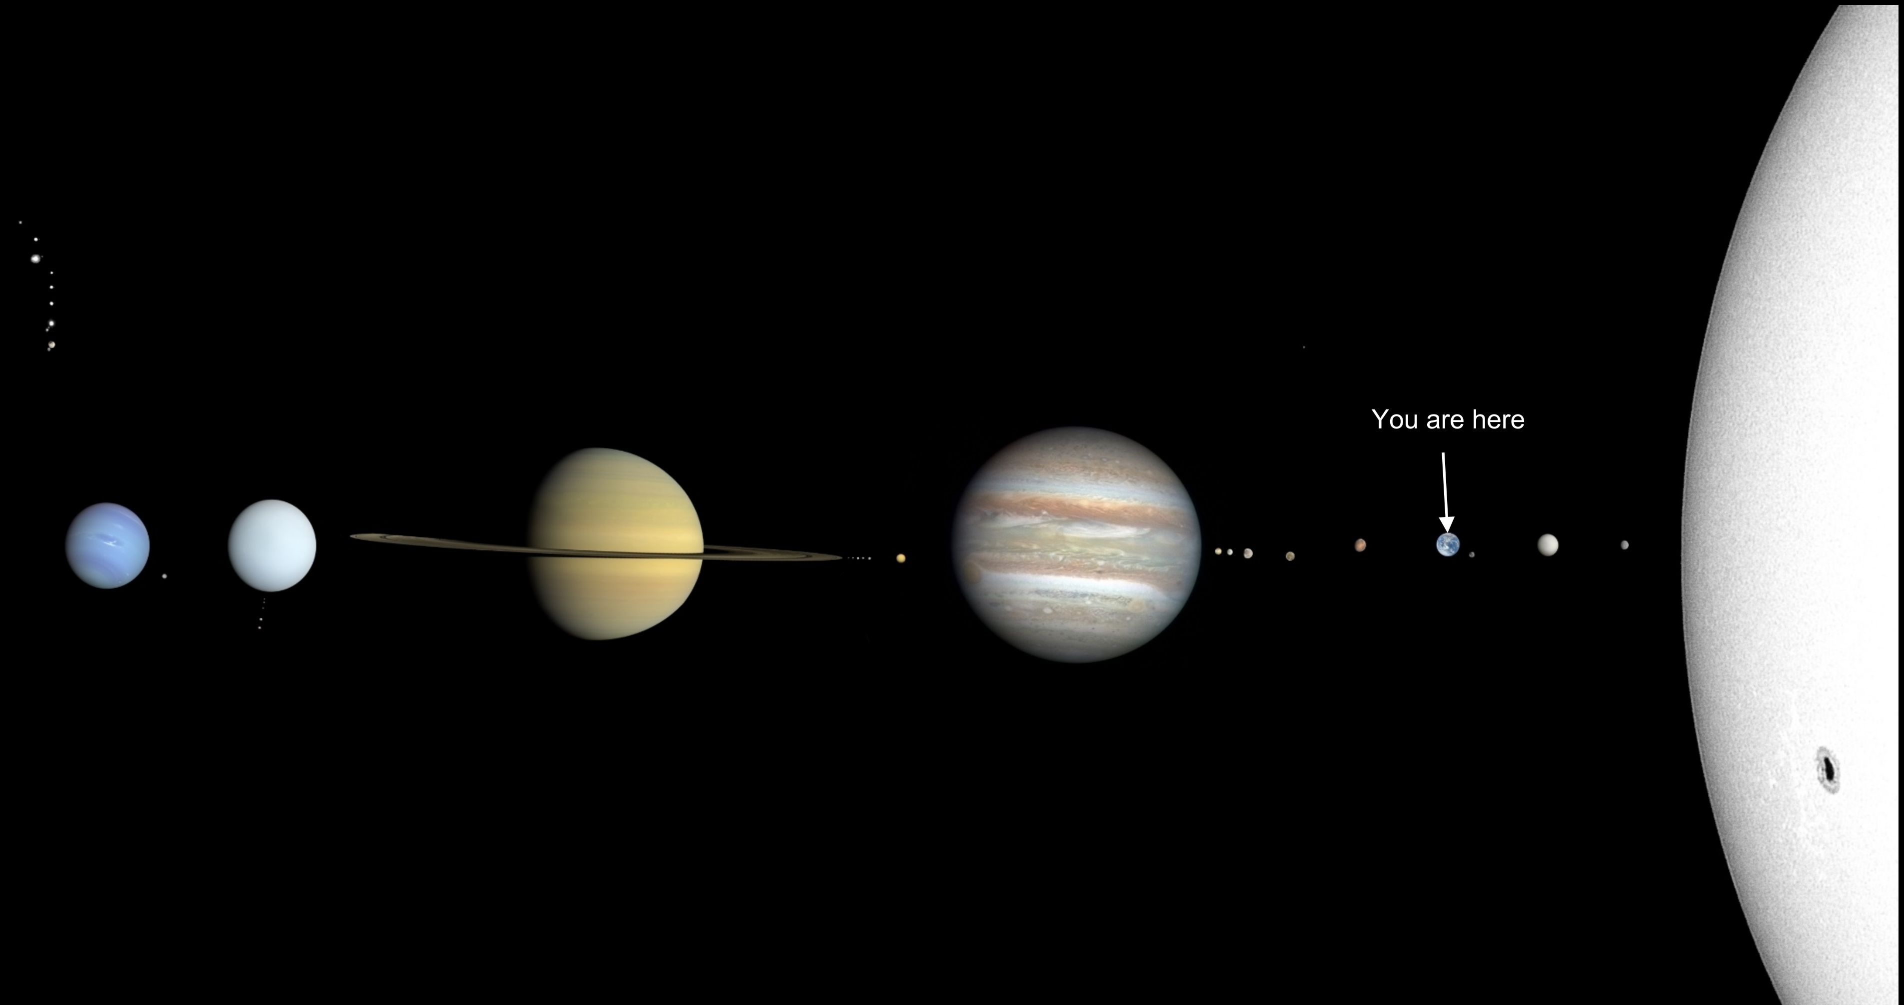

We have small, rocky planets close to the Sun, and large, gaseous planets far away. Why? And how?

The standard model of terrestrial planet formation

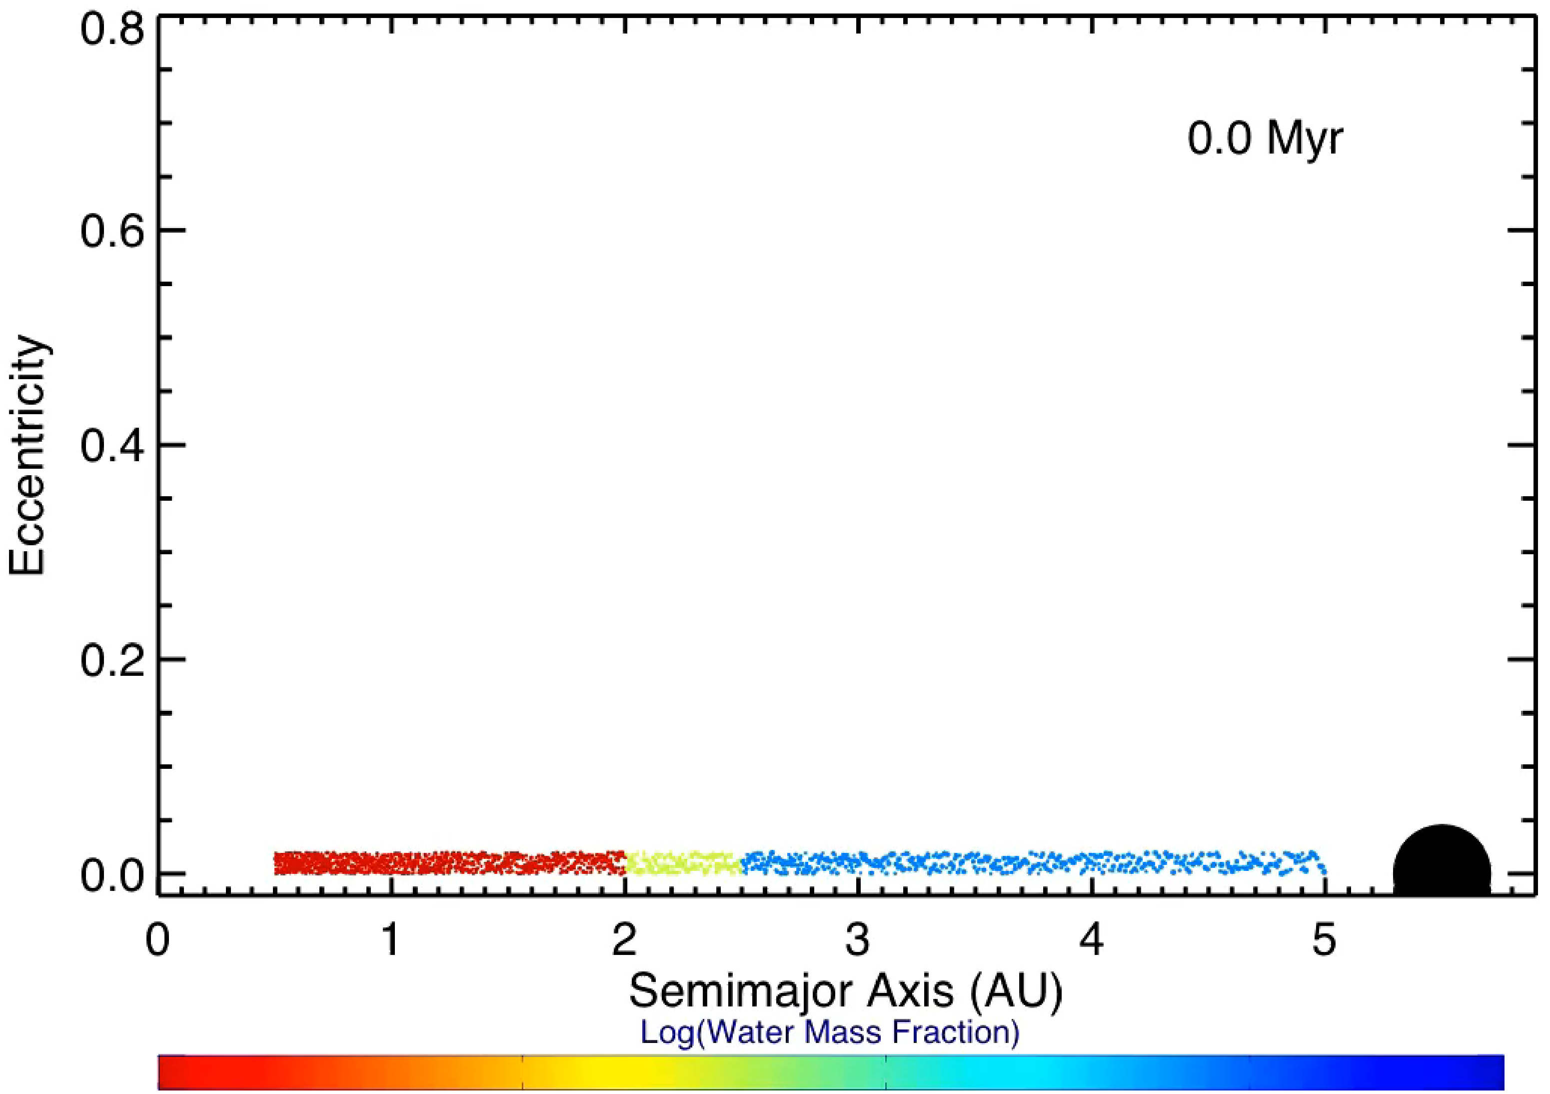

The classical model relies on a disc of planetesimals according to the minimum-mass solar nebula.

The idea is that the planetesimals accrete to form the planets.

Raymond, Quinn & Lunine 2006

Final planets

- The final planets are a mixture of all three colour groups.
- All three planets have a similar composition.
- All three planets have a small amount of water.

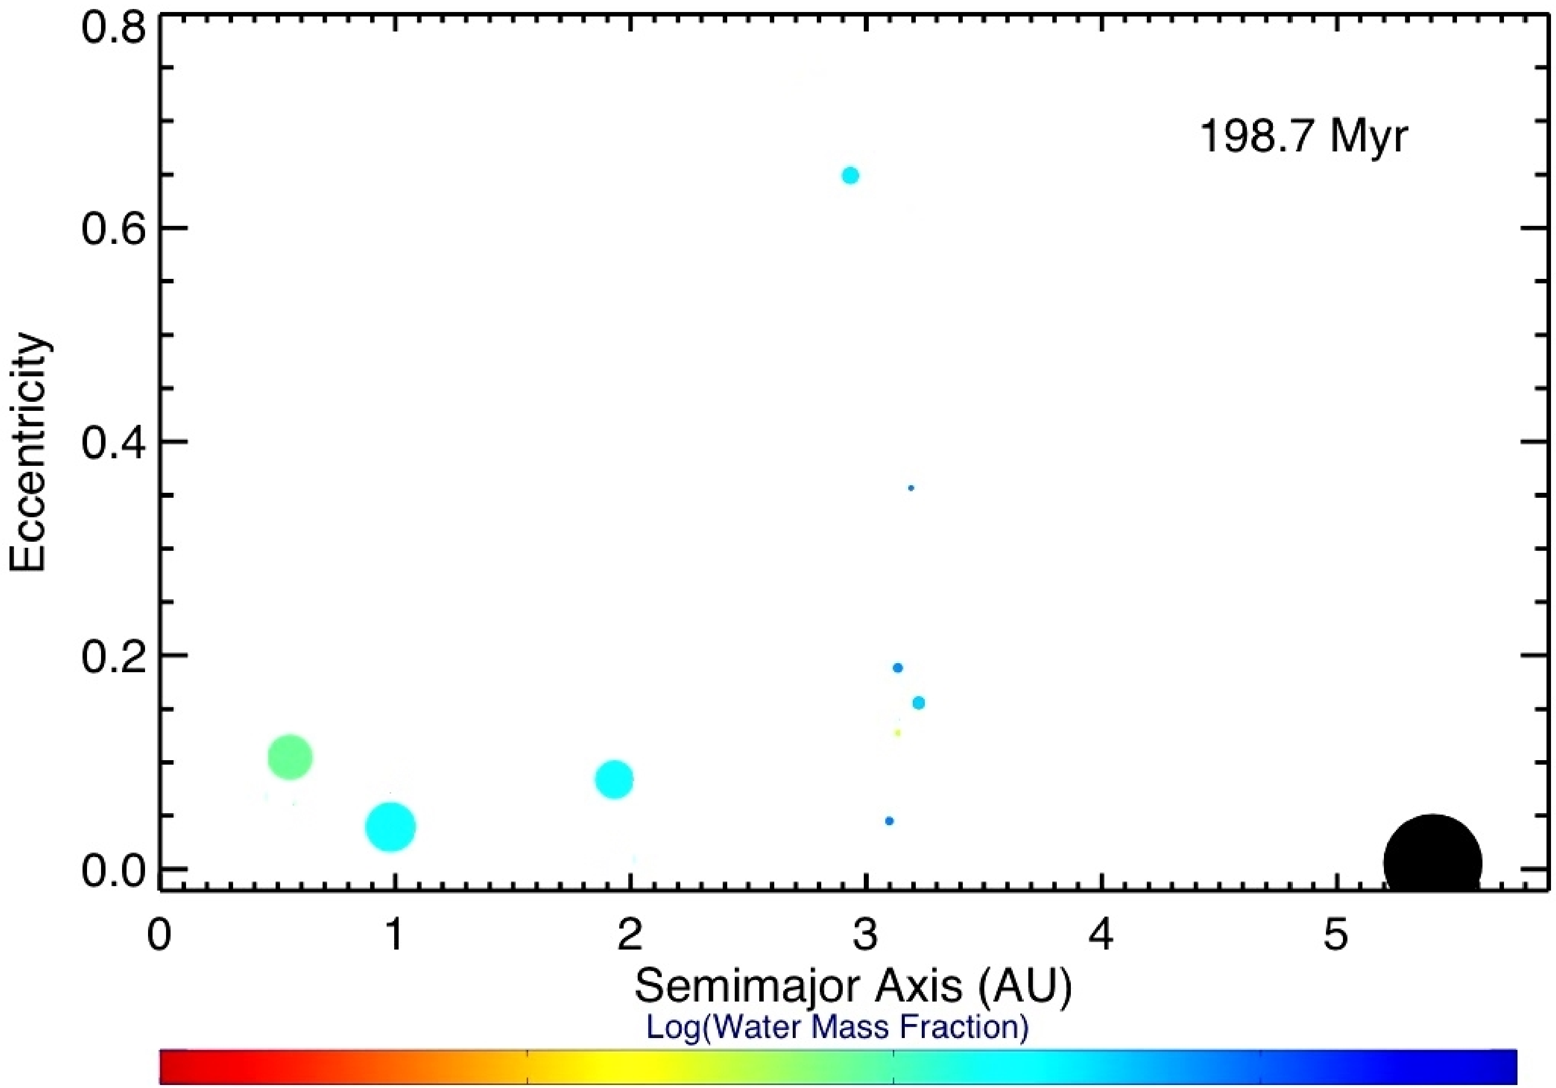

The N^2 problem

Planet formation simulations suffer from the N^2 problem i.e. the number of floating point operations increases with the number of particles N as N^2.

Since a x86 CPU core operates in sequence, the time taken to complete a simulation increases as N^2.

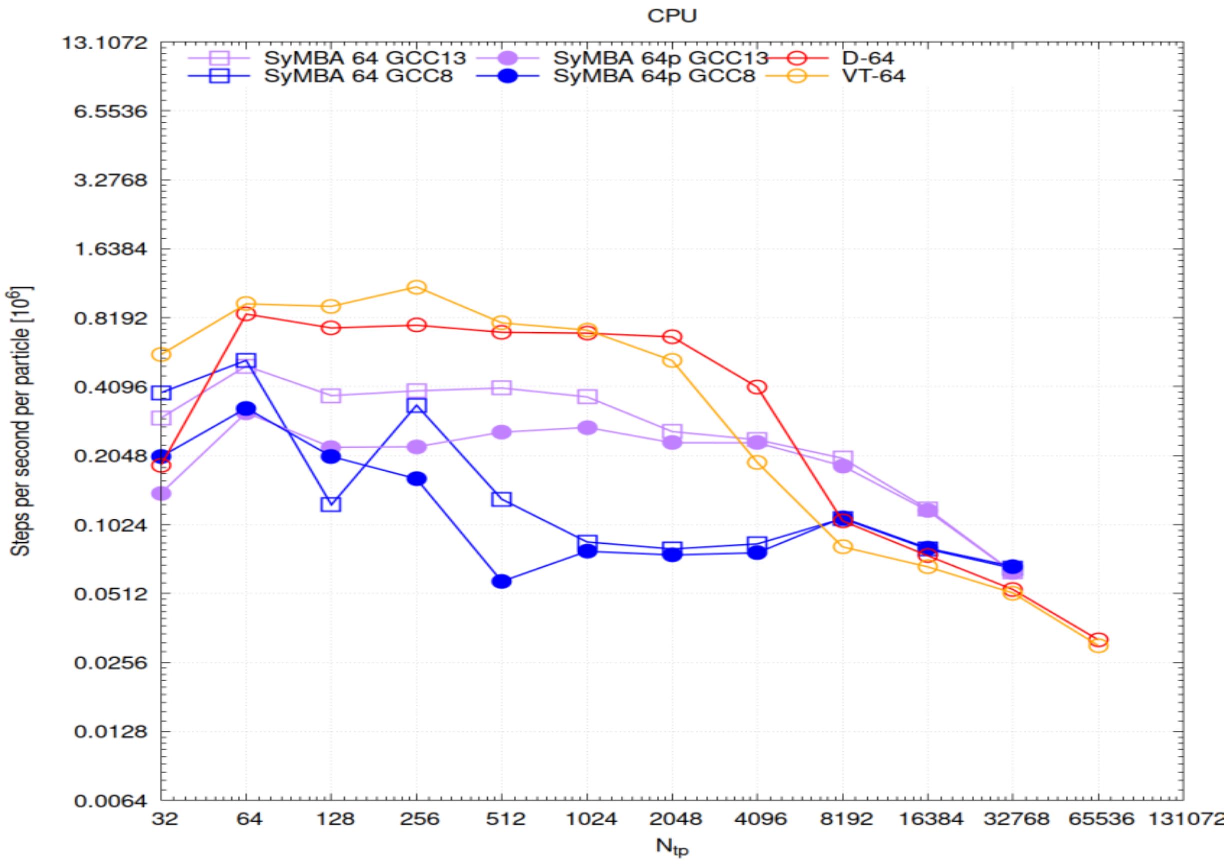

On the CPU the forces are computed as

```
do i=1, N-1
  do j=i+1, N
    !particles i and j interact
  enddo
enddo
```

On modern CPUs at 3 GHz the number of integrations per particle per second is about 4×10^5. Thus doing an integration for 10^{10} steps (100 Myr for the solar system) with 2000 particles would take 8 years assuming no particle loss.

We need something better!

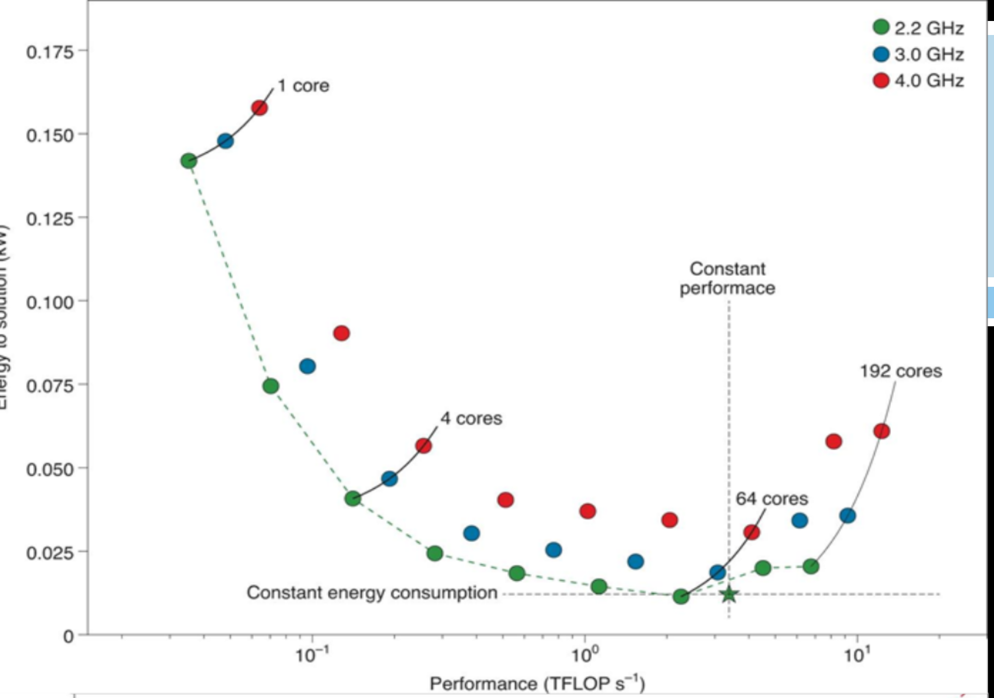

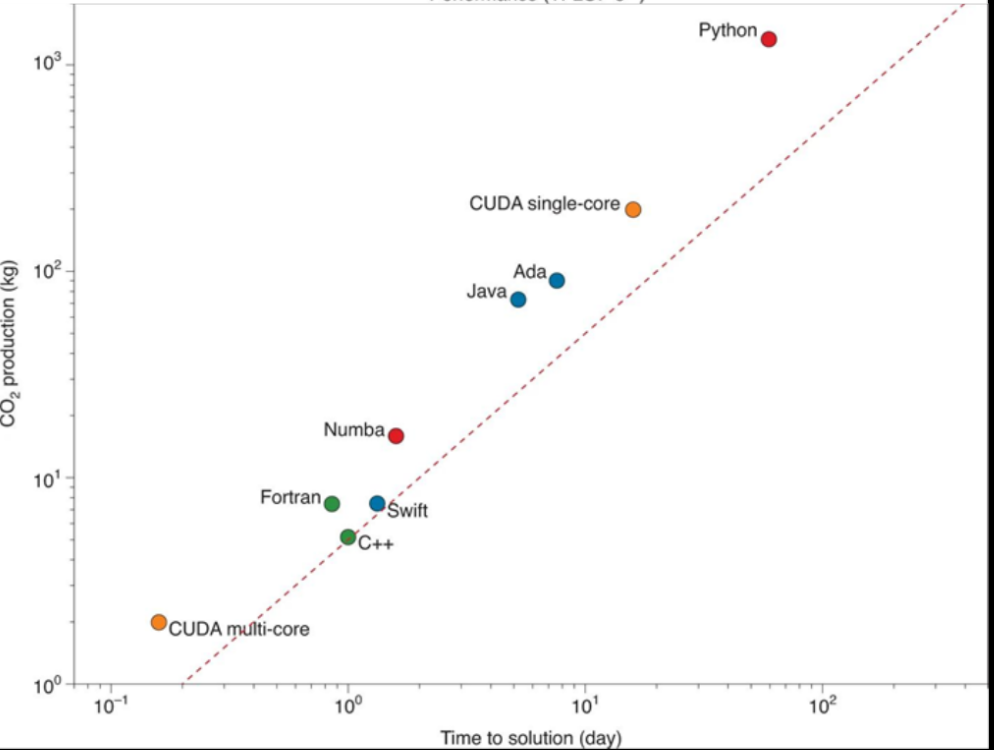

Parallelisation

The integration speed can be improved with parallelisation of the N-body code across multiple CPU cores e.g. with OpenMP. But because certain parts are ill suited to parallelisation and because of Amdahl's law the speedup eventually flattens off, typically around 32 to 64 CPU cores.

The code the N-body integrator is written in also plays a role, with basic Python being much slower than C++ and Fortran. But even this parallelisation will result in long integration times and massive CO_2 production from electricity usage.

Can we do even better?

Enter GPU computing

GPU computing offers a massively parallel infrastructure and enormous computational power. The Nvidia A100 GPU delivers a peak of 9.8 TFlops FP64 precision, while an AMD Epyc 64-core CPU has a peak of 3.5 TFlops in FP64. Also, performance across a CPU has decreasing returns for more than 8 or 16 CPU cores.

The best GPUs can perform a peak 2.5 M steps per particle per second, while this is 0.4 M for a CPU!

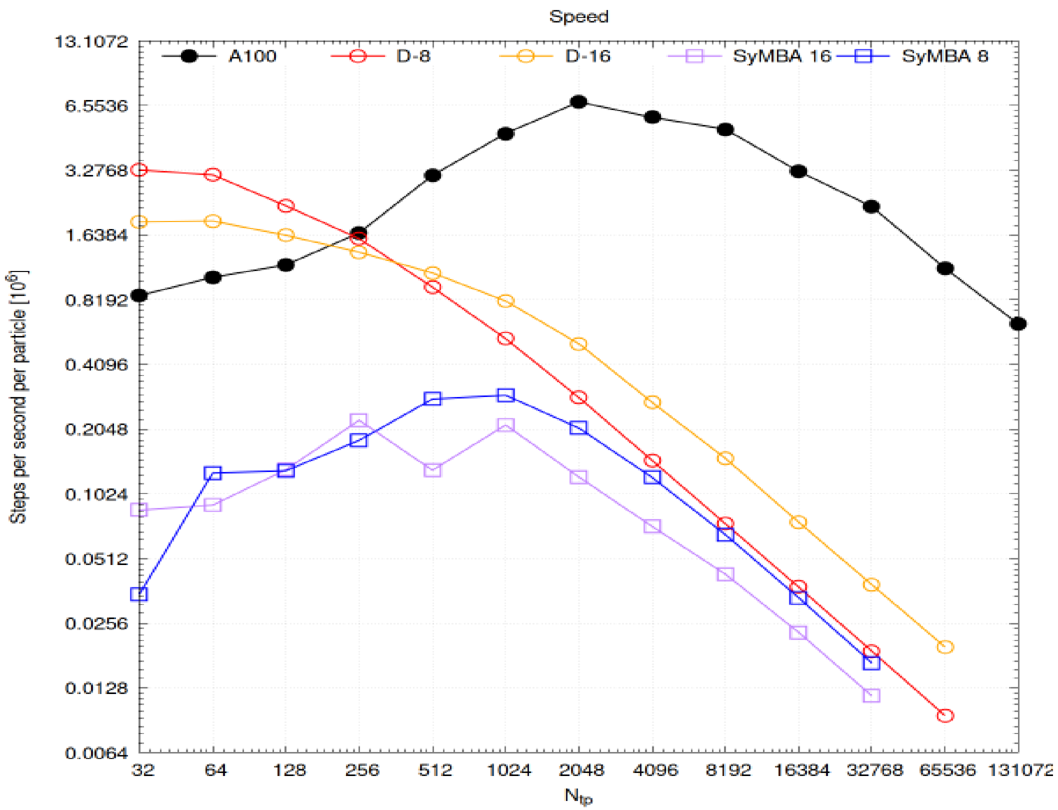

In reality, N-body simulations on the GPU are generally much faster than on the CPU. Using two different N-body codes the peak performance on the Nvidia A100 GPU is about an order of magnitude better than a 32 core Epyc CPU at 3 GHz.

As such, GPU computing is the future choice for planet formation simulations!

The GENGA GPU N-body code

- Open source GPU N-body integrator for planetary problems.
- https://bitbucket.org/sigrimm/genga/src/master/
- requires NVidia GPU to run; one simulation per GPU.
- can be easily compiled with nvcc compiler, CUDA 10+
- can be run on AMD GPUs with HIP
- a CPU-only version is in beta, as well as multi-GPU version
- **no special modules on HPC resources necessary.** Submit executable in SLURM script and you're golden!
- Publication: Grimm et al. (2022). GENGA. II. GPU Planetary N-body Simulations with Non-Newtonian Forces and High Number of Particles, The Astrophysical Journal 932, ID 124. DOI: 10.3847/1538-4357/ac6dd2

Genga CPU benchmarks on Discoverer

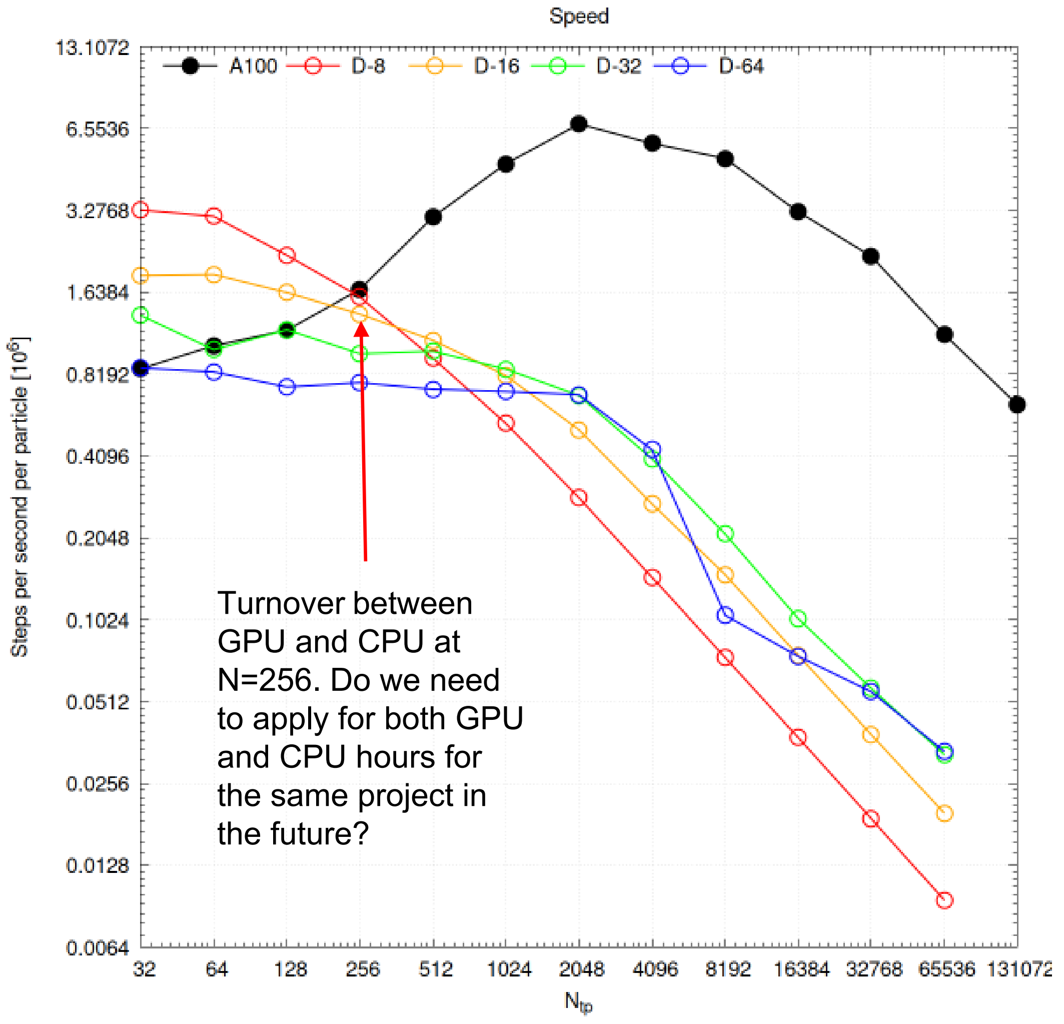

Classical model of terrestrial planet formation with GENGA

With GPU computing, things become really cool!
40k particles

Woo et al., 2021; 2022

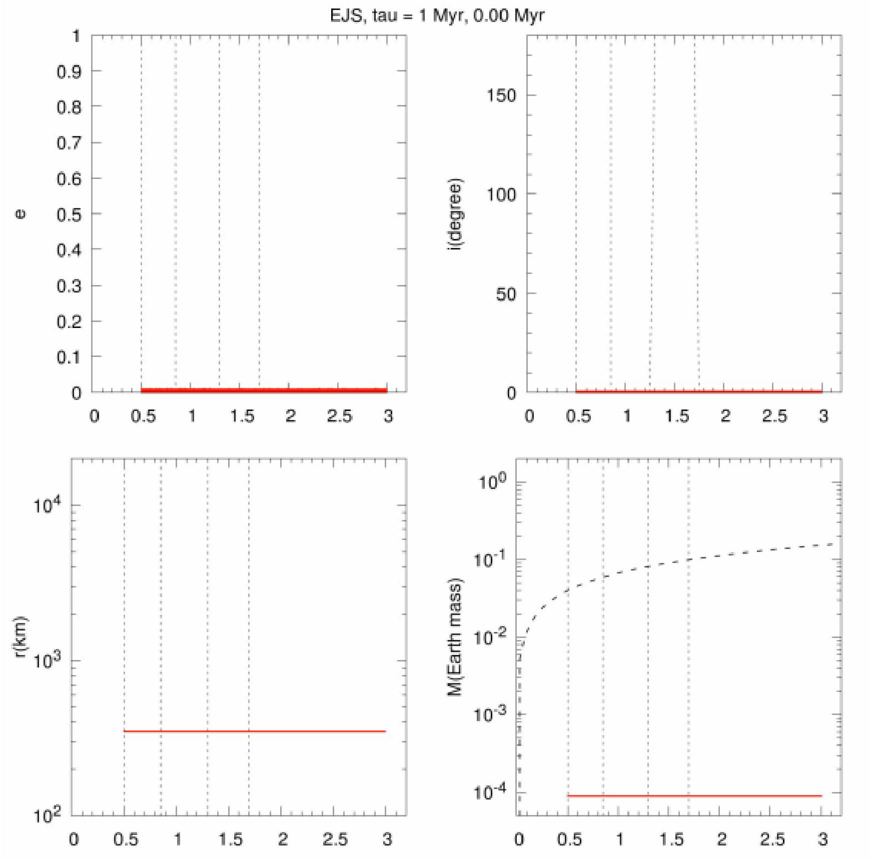

Jupiter and Saturn on their current orbits

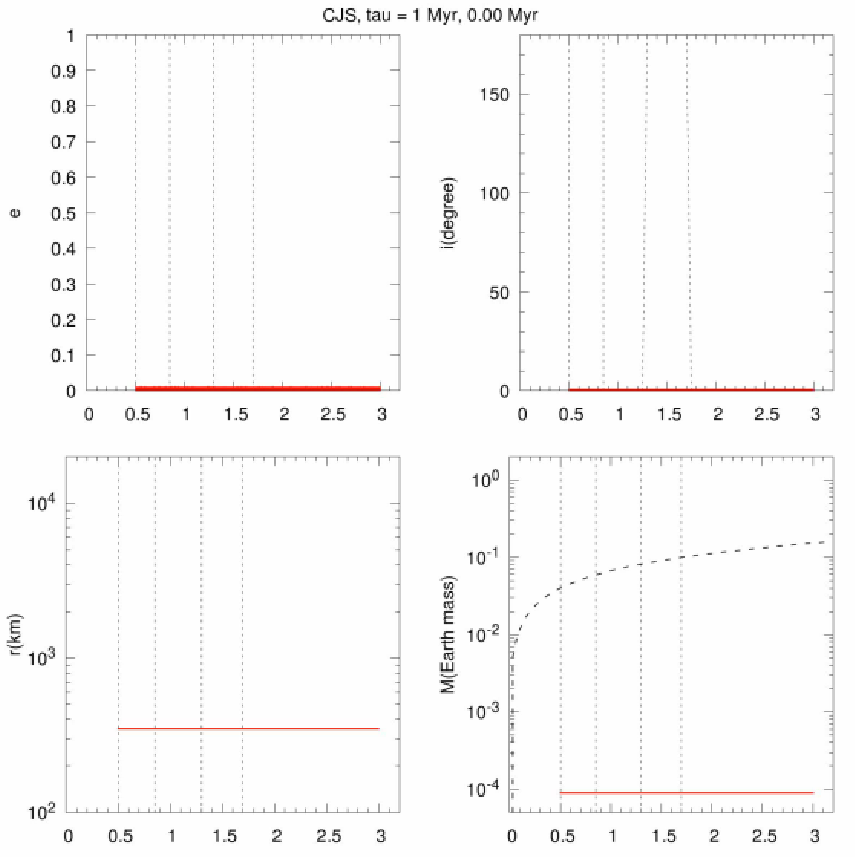

Jupiter and Saturn on circular orbits

Outcome of terrestrial planet systems from GPU runs

Jupiter and Saturn on their current orbits (EJS) seems to work, but not on circular orbits (CJS).

So we need giant planet migration

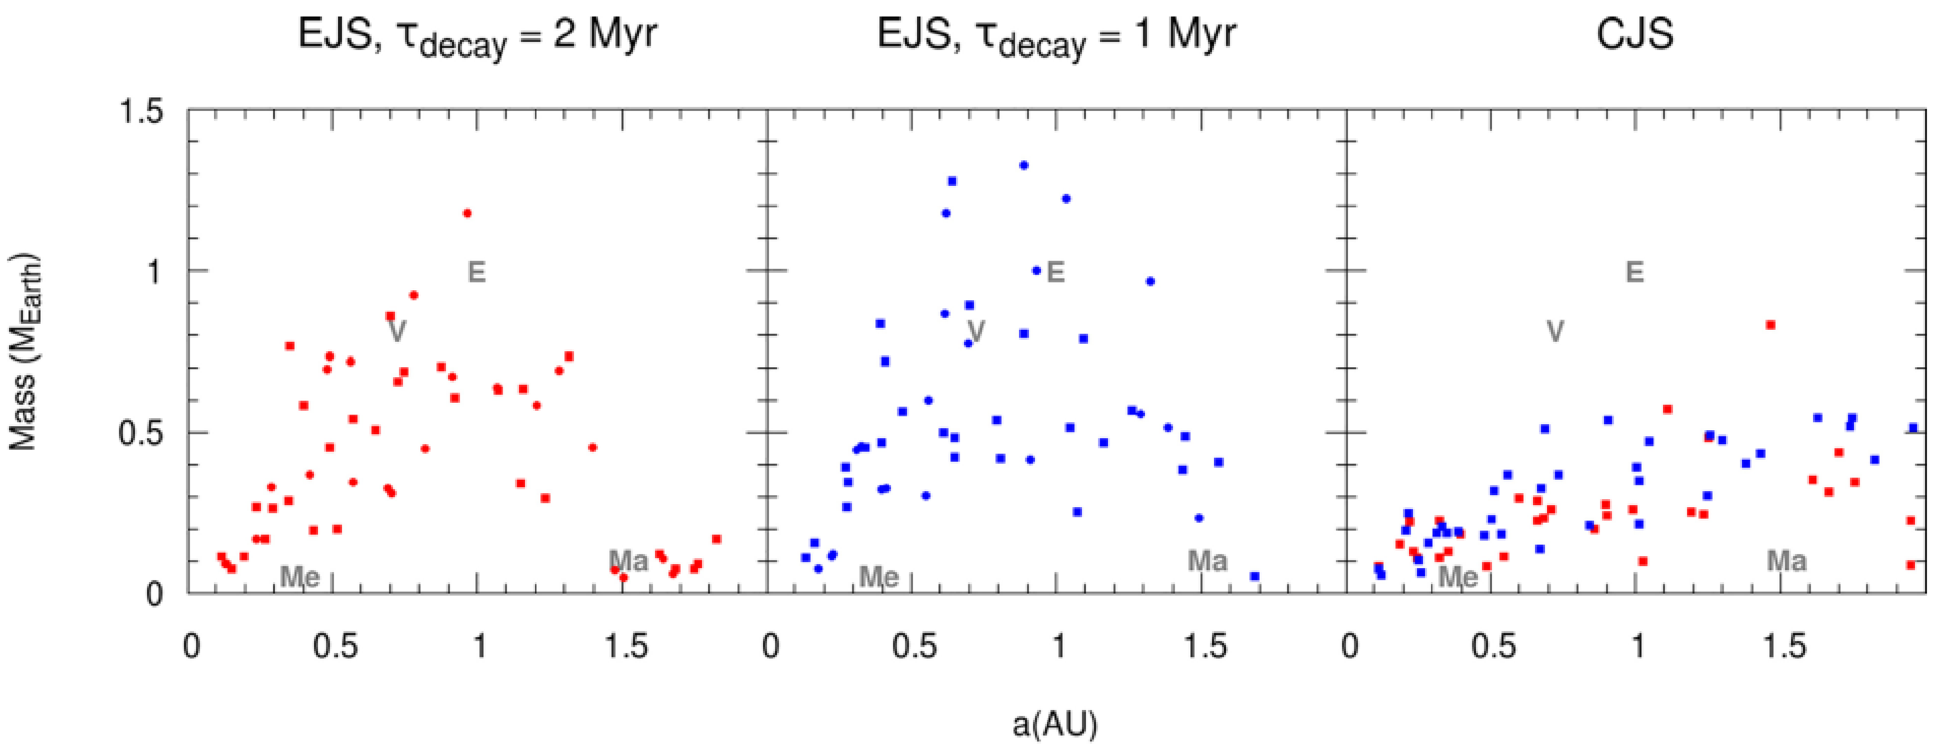

Exoplanet formation on Vega using A100 GPUs

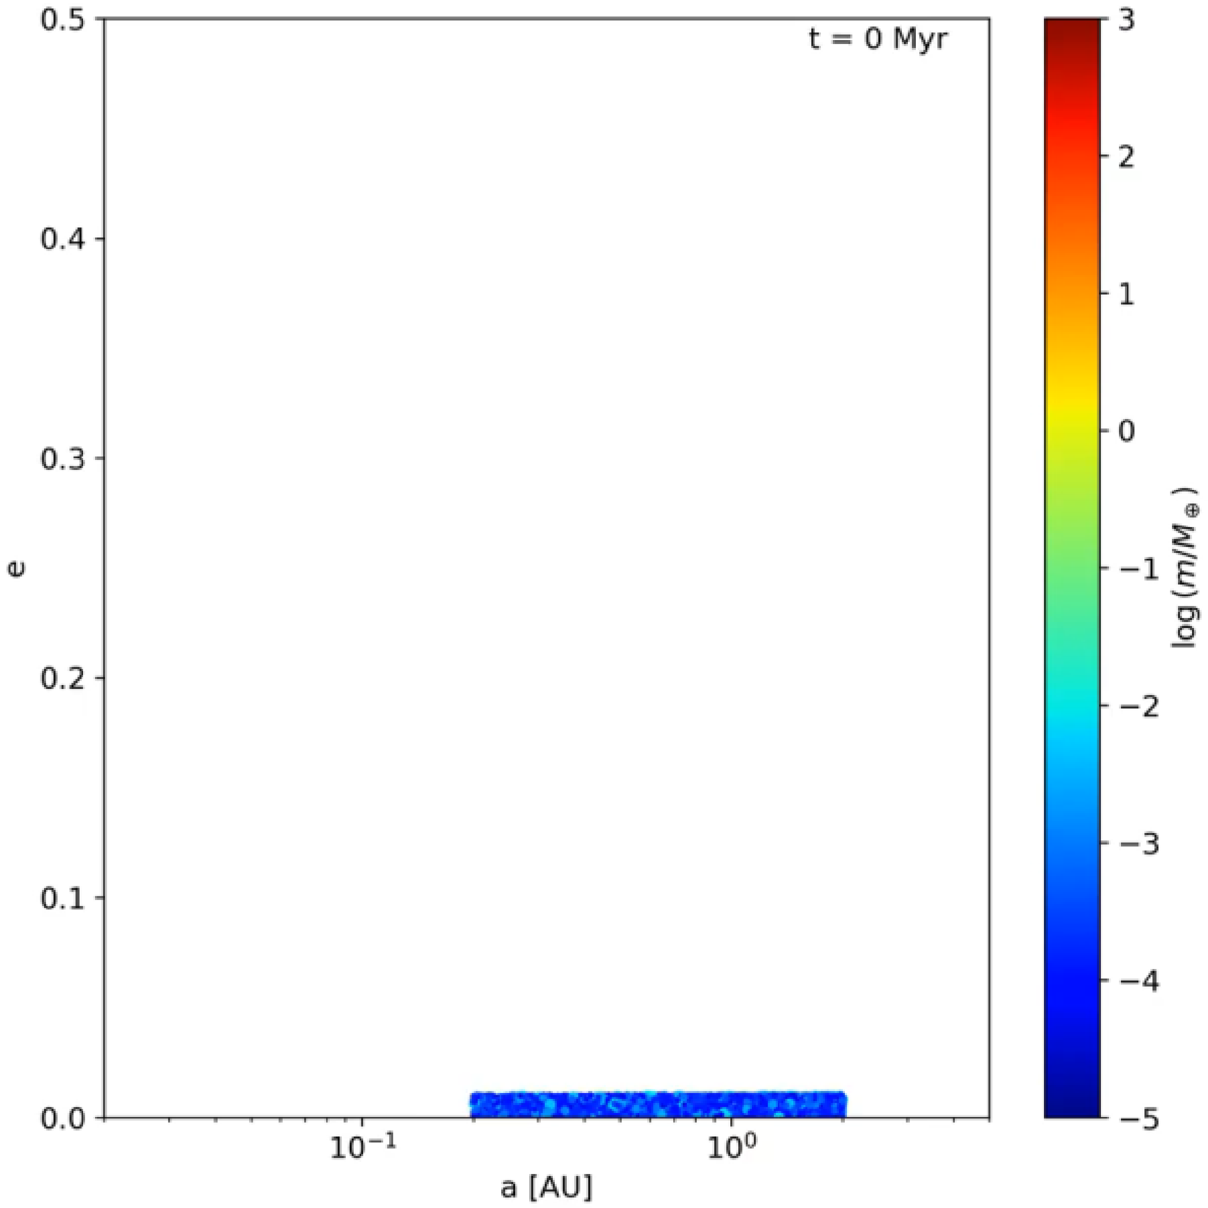

This simulation was run on Vega and took about 40 days. It starts with 28k bodies around a star with 0.7 solar masses, to attempt to form rocky close-in exoplanets.

This simulation is part of a study on rocky exoplanet formation, as well as whether giant impacts occur, and when.

Such impacts may be needed to generate conditions favourable for the origins of life on Earth-mass and larger planets.

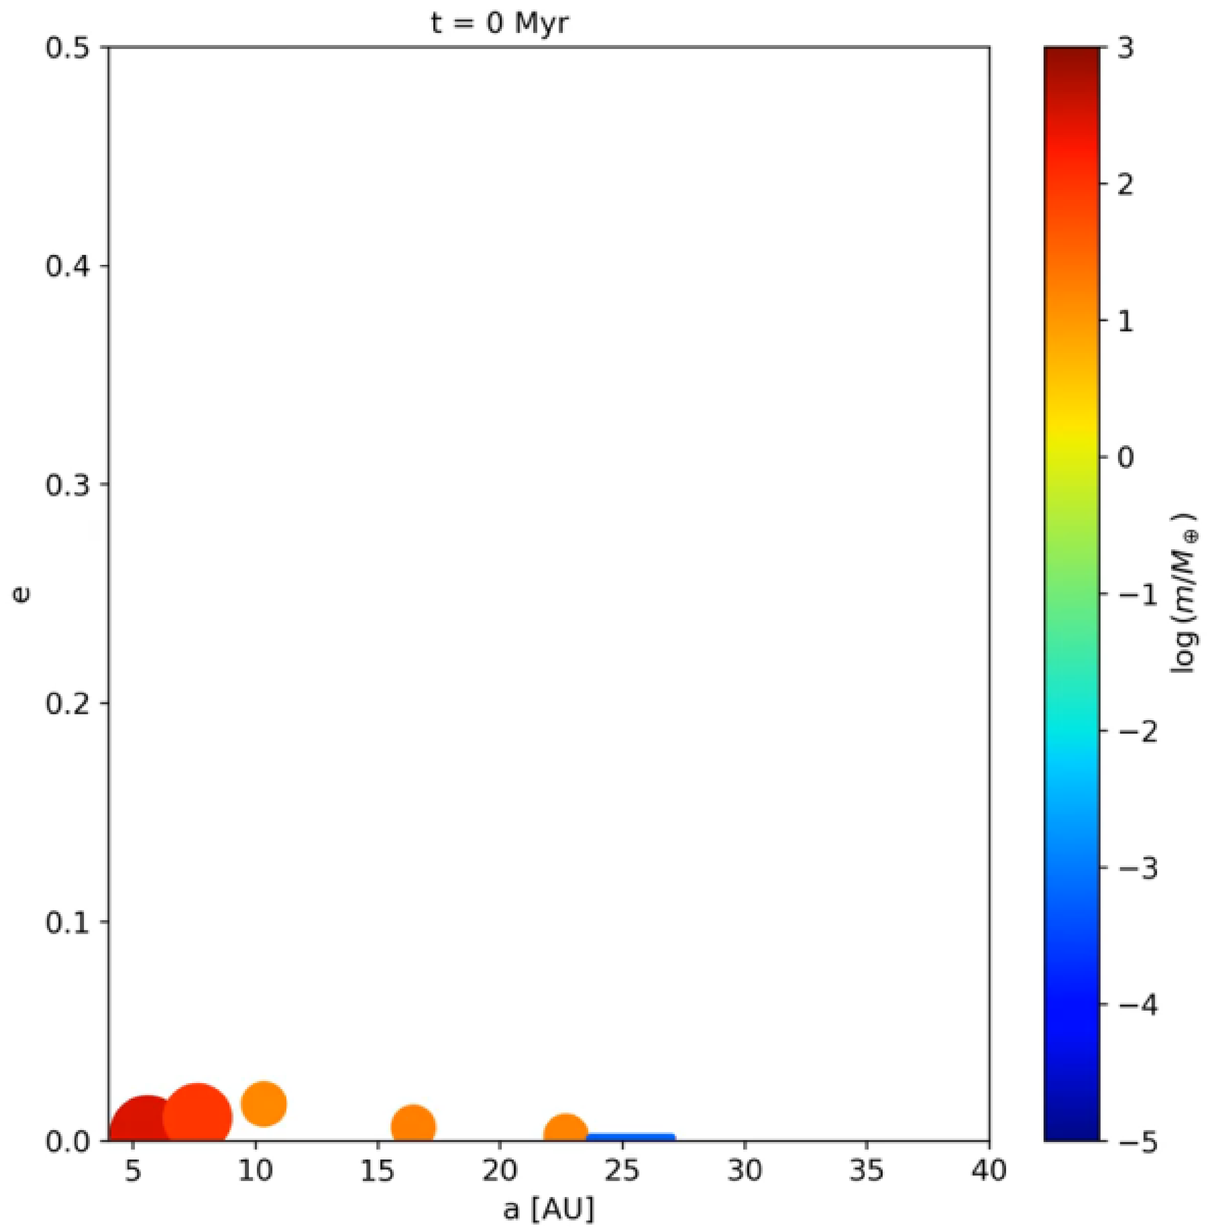

Giant planet migration on Vega

Evidence shows that the giant planets must have migrated. Here is a simulation of that, run on Vega as a test case for understanding how giant planet migration affected the formation of the terrestrial planets. Approximately 30k planetesimals, but these are not self-gravitating (TP2 mode). Fully-self-gravitating forthcoming.

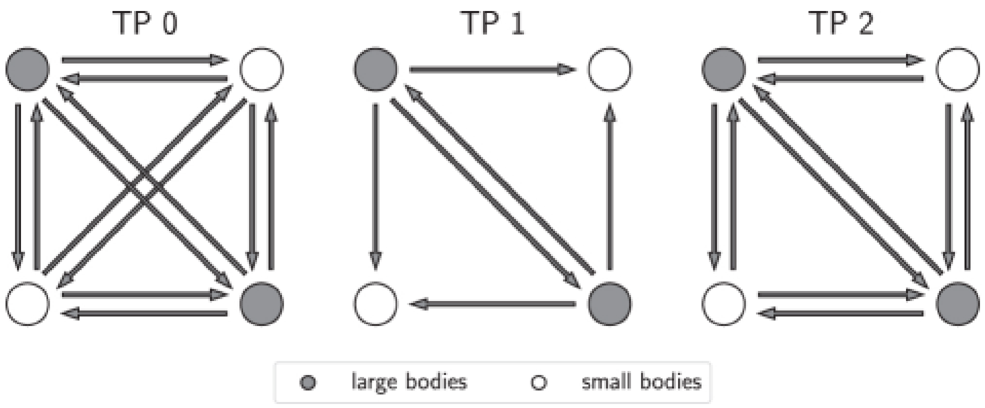

Terrestrial planet formation and giant planet migration on Vega

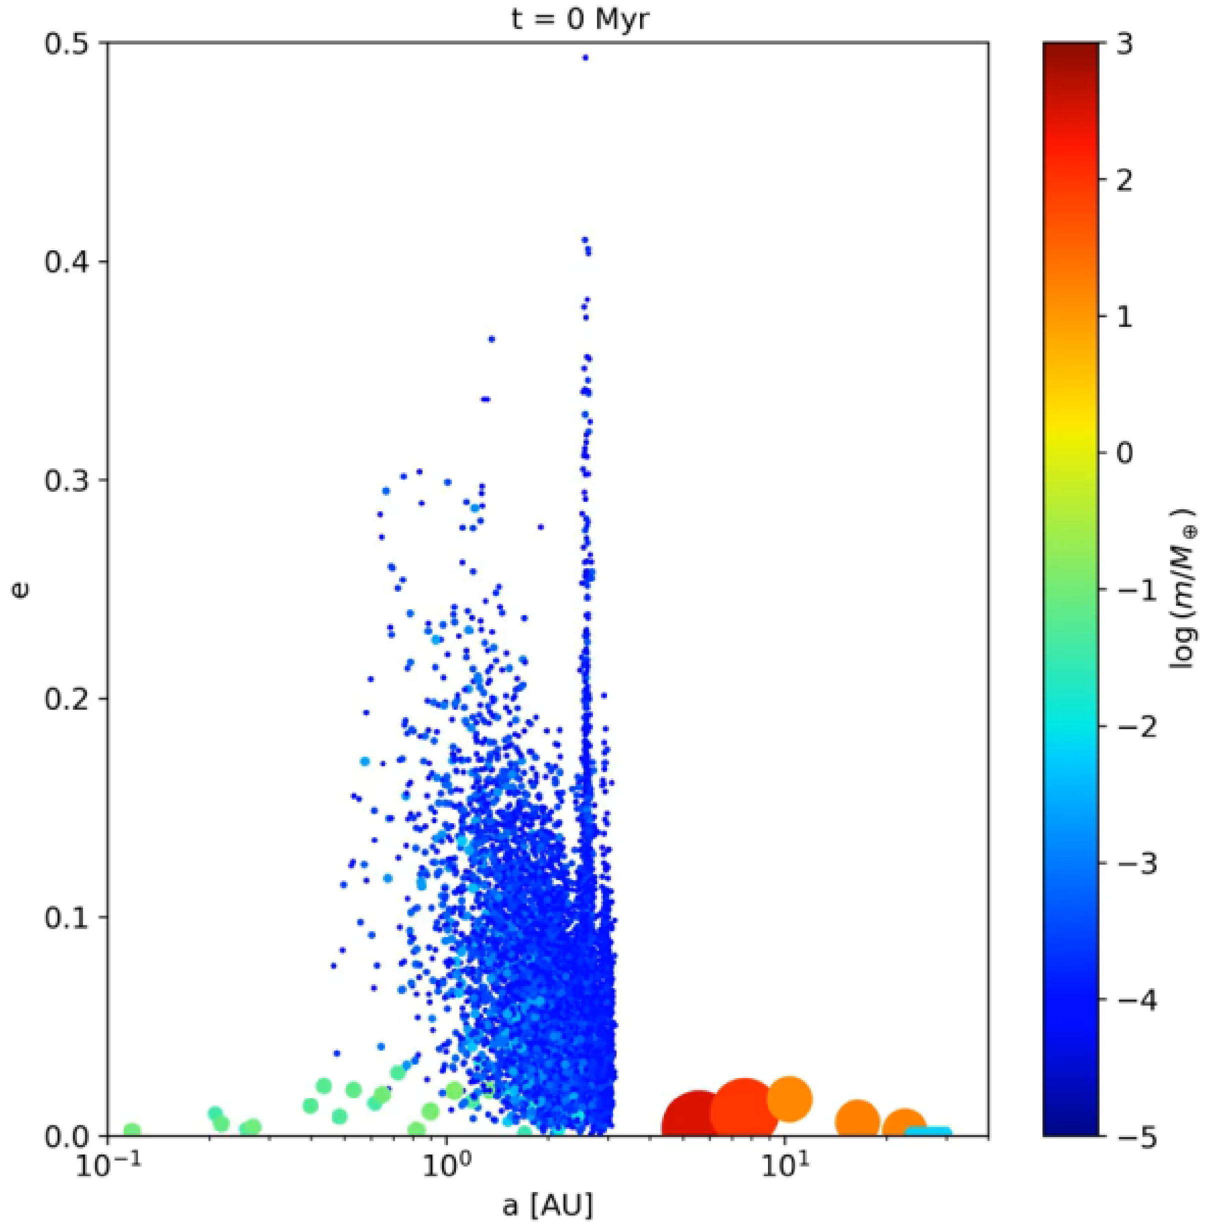

The giant planets probably migrated when the terrestrial planets were still forming. We might be able to use the formation of Venus and Earth and the timing of Earth's core formation to constrain when giant planet migration began in the solar system.

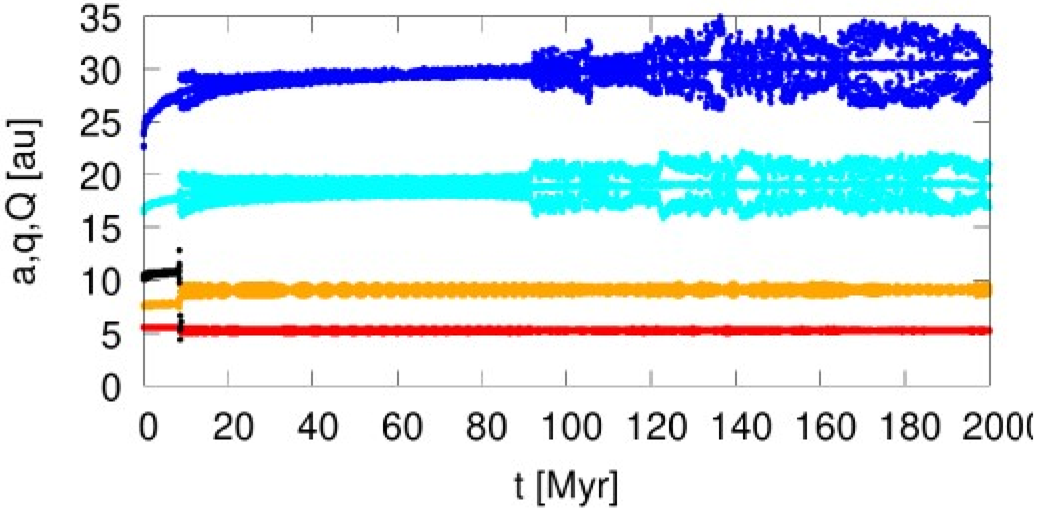

Conclusions and future work

- Planet formation is a thriving and dynamic field.
- N-body planet formation models are chaotic, so we need to run any simulations.
- Due to this and the N^2 problem we get the best results using GPU computing.
- This requirement implies that we have to use HPC computing to push the field forward.
- We gratefully acknowledge EuroHPC JU for awarding this project access to Vega hosted at IZUM, Slovenia.

Thank you for your attention!

Multi-scale microstructure modeling for corrosion and hydrogen embrittlement

EUROHPC USER DAY 2023 Brussels 11.12.23

Project: *Microstructure modeling*

EuroHPC used: LUMI-C

Speaker: Alexander Mavromaras, Materials Design

Executive Summary

- **Industry Challenges:** Corrosion and hydrogen embrittlement pose significant challenges across numerous industries
- **Mitigation Strategies** necessitate a comprehensive understanding and control of the material's microstructure
- **The Value of Simulations:** Multi-scale simulations insights and data where experiments are costly and time-consuming
- **Collaborative Approach:** Materials Design is developing a multi-scale modeling approach together with industrial and academic partners.
- **Case Study:** The evolution of microstructure during the oxidation of metals, such as zirconium alloys

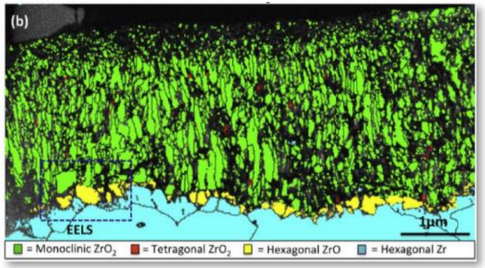

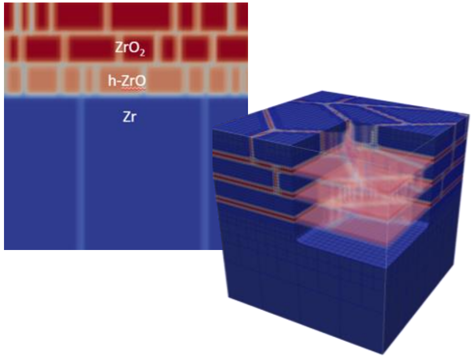

materials design

Overview

- **Company Snapshot**

  Materials Design - A Brief Overview of the Company and Its Technology.

- **Material Challenges**

  Exploring Corrosion and Hydride Formation in Zirconium Alloys.

- **Bridging the Gap**

  Scaling Up from Atomistic Simulations to Microstructure Modeling.

materials design

Materials Design Company Profile

Established Legacy:	Since 1998. Serving 700+ institutions worldwide.
Our Mission:	Creating Engineering Value from Materials Simulations.
Offerings:	MedeA® software, support, consulting, and contract research.
International Presence:	Headquartered in San Diego, USA, and Paris, Europe.
Partnerships:	Collaborating with technology and business partners globally.
Expertise:	Computational materials science, chemistry, chemical engineering.

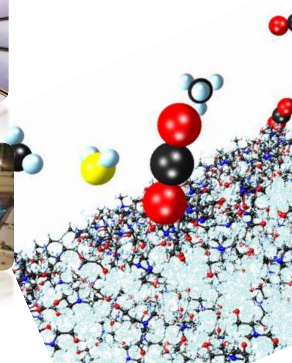

Improved Decision-Making Through Data And Understanding

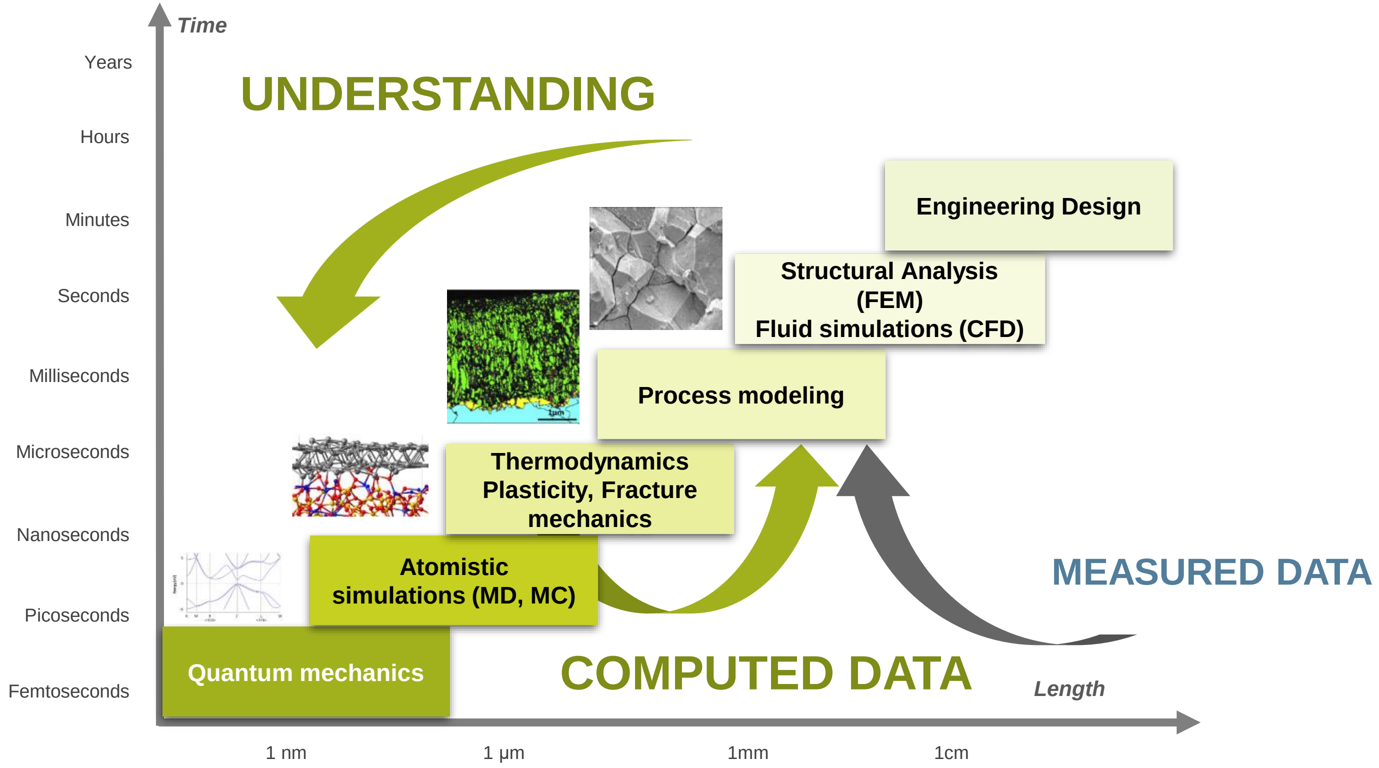

materials design

MedeA Environment

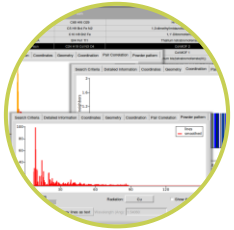

Databases

Experimental structures and properties,
ICSD, Pearson,, COD, Eureka

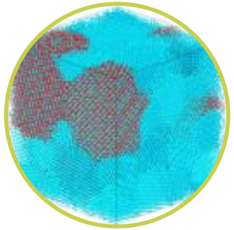

Builders

Models: molecular, fluid, solid,
nano, micro, …

Engines

VASP, LAMMPS, GAUSSIAN, GIBBS, MOPAC …

Automation

Properties: thermal, mechanical, transport, phase stability, …

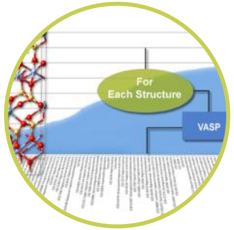

High Throughput

Workflows, structure lists,
Job control,
local, cloud

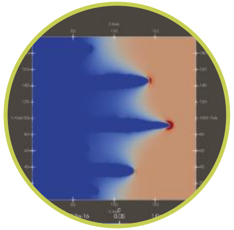

Analysis

Structure-property, graphical,
machine learning

Combines experimental data, builders and compute engines with automation high-throughput, and analysis

materials design

Corrosion

The cost of corrosion is estimated at 3.4% of the global GDP

- Water-cooled nuclear reactors: Corrosion of cladding of fuel rods is a life-limiting degradation mechanism
- Oil & Gas: Corroding pipelines drive up costs
- Automotive: Poor corrosion resistance of Magnesium

  alloys prevent lighter cars

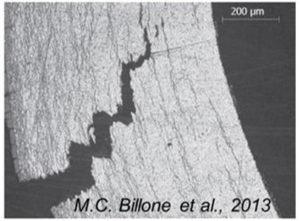

Can we simulate corrosive processes over days/months/years and make predictions?

Zirconium Corrosion And Hydride Formation In Reactors

- Water reacts with Zr cladding of fuel rods, freeing hydrogen

Corrosion reaction:

$$Zr + 2H_2O \rightarrow ZrO_2 + 2H_2$$

- Formation of Zr hydride
- Hydrogen embrittlement and mechanical weakening

How does the hydride microstructure look for a given H influx, temperature, and material?

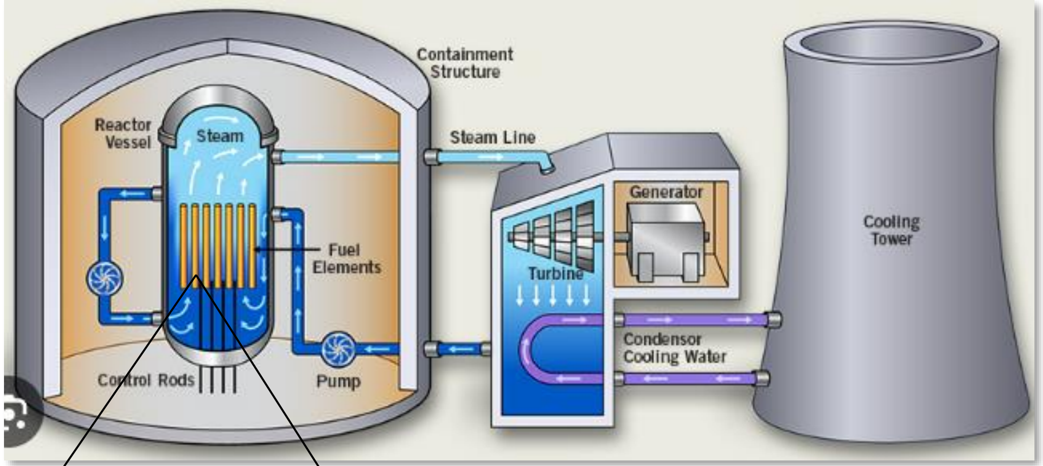

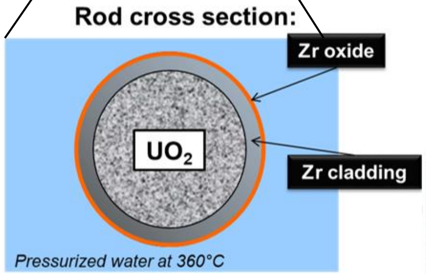

Hydrided Zr cladding

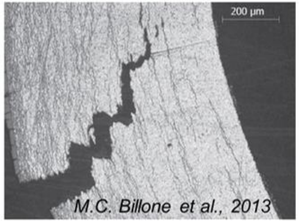

[Modeling Corrosion of Zirconium Alloys Fuel Cladding – The UW-Madison Materials Degradation under COrrosion and Radiation (MADCOR) – UW–Madison (wisc.edu)]

materials design

Hydride growth

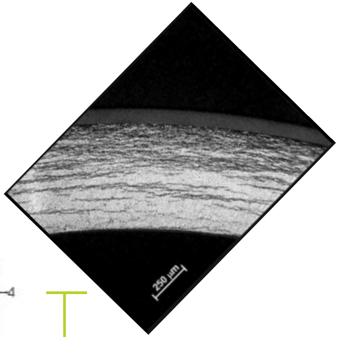

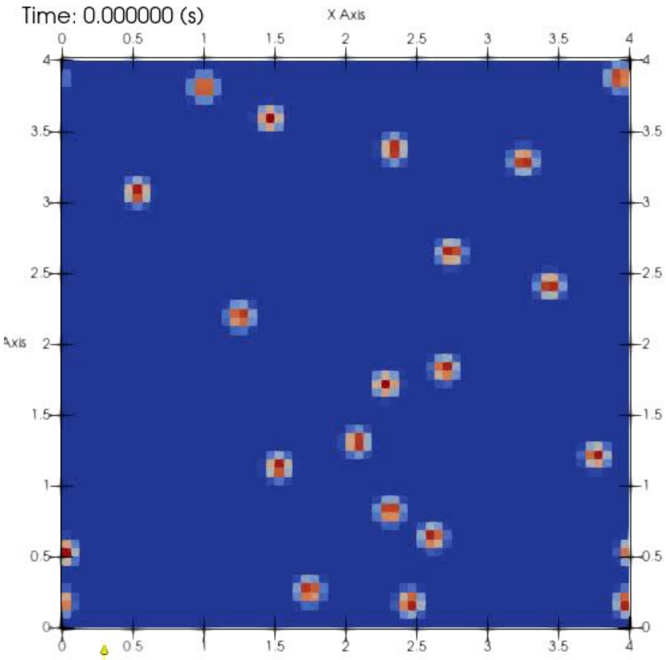

4µm

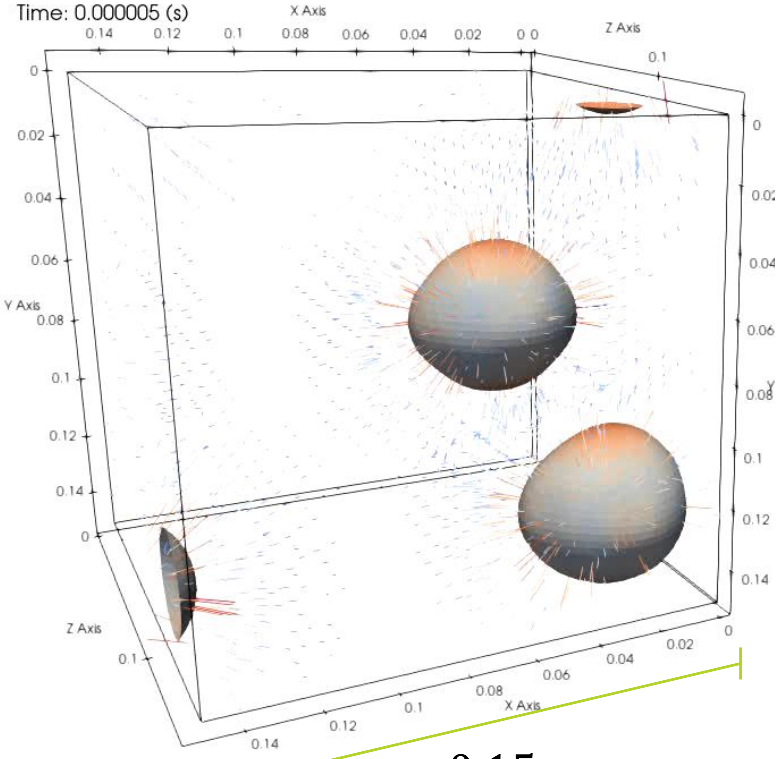

0.15µm

Zr and hydride phase evolution (for 0.07s at 600K) shows effect of elastic anisotropy

3D simulation with H fluxes and elastic energy at interfaces

materials design

From Nanoseconds To Years

Atomistic simulations

Machine Learning Potentials

Microstructure evolution

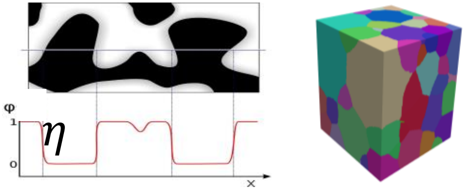

$$F = \int dV\,[f_{chemical} + f_{interface} + felastic]$$

$$\frac{\partial c}{\partial t} = \nabla\,\mathrm{M}(\mathrm{c},\eta,\sigma)\nabla\frac{\delta F}{\delta c}$$
$$\frac{\partial \eta}{\partial t} = -L\nabla\frac{\delta F}{\delta \eta}$$

Diffusion

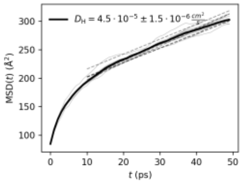

Free energy

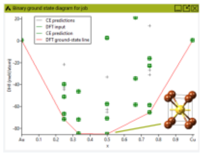

Interfaces

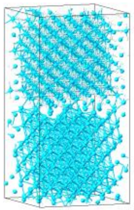

Elastic energy

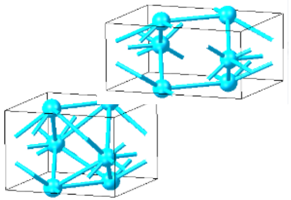

materials design

Summary

- Safety, economy, and sustainability of primary energy generation are major challenges of this century
- Understanding corrosion and material degradation requires multi-scale modeling
- Multi-scale simulations require
  - Accurate and efficient computations of chemical, thermo-mechanical, interfacial, and transport properties
  - Machine-learned potential for scale-up and complexity
  - Realistic 3D Microstructure simulations with the above input
  - Significant parallel computing resources

Phase field

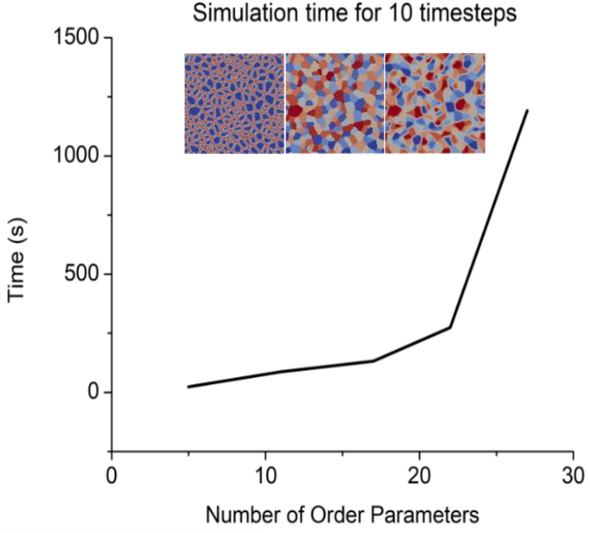

materials design

More Examples

Crack formation in α-Zr – m-ZrO_2 and Mg – MgO

Cracking In Materials With Different Pilling-Bedworth Ratios

Zirconium oxide PBR: 1.56

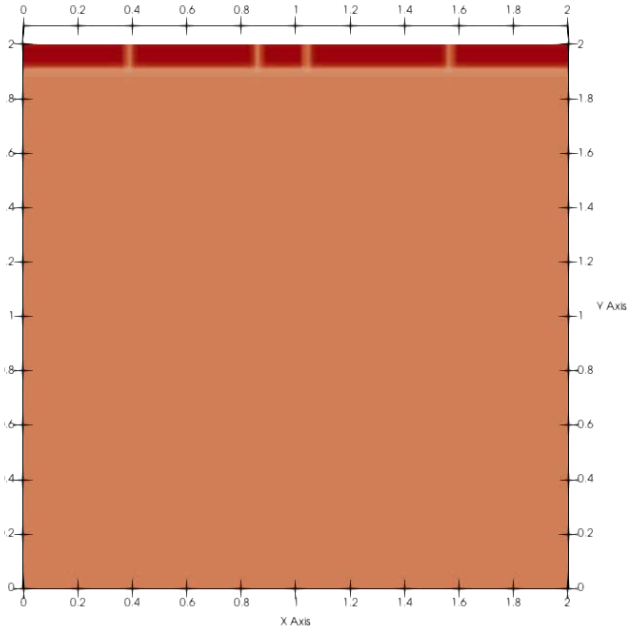

Cracks forming in a α-Zr – m-ZrO_2 multi-grain simulation at high O flux

Magnesium oxide PBR: 0.81

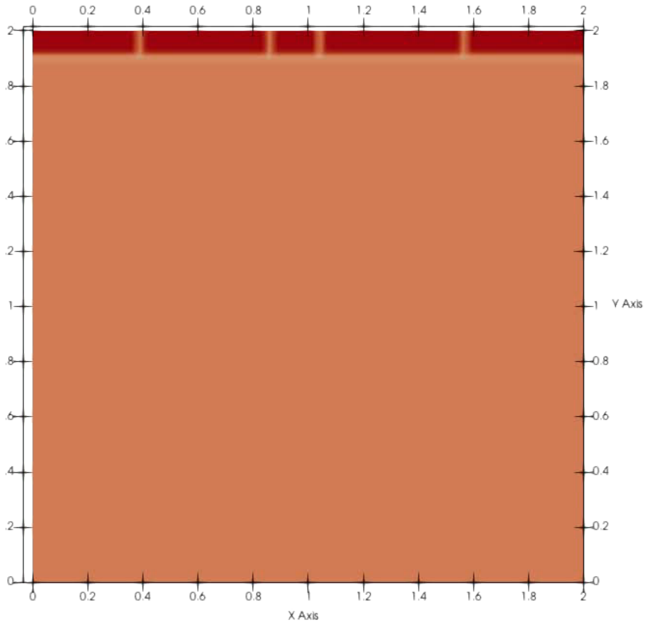

Cracks forming in a Mg – MgO multi-grain simulation at high O flux

$$PBR = \frac{V_{oxide}}{V_{metal}}$$

materials design

Modelling the formation of multiple stellar populations in globular clusters

EUROHPC USER DAY 2023 Brussels 11.12.23

EuroHPC Joint Undertaking

Project: "EHPC-REG-2021R0052"

EuroHPC used: Discoverer

Speaker: *Elena LACCHIN - University of Padua*

Globular clusters

NGC 1866

Credit: NASA / ESA / Hubble

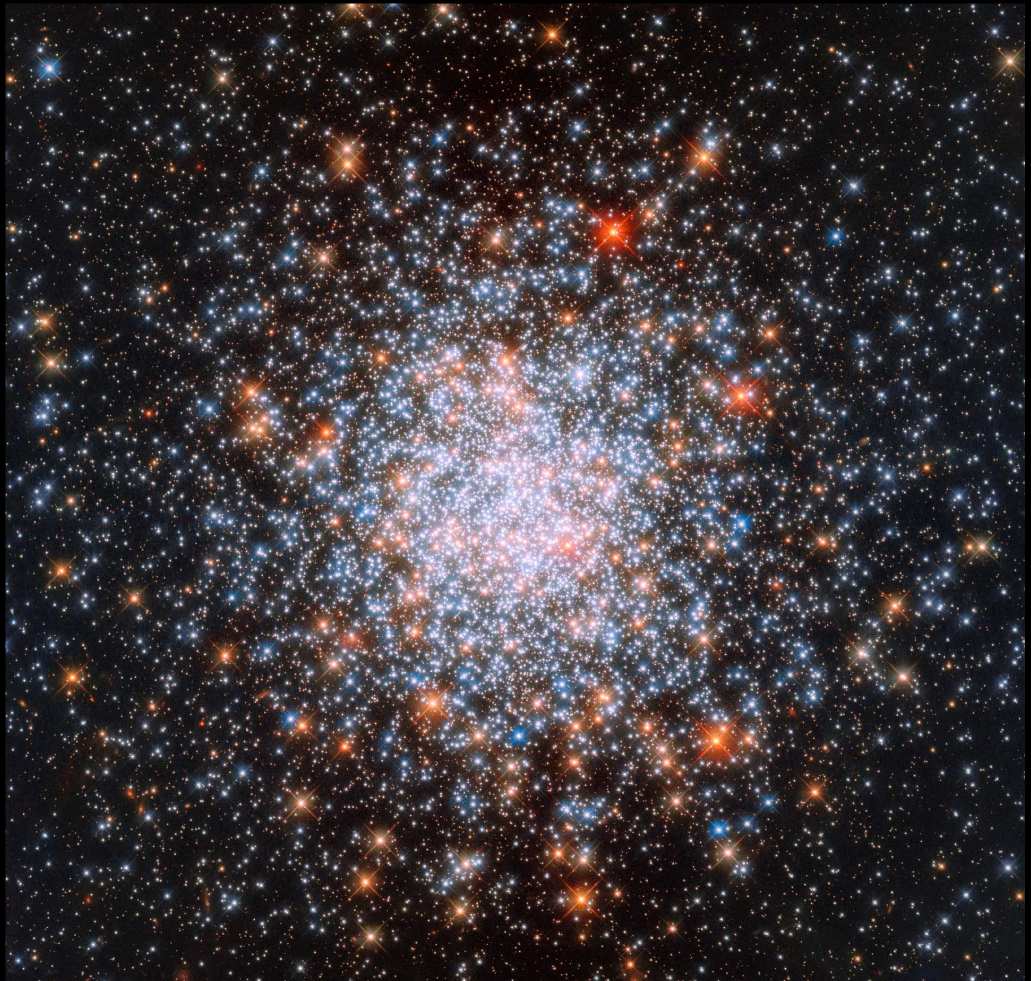

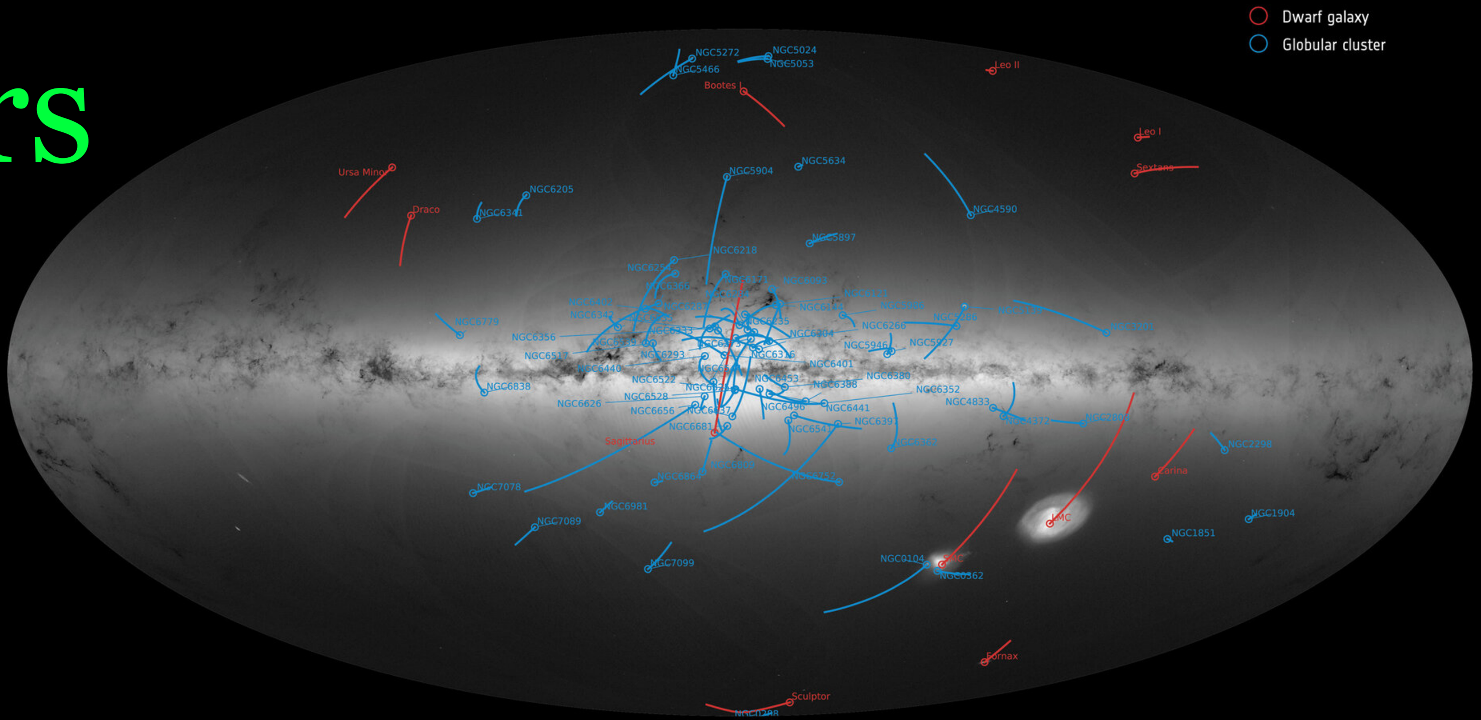

- ★ Gravitationally bound stars
- ★ All galaxies (in bulges, disks and halos)
- ★ $N_\star = 10^5 - 10^6$
- ★ Age $\sim 10 - 12$ Gyr
- ★ Devoid of gas
- ★ $N_{GC}^{MW} \sim 150$

Why globular clusters are so important ?

Validate the age of the Universe

Constrain dark matter properties

Understand star formation in the early and dense Universe

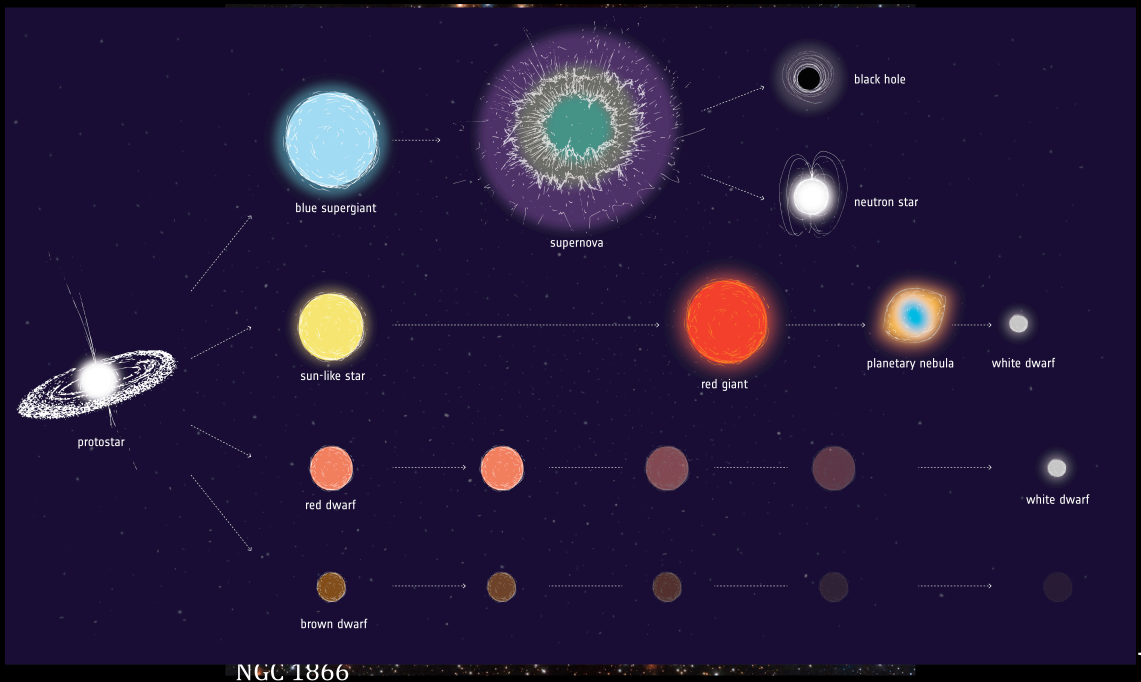

Constrain stellar evolution and nucleosinthetic models

May played a role in cosmic reionization and galaxy formation

Reservoir of exotic objects, e.g. binary black holes

Reconstuct the assembly of the Milky Way

Globulars host multiple stellar populations!

From spectroscopy…

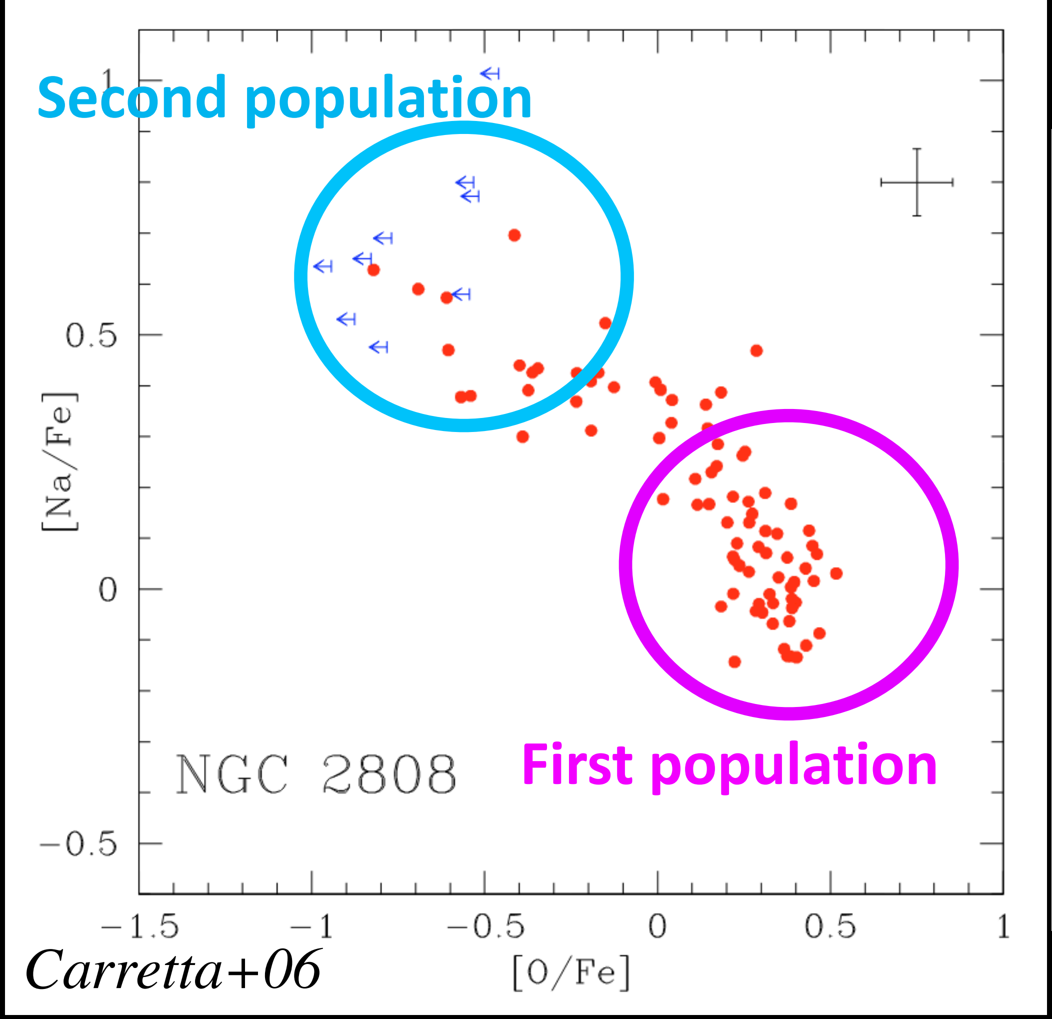

…and photometry.

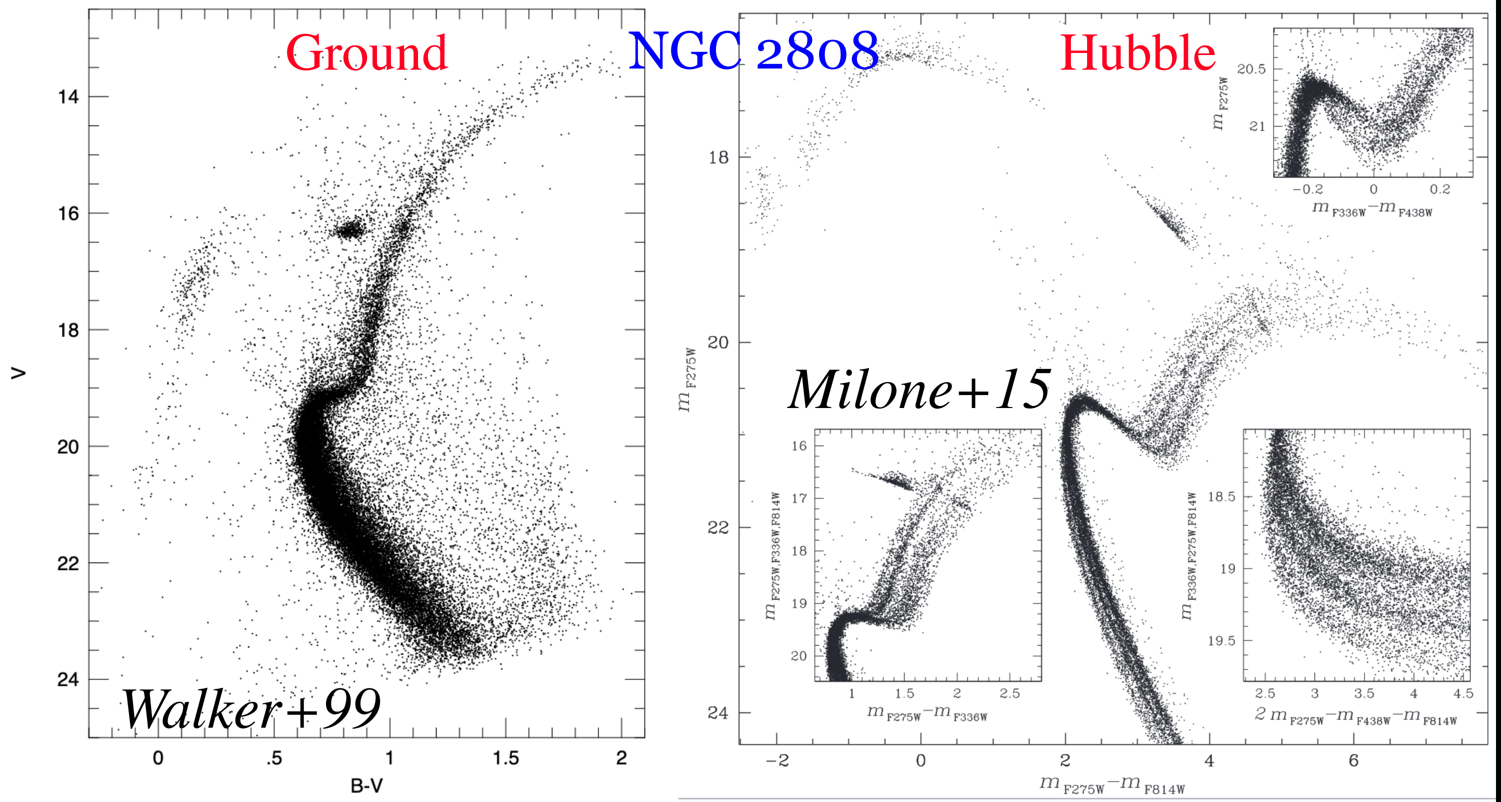

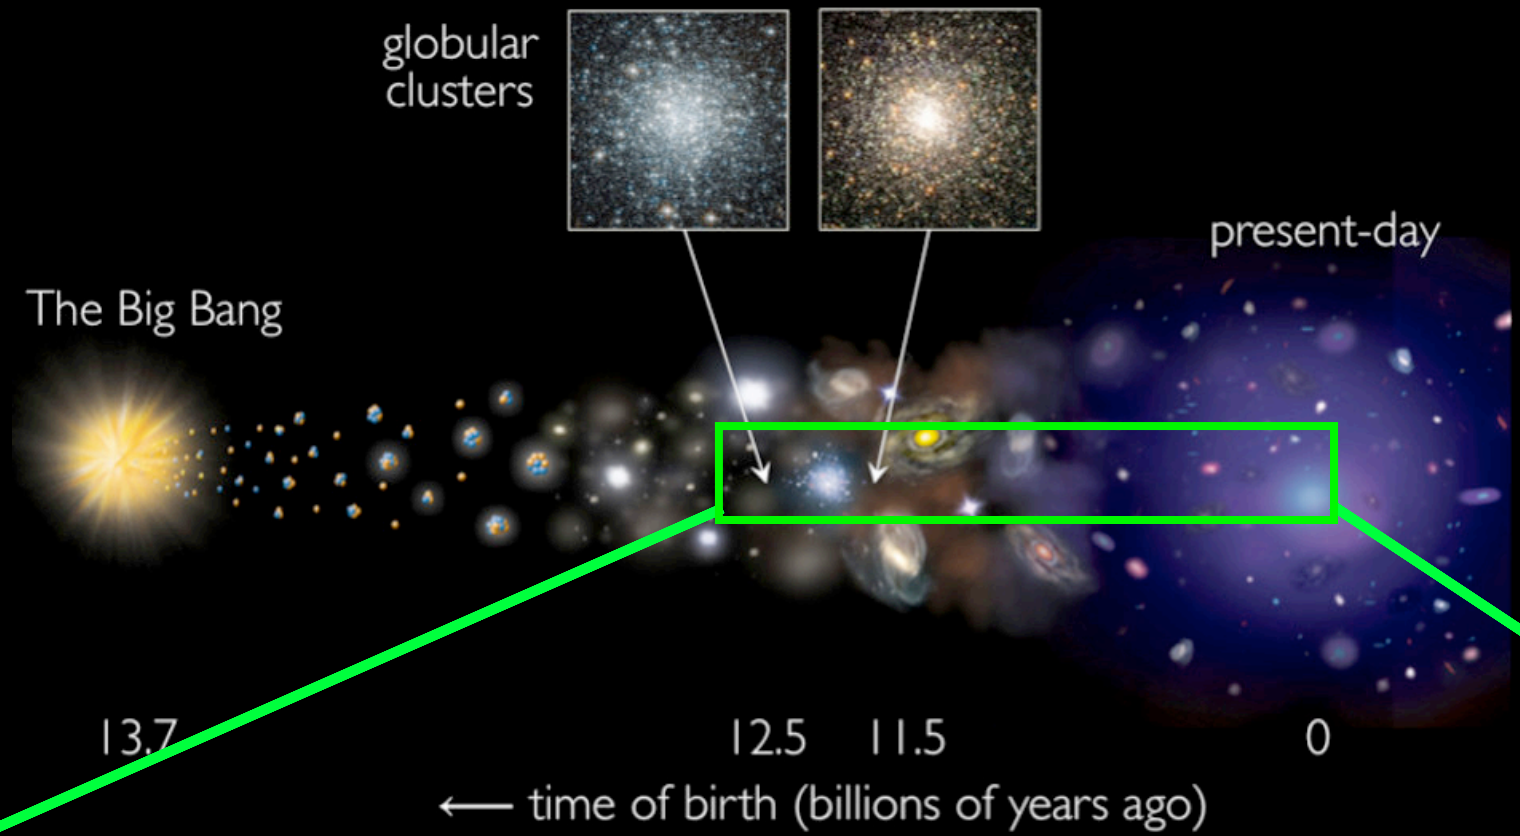

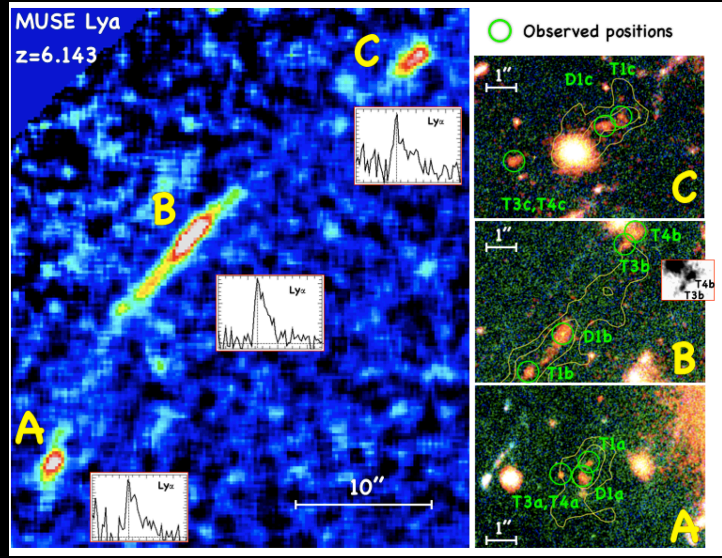

Vanzella+19

JWST and ELT will tell us more !

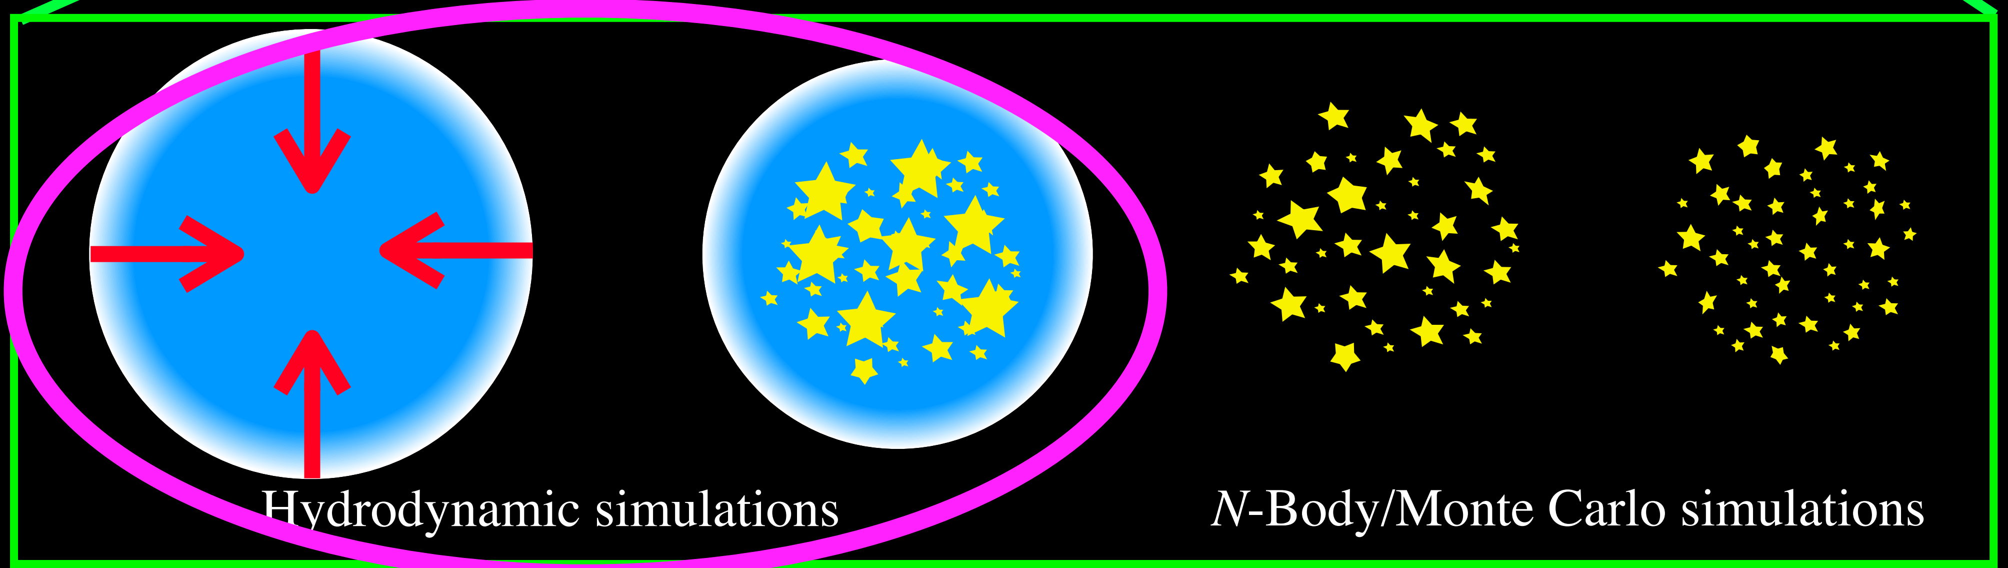

Modelling a realistic physical environment for the formation of MPs in globular clusters

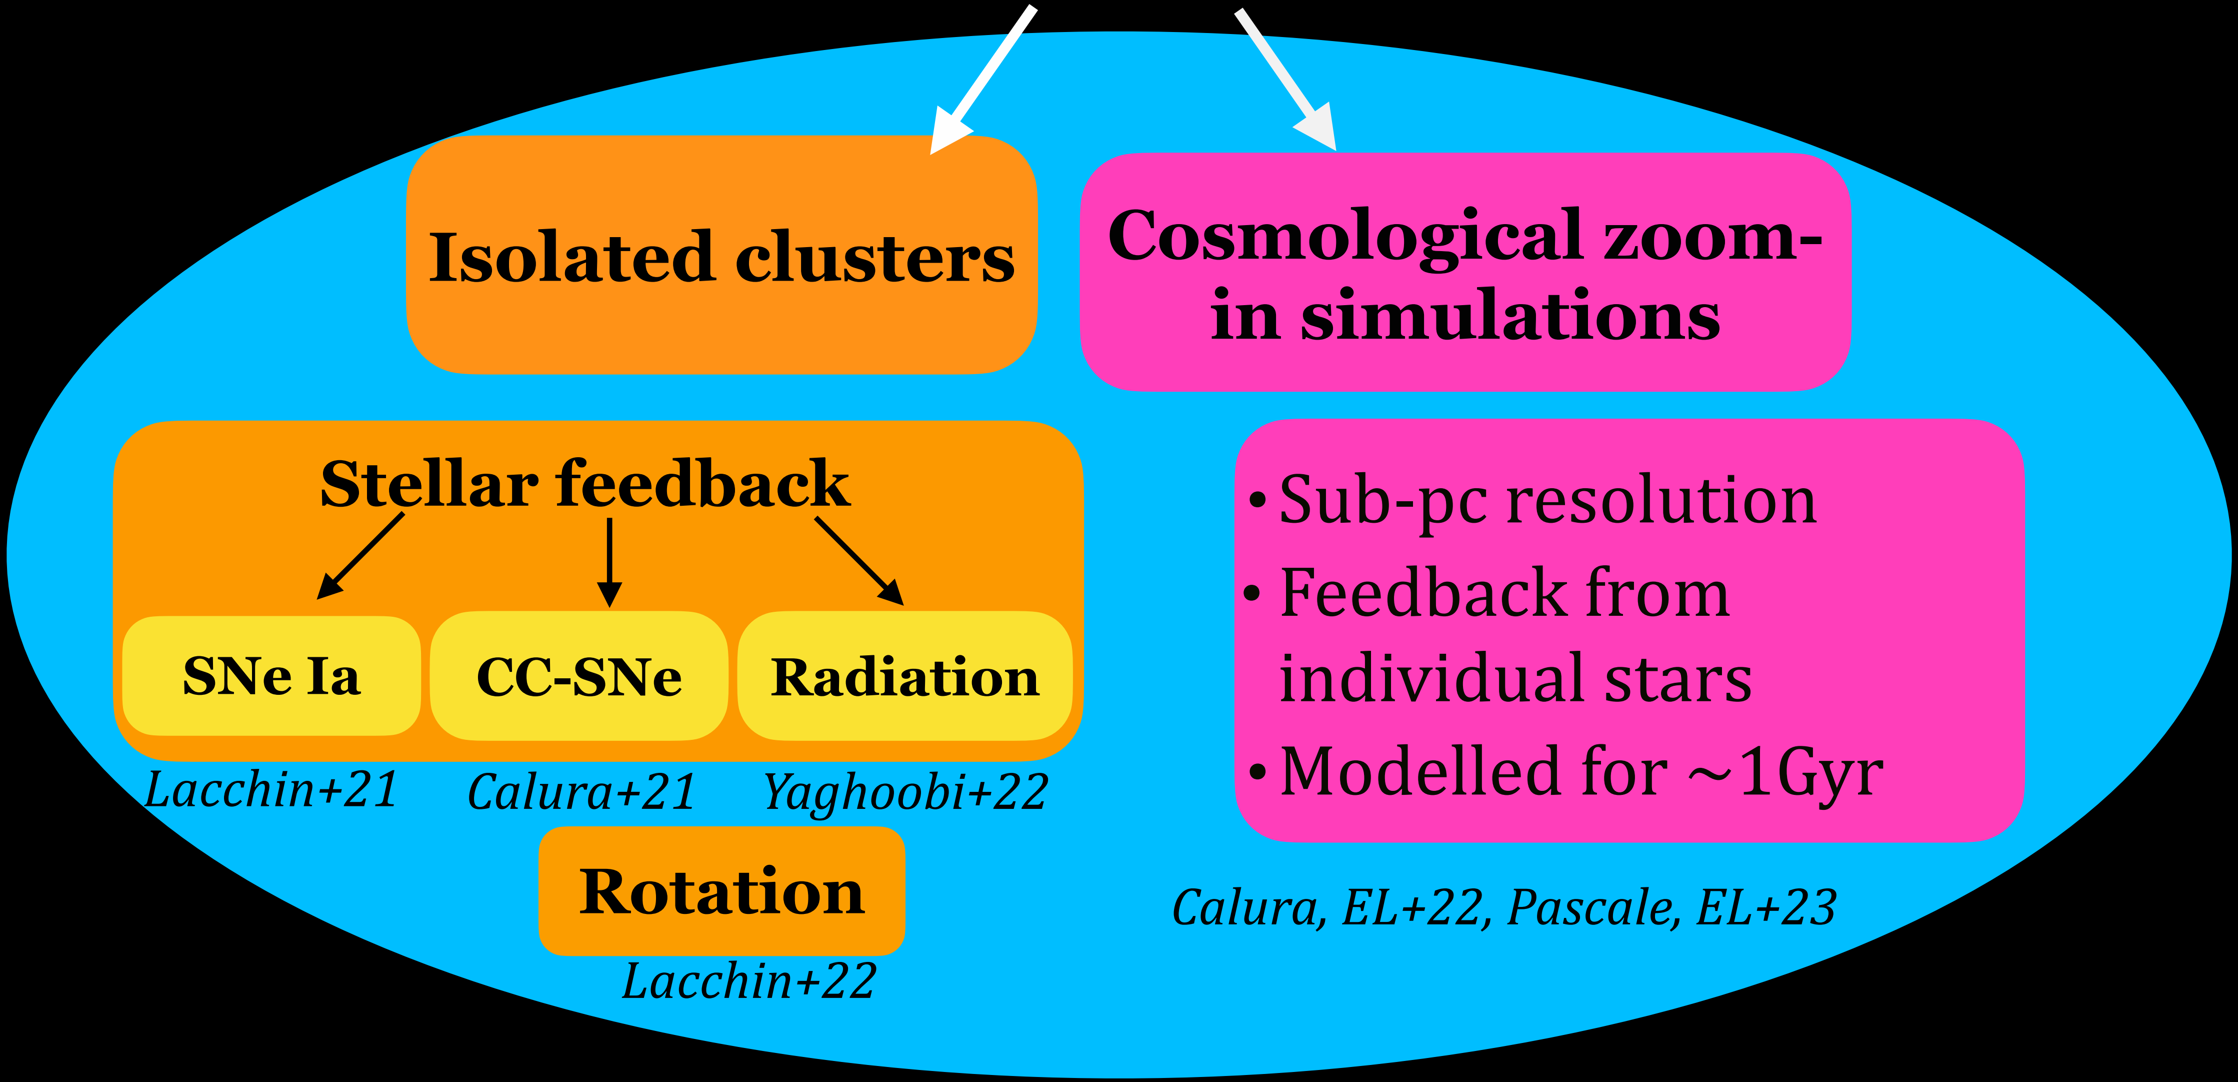

RAMSES

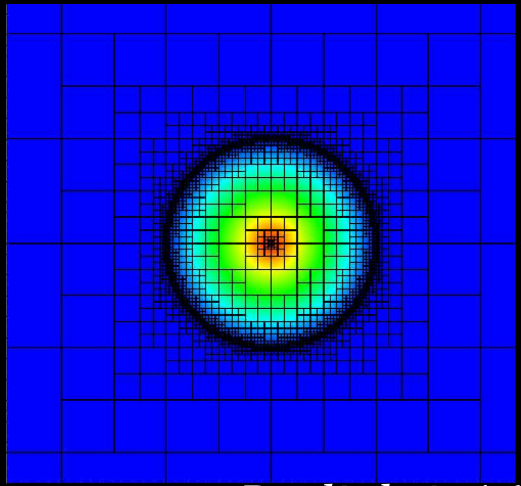

Peplinksi+16

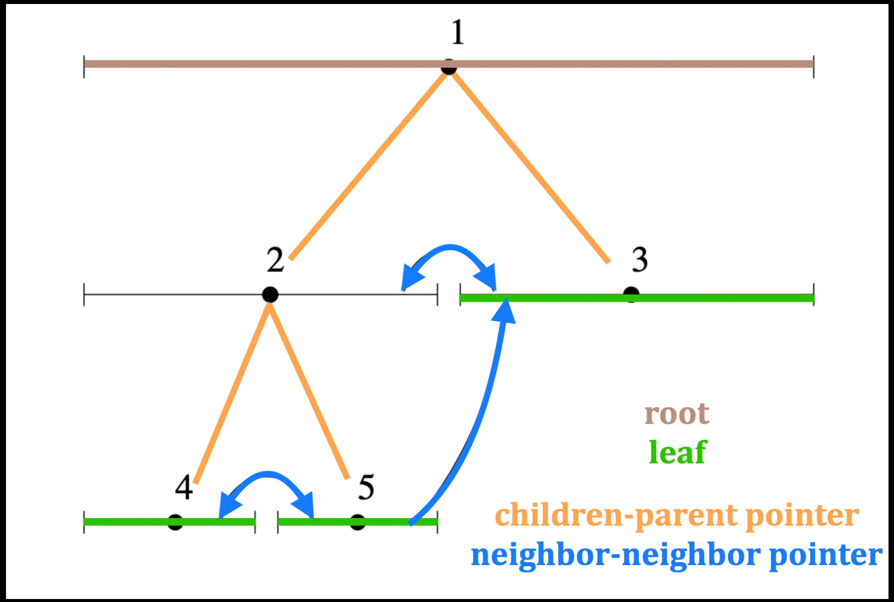

Khokhlov+98

Hydrodynamical + N-body code

- ★ Adaptive Mesh Refinement technique using the Fully Threaded Tree data structure
- ★ N-body solver: Particle-Mesh method on AMR grids.
- ★ Unsplit second order Godunov method (MUSCL) with various Riemann solvers and slope limiters.
- ★ MPI-based parallel computing using time-dependent domain decomposition based on Peano-Hilbert cell ordering

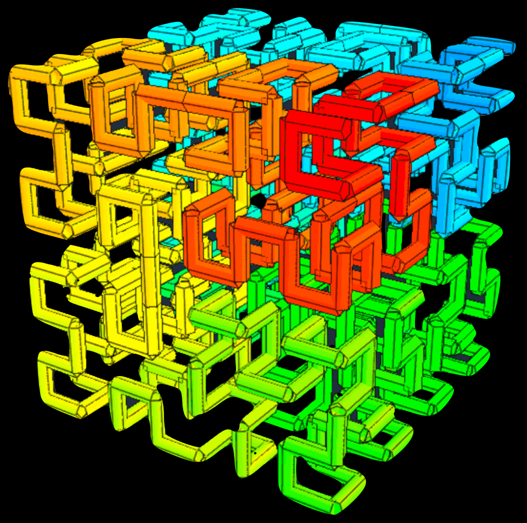

We acknowledge EuroHPC JU for awarding this project access to Discoverer

1. Isolated clusters

Asymptotic Giant Branch scenario

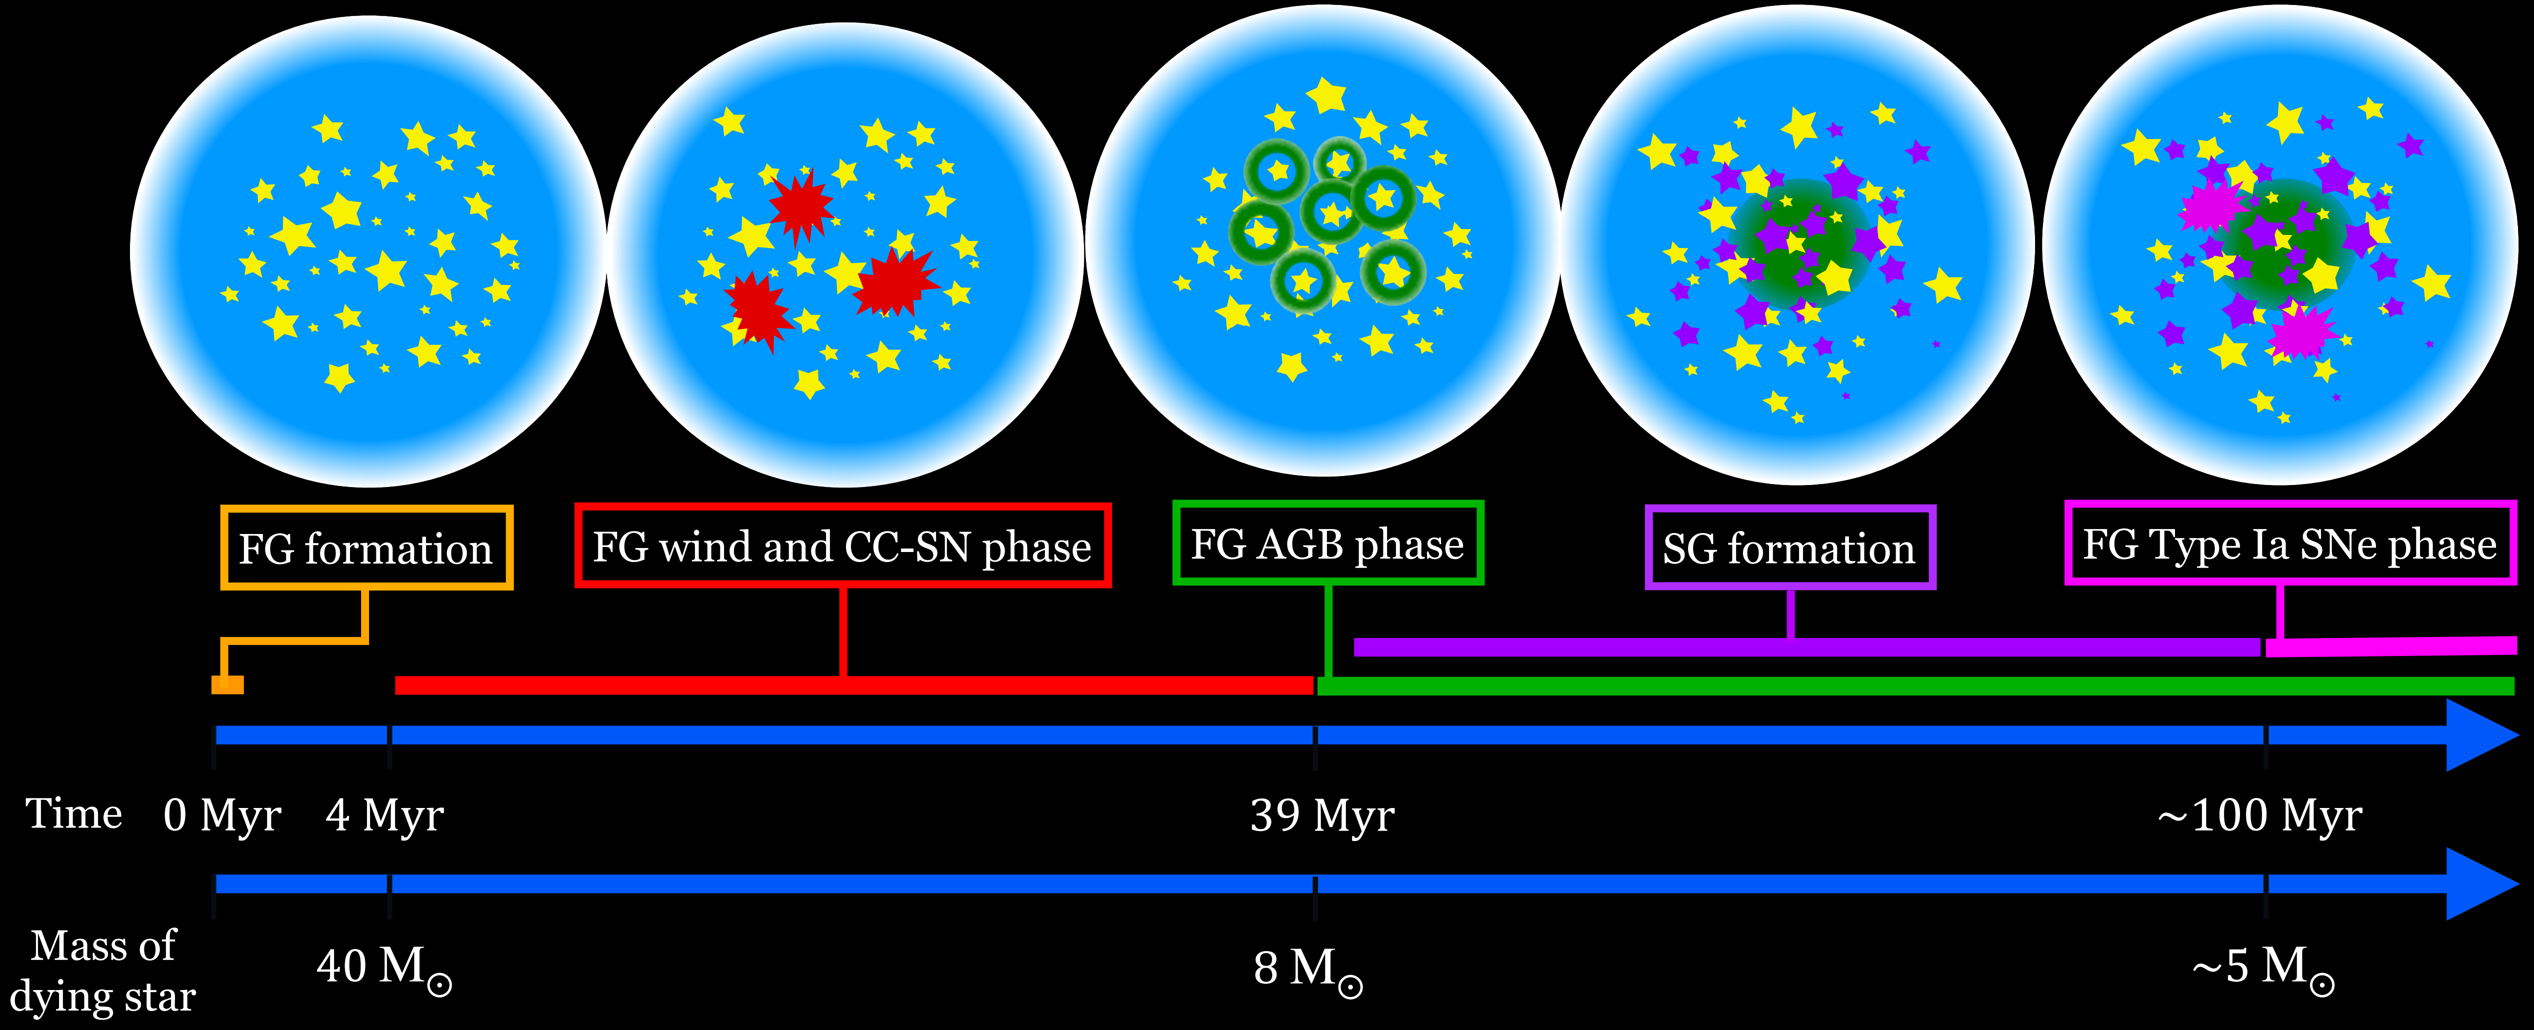

D'Ercole+08,+16

Why cluster rotation?

Increasing number of GCs has been found to rotate

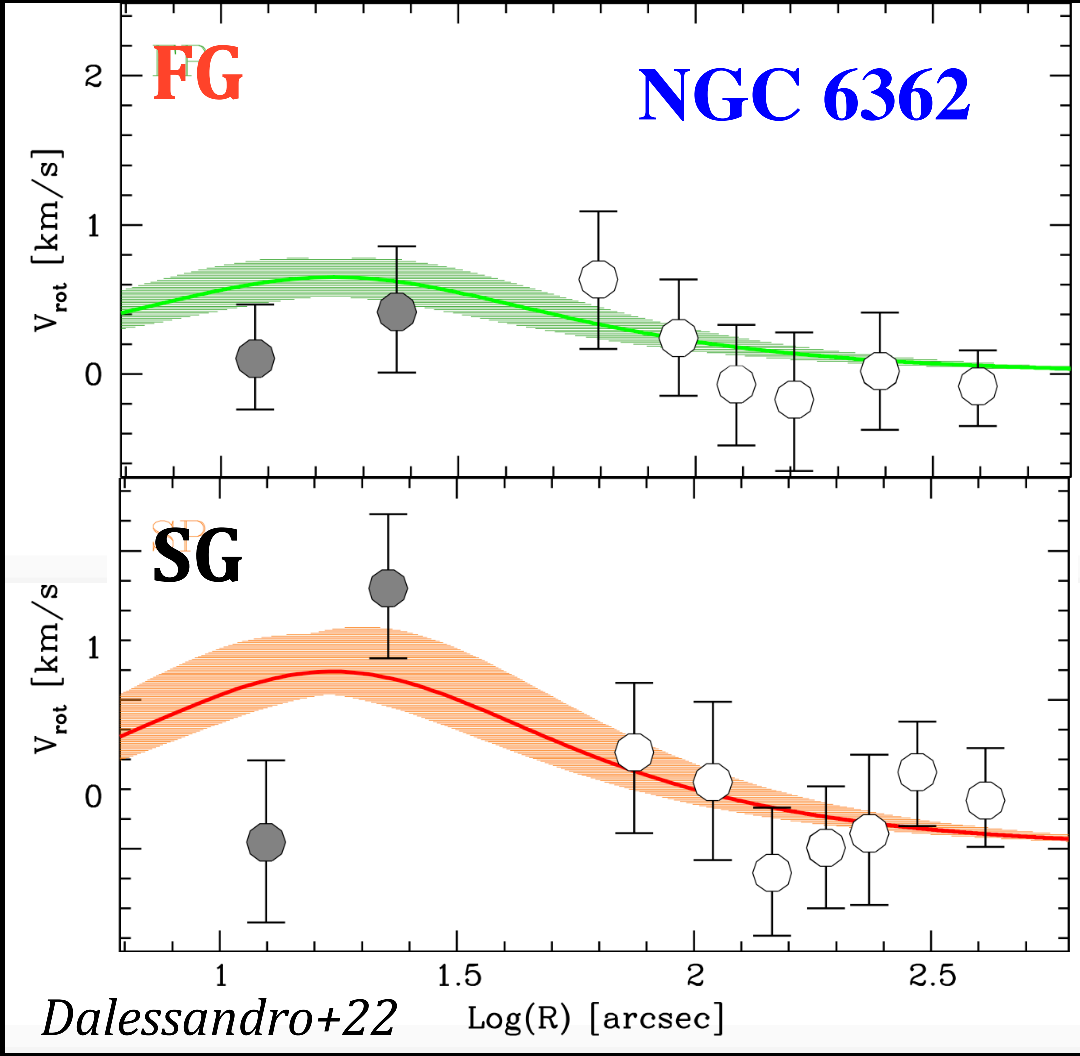

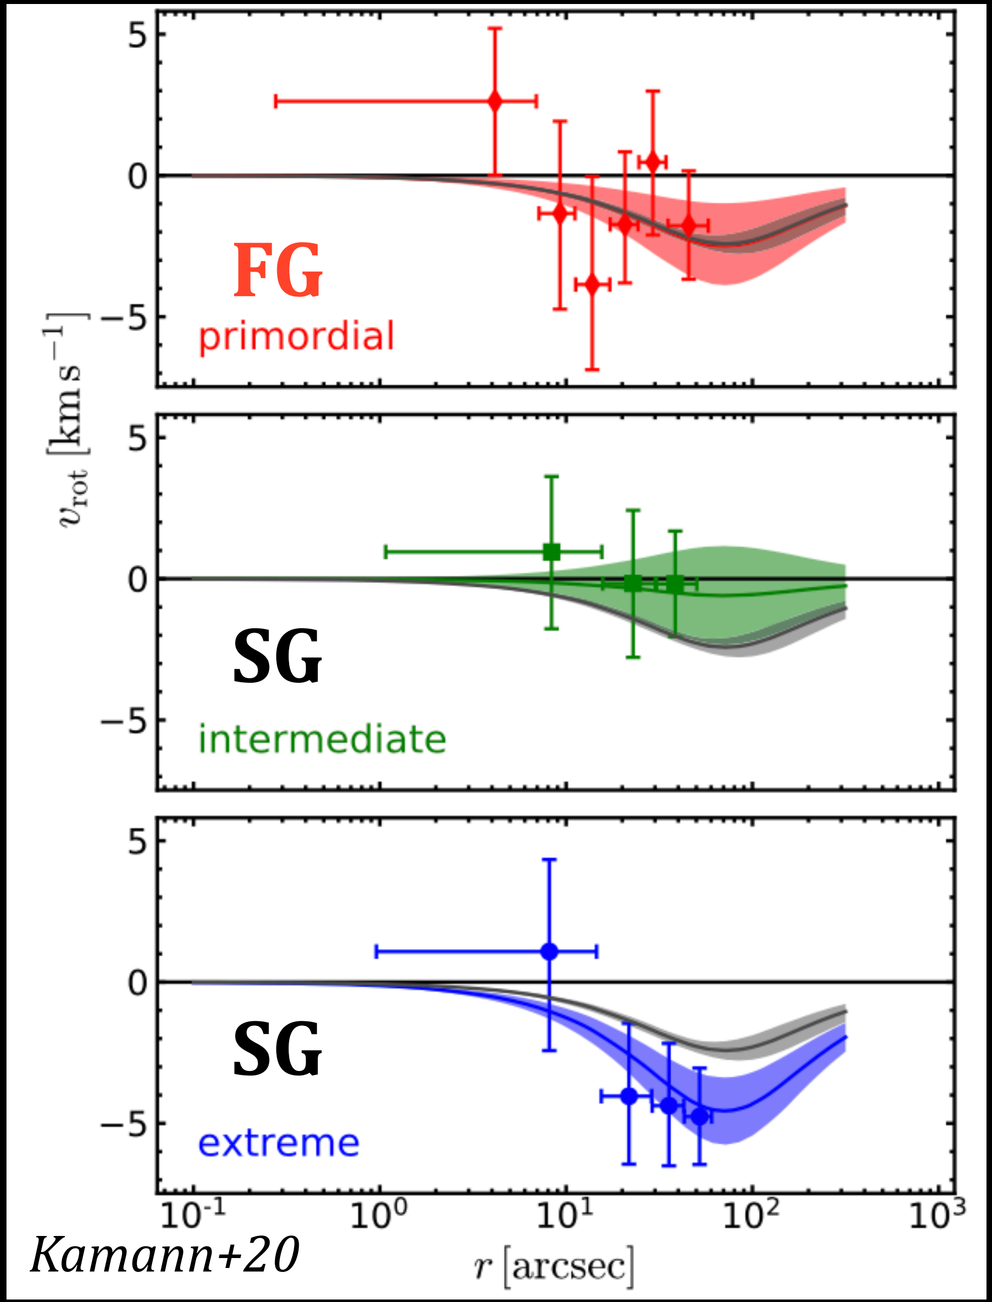

Simulation setup

$M_{FG} = 10^7 M_\odot$ in rotation

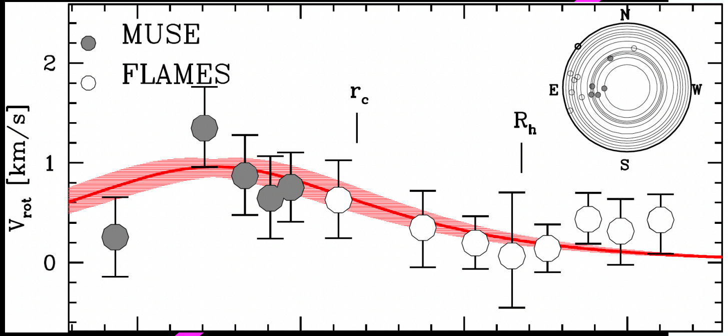

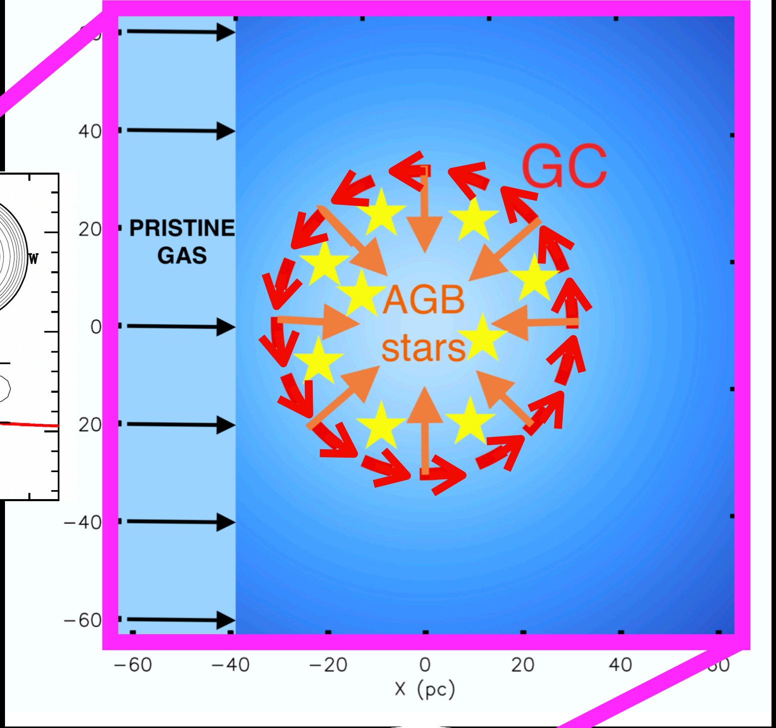

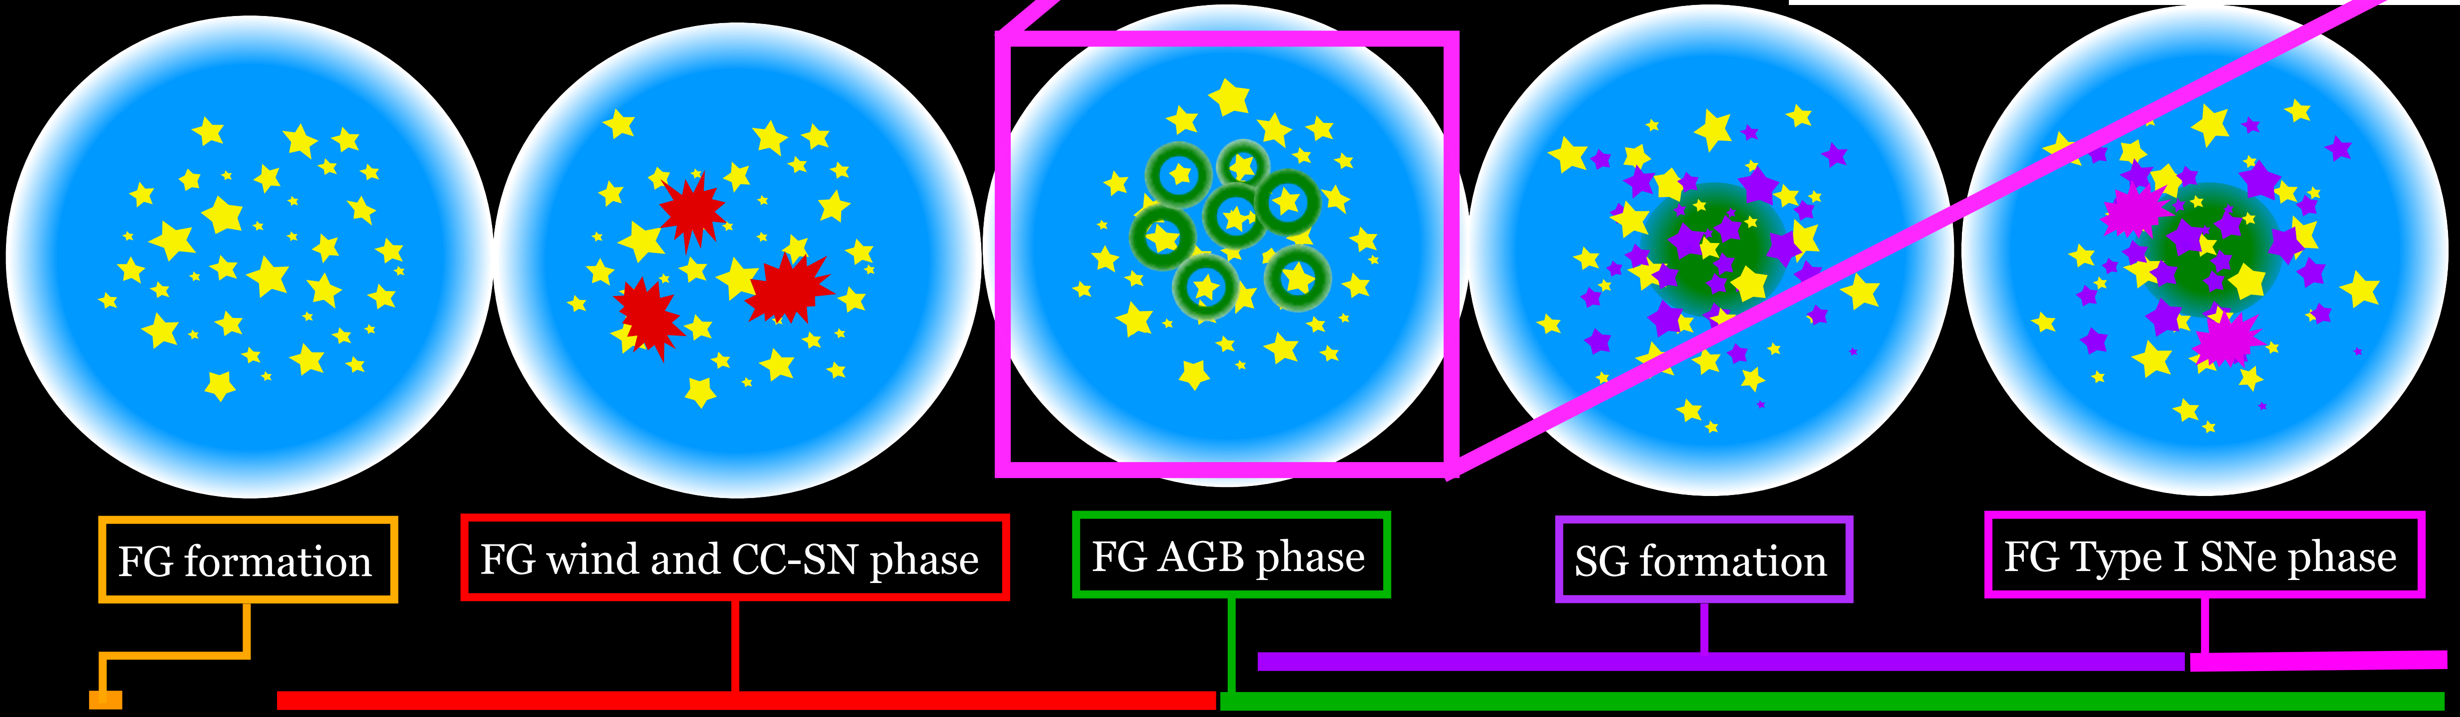

SG evolution - xy plane

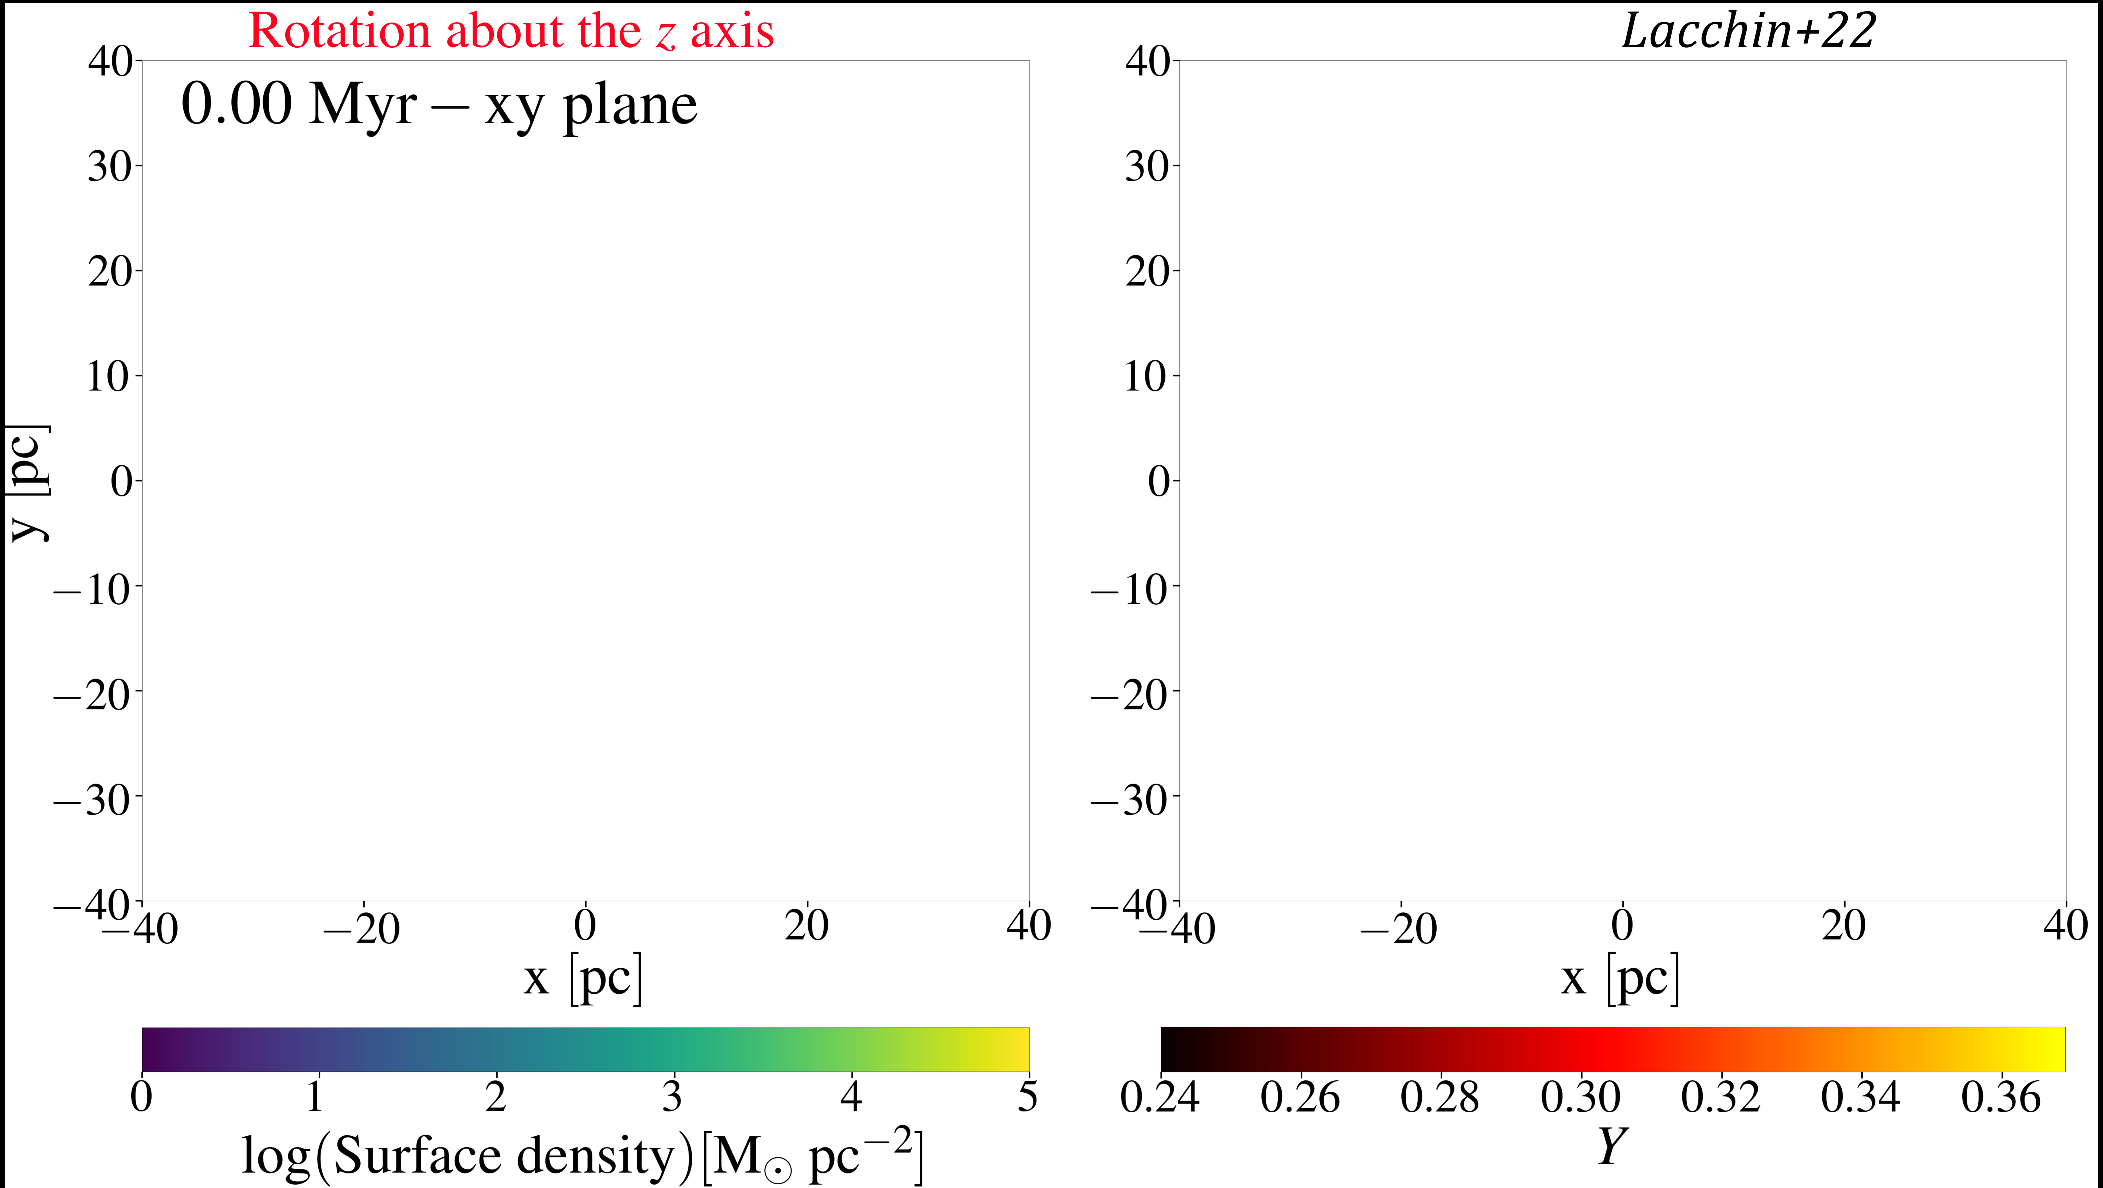

SG evolution - yz plane

A stellar disk is formed which is mainly composed by helium-enhanced stars

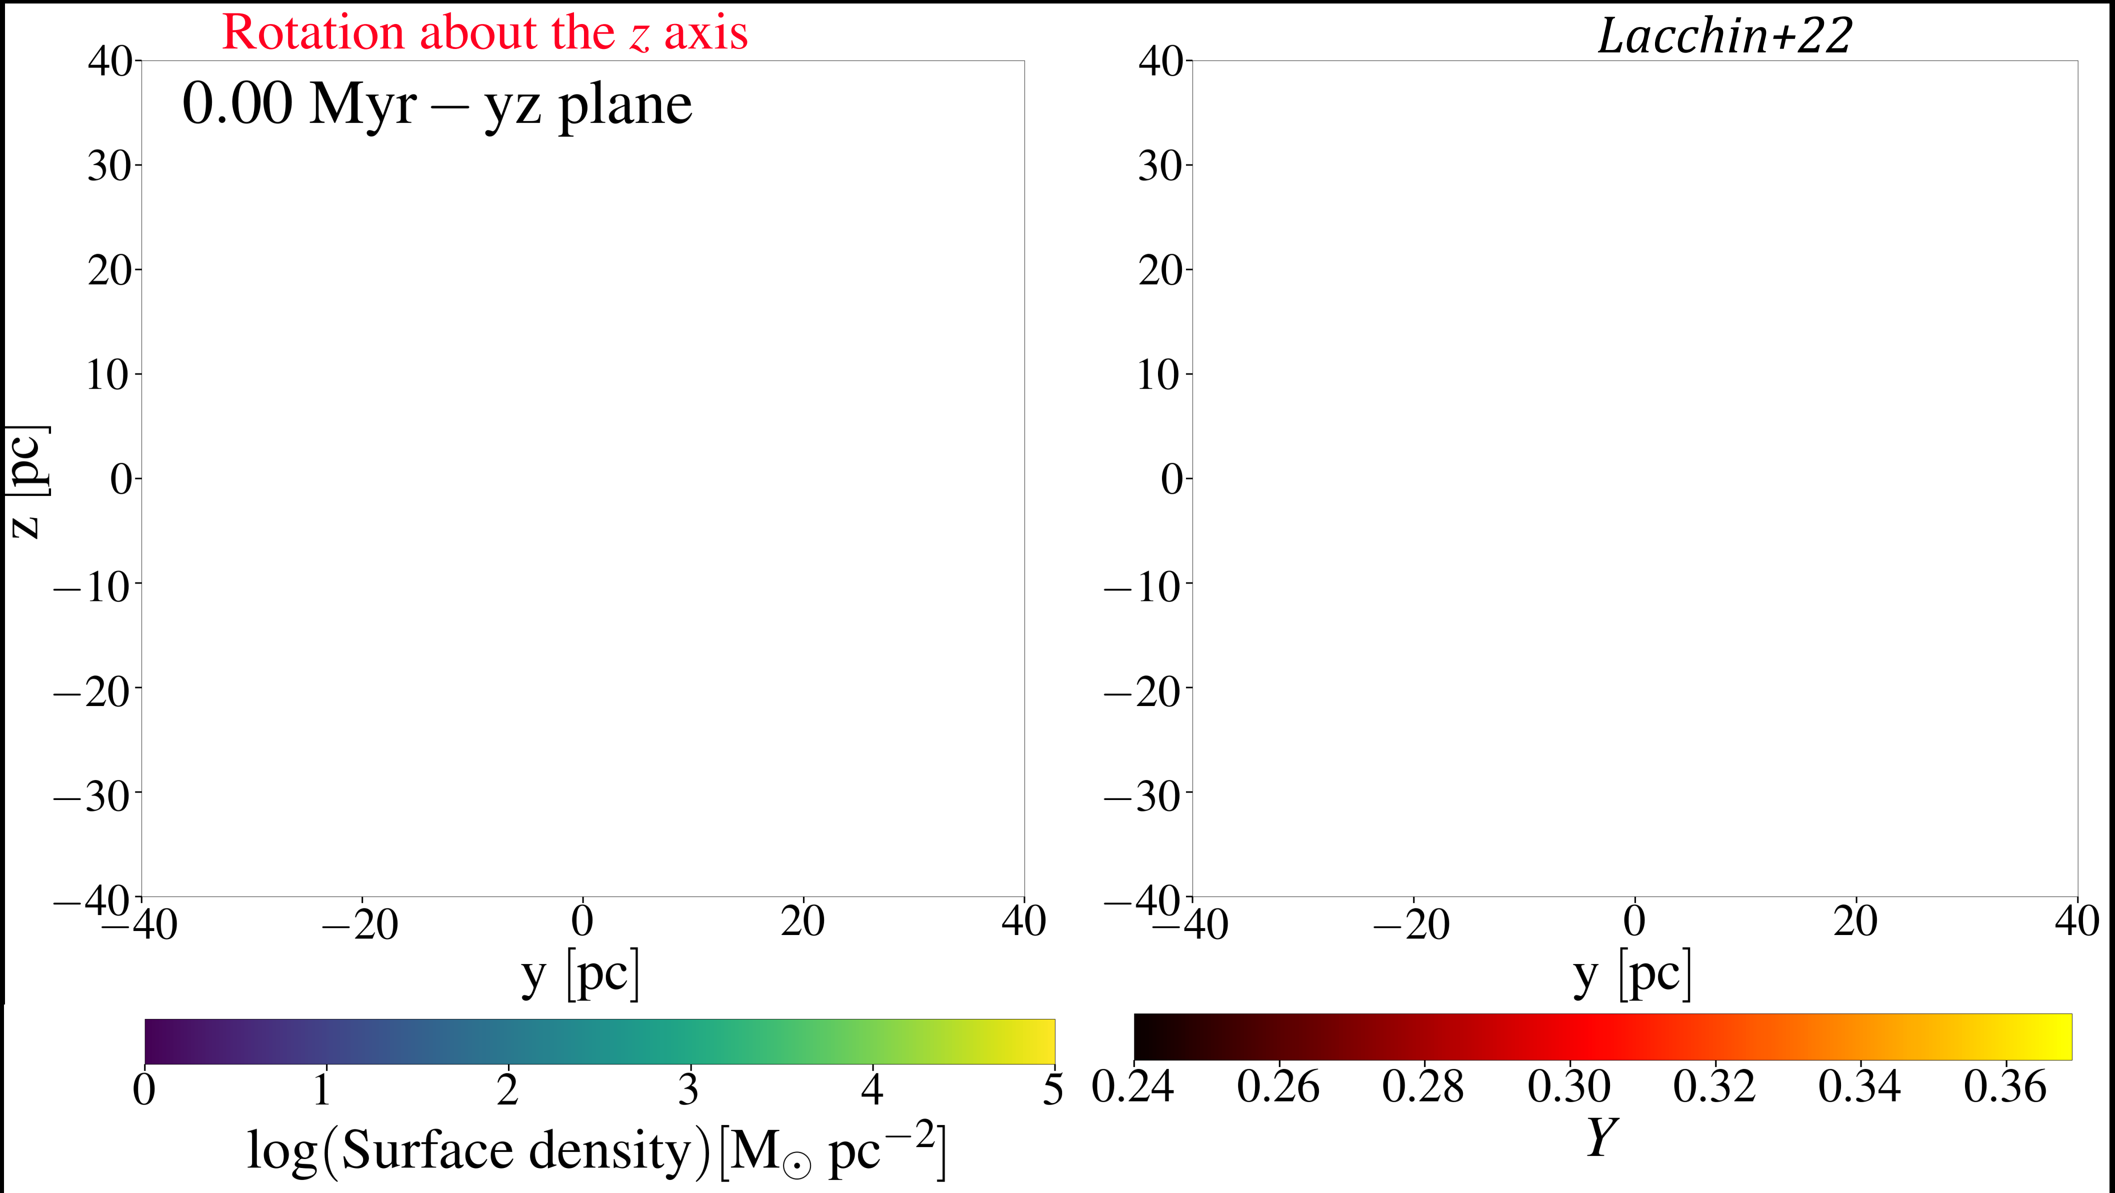

SG rotation profiles

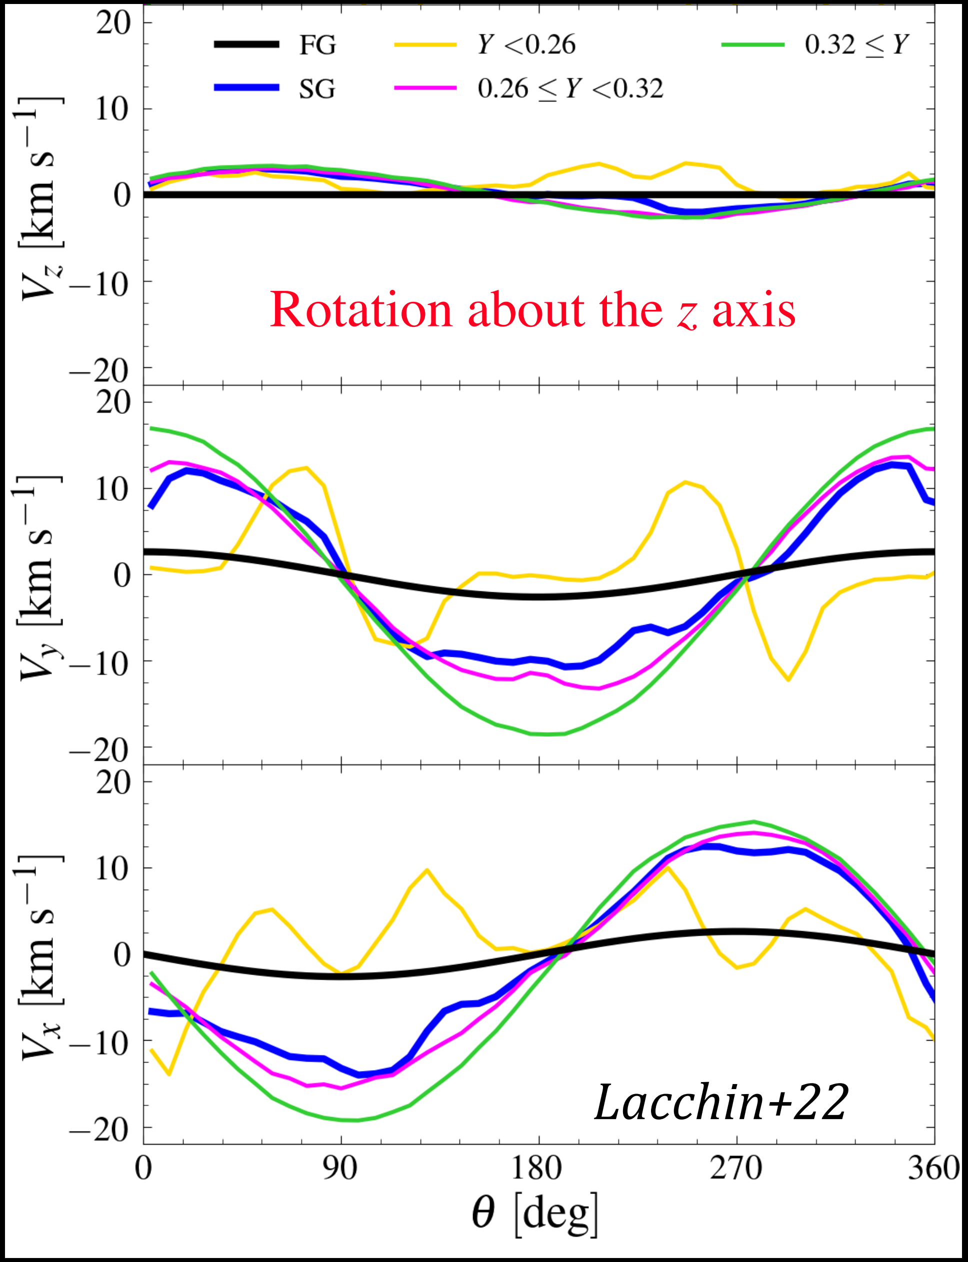

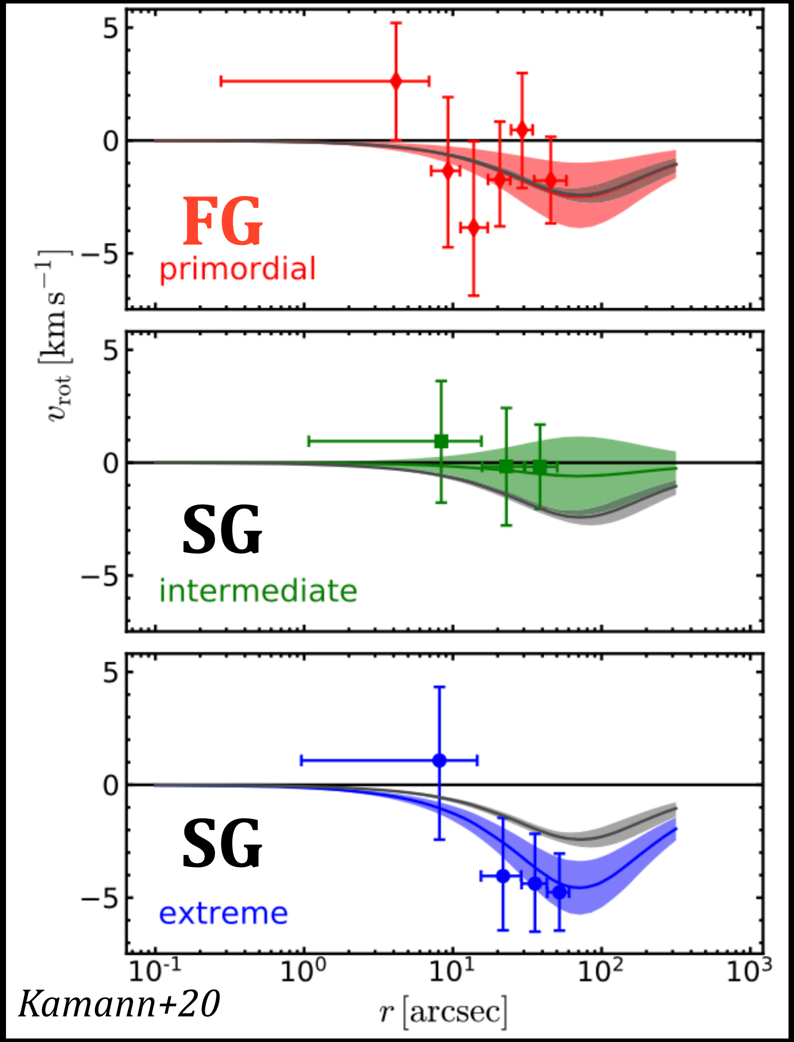

★ the SG rotational amplitude is significantly greater than the FG one

★ the more helium enhanced SG stars rotate faster than the helium poor ones

Type Ia SNe - feedback

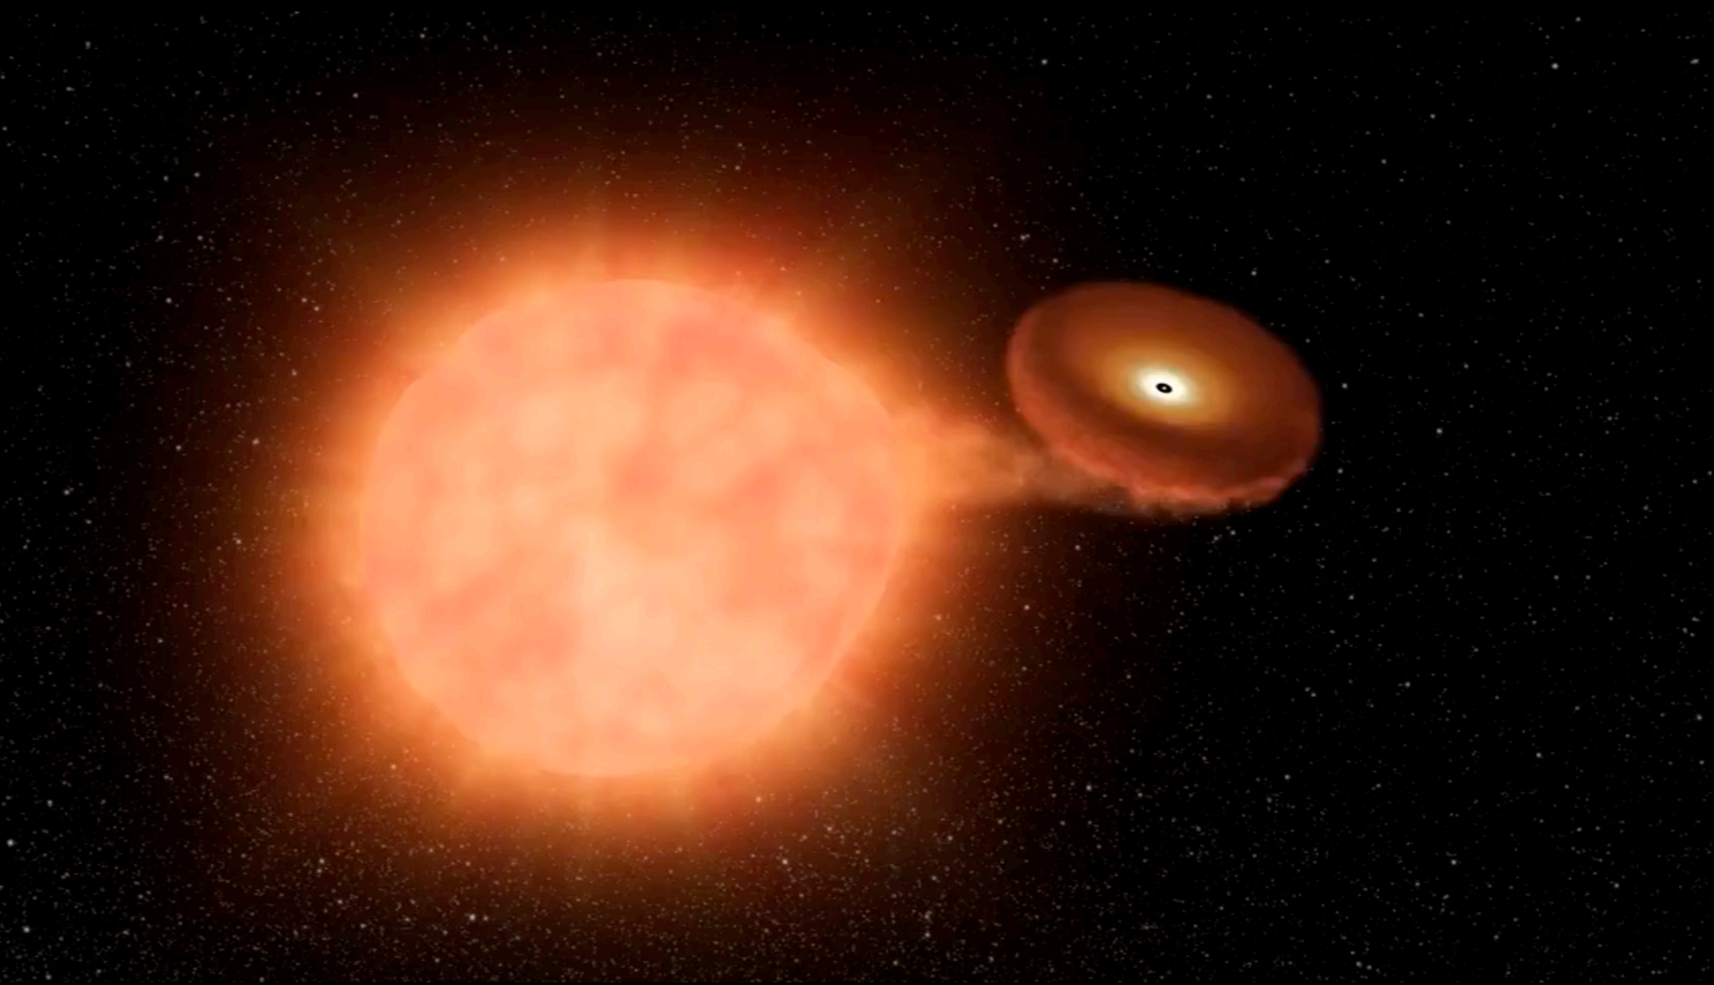

Credit: NASA/JPL-Caltech

Each SN injects:

- ★ 1.44 $M_\odot$ of ejecta, all metal
- ★ 0.5 $M_\odot$ of Fe *Scalzo+14*
- ★ 10^{51} erg of thermal energy

SNR 0509-67.5

Image: NASA, ESA, CXC, SAO, the Hubble Heritage Team

Stellar mass vs time

$M_{\rm FG} = 10^7 {\rm M}_\odot$

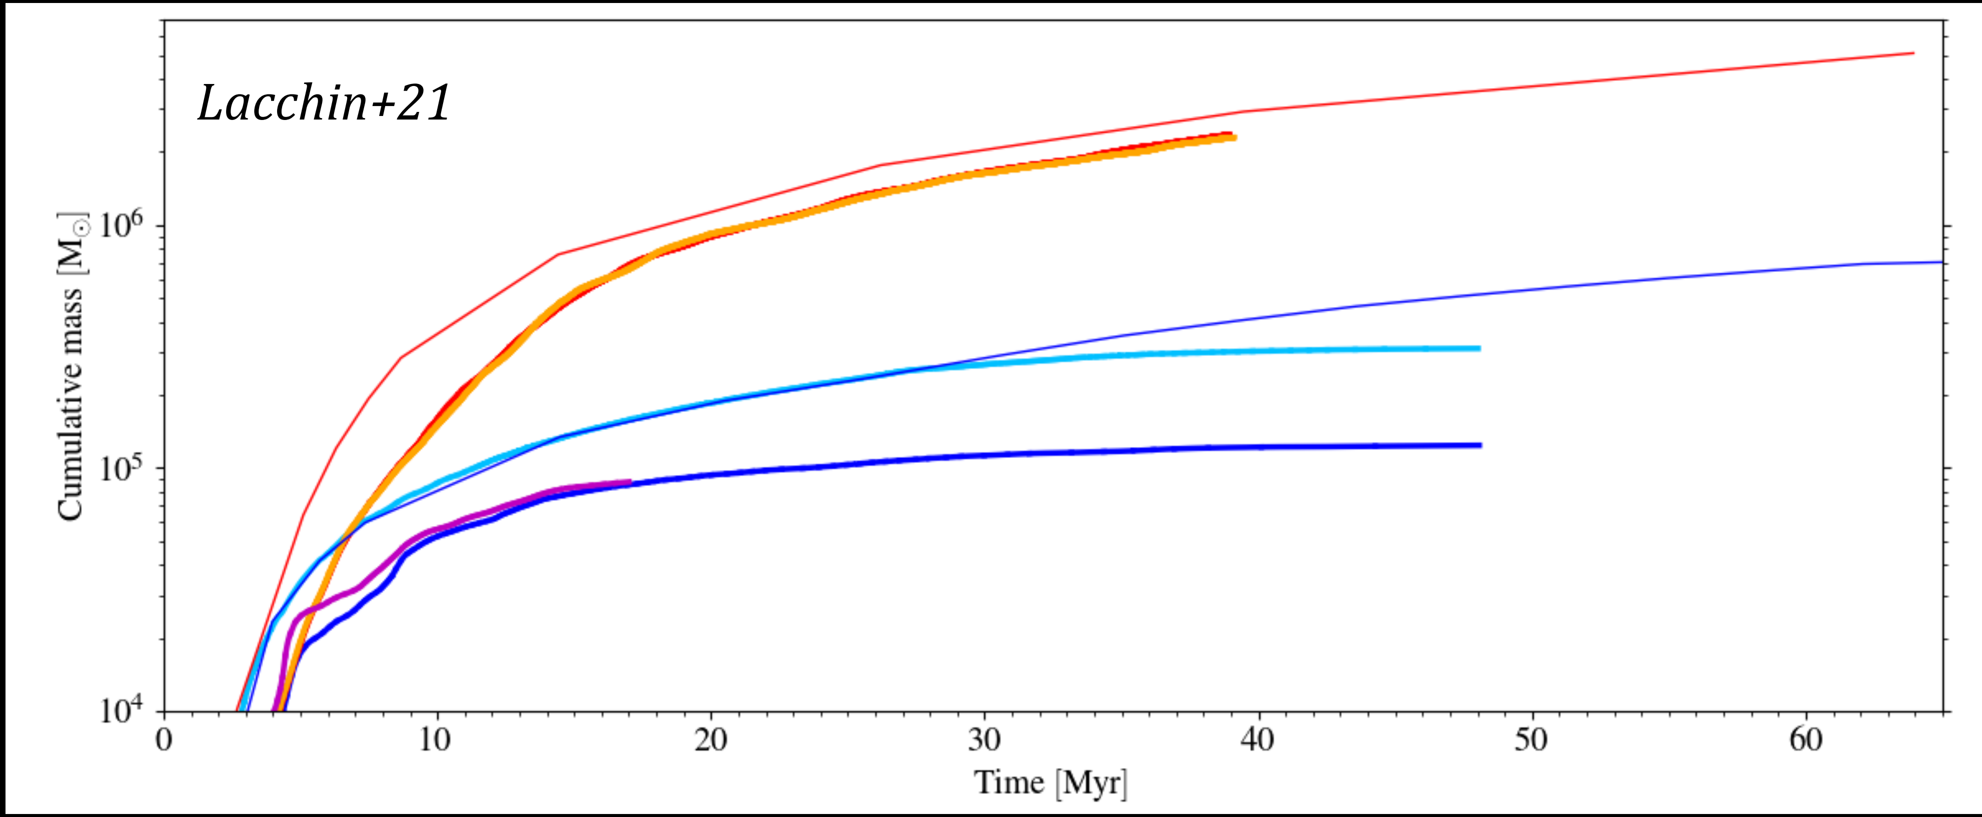

Without SNe Ia

With SNe Ia

High density

Low density

Low density:
$M_\star^{final}$ reduced by 80%

High density:
$M_\star^{final}$ reduced by 20%

Clusters of lower mass

$M_{\mathrm{FG}} = 10^6 \mathrm{M}_{\odot}$

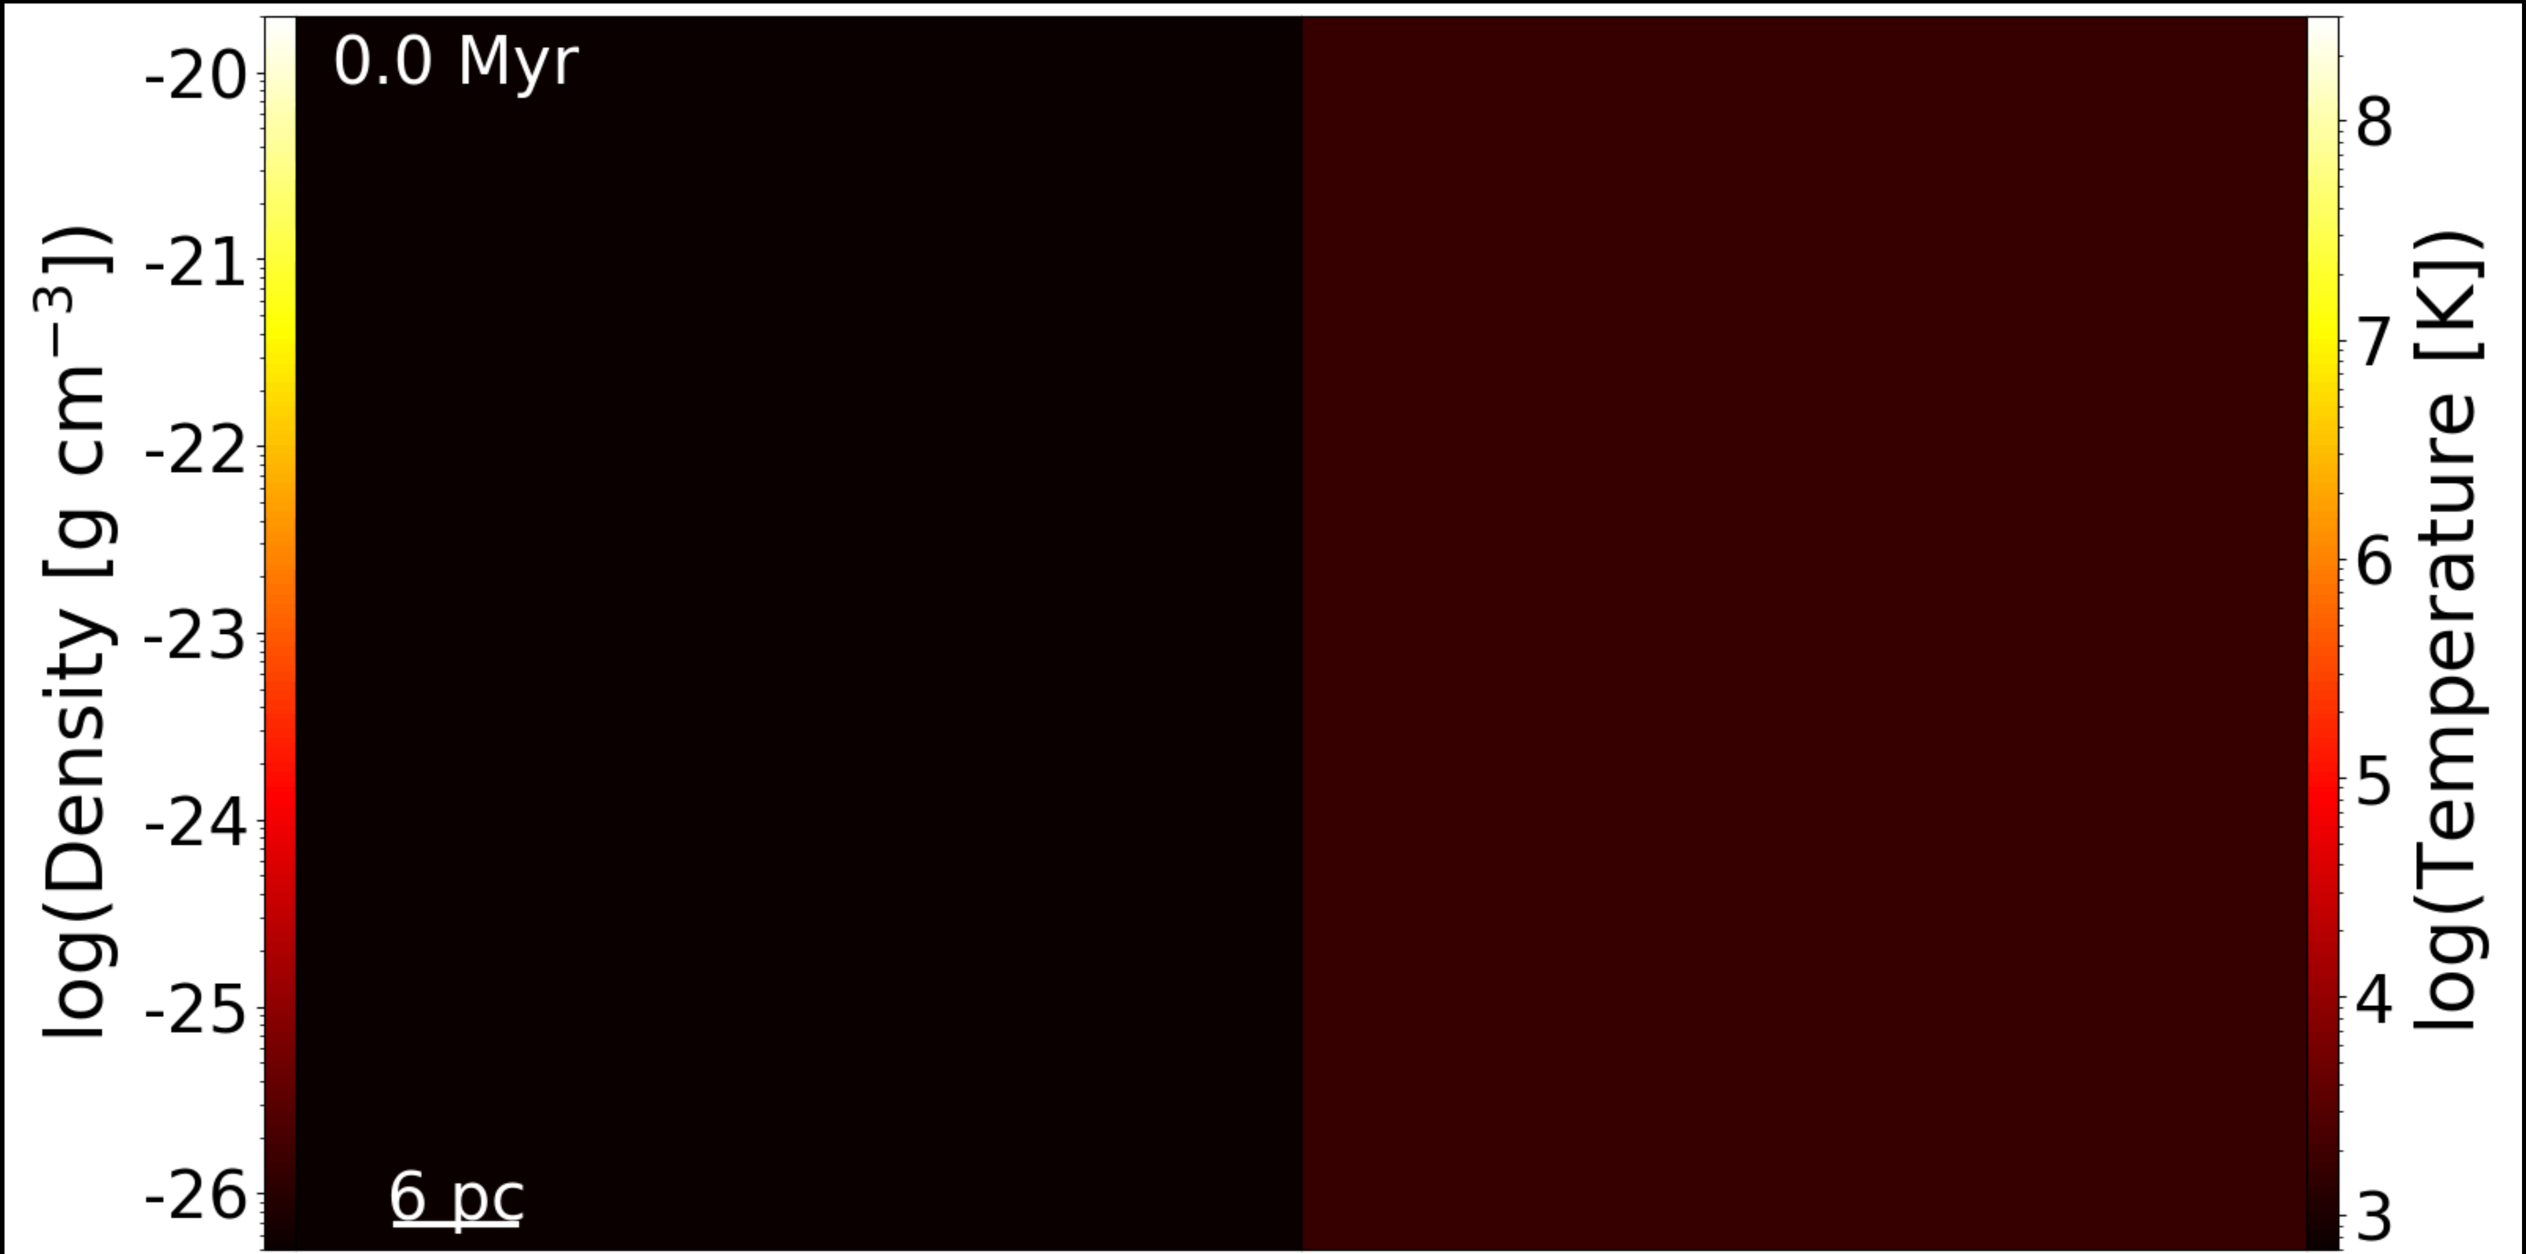

2. Cosmological framework

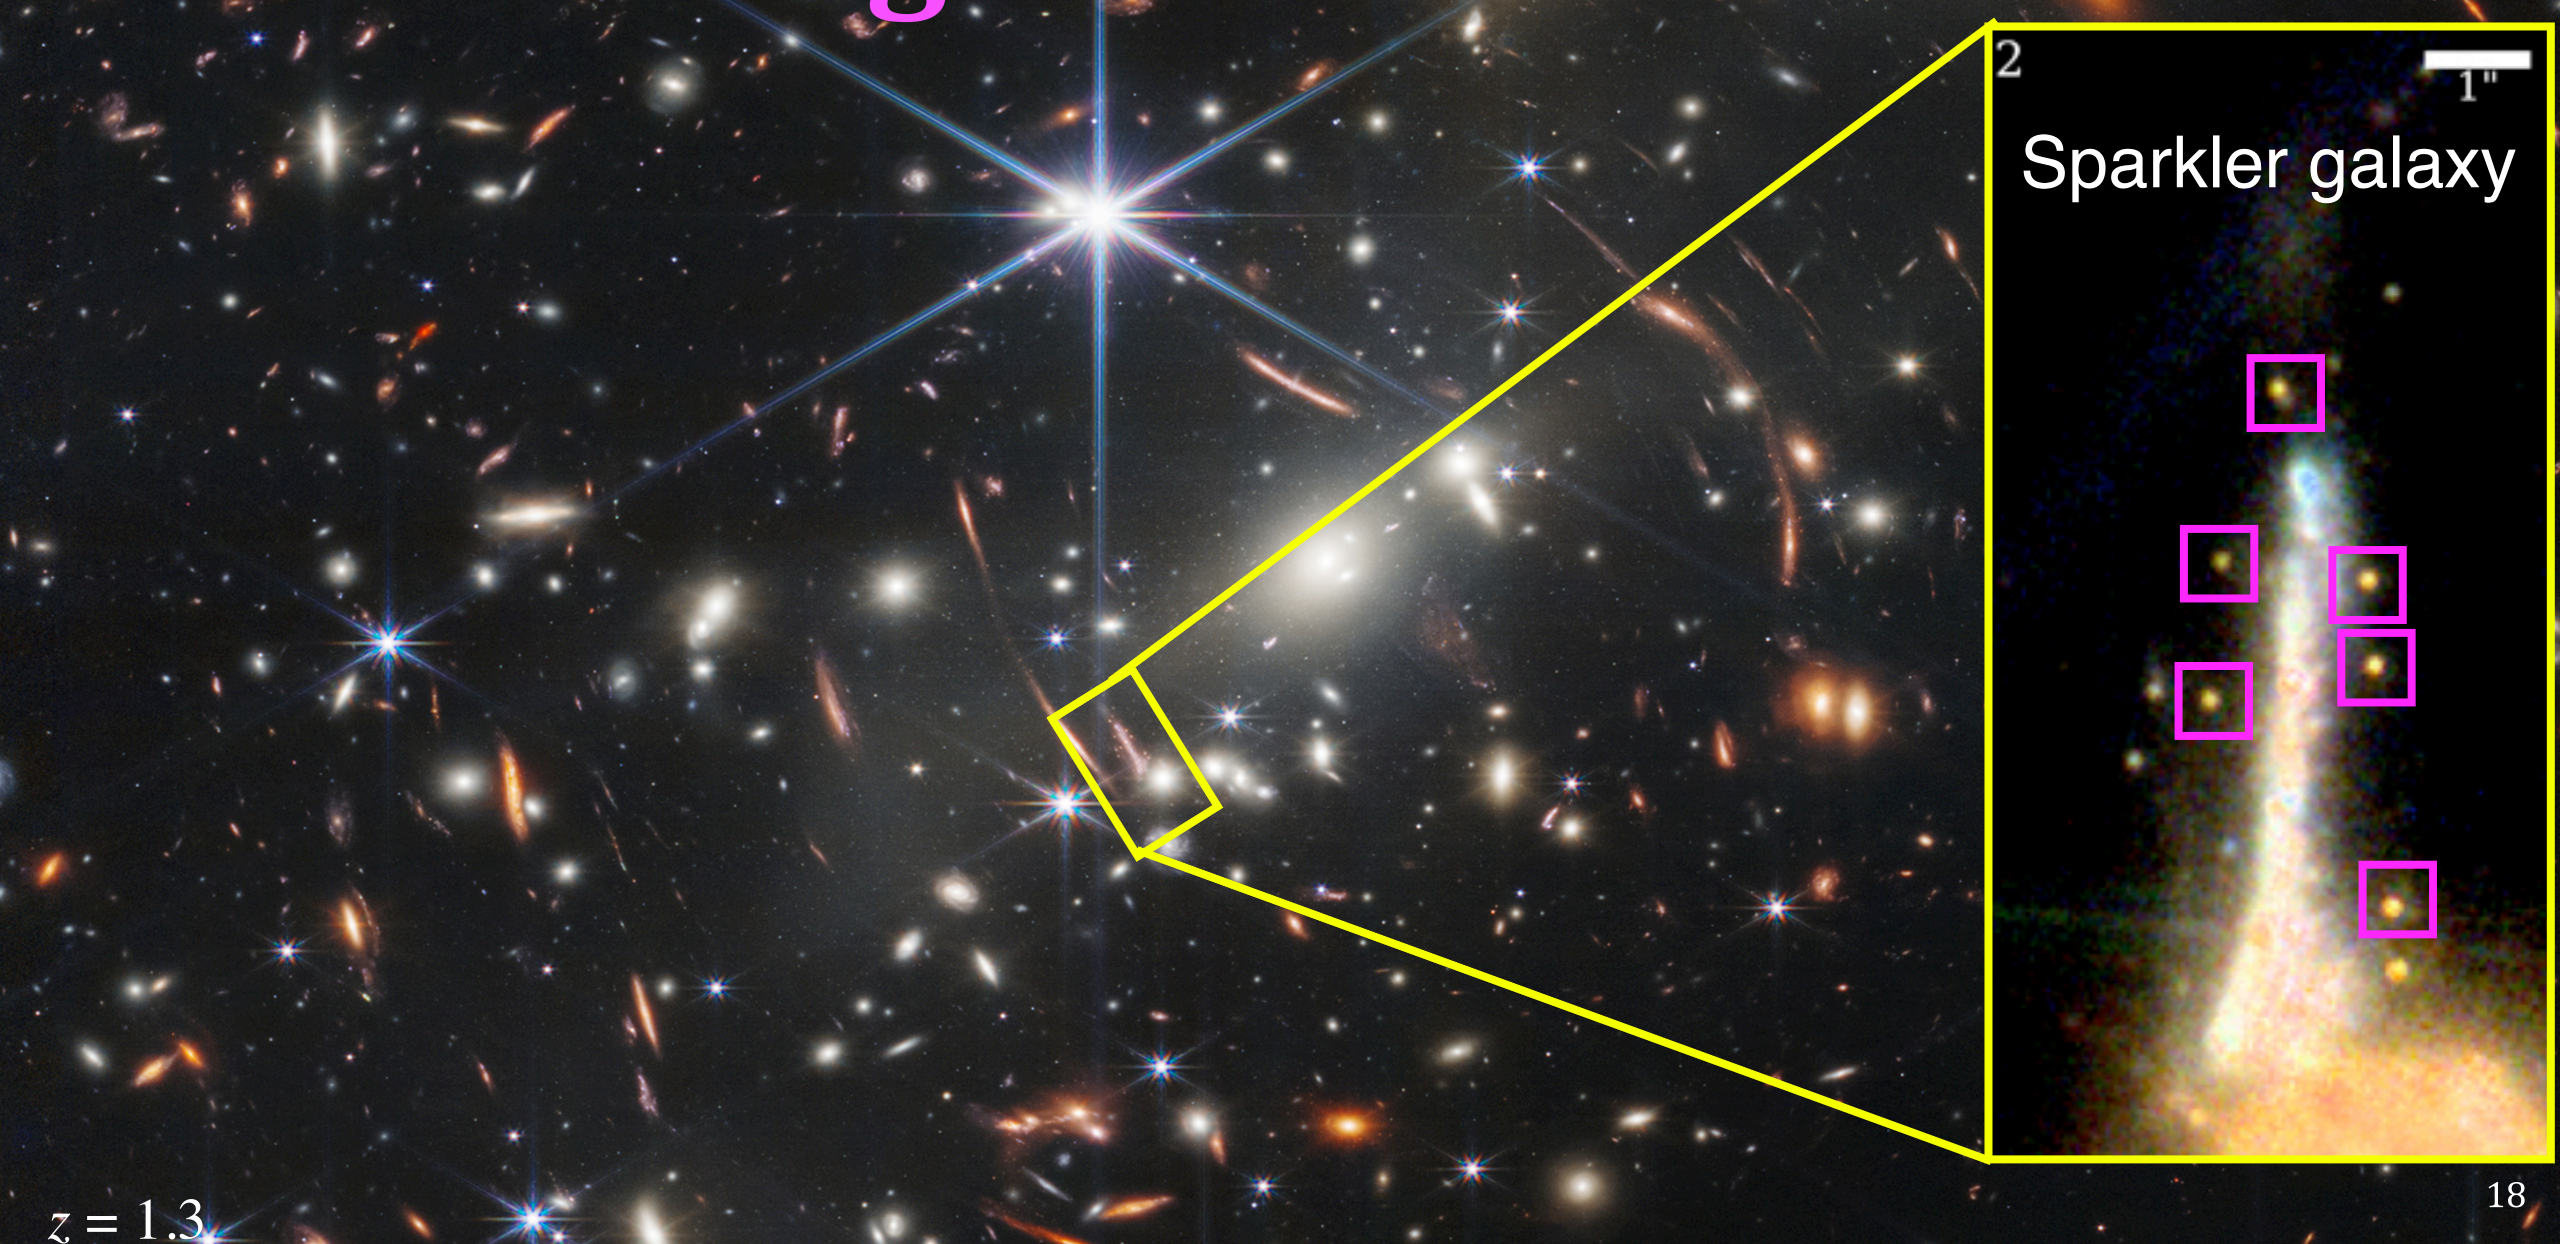

Star forming clumps at high-z

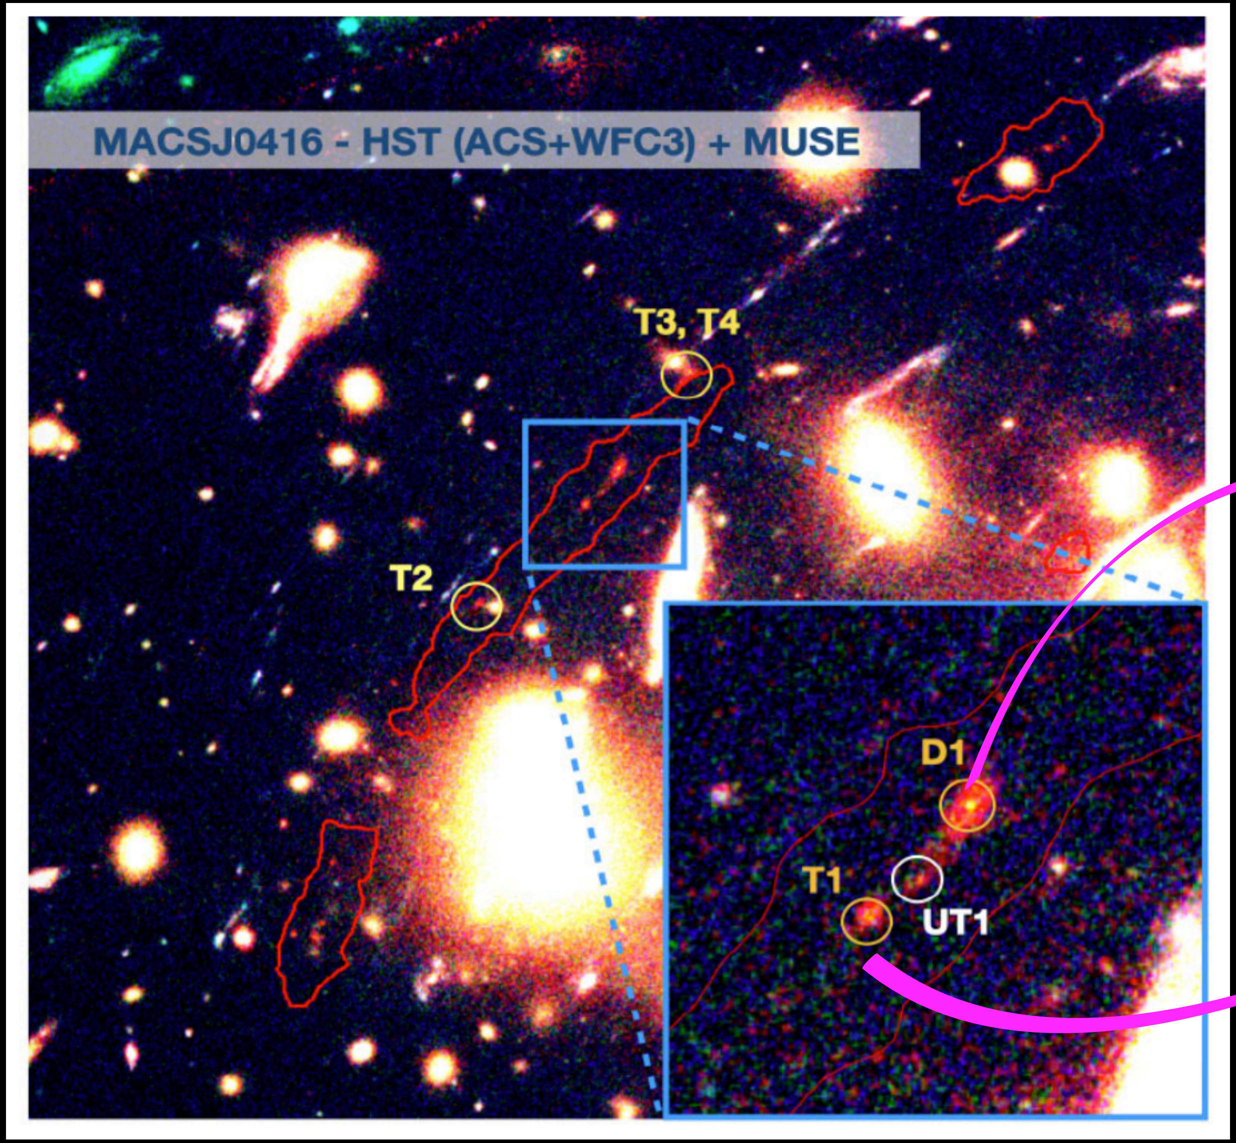

Strongly magnified star forming region at z = 6.14 (t=1Gyr)

D1

- ⋆ $M_\star = 2 \times 10^7 \mathrm{M}_\odot$
- ⋆ $r = 44$ pc

T1

- ⋆ $M_\star = 2 \times 10^6 \mathrm{M}_\odot$
- ⋆ $r < 30$ pc

Vanzella+19

Cosmological view

Calura, EL+22

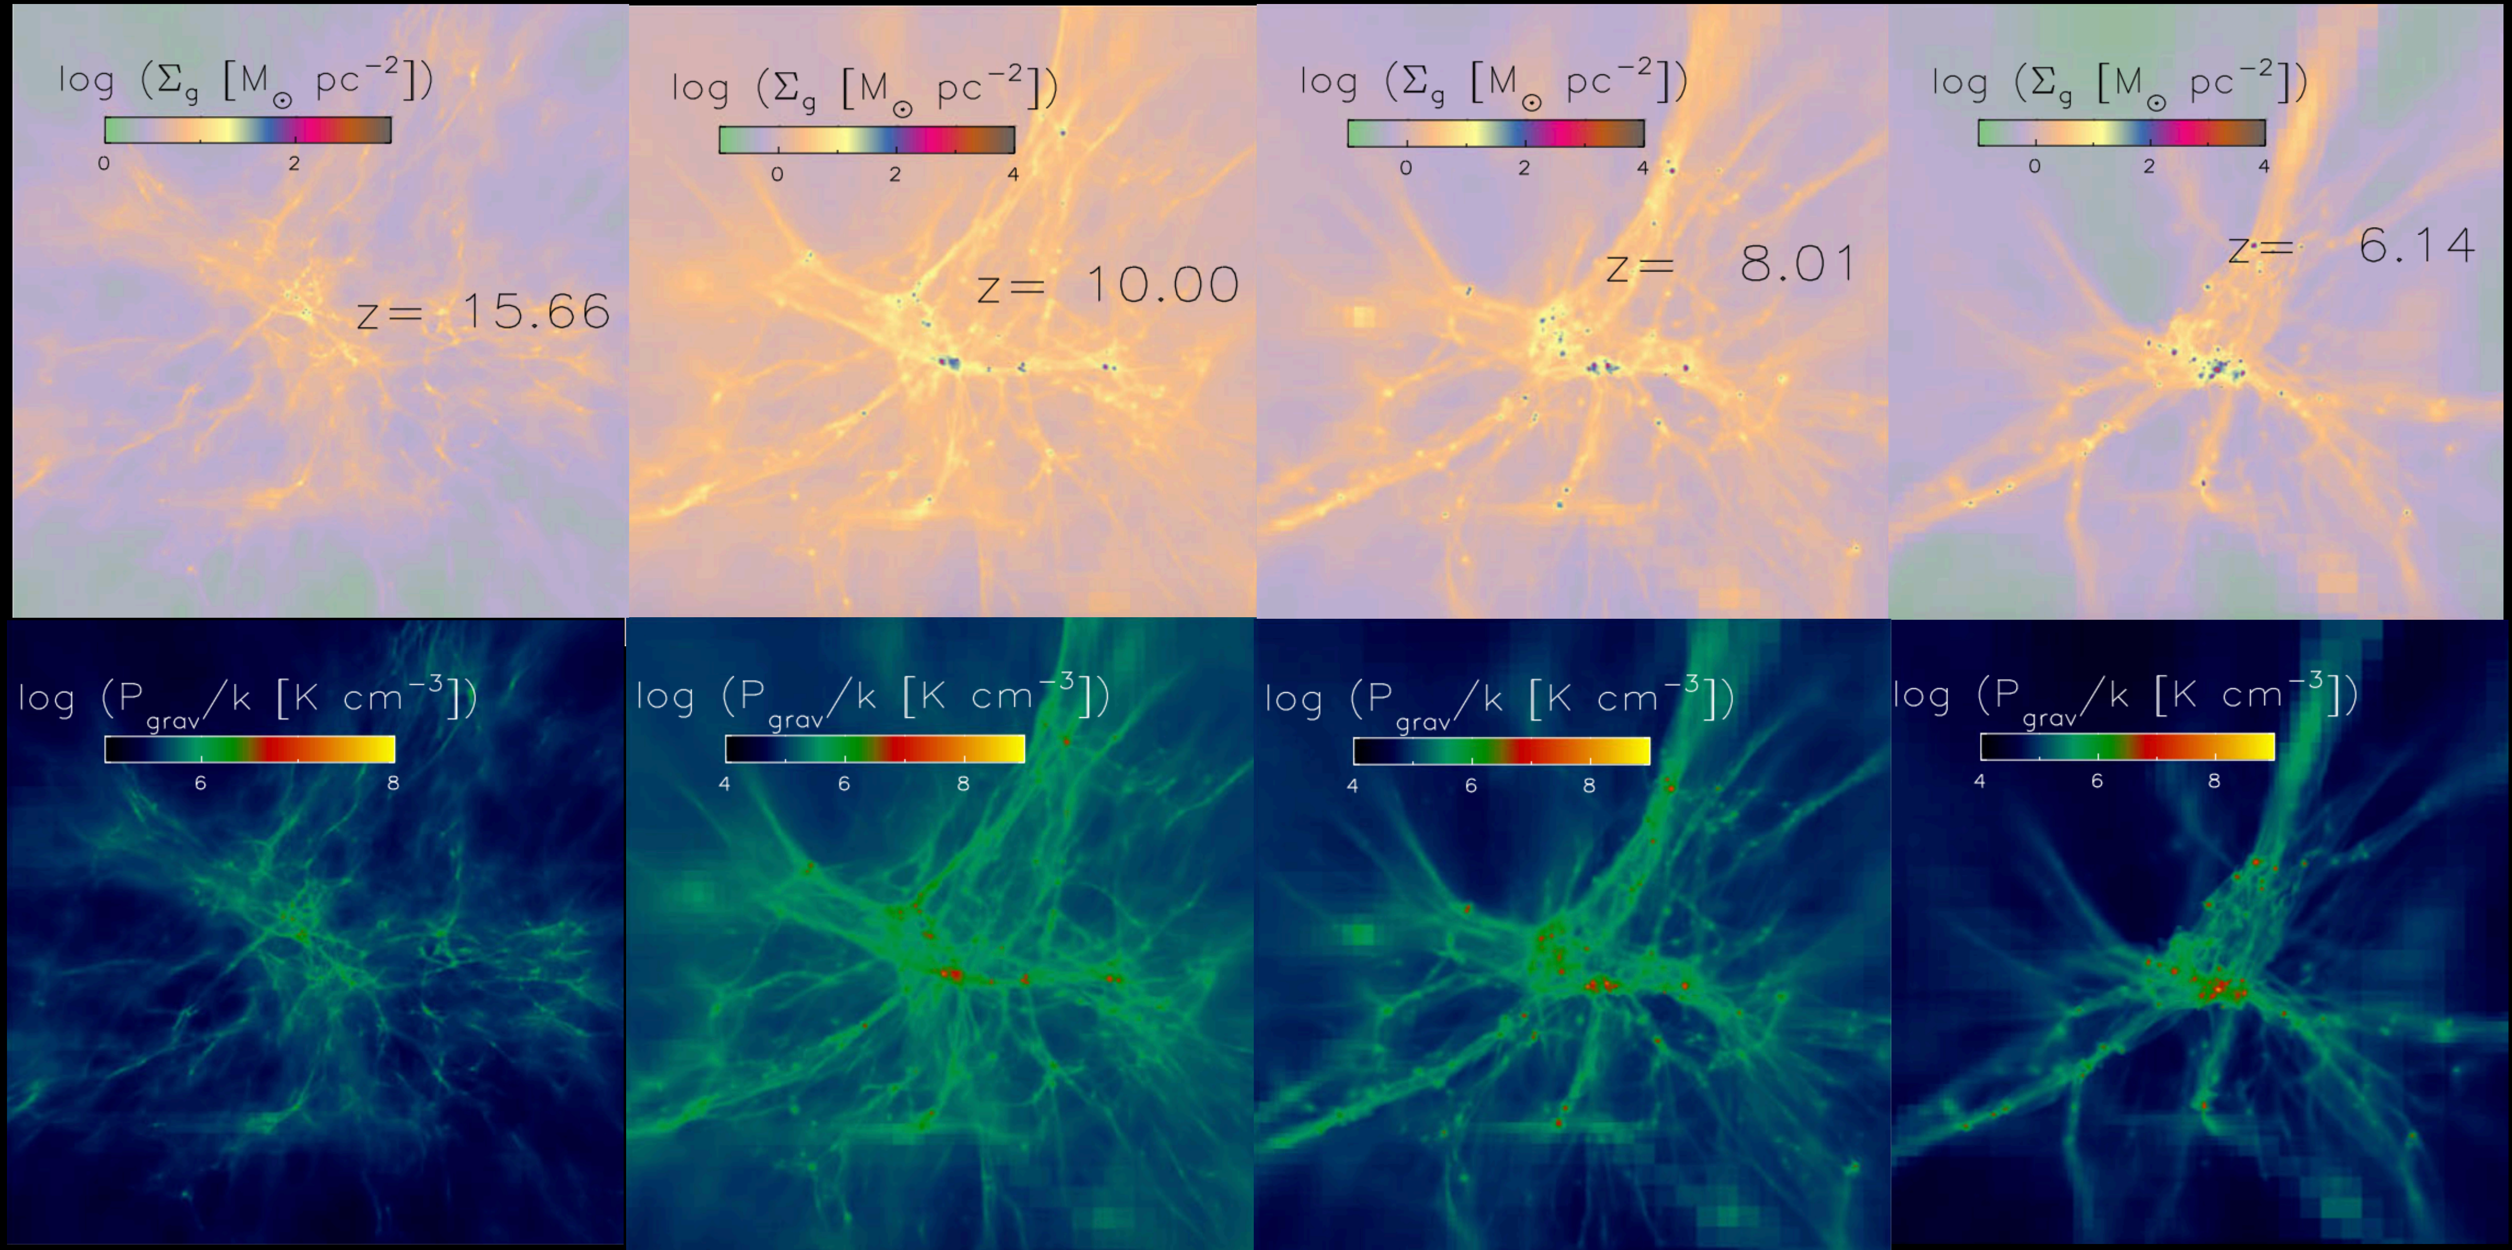

Cosmological view

Pascale, EL+23

Dark matter | Stellar clumps

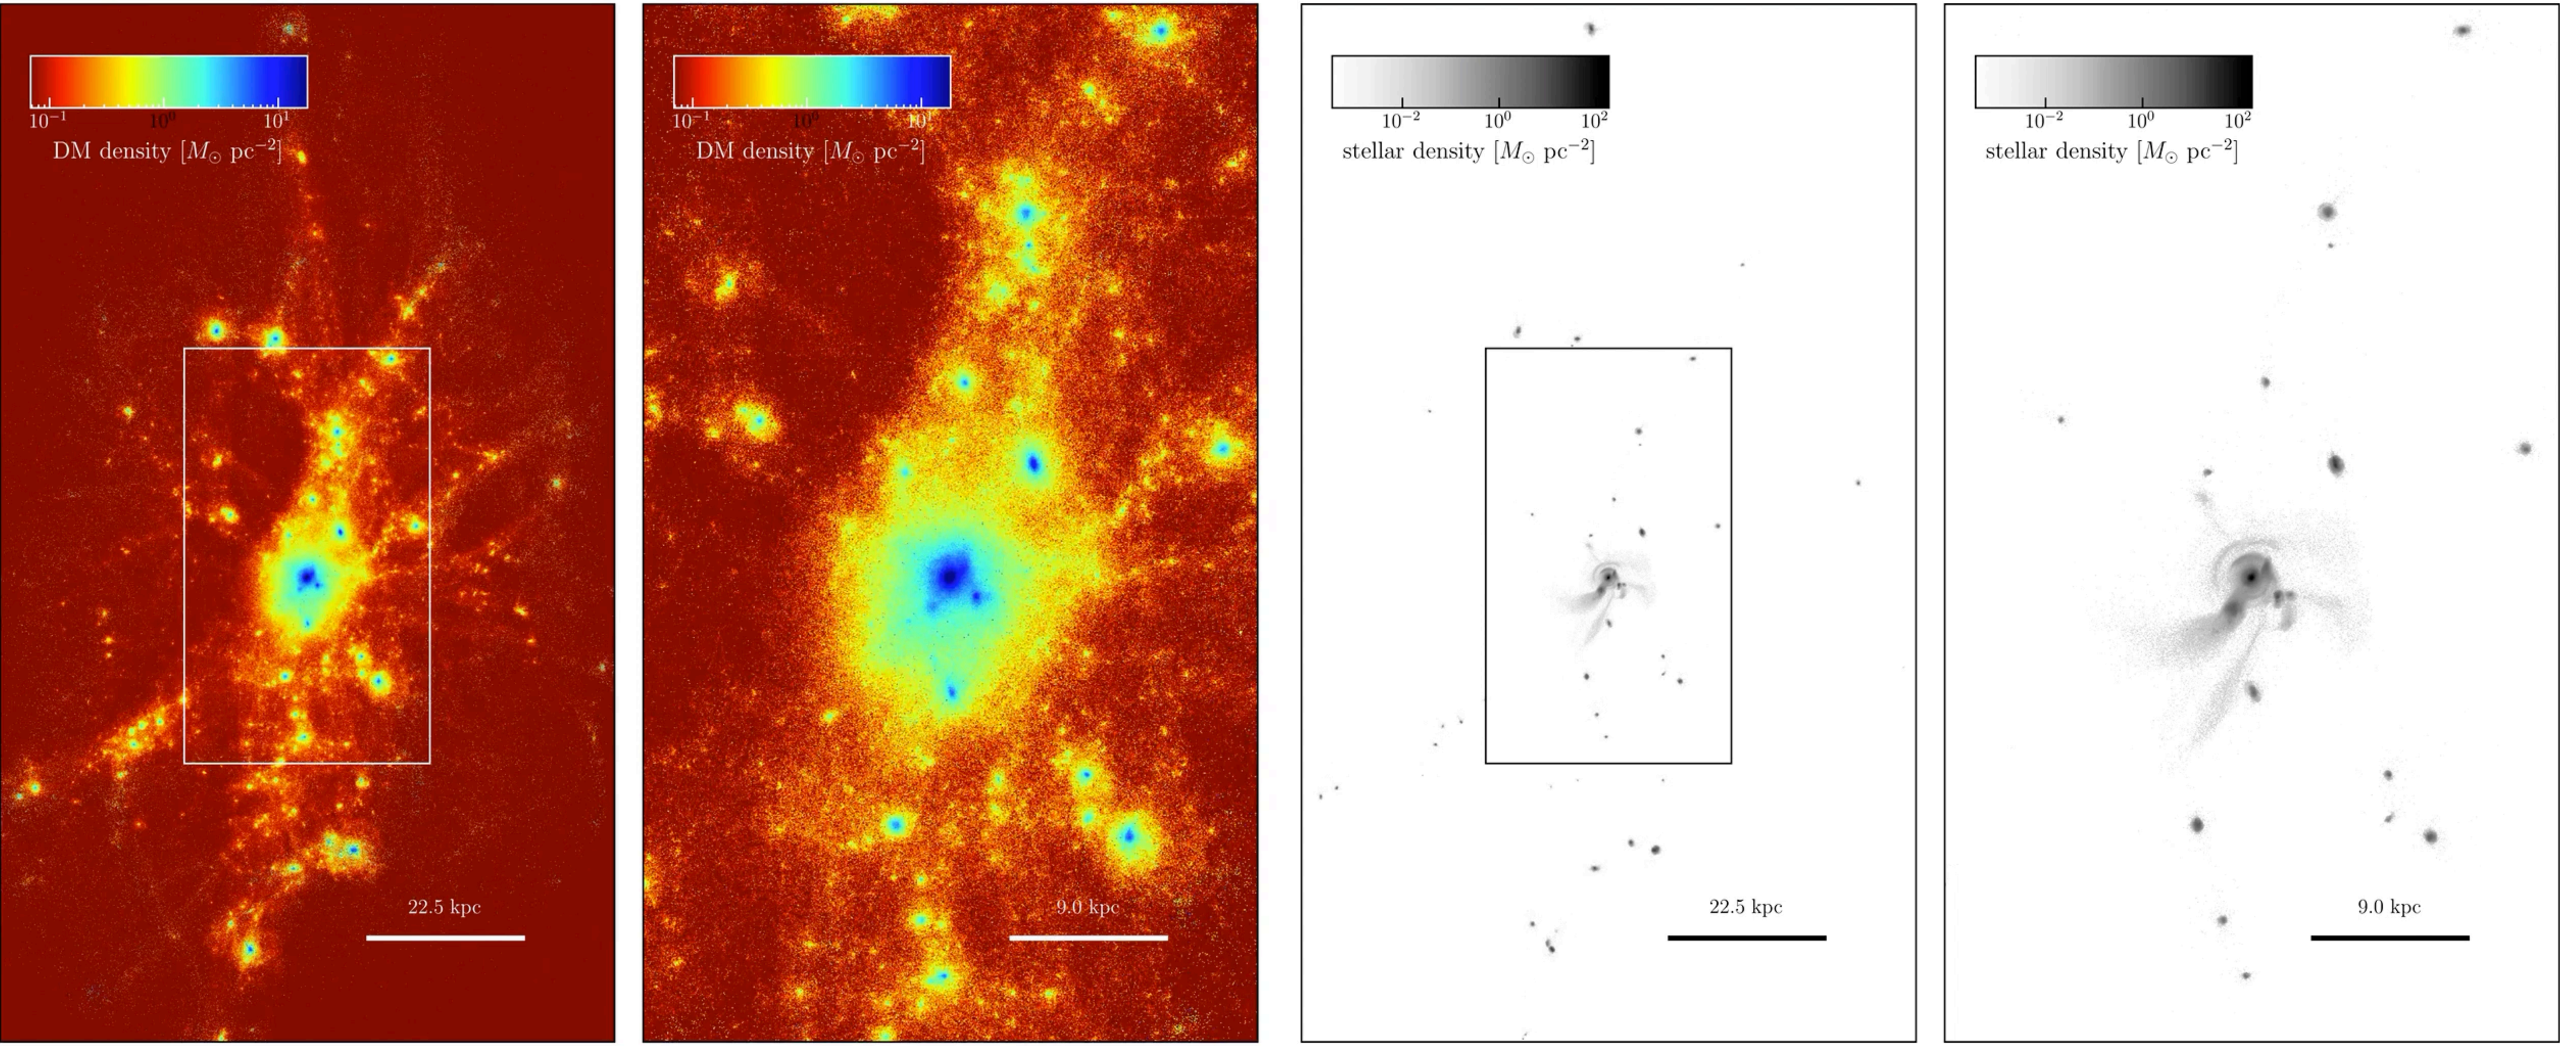

Stellar clumps

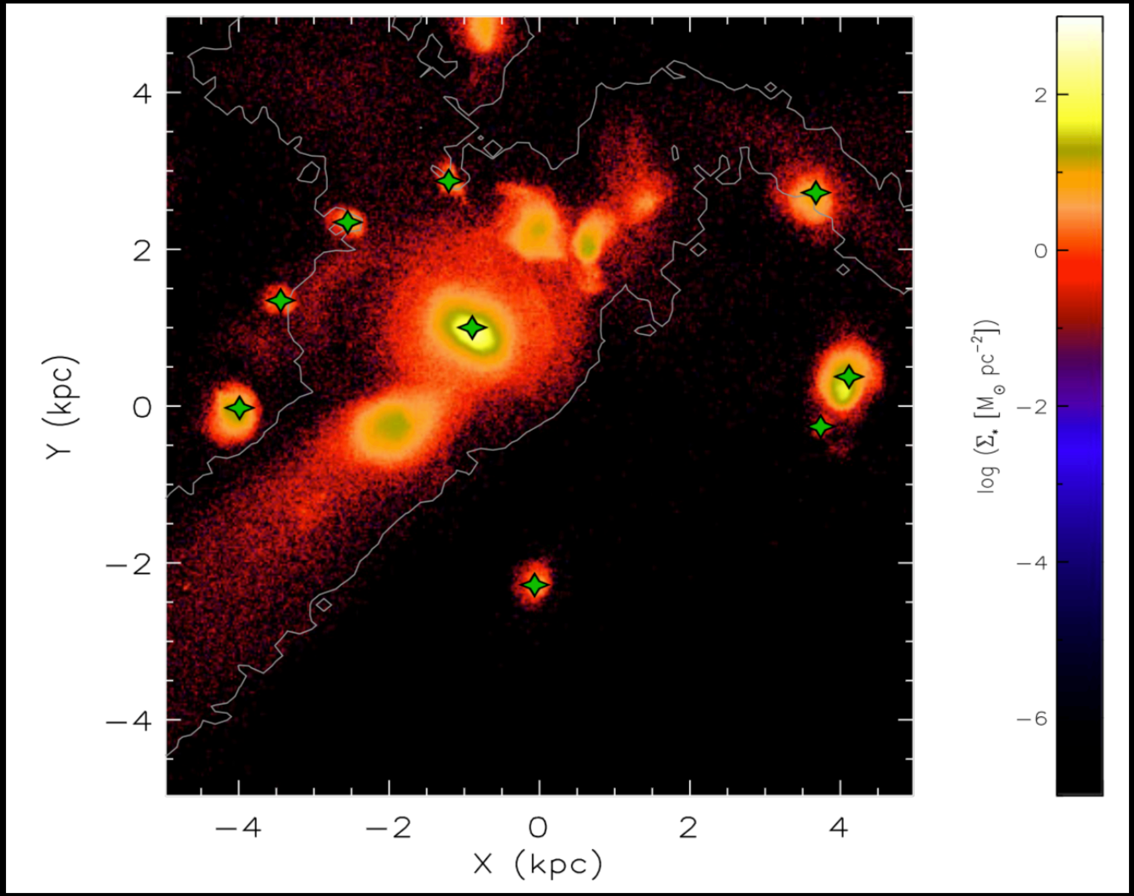

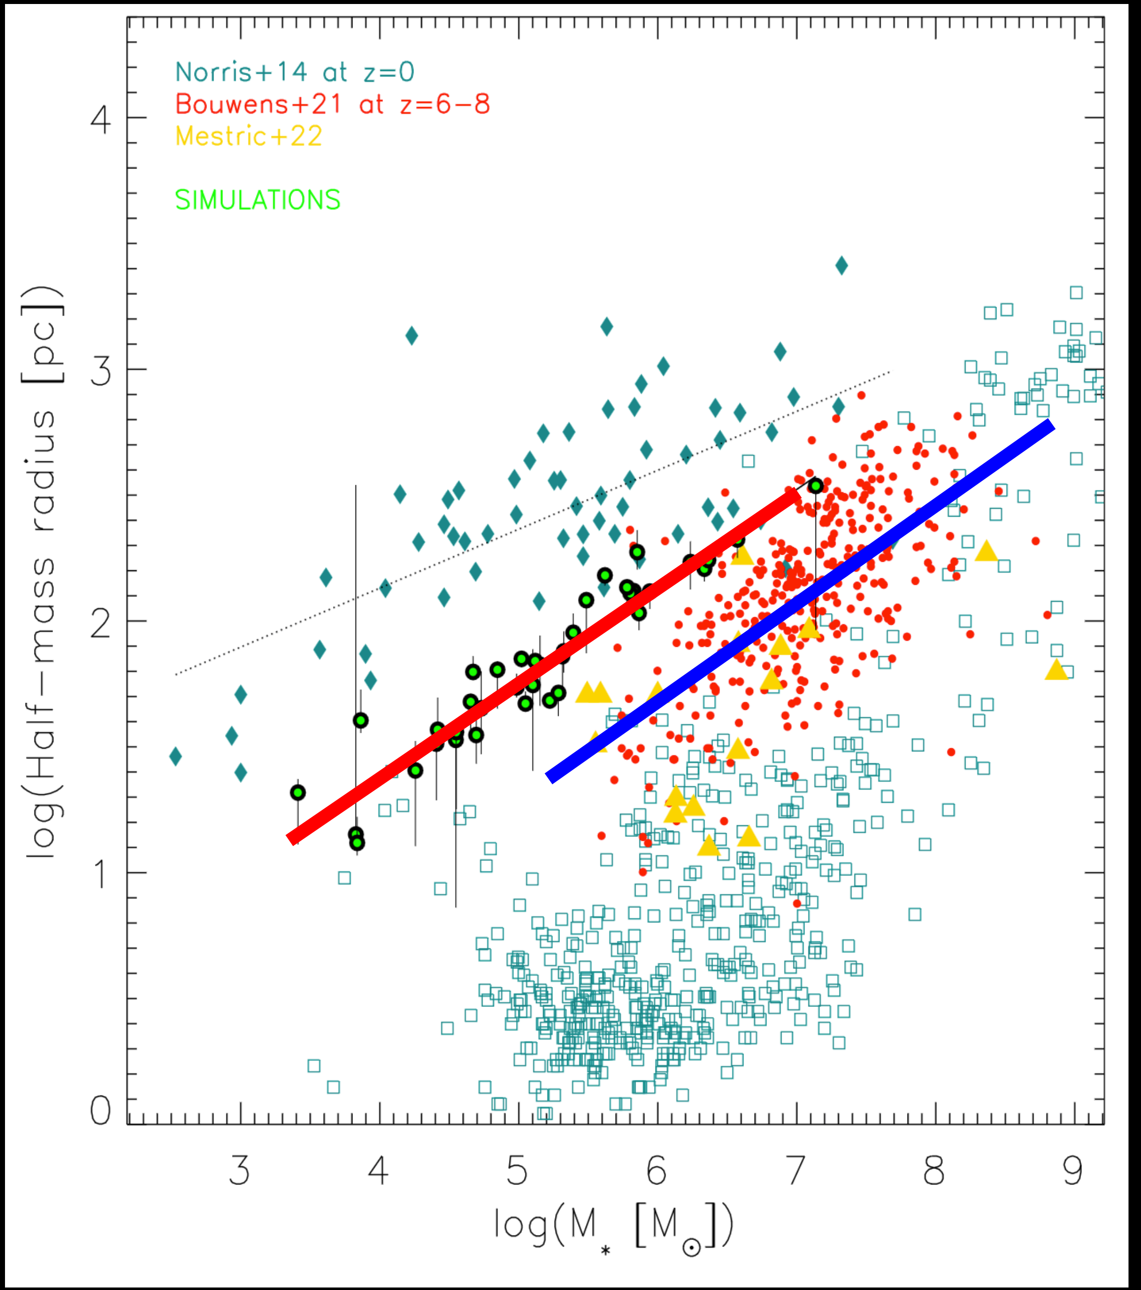

Masses comparable with the D1+T1 system but too large sizes!

New feedback and star formation recepies

Calura, EL+22

Stellar clumps shapes

High mass halos: formed stars

Low mass halos: no stars formed

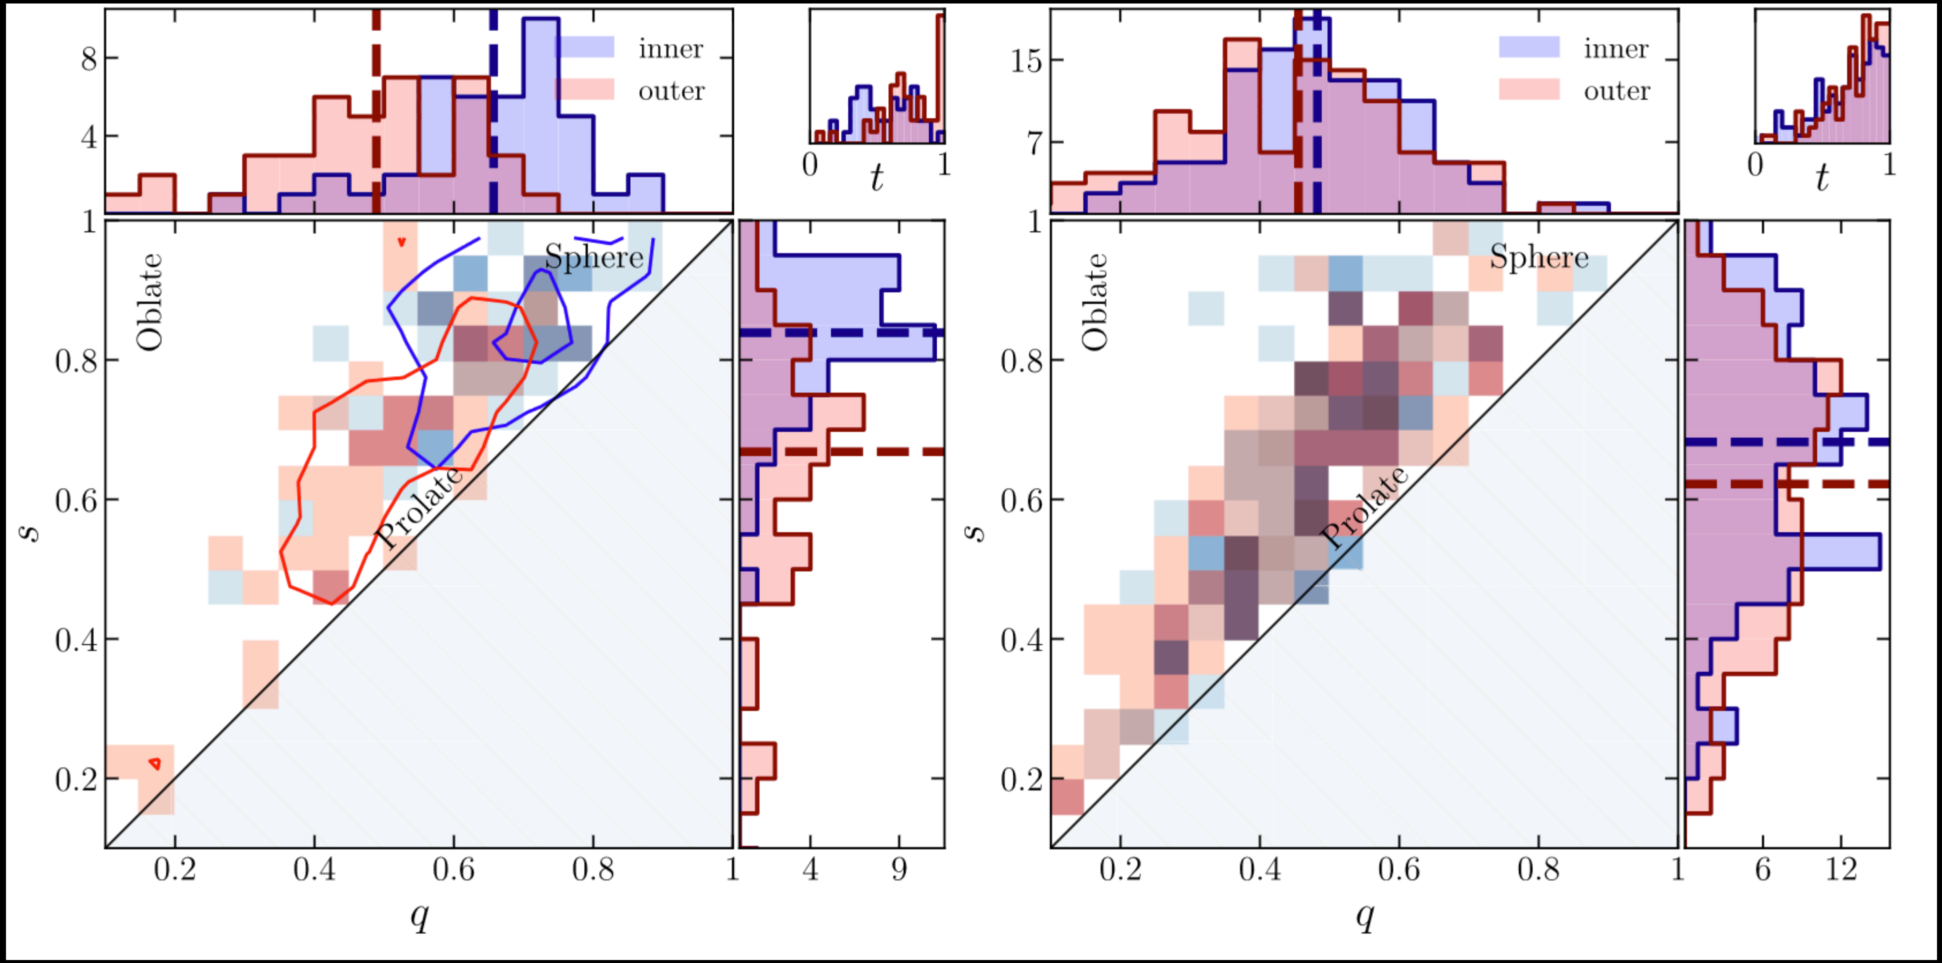

Pascale, EL+23

★ Spherical at the centre

★ Prolate in the outskirts

★ Prolate both in the centre and in the external part the outskirts

Future steps..

Isolated clusters

★Type Ia Supernovae effects on clusters of different masses

★Contribution of first generation massive stars to the iron spread

Cosmological simulations

★New and more realistic feedback recepies to form denser stellar clumps and study their shape and star formation histories

5M core hours used on Discoverer