


Phoebos Simulations into the Cosmic Dawn

Darren Reed (presenter), Pedro Capelo, Floor van Donkelaar, Lucio Mayer (UZH); Tom Quinn (UW Seattle)

- **Formation of first generations of galaxies**
- Cosmological Smoothed Particle Hydrodynamical simulations
- Allocation on Piz Daint (GPU Nodes) at CSCS
 - SKACH Computing Platforms and Infrastructure Program
 - July 2022 – June 2023
- (very) preliminary results & status update



Phoebos Simulations into the Cosmic Dawn

- Phoebos MR (“medium resolution”) simulation using **ChaNGa**
- 32 billion particles, 100Mpc volume, $M_{\text{gas}} < 10^6 M_{\text{sun}}$
 - Need *large volume* for ***statistics***. Small particle masses for ***resolution***.
- Currently at redshift 14. plan to reach $z \sim 5$.
- 100K Node hours used.
 - Job size: 2048 Nodes
 - Each Node. 12 cores/64GB + Nvidia Pascal P100 GPU

The PHOEBOS simulations roadmap:

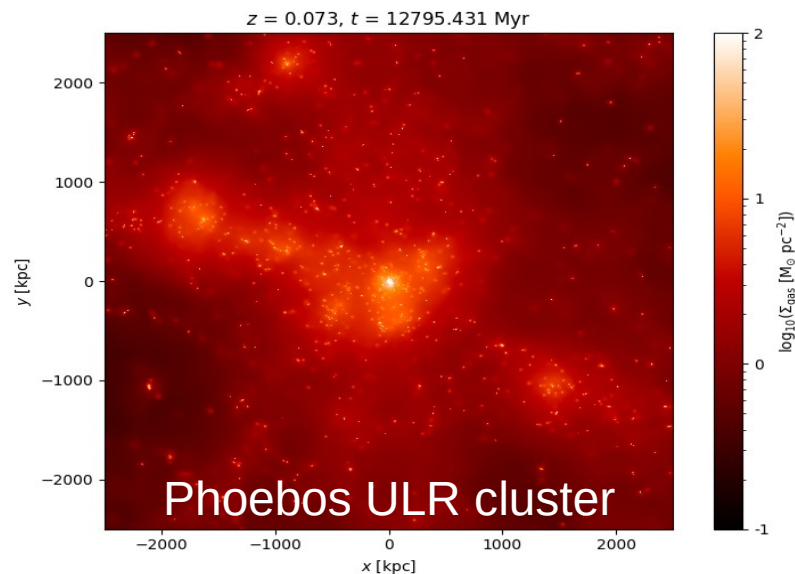
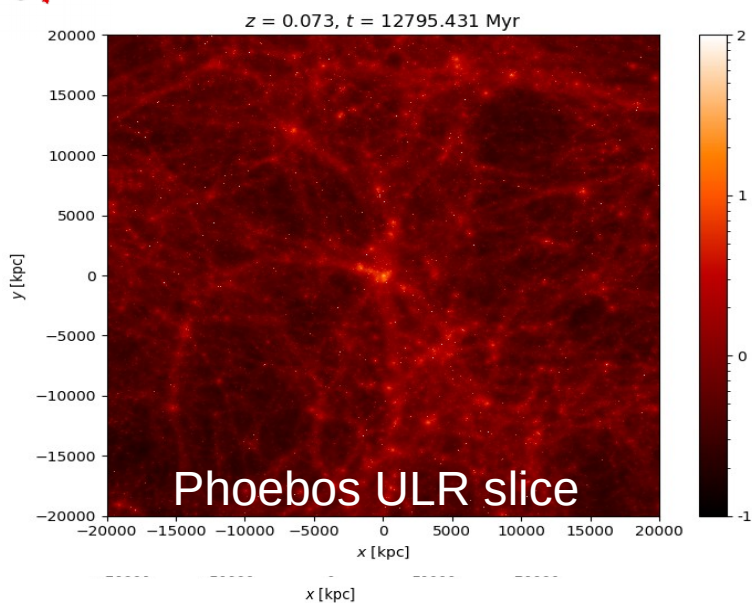
Large cosmological volume (100 comoving Mpc)

Run:	# DM	# gas	# tot	$m_{\text{DM}} [M_{\odot}]$	$m_{\text{gas}} [M_{\odot}]$	ϵ [kpc]	memory [kB]	# nodes
PhoebosHR	5808^3	3888^3	2.547×10^{11}	1.699×10^5	1.059×10^5	0.15	3.135×10^{11}	
PhoebosMR	2904^3	1944^3	3.184×10^{10}	1.360×10^6	8.473×10^5	0.30	3.918×10^{10}	2048
PhoebosLR	1452^3	972^3	3.980×10^9	1.088×10^7	6.778×10^6	0.60	4.898×10^9	
PhoebosULR	726^3	486^3	4.974×10^8	8.701×10^7	5.423×10^7	1.20	6.122×10^8	

ChaNGa {
 ✓
 ✓
 ✓

~~SPH~~
 EXA

→ **Eventual Goal:** EXA-PHOEBOS simulation **1 trillion** particles



Simulation code

- **ChaNGa** (Jetley et al. 2008, 2010; Menon et al. 2015)
 - gravity + hydrodynamical Smooth Particle Hydrodynamical (SPH) + astrophysics
 - Astrophysics: gas cooling; star formation; stellar feedback; ultra-violet background heating
 - Charm++ parallel programming framework
 - Scales efficiently with number of compute nodes
- Future bigger simulations plan to use **SPH-EXA**
 - PASC project: SPH-EXA2 (PI: Ciorba, CoPIs: Mayer, Cabezón)
 - see talks by Sebastian Keller + Osman Seckin Simsek; Axel Sanz
 - designed for exascale → more particles + faster
 - in development to required physics for cosmological simulations with star formation

Gas Density slice

100 x 100 x 0.4 Mpc


redshift=22



Gas Density slice

100 x 100 x 0.4 Mpc


redshift=15



Gas Density slice

100 x 100 x 0.4 Mpc

redshift=13.7



Some technical challenges so far

- Higher memory usage than expected
 - due to memory allocated for GPU communication and memory associated with file writing
 - → **crashes** in messy way if 1 node exceeds available memory
 - Workarounds: shorter jobs. Stop and restart checkpoints written & disable core dumps
- Data Storage (>~ Petabyte)
- Data Analysis (tools not optimized for huge data)

Analysis pipeline

- **Analysis pipeline products:**
 - *Images* (density maps, temperature maps)
 - *Galaxy identification* using group (halo) finders (e.g. Amiga Halo Finder, AHF)
 - *Galaxy properties*: formation, evolution, statistics (e.g. n vs time)
 - *Cosmic web*: Evolution of *gas distribution*, *gas inflows* to galaxies
- **Strategies**
 - “Divide and conquer” strategies needed for all analysis for standard tools (e.g. pynbody)
 - **Data too big** (2TB+/snapshot) to fit into RAM on available shared memory machines
 - *Simplest*: Cut simulation into subvolume files (e.g. slices + boundary region)
 - analyze each subvolume file independently
 - Inefficient but “always” possible
 - Better way: Efficiently parallelized analysis tools that output all needed statistics
 - Efficient but usually impractical (need rewrite most analysis tools)
 - *Ideal/perfect*: On the fly analysis tools integrated into the simulation code

1st galaxies region:

redshift 22

redshift 15


redshift 13.7

2Mpc



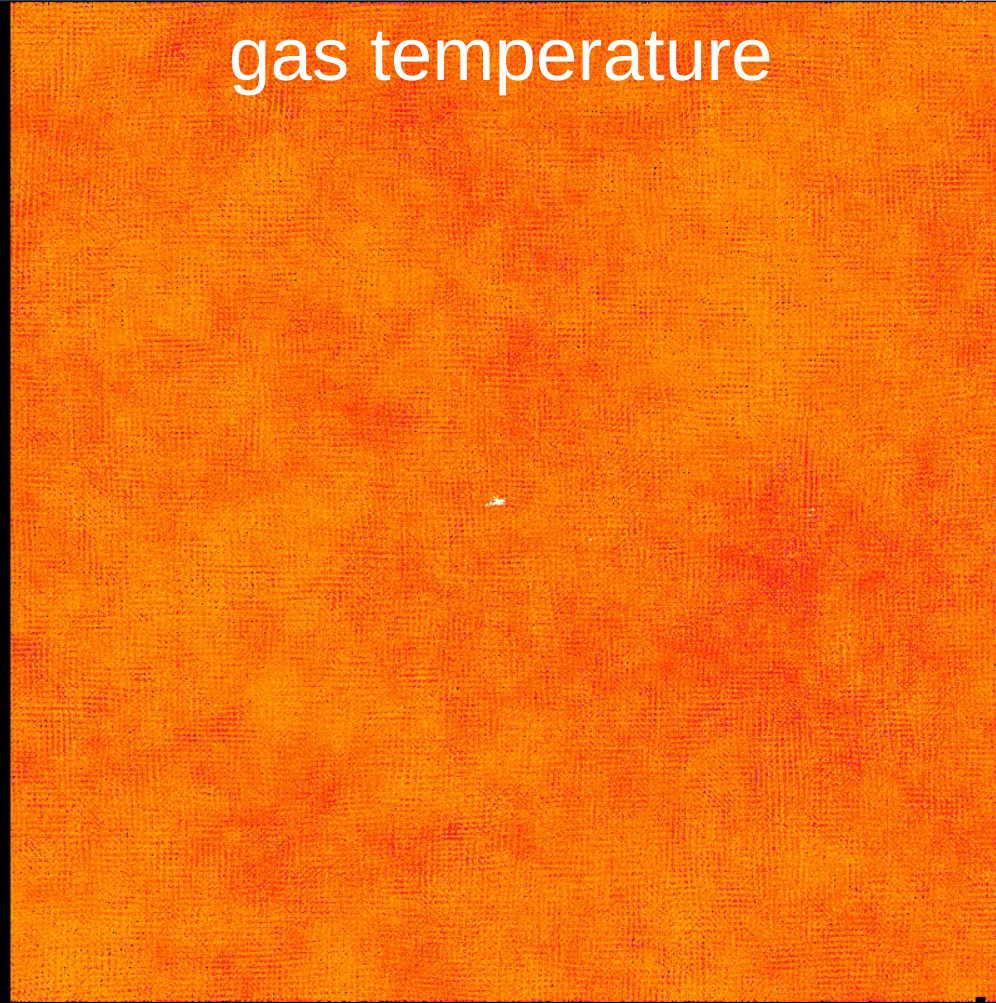
1st galaxies region – redshift 13.7

gas density



2Mpc


gas temperature




1st galaxies: simulation mass resolution

redshift 13.7


ULR: N=0.5 billion



LR: N=4 billion



MR: N=32 billion



Summary:

Phoebos Simulations into the Cosmic Dawn

- **Formation of first generations of galaxies**
 - galaxies resolved at redshift 14
- 3 quarters remain of Piz Daint allocation
- Work continues on analysis pipeline

End:
