Institut d'astrophysique de Paris

Potentially brilliant: a history of boundaries

The Random Universe conference, Imperial College London

Florent Leclercq

www.florent-leclercq.eu

Institut d'Astrophysique de Paris
CNRS & Sorbonne Université

2 July 2026

CNRS

SORBONNE UNIVERSITÉ

HOW A PAPER COMES INTO BEING

Simbelmynë: an end-to-end code for synthetic galaxy surveys

- Our code Simbelmynë is publicly available at https://simbelmyne.florent-leclercq.eu/.

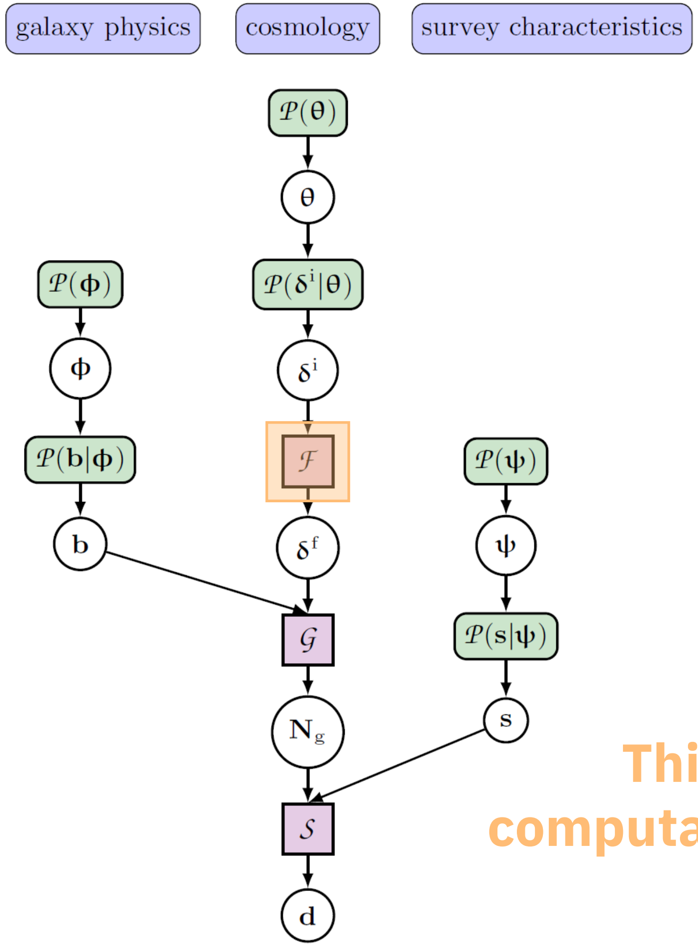

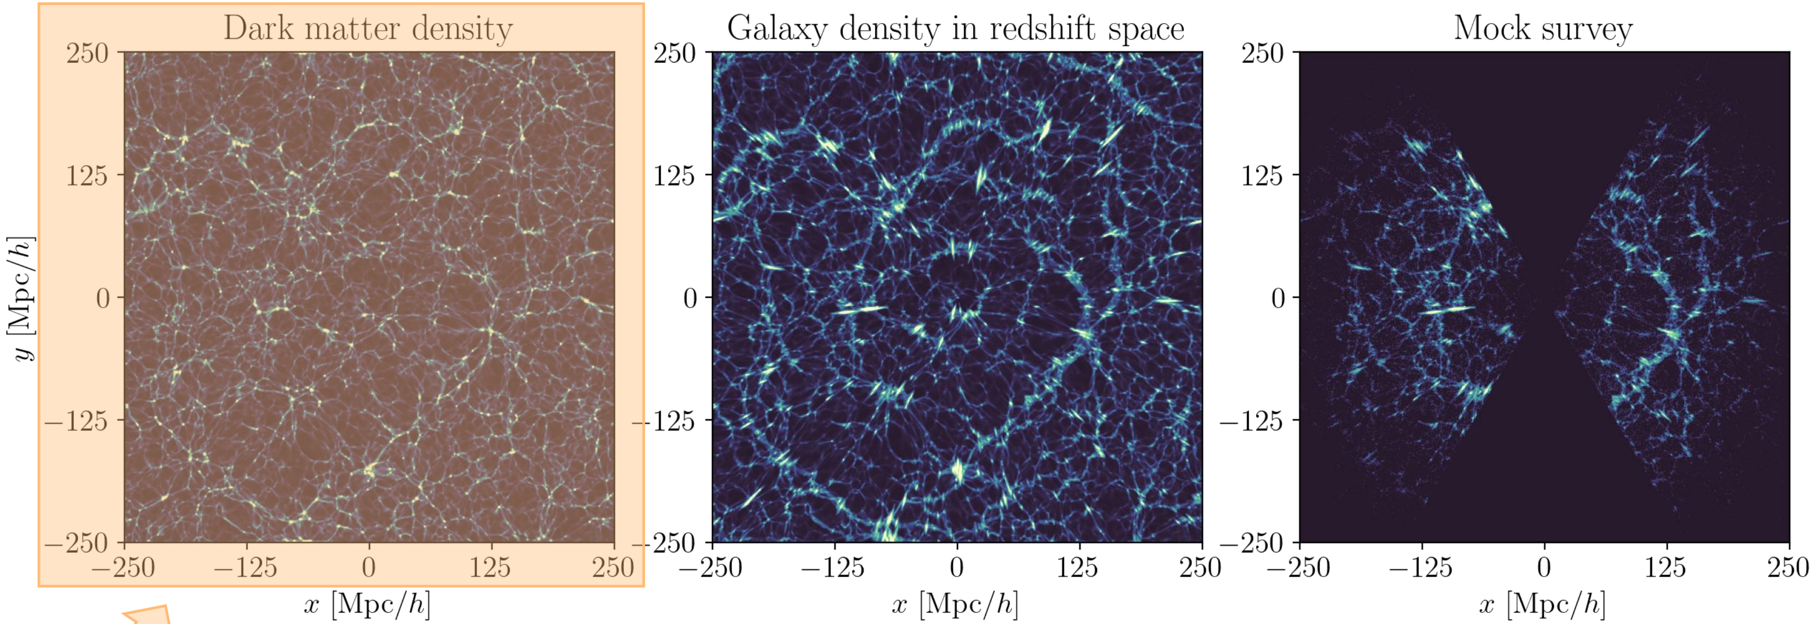

This step is the computational bottleneck

FL, Jasche & Wandelt, 1502.02690

"How fair are the bright eyes in the grass! Evermind they are called, simbelmynë in this land of Men, for they blossom in all the seasons of the year, and grow where dead men rest."

The tCOLA framework: temporal COmoving Lagrangian Acceleration (2013)

- Idea behind tCOLA: we can make use of the analytical solution at large scales and early times: Lagrangian perturbation theory (LPT).

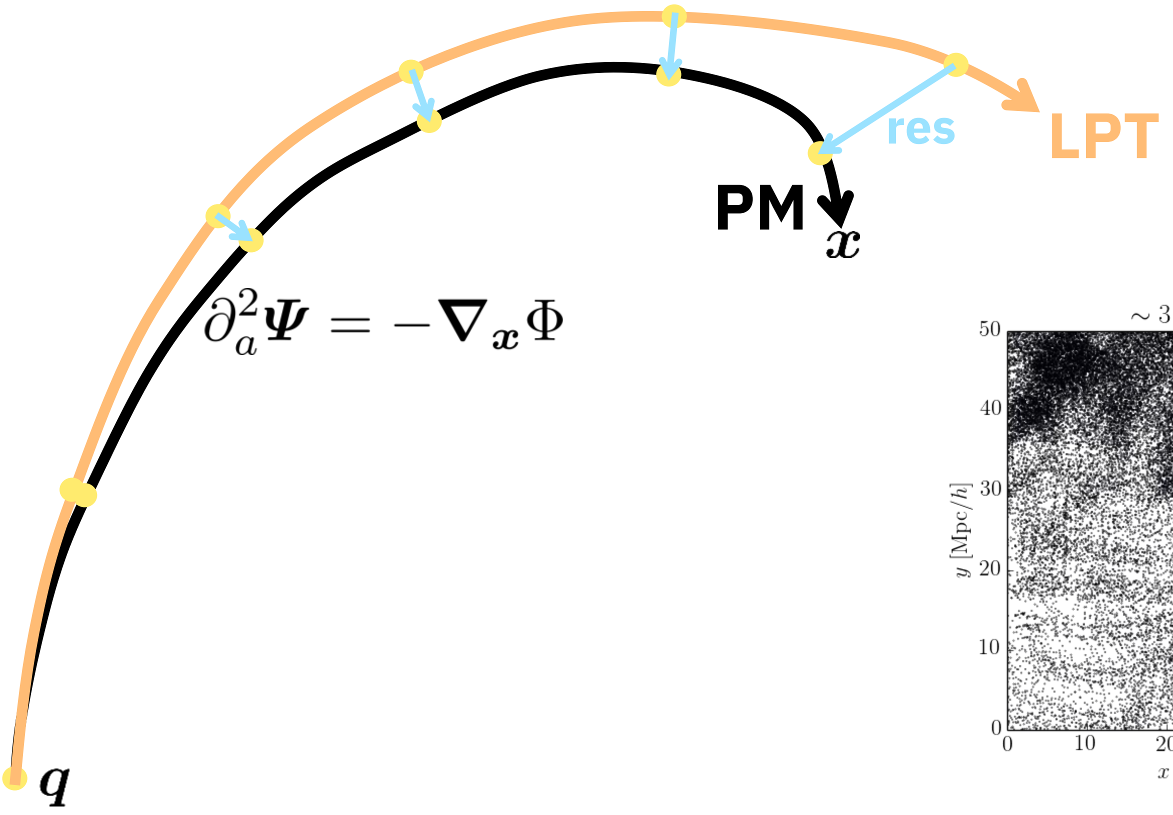

- Write the displacement vector as:

$$\boldsymbol{\Psi} = \boldsymbol{\Psi}_{\mathrm{LPT}} + \boldsymbol{\Psi}_{\mathrm{res}}^{\mathrm{COLA}} \qquad (\boldsymbol{x} = \boldsymbol{q} + \boldsymbol{\Psi})$$

Tassev & Zaldarriaga, 1203.5785

- Equation of motion (omitted constants and Hubble expansion):

Analytical solutions!

$$\partial_a^2 \boldsymbol{\Psi}_{\mathrm{res}}^{\mathrm{COLA}} = \partial_a^2 (\boldsymbol{\Psi} - \boldsymbol{\Psi}_{\mathrm{LPT}}) = -\boldsymbol{\nabla}_{\boldsymbol{x}} \Phi - \partial_a^2 \boldsymbol{\Psi}_{\mathrm{LPT}}$$

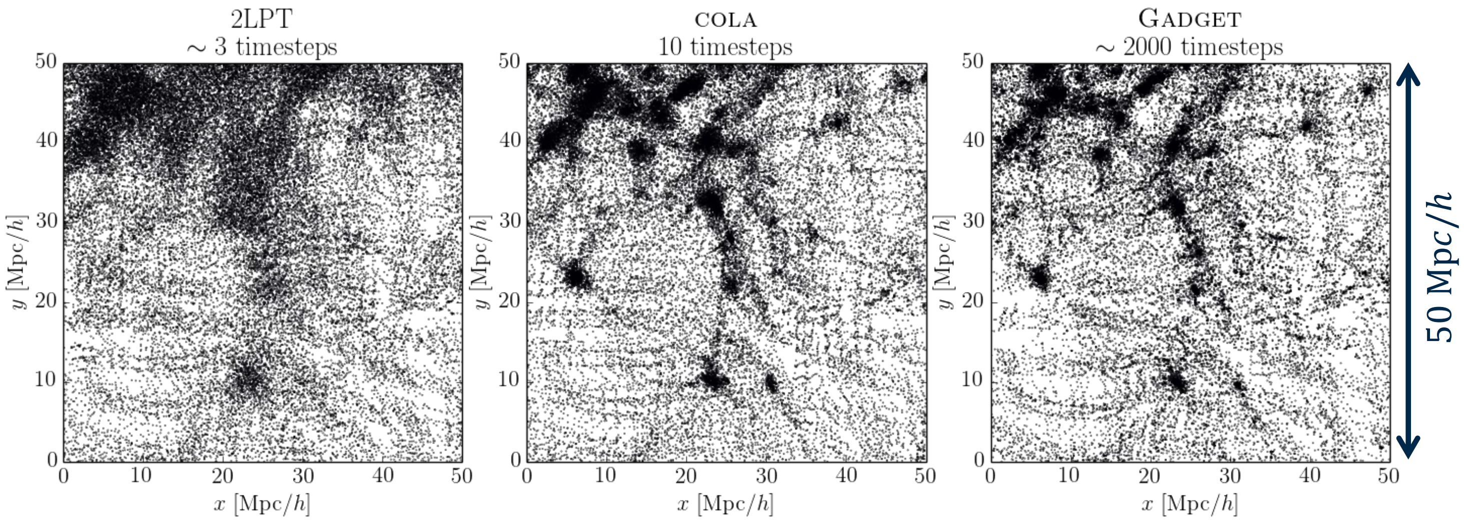

Tassev, Zaldarriaga & Einsenstein, 1301.0322

The real challenge of parallelisation: long-range gravity

- tCOLA provides a pleasant reduction in the number of timesteps...
- ...but: the real challenge, preventing the easy parallelisation of *N*-body codes, is the long-range nature of gravitational interactions:
  - It is a communication-dominated problem.
  - It requires sophisticated memory management to maximise locality in the storage hierarchy (cache, RAM, disk I/O).
- Amdahl's law: latency kills the gains of parallelisation.
  - For example, Gadget-4 loses strong scaling before 1,000 cores due to domain decomposition and long-range gravity.

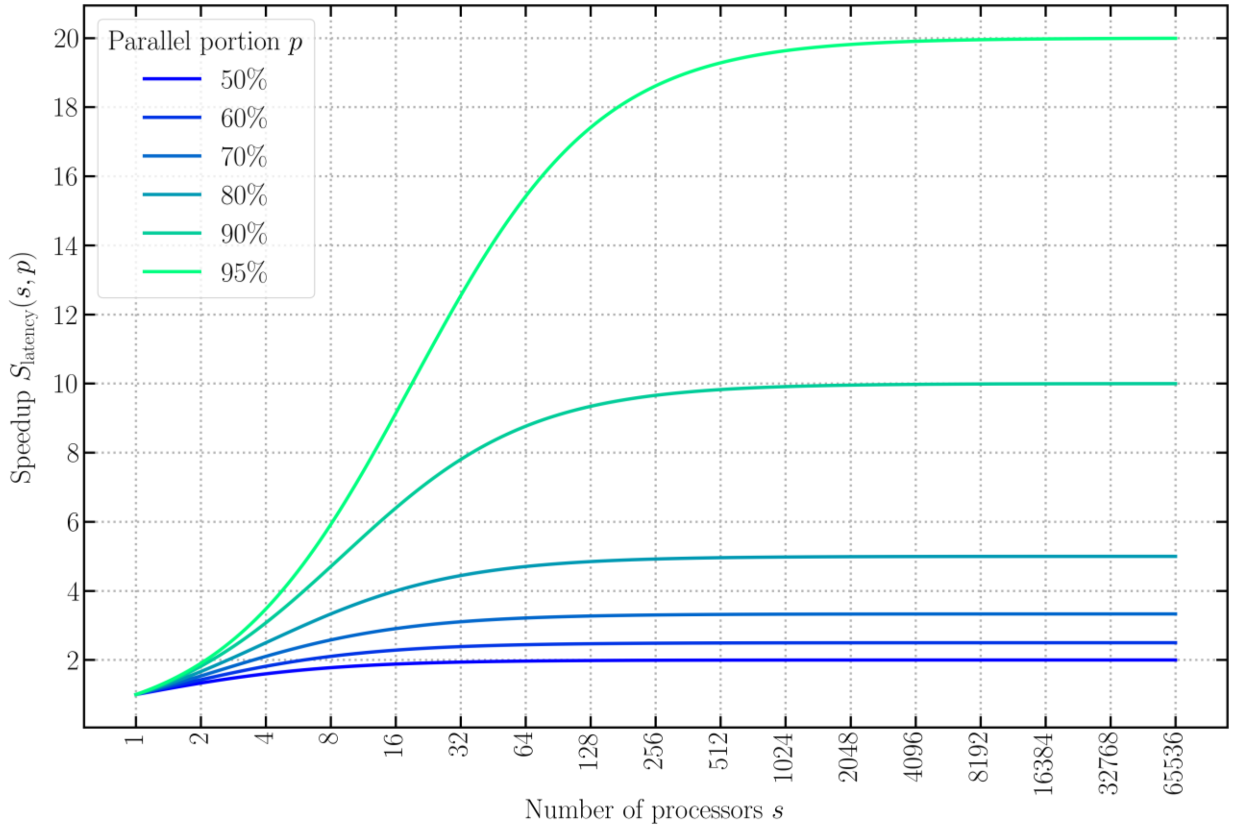

Are high-performance communications really needed to solve a quasi-linear problem at large scales? Shouldn't locality in memory come by construction?

Amdahl 1967, doi:10.1145/1465482.1465560, Springel *et al.*, 2010.03567 (Figure 63)

Let's define some concepts

- Perfect parallelism: dividing a computational task into independent sub-tasks with no communication between them, allowing for highly efficient parallel processing.
  - It is a.k.a. an "embarrassingly parallel workload" — but there's nothing to be embarrassed about, really.
  - Examples: computer simulations comparing many independent scenarios, ensemble calculation of i.i.d. numerical model predictions (e.g. for covariance matrix estimation).
  - Counter-examples: usual *N*-body simulation codes, recursive algorithms, Markov Chain Monte Carlo.

- Machine-Learning safety: applying machine learning (ML) — in particular neural networks (NNs) — in a way that ensures the results are either correct by construction or, at worst, suboptimal.
  - Safe uses of ML eliminates the requirement for explainability.
  - Examples: data compression, e.g. denoising autoencoders (DAE) to build summaries, information-maximising neural networks (IMNN) for implicit likelihood inference; Learning the Universe by Learning to Optimise (LULO)

Charnock *et al.*, 1802.03537, Makinen *et al.*, 2107.07405, Makinen *et al.*, 2410.07548, Doeser *et al.*, 2502.13243

  - Counter-example: emulation of *N*-body simulations. There remains an emulation error [up to $\mathcal{O}(10\%)$] that we cannot ever correct for.

He *et al.*, 1811.06533, Lucie-Smith *et al.*, 1802.04271, Jamieson *et al.*, 2206.04594, Conceição *et al.*, 2304.06099, Doeser *et al.*, 2312.09271, Jamieson *et al.*, 2408.07699

A first attempt at spatial decoupling (2015)

- Computing the LPT reference frame suggests a new strategy: using the change of frame of reference to split simulations spatially.
- The proof of concept used one sub-box embedded into larger simulation box.
- The sub-box is evolved using fast Fourier transforms (FFTs) to compute forces, which incorrectly assumes periodic boundary conditions.

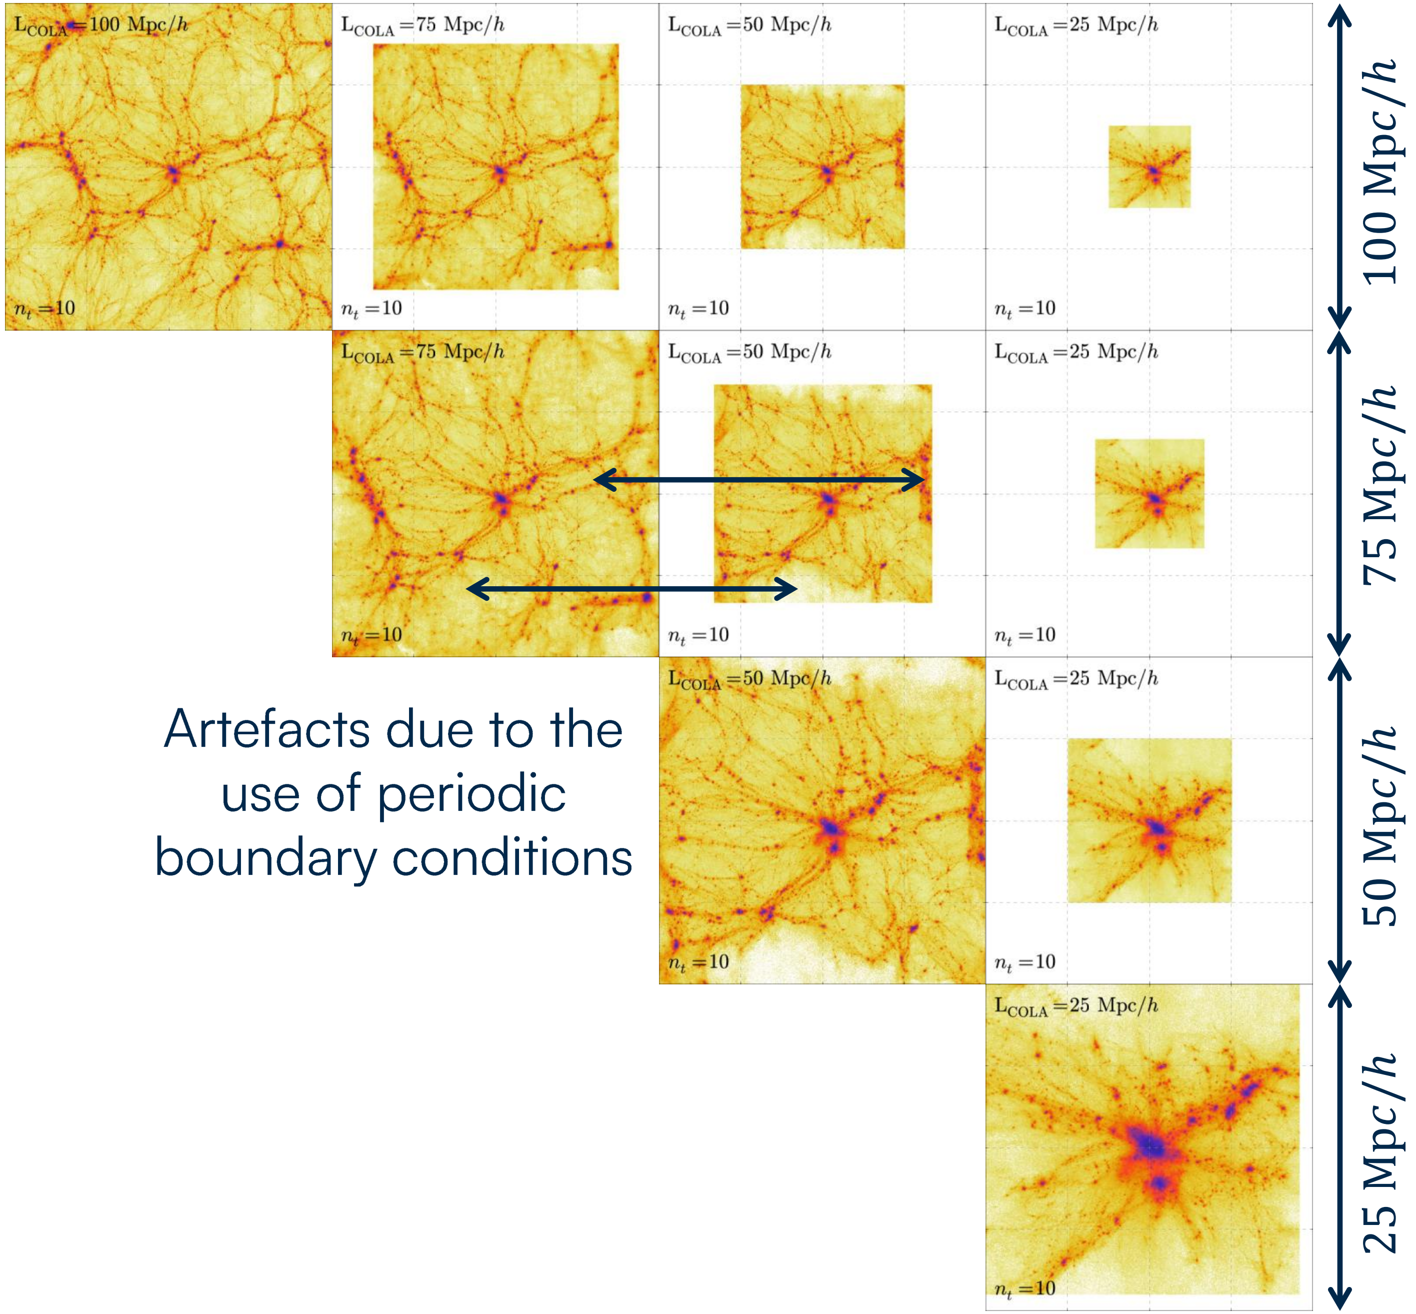

Tassev, Eisenstein, Wandelt & Zaldarriaga, 1502.07751

Perfectly parallel cosmological simulations using spatial comoving Lagrangian acceleration (sCOLA)

- Can we decouple sub-volumes by using the large-scale solution?

LPT (analytical solution) → sCOLA

$$\partial_a^2 \boldsymbol{\Psi} = -\boldsymbol{\nabla}_{\boldsymbol{x}} \underbrace{\left[\Delta^{-1}\delta\right]}_{\text{non-local}} \Leftrightarrow \partial_a^2 \left(\boldsymbol{\Psi} - \boldsymbol{\Psi}_{\mathrm{l.s.}}\right) = -\boldsymbol{\nabla}_{\boldsymbol{x}} \underbrace{\left[\Delta^{-1}\left(\delta - \delta_{\mathrm{l.s.}}\right)\right]}_{\text{local}}$$

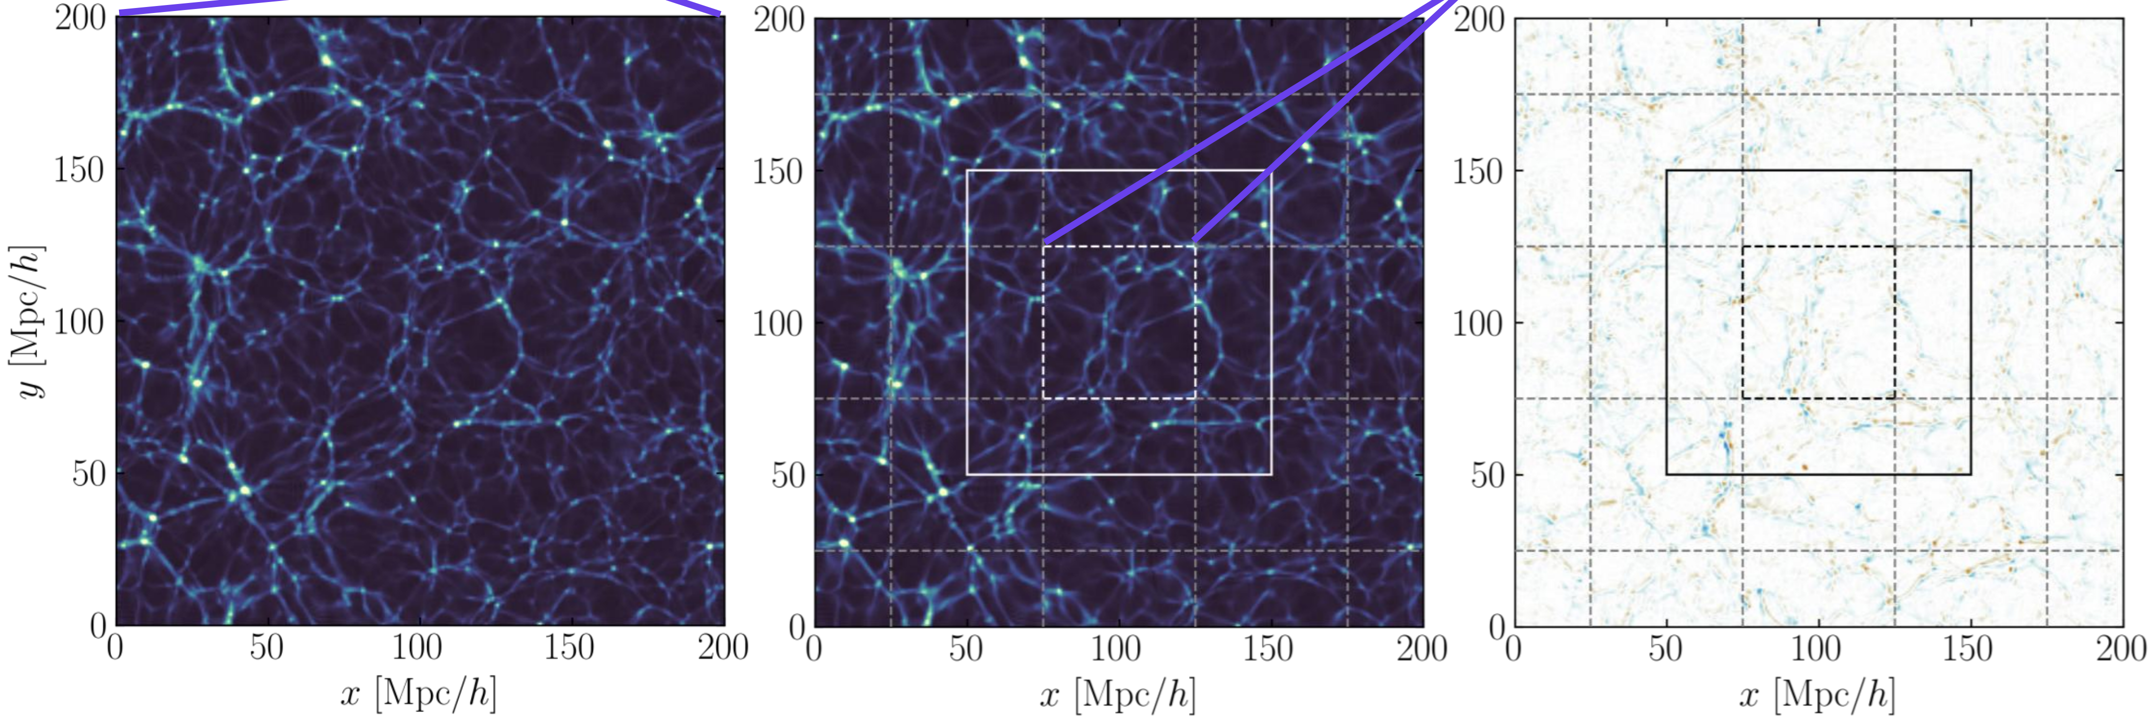

tCOLA (reference)

sCOLA

Difference

FL, Faure, Lavaux, Wandelt, Jaffe, Heavens, Percival & Noûs, 2003.04925

The ideas that didn't work (2016-2017)

- Using a Eulerian patching of the density field (does not conserve mass)

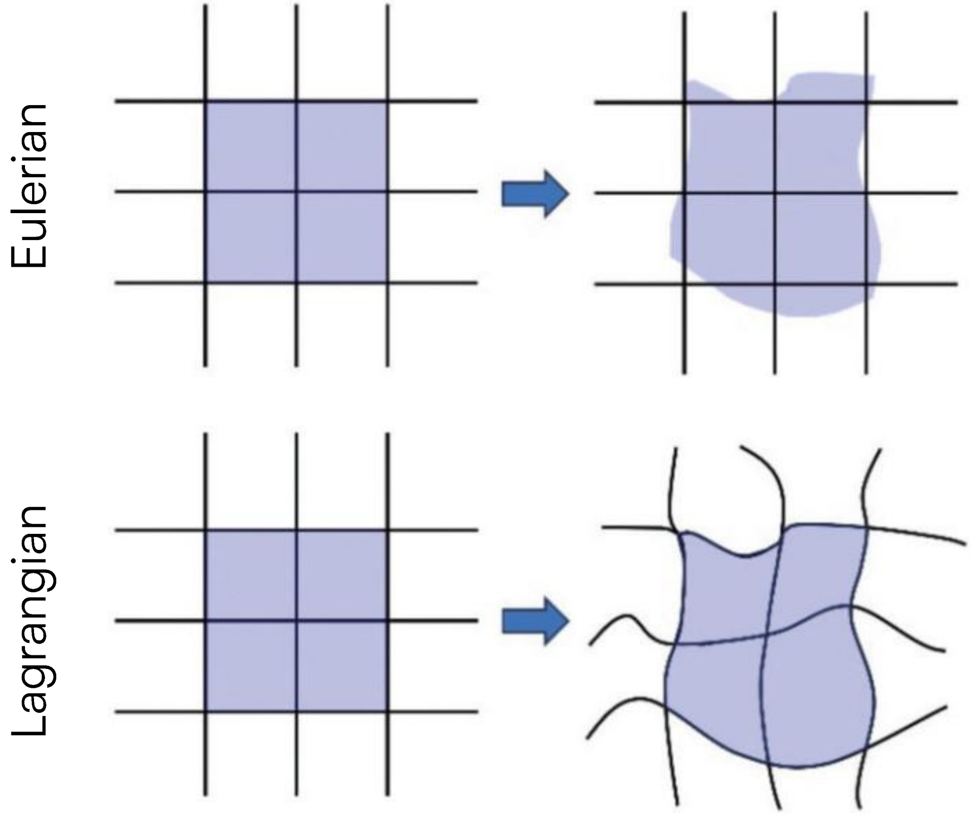

- Keeping periodicity (through FFTs), even with large buffer regions and some cooking recipes there (not accurate enough)

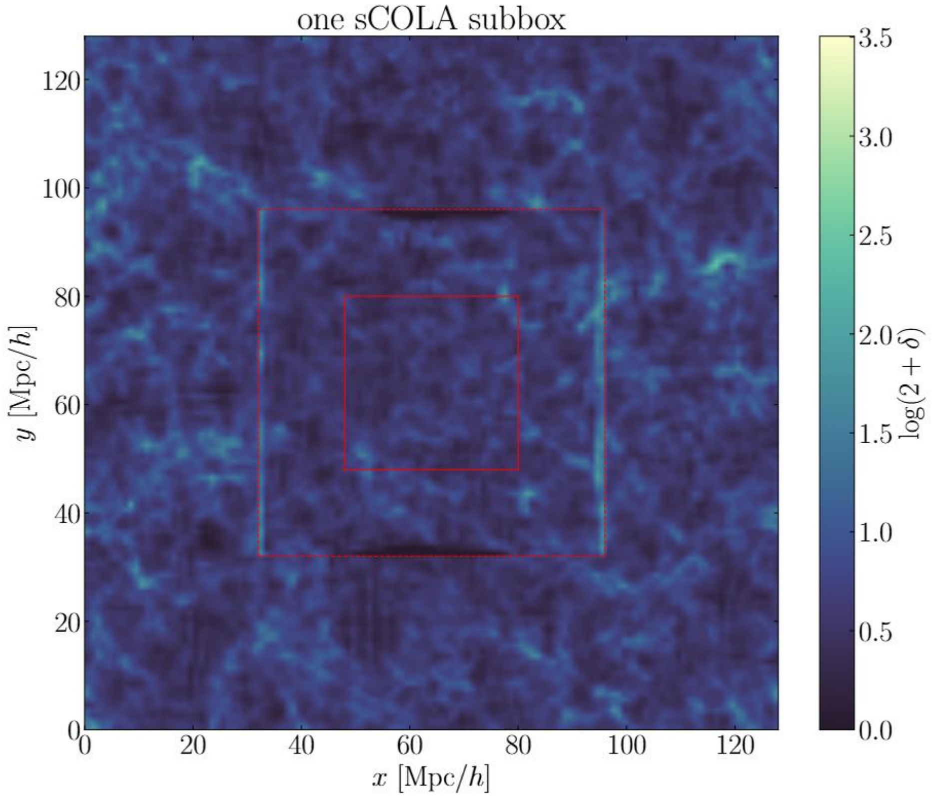

Baptiste Faure's internship report (2017)

Andrew's suggestion (2018): why don't you use the *exact* boundary conditions?

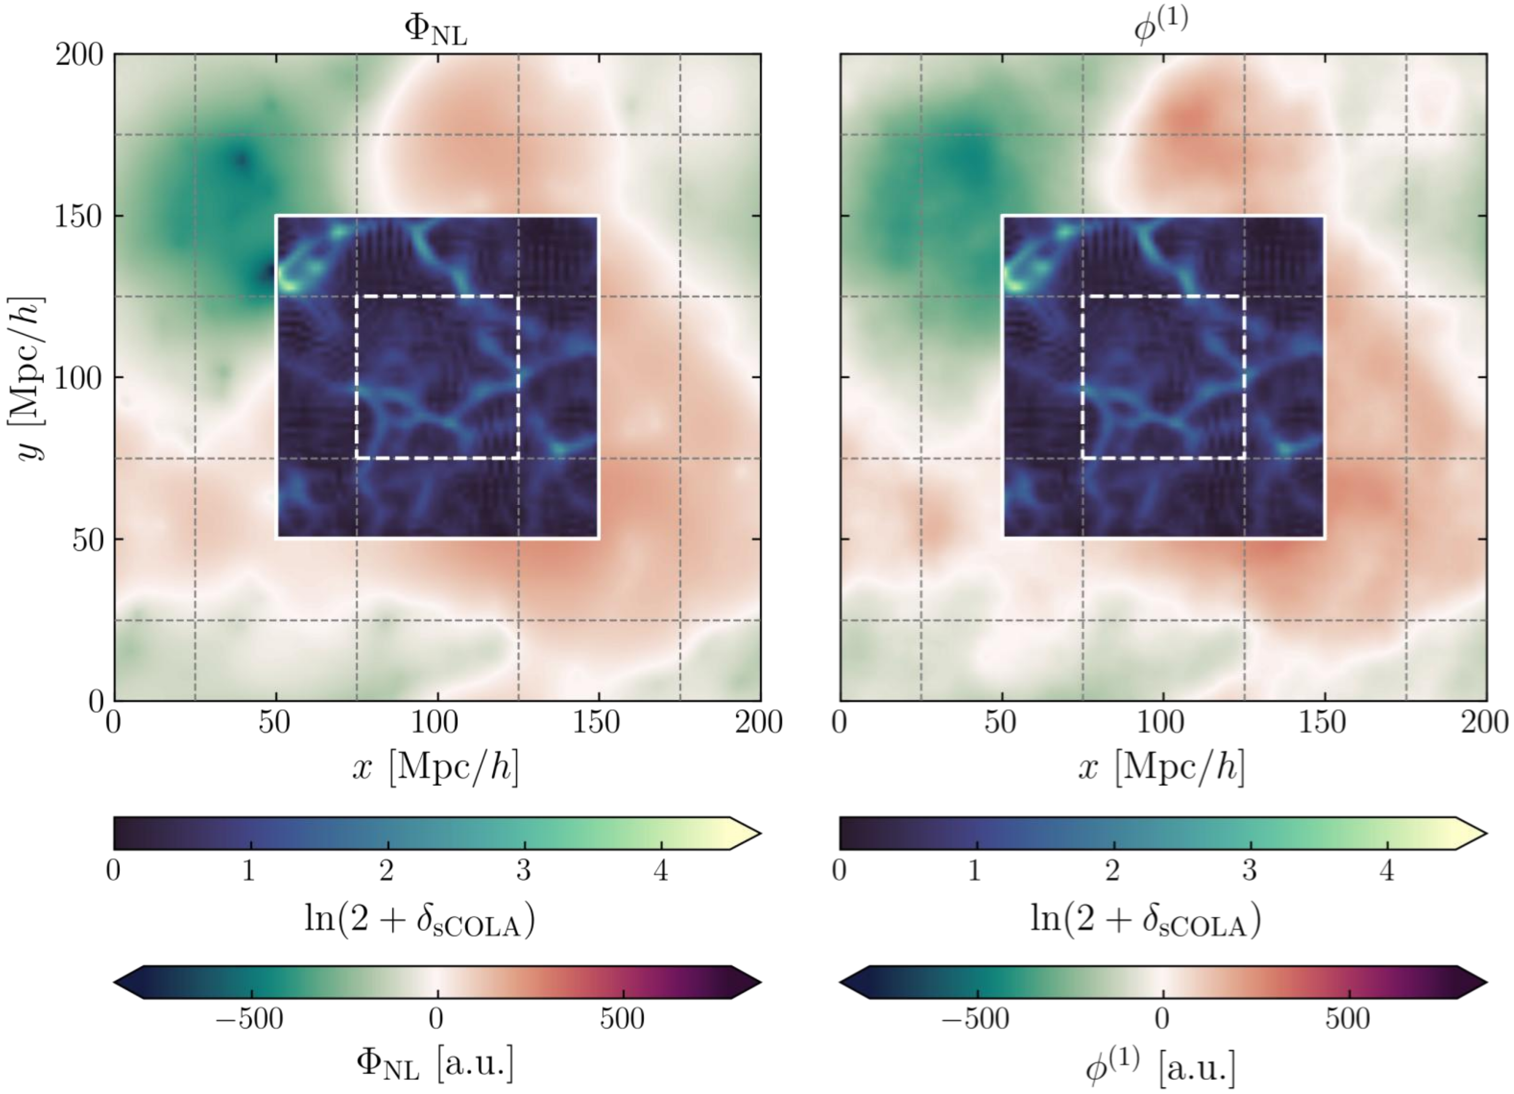

- Of course, we don't know the exact boundary conditions in the full volume.
- But we can use the linearly-evolving potential approximation (LEP):

We need it already!

$$\Phi_{\mathrm{BCs}}(\boldsymbol{x}, a) \approx D_1(a)\phi^{(1)}(\boldsymbol{x})$$

Brainerd *et al.* 1993; Bagla & Padmanabhan 1994

The idea was already old back then!

- Then we are faced with a Dirichlet problem: solving the Poisson equation with fixed (non-periodic) boundary conditions. There exists a solution since the 1970s.

The solution to boundary artefacts and the perfectly parallel algorithm (2020)

1. A buffer region around each tile

2. Appropriate Dirichlet boundary conditions for the potential

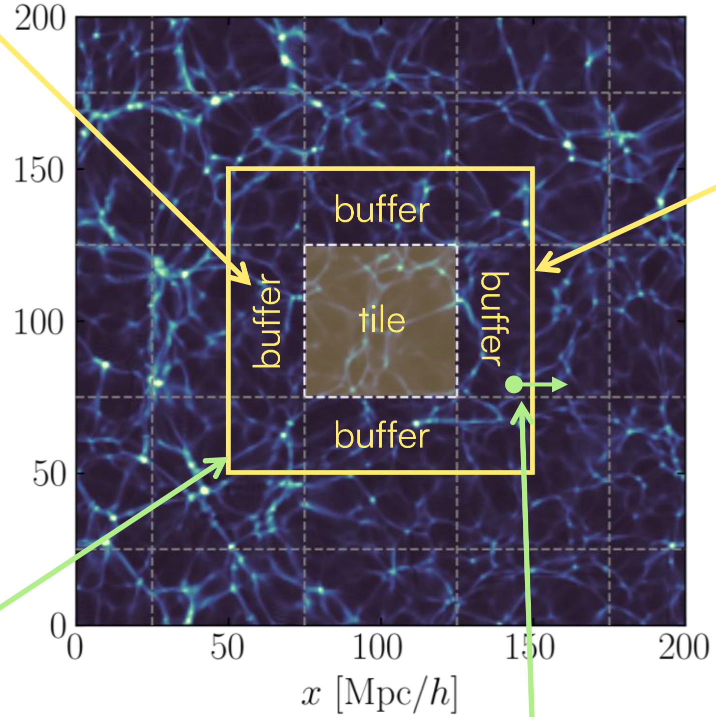

The Poisson solver uses discrete sine transforms (DSTs) instead of FFTs.

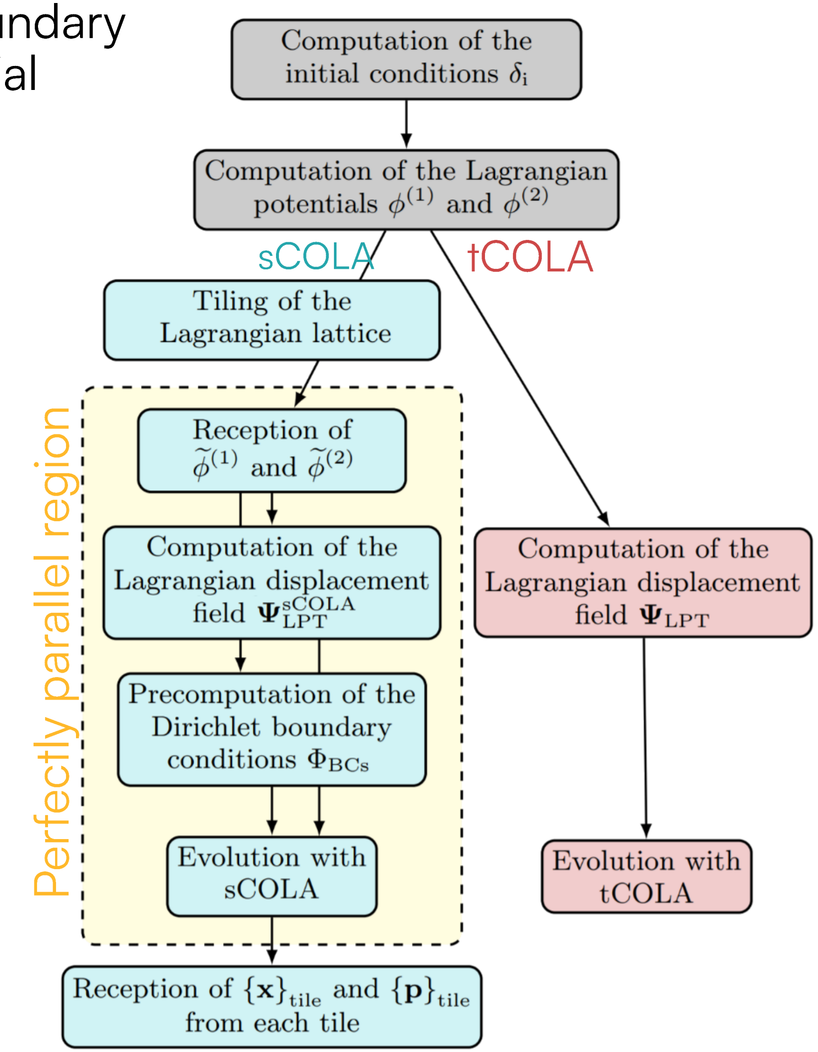

Two approximations (to ensure no communication between tiles):

1. Linearly-evolving potential (LEP) at the boundaries: $\Phi_\mathrm{BCs}(\boldsymbol{x}, a) \approx D_1(a)\phi^{(1)}(\boldsymbol{x})$

2. Outgoing particles do not deposit mass

FL, Faure, Lavaux, Wandelt, Jaffe, Heavens, Percival & Noûs, 2003.04925

Along these lines, we can do better in terms of accuracy! (2024-2026)

1. A second buffer of particles to approximate the inflow: the "LPT buffer".

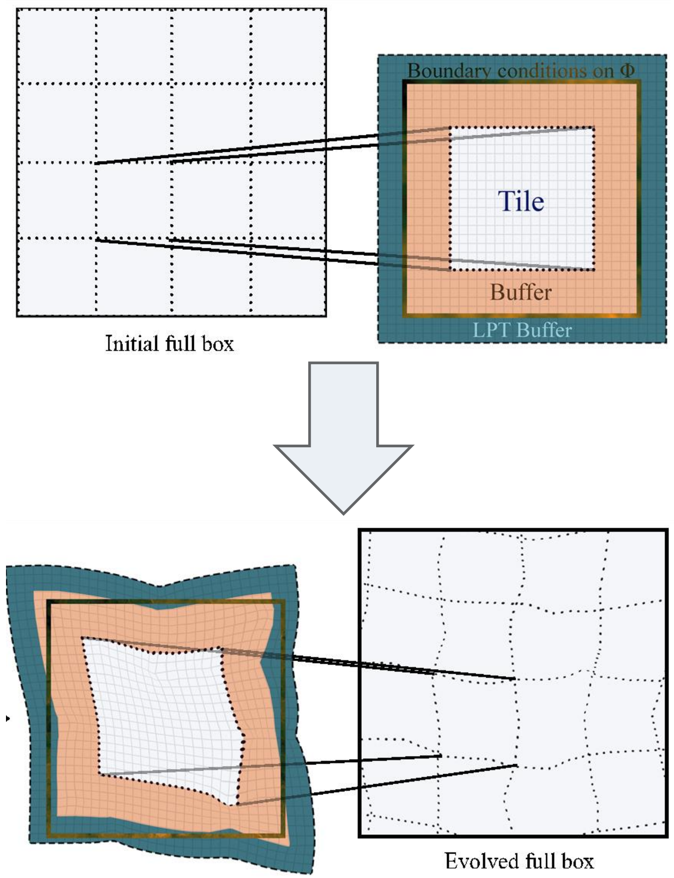

Aubin, FL & Lavaux, to be submitted

2. A second-order approximation for the gravitational potential at the boundaries.

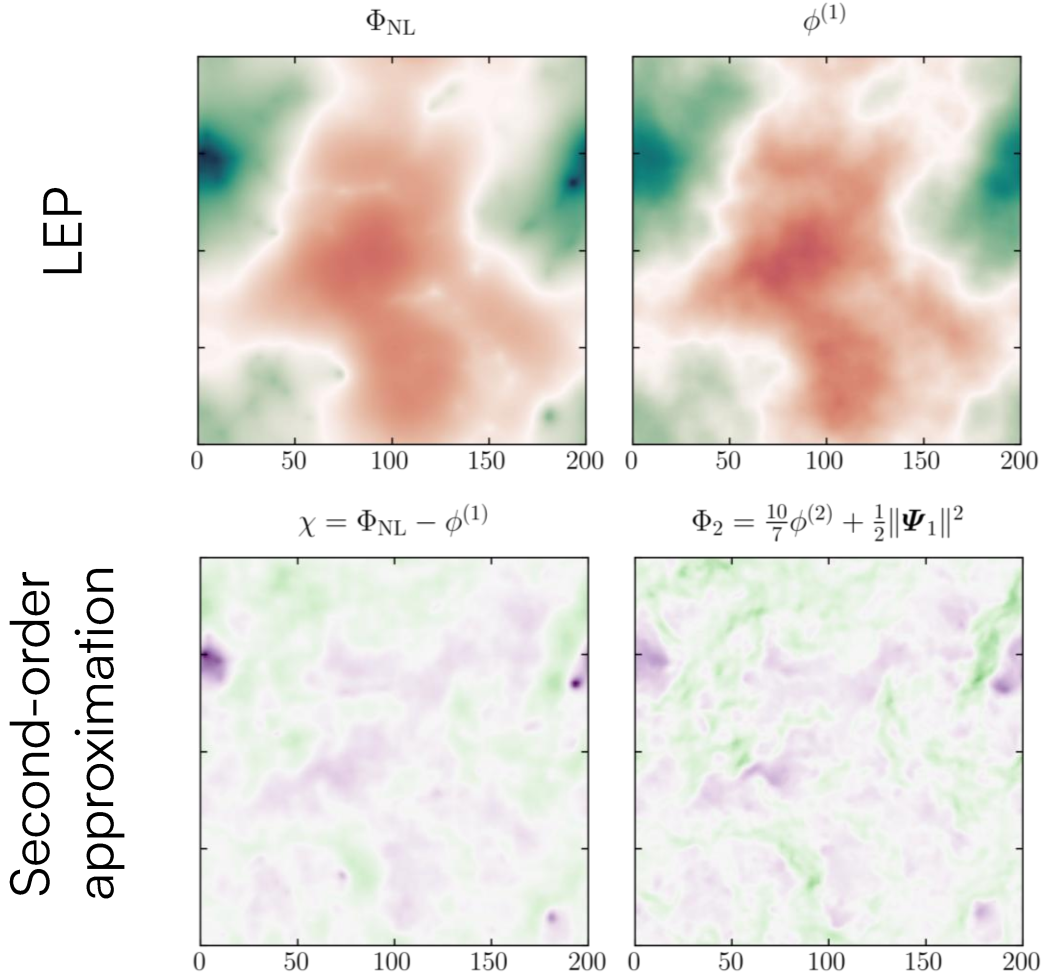

THE STORY CONTINUES: sCOLA LIGHTCONES

Lightcones and mock catalogues with sCOLA

- The workload in sCOLA is perfectly parallel, with a parallelisation potential factor:

$$\boxed{p = s\left(\frac{L}{L_\mathrm{sCOLA}}\right)^3}$$

hardware "boost" factor due to small memory requirements

- Generation of lightcones and mock catalogues:
  - sCOLA boxes only need to run until they intersect the observer's past lightcone.
  - Most of the high-z volume will run faster than $z = 0$.
  - Many unobserved sCOLA boxes do not even have to run!
  - The wall-clock time limit is the time for running a single sCOLA box to $z = 0$ at the observer's position.

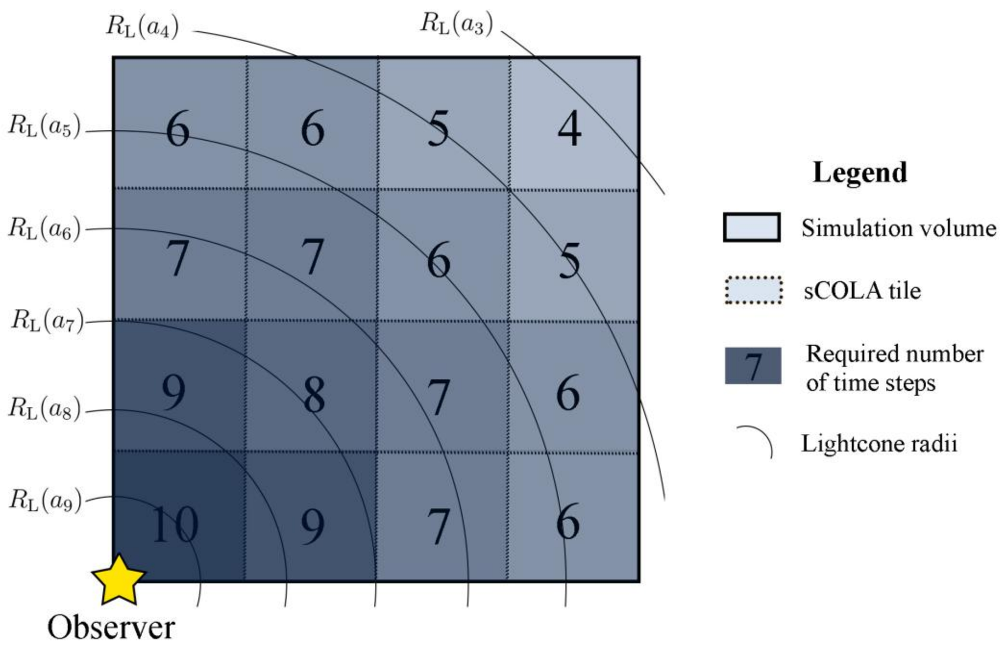

- Additional benefits:
  - Grid computing: the algorithm is suitable for inexpensive, strongly asynchronous networks
  - Robustness to node failure

sCOLA versus Gadget-4 lightcone simulations (2026)

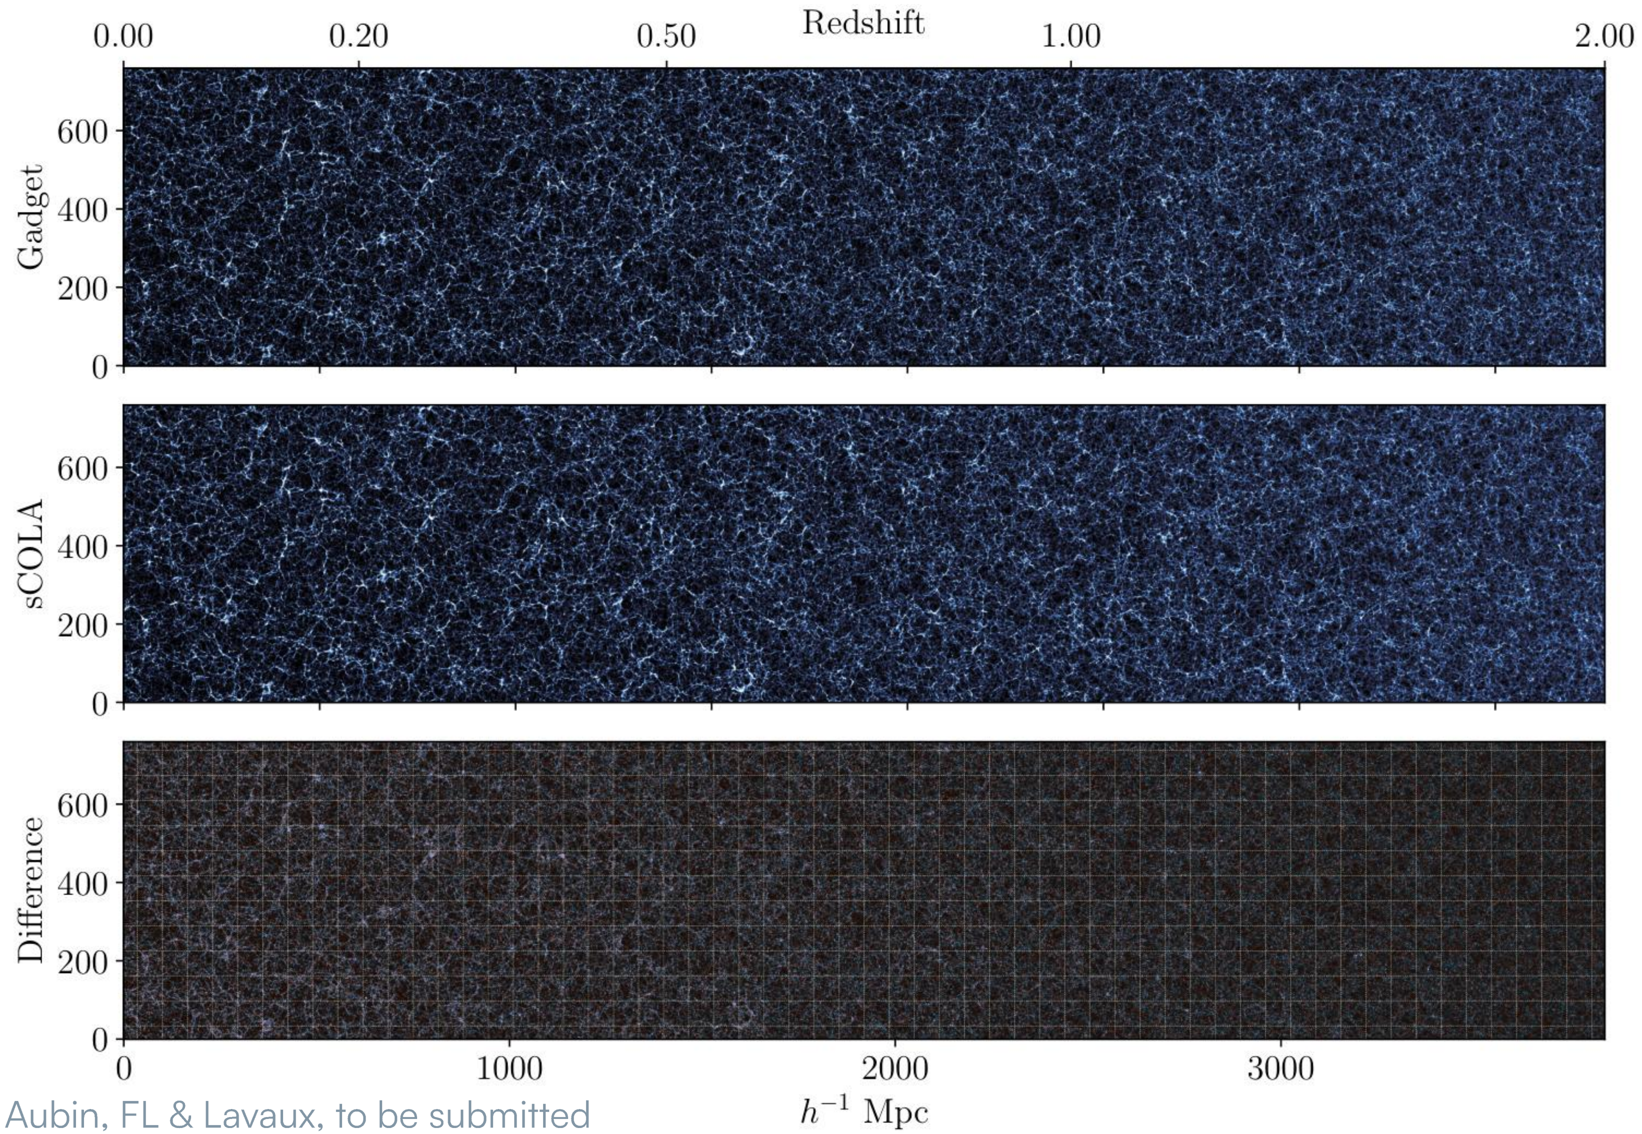

Aubin, FL & Lavaux, to be submitted

(PRELIMINARY)

Mayeul Aubin
(doctoral researcher at IAP)

The Silmarilli simulations (2026+)

(PRELIMINARY)

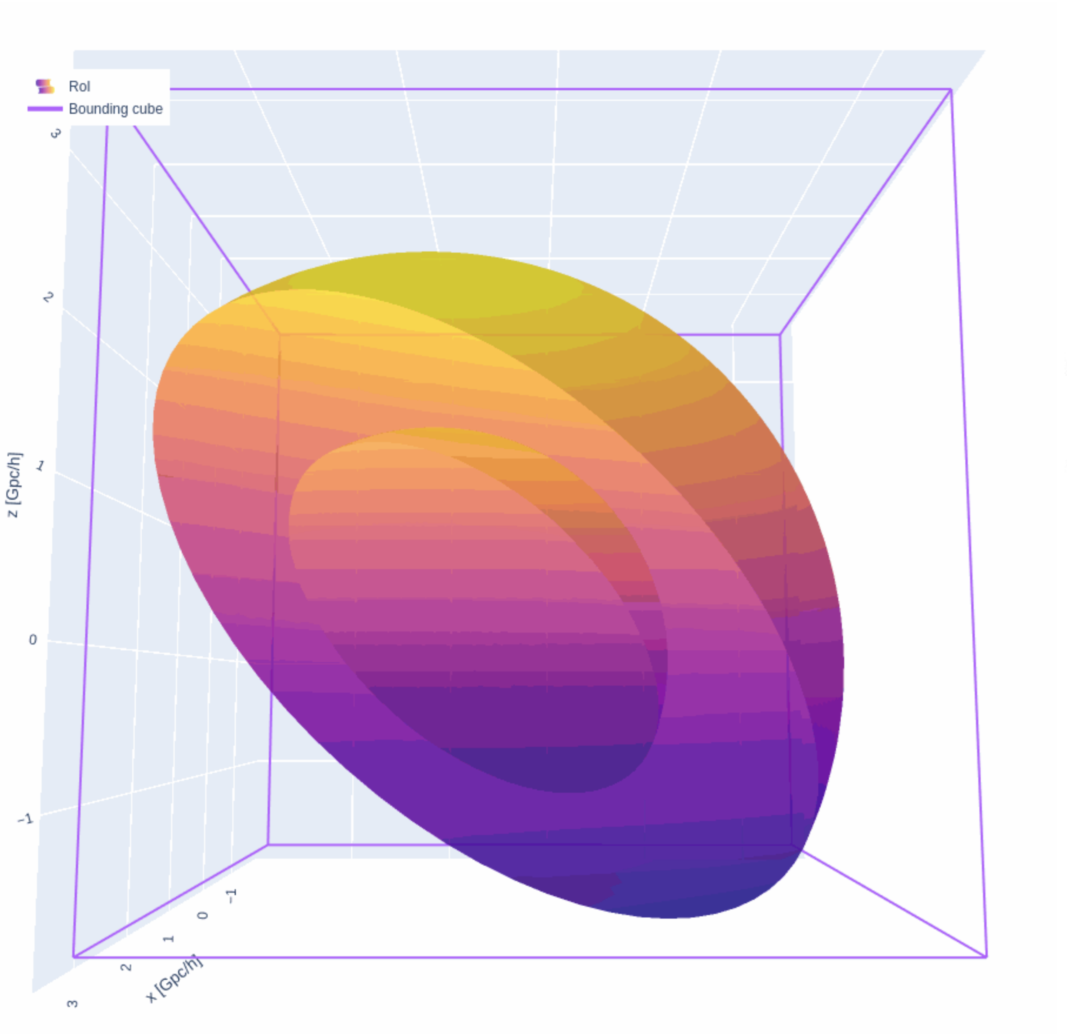

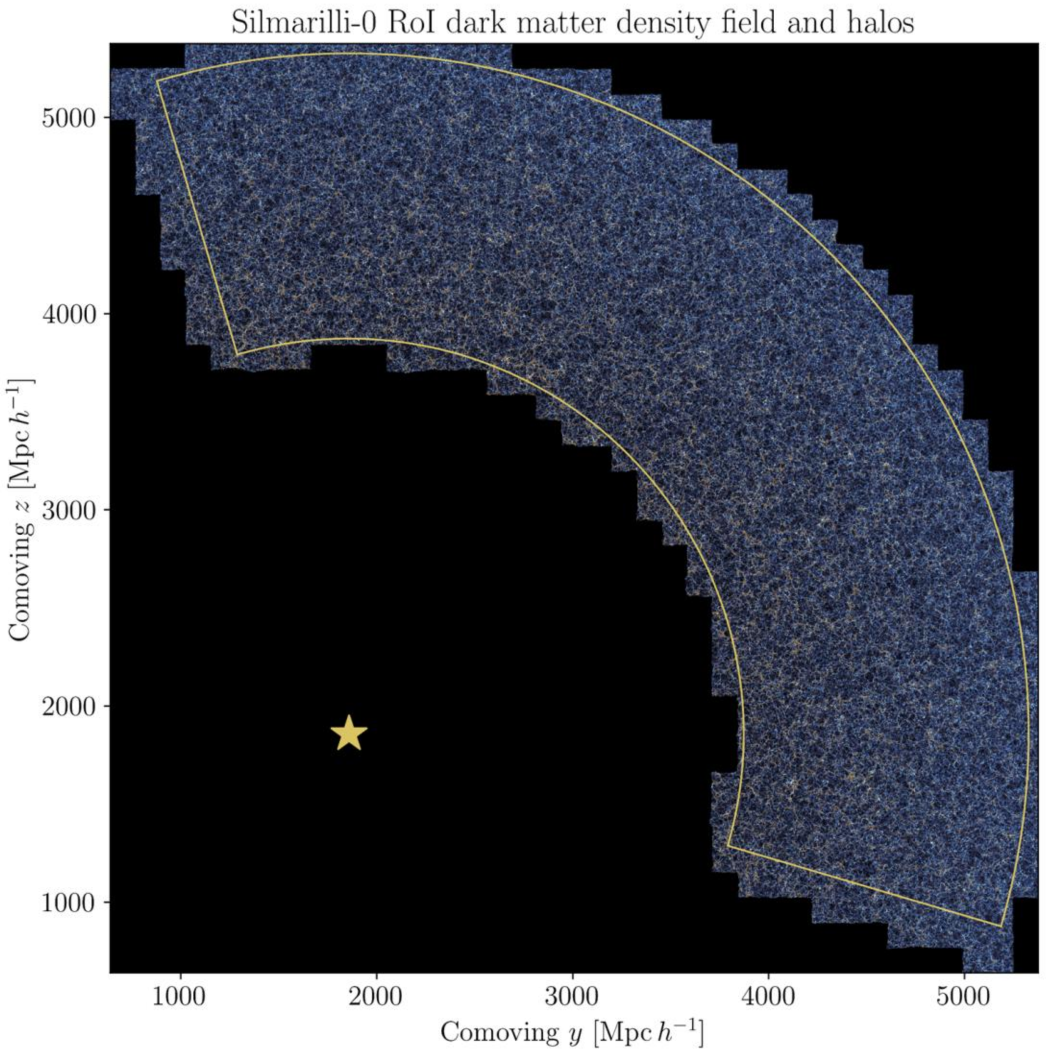

- Silmarilli-0 test completed
- Took in total ~30k CPUh, ~5 000x speed up w.r.t. Gadget-4

Silmarilli-0 setup: L = 5 376 Mpc/*h*; tiles of 128 Mpc/*h*, 256^3 particles, and 512^3 grid cells; 25 718 tiles simulated (34.7%); particles of mass 1.08×10^{10} $M_\odot$; redshift range of $0.84 < z < 1.88$; safe buffers of 40 Mpc/*h* around the RoI; 40 time steps from $z=19$ to $z=0$ (but with early stopping)

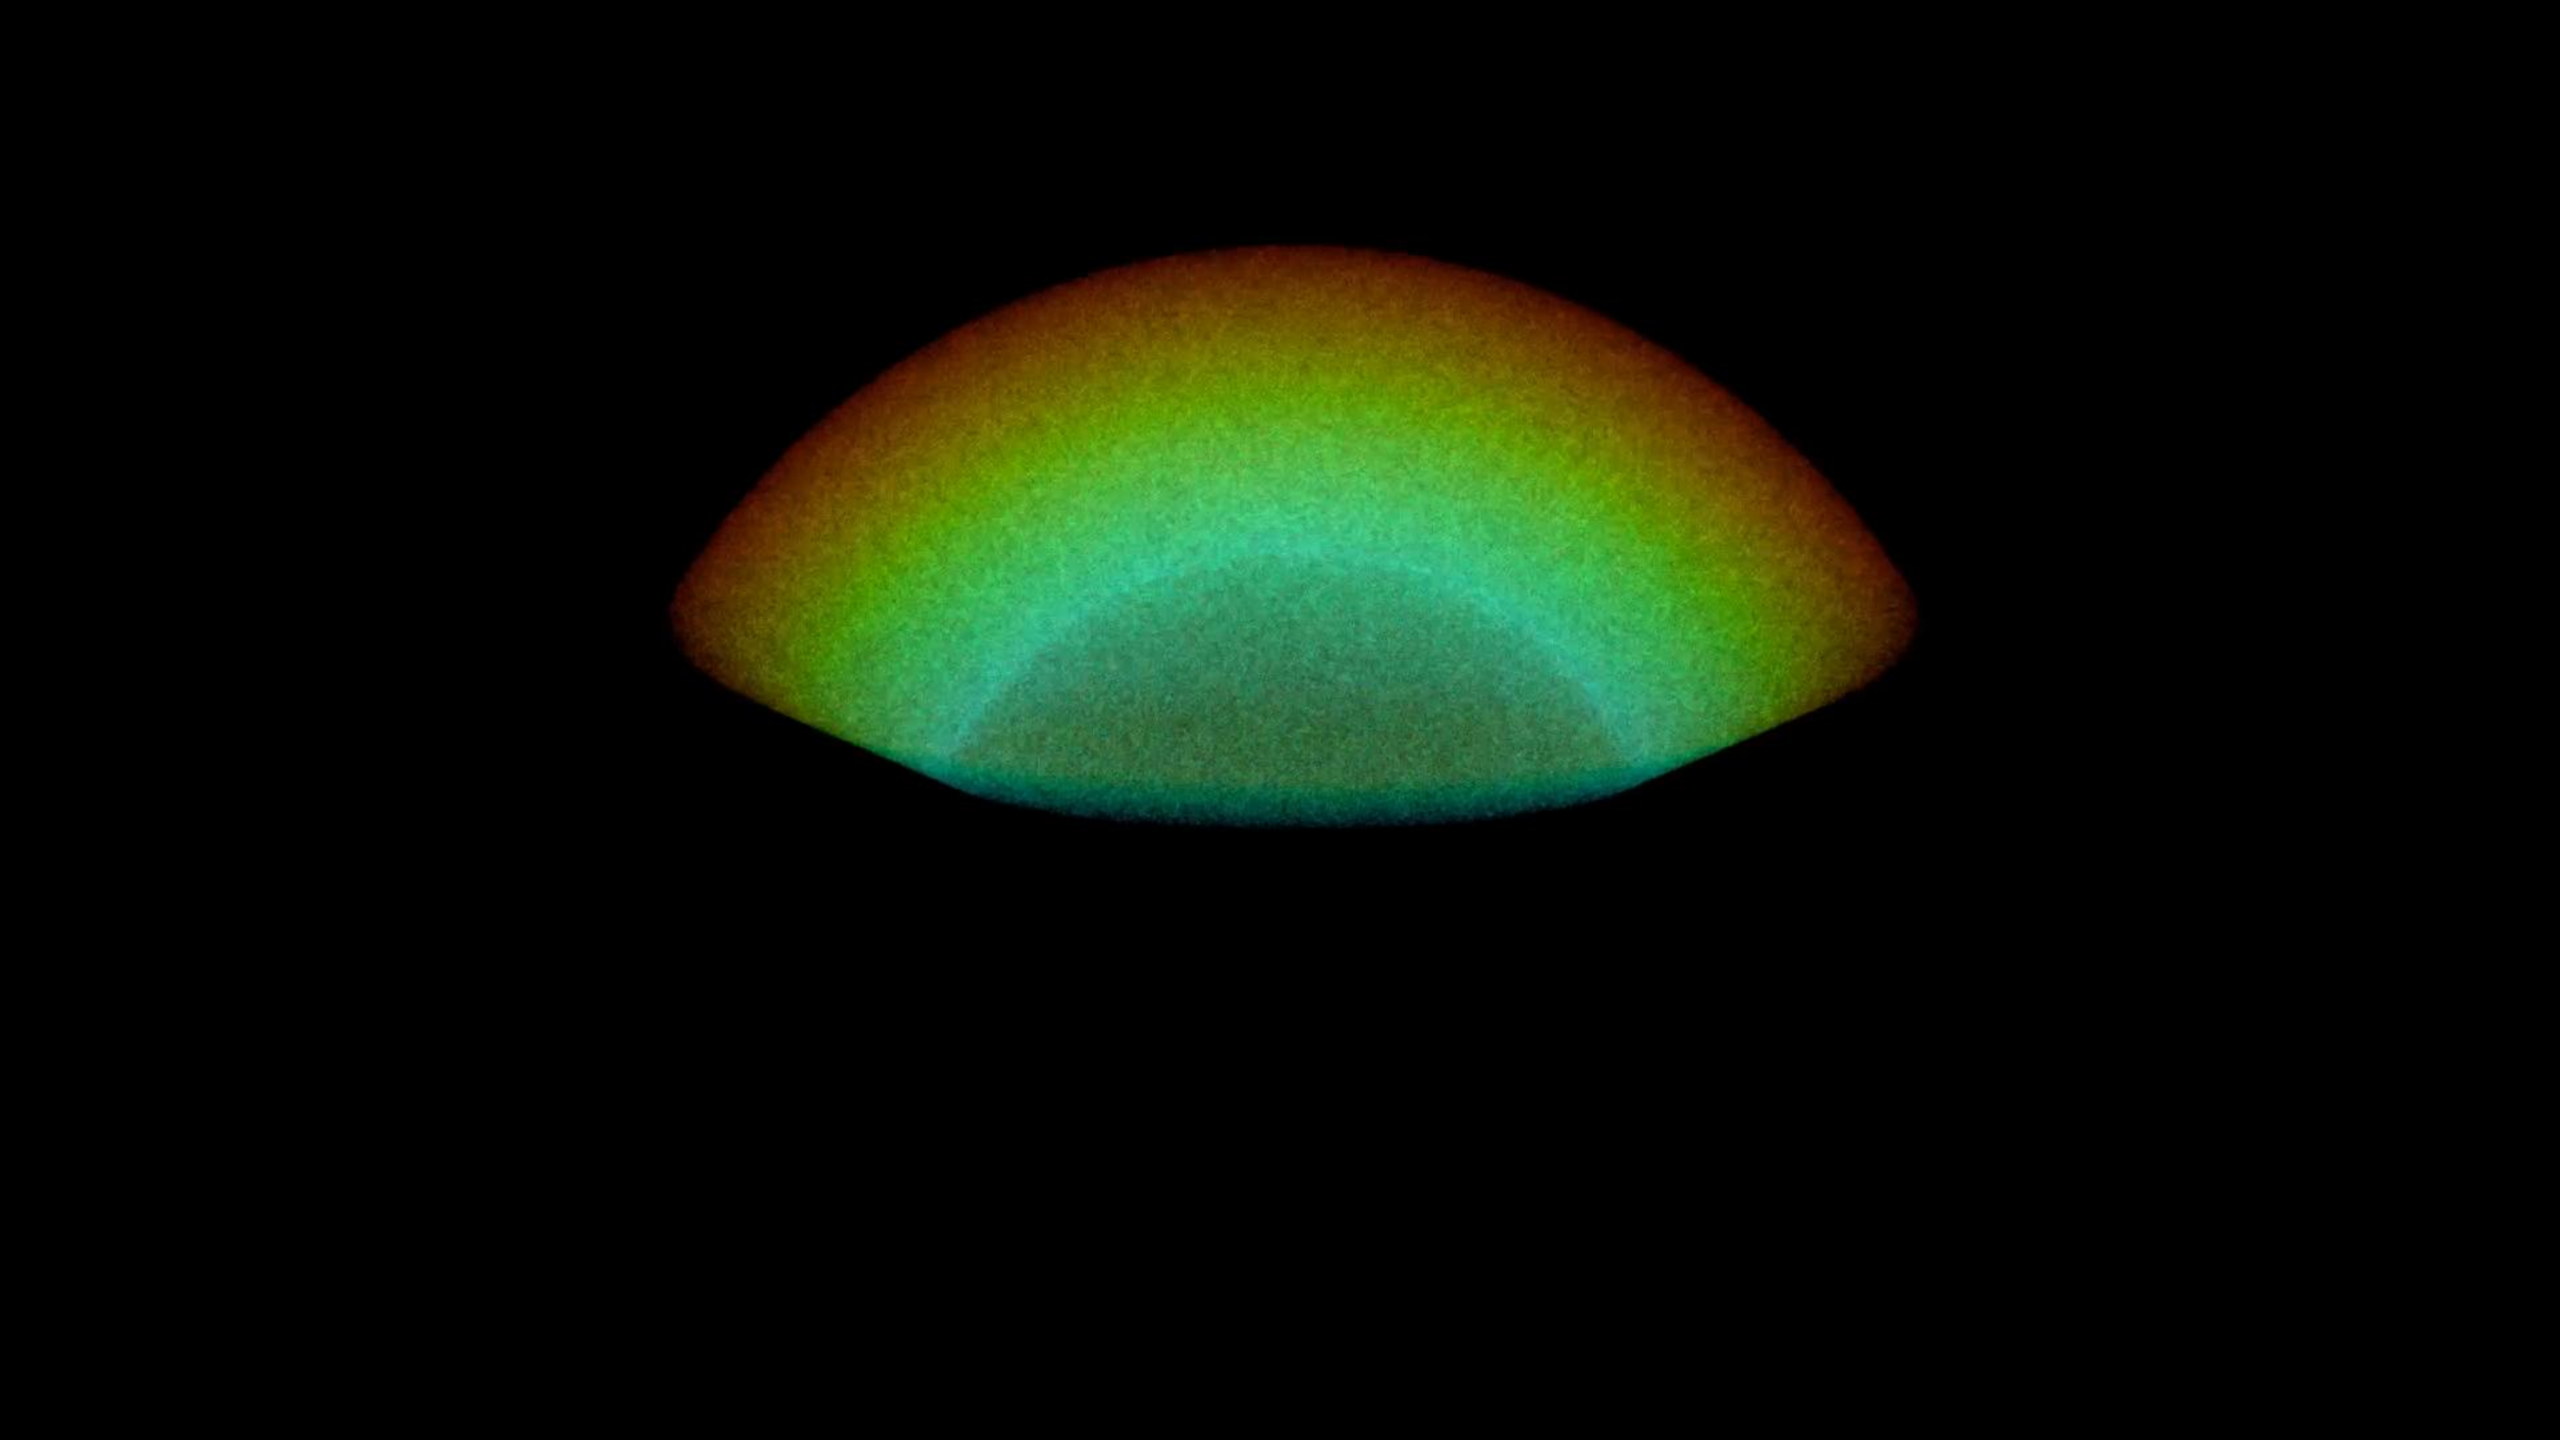

WAIT, DIDN'T YOU PROMISE SOME ML IN THIS STORY?

The tCOCA framework: temporal COmoving Computer Acceleration (2024)

- Idea behind tCOCA: the easiest simulation to run is the one where nothing moves!

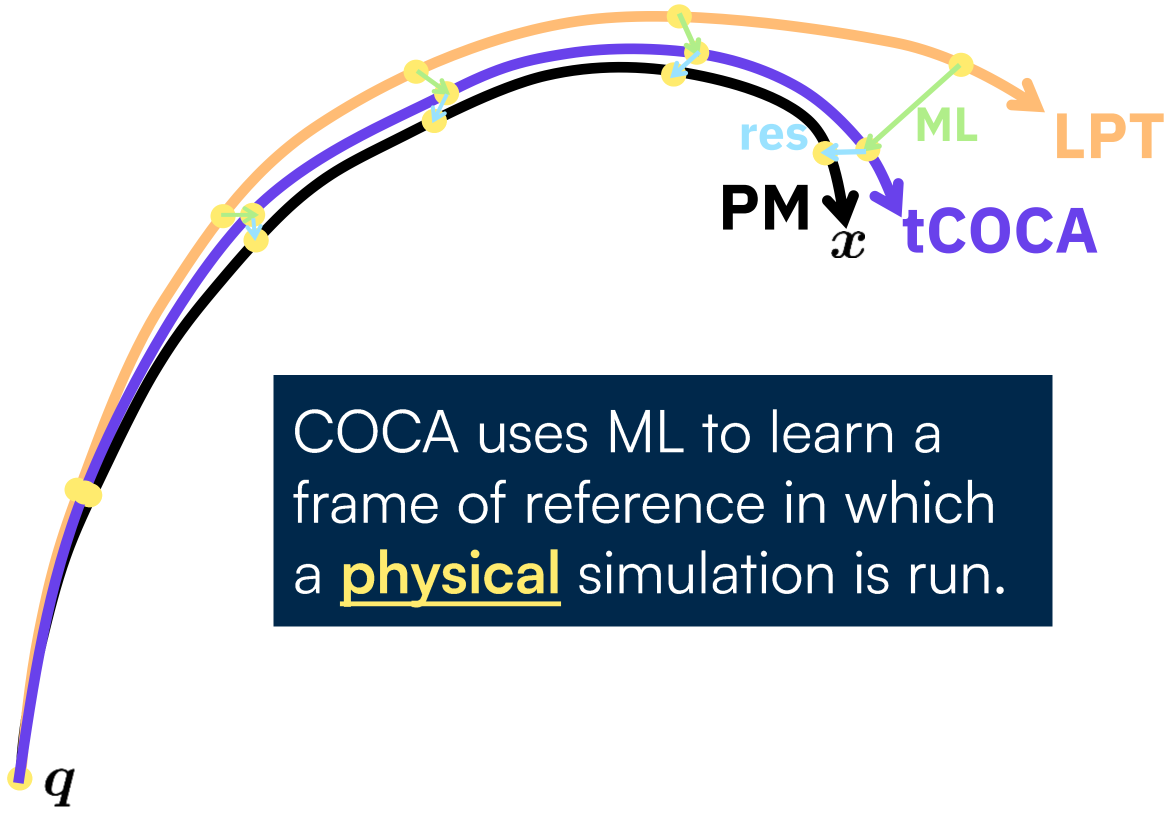

COCA uses ML to learn a frame of reference in which a **physical** simulation is run.

Bartlett, Chiarenza, Doeser & FL, 2409.02154

- Write the displacement vector as:

$$\boldsymbol{\Psi} = \boldsymbol{\Psi}_{\mathrm{LPT}} + \boldsymbol{\Psi}_{\mathrm{ML}} + \boldsymbol{\Psi}^{\mathrm{COCA}}_{\mathrm{res}} \qquad (\boldsymbol{x} = \boldsymbol{q} + \boldsymbol{\Psi})$$

- Equation of motion (omitted constants and Hubble expansion):

$$\partial_a^2 \boldsymbol{\Psi}^{\mathrm{COCA}}_{\mathrm{res}} = -\boldsymbol{\nabla}_{\boldsymbol{x}} \Phi - \partial_a^2 \boldsymbol{\Psi}_{\mathrm{LPT}} - \partial_a^2 \boldsymbol{\Psi}_{\mathrm{ML}}$$

$$\Leftrightarrow \quad \partial_a^2 \boldsymbol{\Psi} = -\boldsymbol{\nabla}_{\boldsymbol{x}} \Phi$$

- With tCOCA:
  - Any emulation error will be corrected by solving the *correct* physical equation of motion.
  - Any ML algorithm can do the job!
  - Building a data model is a safe use of ML.

Time stepping and force calculations in tCOCA

- Our implementation of tCOCA in the Simbelmynë code uses the standard Kick-Drift-Kick (leapfrog) discretisation of the equation of motion.

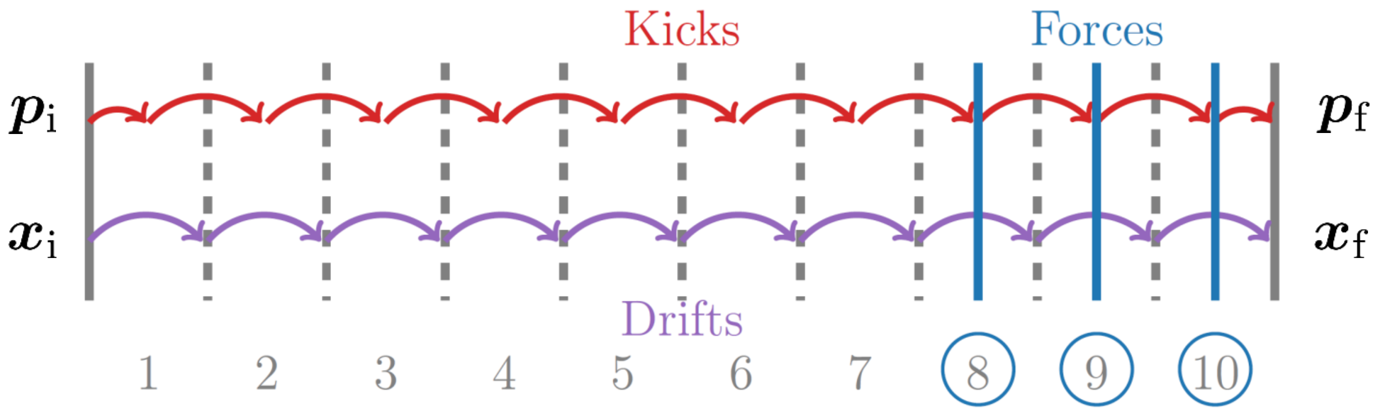

- Learning the new frame of reference means emulating the tCOLA residual momenta at every time step: $\boldsymbol{p}_\mathrm{res}^\mathrm{COLA} = \boldsymbol{p} - \boldsymbol{p}_\mathrm{LPT}$.

- When the emulation error is small $\left(\boldsymbol{p}_\mathrm{ML} \approx \boldsymbol{p}_\mathrm{res}^\mathrm{COLA}\right)$, particles are already at rest in the tCOCA frame of reference, so it is unnecessary to compute forces at every step.

- A good frame-of-reference emulator therefore makes tCOCA cheaper than tCOLA.

Bartlett, Chiarenza, Doeser & FL, 2409.02154

ML-safe use of neural networks in N-body simulations

Deaglan Bartlett
(PDRA at IAP → Oxford)

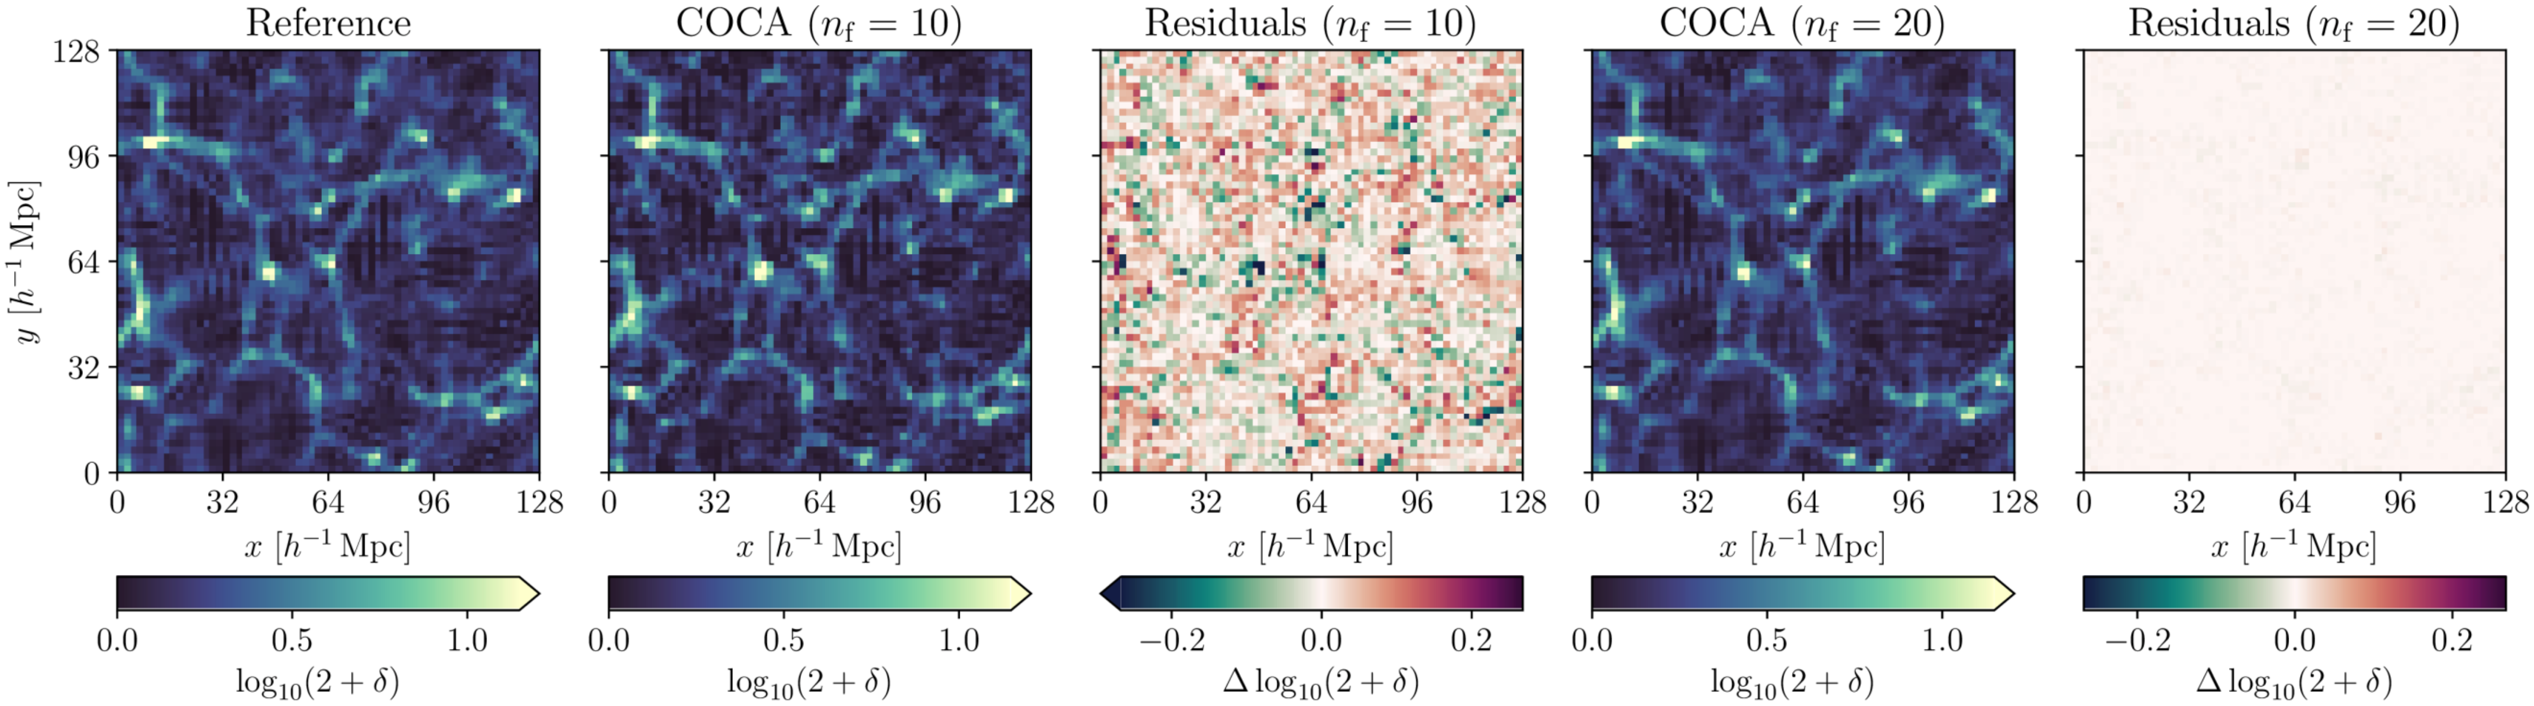

We reach percent-level accuracy up to $k = 1\ h/\mathrm{Mpc}$ on standard correlations functions, using only 8 to 10 particle-mesh (PM) force evaluations.

Bartlett, Chiarenza, Doeser & FL, 2409.02154

Putting it all together (2026+): perfect parallelism & ML-safety
spatial comoving COmputer acceleration (sCOCA)

- Can we decouple sub-volumes by using the large-scale solution?

LPT so far
(analytical solution) → sCOLA;
soon **ML** solution → sCOCA

$$\partial_a^2 \boldsymbol{\Psi} = -\boldsymbol{\nabla}_{\boldsymbol{x}} \left[\underbrace{\Delta^{-1}\delta}_{\text{non-local}} \right] \Leftrightarrow \partial_a^2 \left(\boldsymbol{\Psi} - \boldsymbol{\Psi}_{\mathrm{l.s.}} \right) = -\boldsymbol{\nabla}_{\boldsymbol{x}} \left[\underbrace{\Delta^{-1} \left(\delta - \delta_{\mathrm{l.s.}} \right)}_{\text{local}} \right]$$

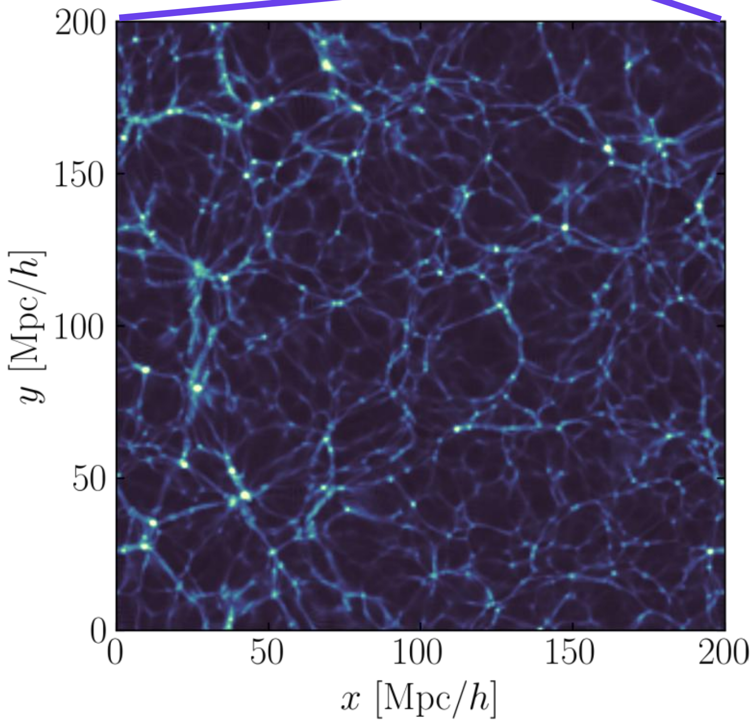

tCOLA (reference)

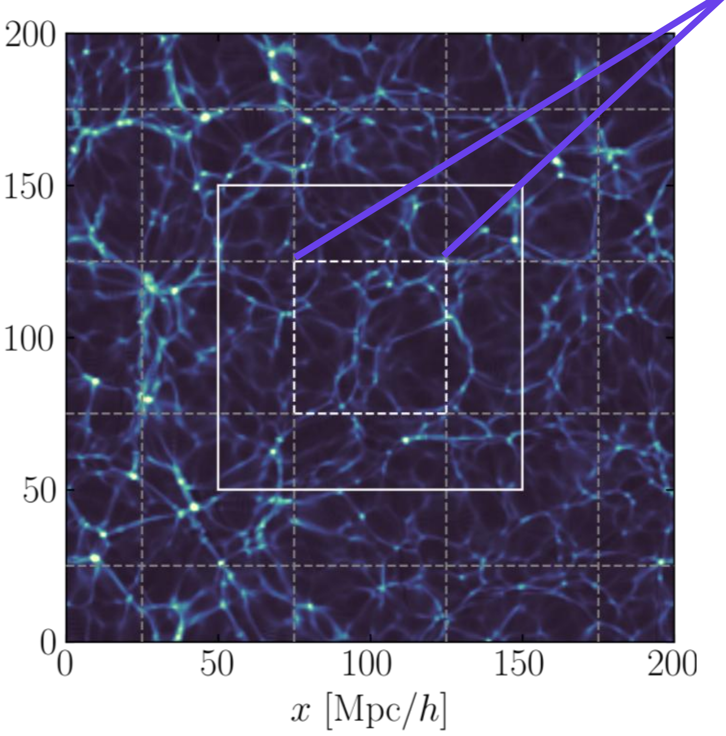

sCOCA

Stay tuned!

Tiling for scaling field-level inference of initial conditions (BORG)

- The **Manticore-Deep** project: a Bayesian digital twin of the Universe in a 4096 Mpc/*h* box constrained by 2M++, 6dFGS, 2dFGRS, SDSS, BOSS.

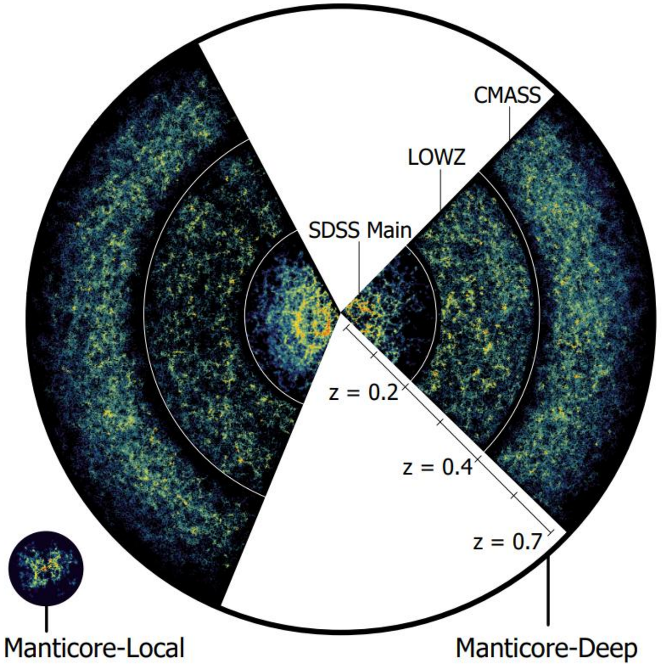

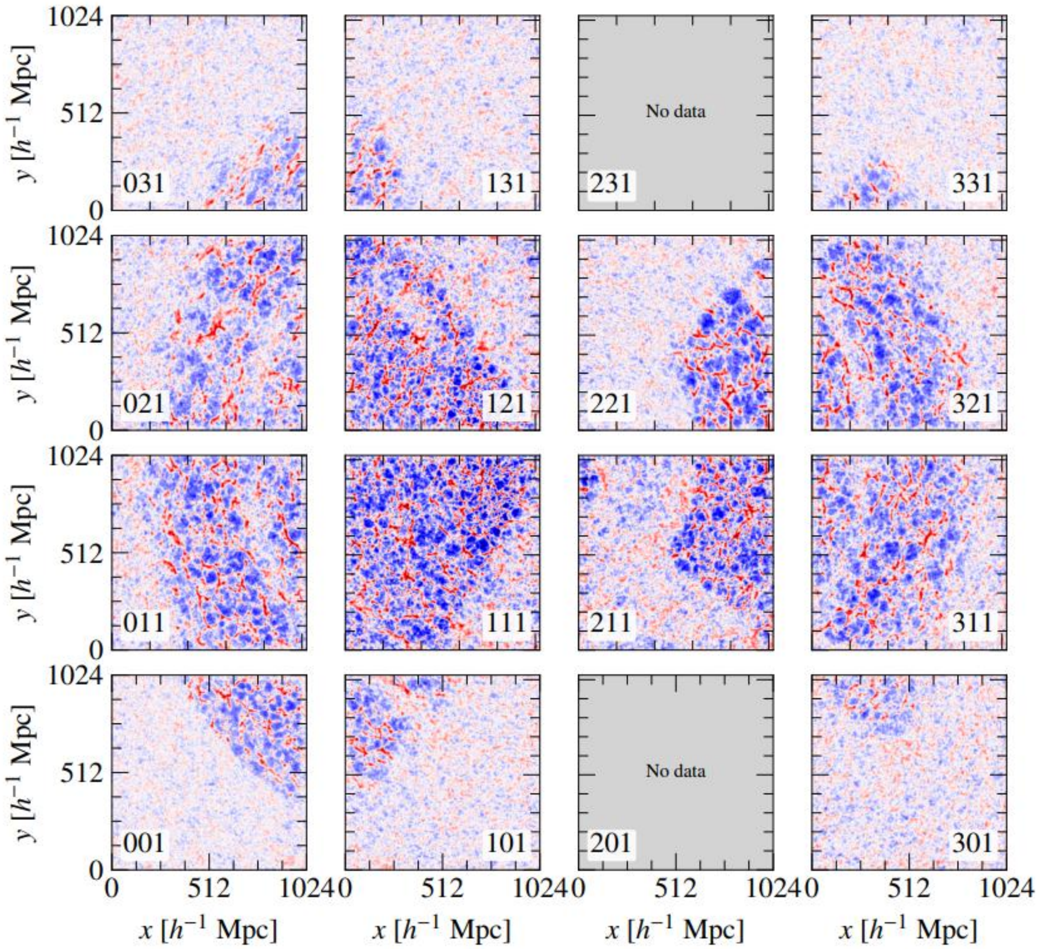

- The reconstruction (effectively) has 1024^3 parameters in 64 tiles (27 with data).

Jasche & Wandelt, 1203.3639; Jasche, FL & Wandelt, 1409.6308; Jasche & Lavaux, 1806.11117; Lavaux, Jasche & FL, 1909.06396; McAlpine *et al.* 2505.10682; McAlpine *et al.* 2606.10020

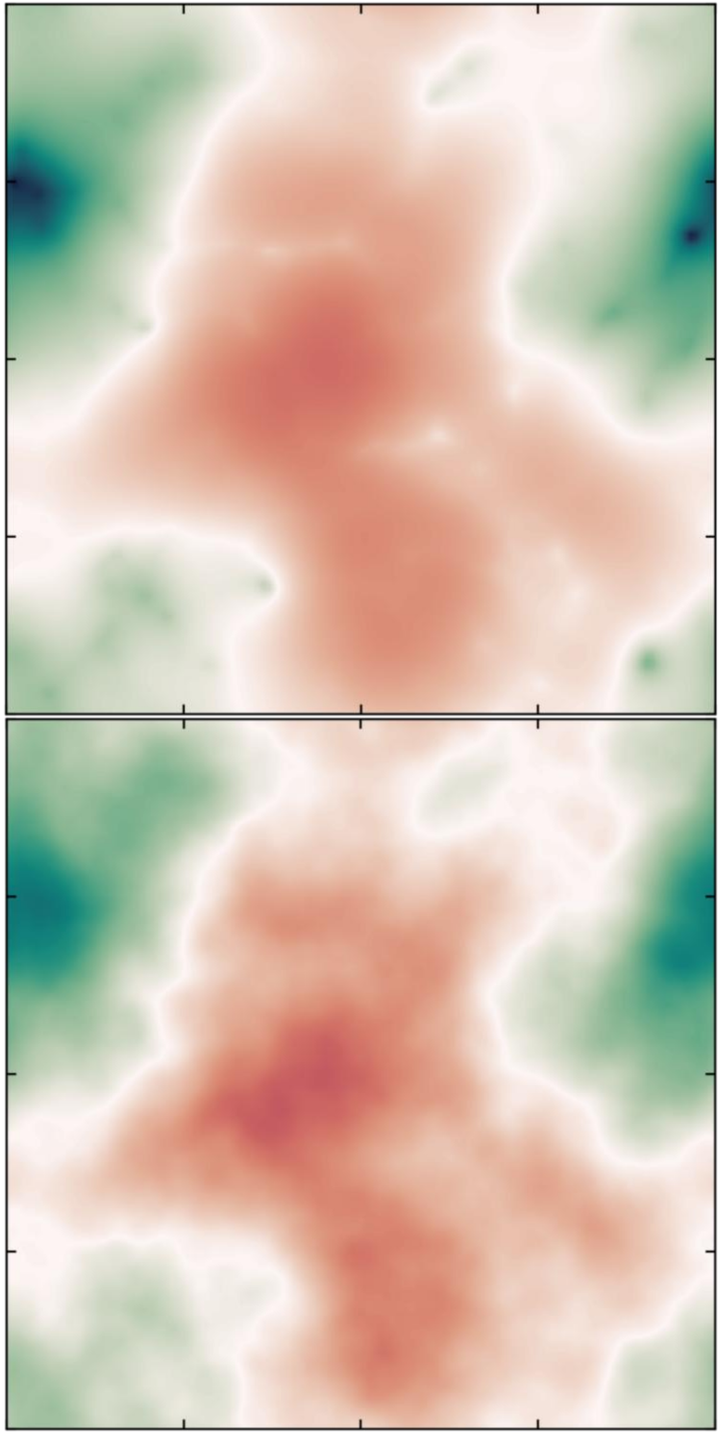

Conclusions: Potentially brilliant, a history of boundaries

- sCOLA uses the large-scale approximate solution to spatially split simulations in independent tiles:
  - It achieves perfect parallelism by fully removing the need for communications across the full computational volume.
  - It allows for fast lightcone and mock catalogue generation.
- tCOCA reimagines the use of neural networks for emulating *N*-body simulations:
  - It generalises the idea of tCOLA: running simulations in a new frame of reference. But it is not an emulator!
  - It solves the correct equations of motion, so it is a ML-safe use of neural networks. Explainability is not needed!
  - It makes simulations cheaper by skipping unnecessary force evaluations. But any force calculation technique (e.g. P3M) can be used.