Perfectly parallel cosmological simulations using spatial comoving Lagrangian acceleration

Colloque national action Dark Energy 2022,
Marseille, France

Florent Leclercq

www.florent-leclercq.eu

Institut d'Astrophysique de Paris
CNRS & Sorbonne Université

In collaboration with:
Baptiste Faure (CEA), Guilhem Lavaux (IAP), Benjamin Wandelt (IAP),Alan Heavens (Imperial College), Andrew Jaffe (Imperial College), Will Percival (U. Waterloo), Camille Noûs (Laboratoire Cogitamus)

and the Aquila Consortium

www.aquila-consortium.org

17 November 2022

Why do we (still) need *N*-body codes?

- *N*-body simulations remain a basic ingredient for many cosmological modelling problems: galaxy clustering, ray-tracing, 21cm intensity mapping, Lyman-α.
- Frequentist approach: mock surveys (e.g. DESI, Euclid, LSST) are used for measurements of summaries and their covariances.
- Bayesian approach: forward numerical data models are the new way to express the theory:
  - … embedded into a field-level likelihood: Bayesian large-scale structure inference (BORG),
  - … or in a simulator-based approach: likelihood-free inference (ABC, DELFI, BOLFI, SELFI, etc.).

- Traditional hardware architectures are reaching their physical limit: per-core compute performance is slowing down.

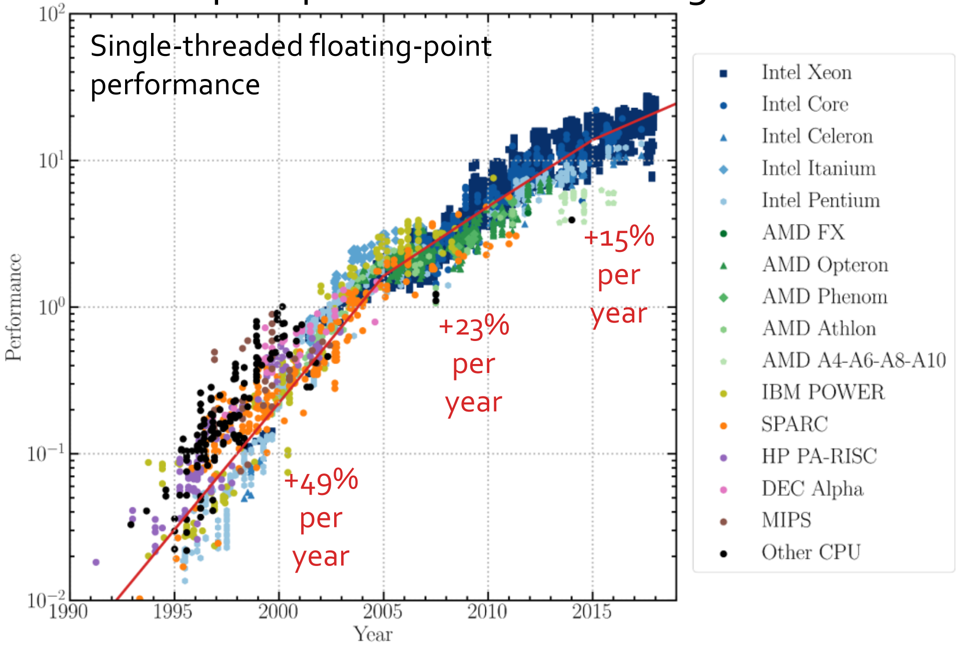

Based on adjusted SPECfp® results, http://spec.org

- Current hardware development focuses on:
  - Packing a larger number of cores into each CPU: currently $\mathcal{O}(10^5)$, soon $\mathcal{O}(10^{6-7})$ in systems that are currently being built.
  - Developing hybrid architectures with cores + accelerators: GPUs, reconfigurable or dedicated chips (FPGAs/ASICs).

Parallelisation of N-body codes: the challenge

- Compute cycles are no longer the scarce resource. The cost is driven by interconnections.
- Amdahl's law: latency kills the gains of parallelisation.

Amdahl 1967, doi:10.1145/1465482.1465560

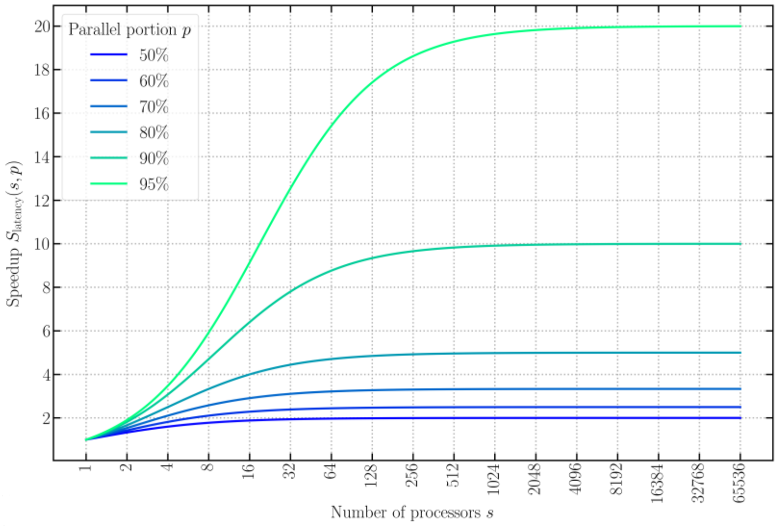

- The main issue preventing the easy parallelisation of N-body codes is the long-range nature of gravitational interactions.
- "Exact" gravity requires $\mathcal{O}(N^2)$ all-to-all communications between N particles.

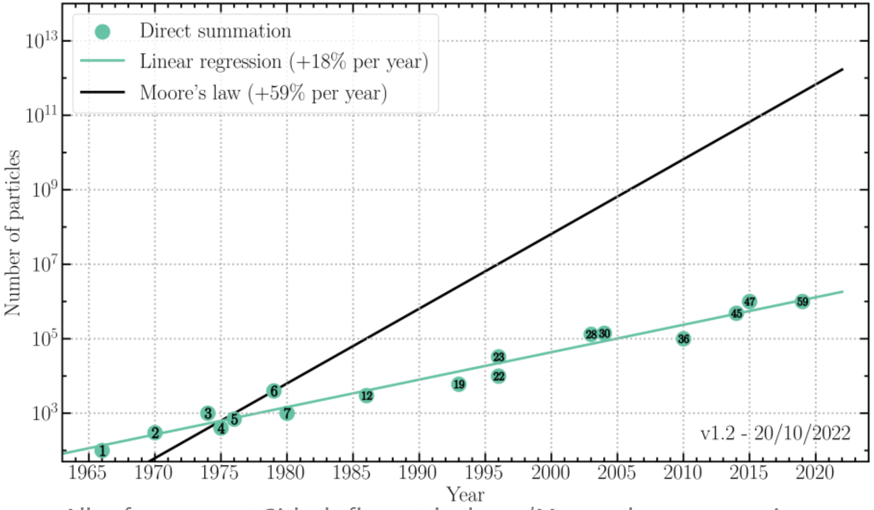

All references at Github:florent-leclercq/Moore_law_cosmosims

Parallelisation of *N*-body codes: the challenge

- Most of the work on numerical cosmology so far has focused on algorithms (such as tree, multipole, and mesh methods) that reduce the need for communications across the full computational volume.

- Since 1990, we observe a super-exponential trend that cannot be explained only by increase in computer speed.

- *N*-body codes cannot merely rely on computers becoming faster to reduce the computational time in the future.

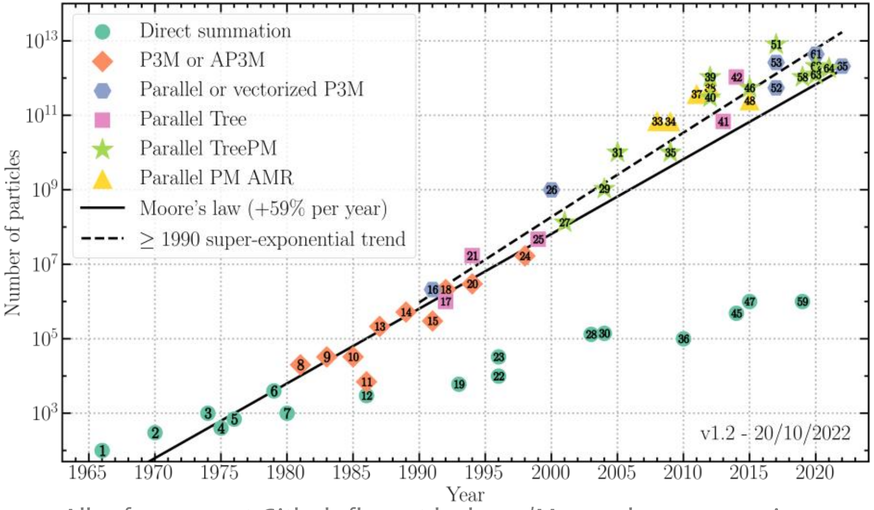

All references at Github:florent-leclercq/Moore_law_cosmosims

tCOLA: Comoving Lagrangian Acceleration (temporal domain)

- Write the displacement vector as: $\boldsymbol{\Psi} = \boldsymbol{\Psi}_\mathrm{LPT} + \boldsymbol{\Psi}_\mathrm{res} \quad (\mathbf{x} = \mathbf{q} + \boldsymbol{\Psi})$

Tassev & Zaldarriaga, 1203.5785

Analytical solutions!

- Time-stepping (omitted constants and Hubble expansion):

Standard: $\partial_a^2 \boldsymbol{\Psi} = -\boldsymbol{\nabla}_\mathbf{x}\Phi$ ⇒ Modified: $\partial_a^2 \boldsymbol{\Psi}_\mathrm{res} = \partial_a^2(\boldsymbol{\Psi} - \boldsymbol{\Psi}_\mathrm{LPT}) = -\boldsymbol{\nabla}_\mathbf{x}\Phi - \partial_a^2 \boldsymbol{\Psi}_\mathrm{LPT}$

Tassev, Zaldarriaga & Einsenstein, 1301.0322

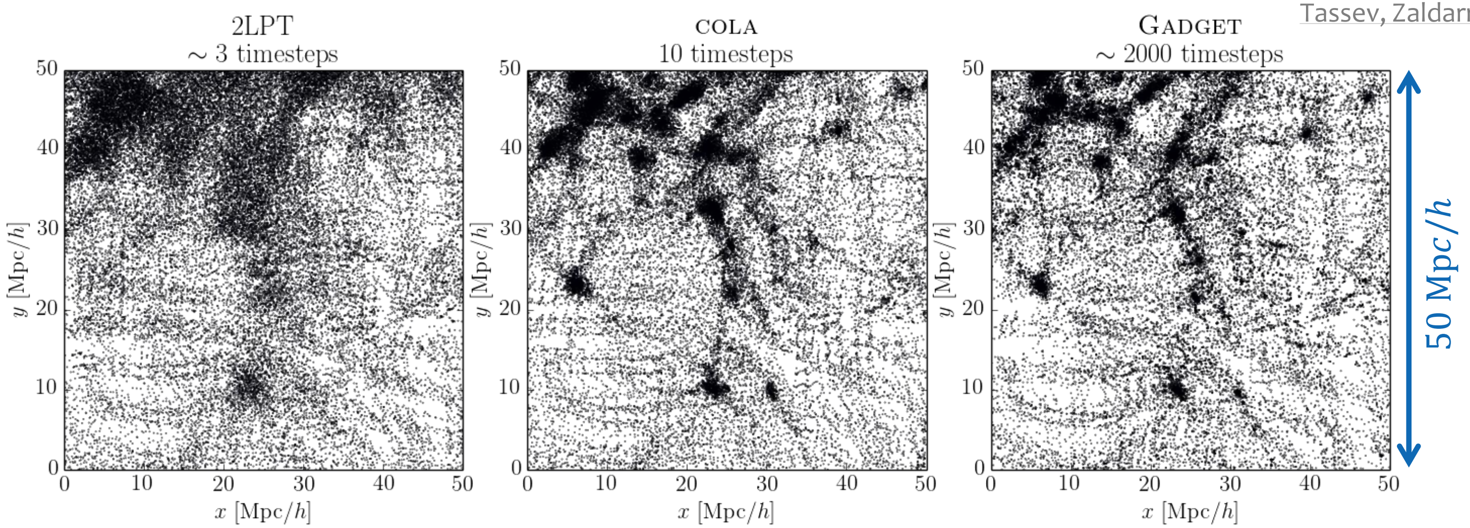

Beneficial gain of efficiency… but the real problem is not CPU-hours, but the inability to run on a very large number of cores due to latencies/parallelisation overhead.

sCOLA: Extension to the spatial domain

- Computing the LPT reference frame suggests a new strategy:

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

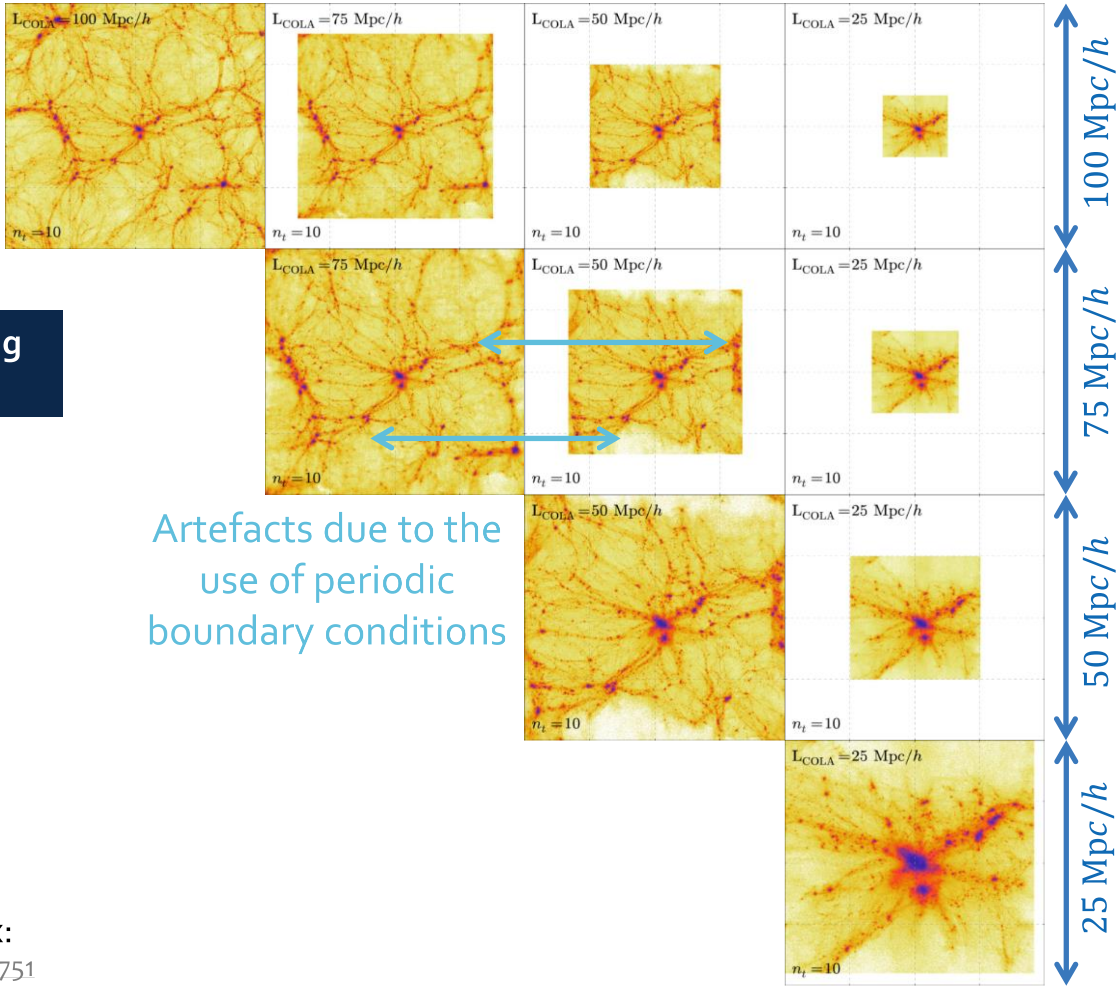

Proof of concept using one sub-box embedded into a larger simulation box:
Tassev, Eisenstein, Wandelt & Zaldarriaga, 1502.07751

The solution to boundary artefacts

1. A buffer region around each tile
2. Appropriate Dirichlet boundary conditions for the potential

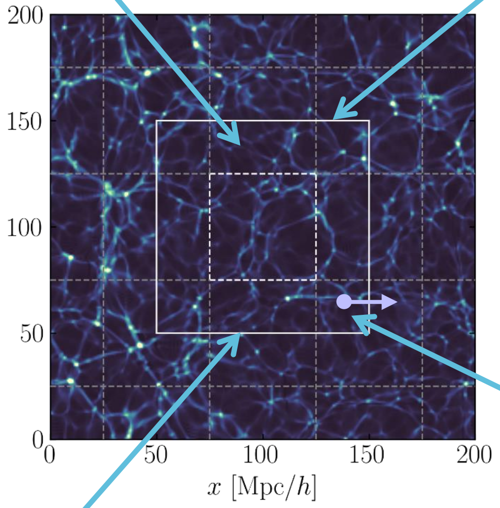

- It is necessary to oversimulate some of the volume.
- The Poisson solver uses discrete sine transforms (DSTs) instead of FFTs.
- Overall, two approximations (to ensure no communication between tiles):

1. Linearly-evolving potential (LEP) at the boundaries: $\Phi_{\mathrm{BCs}}(\mathbf{x}, a) \approx D_1(a)\phi^{(1)}(\mathbf{x})$
2. Outgoing particles do not deposit mass

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

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

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

$$\frac{\partial^2 \mathbf{x}}{\partial t^2} = -\nabla \left[\Delta^{-1}\delta\right] \qquad \frac{\partial^2}{\partial t^2}\left(\mathbf{x} - \mathbf{x}_{\mathrm{l.s.}}\right) = -\nabla \left[\Delta^{-1}\left(\delta - \delta_{\mathrm{l.s.}}\right)\right]$$

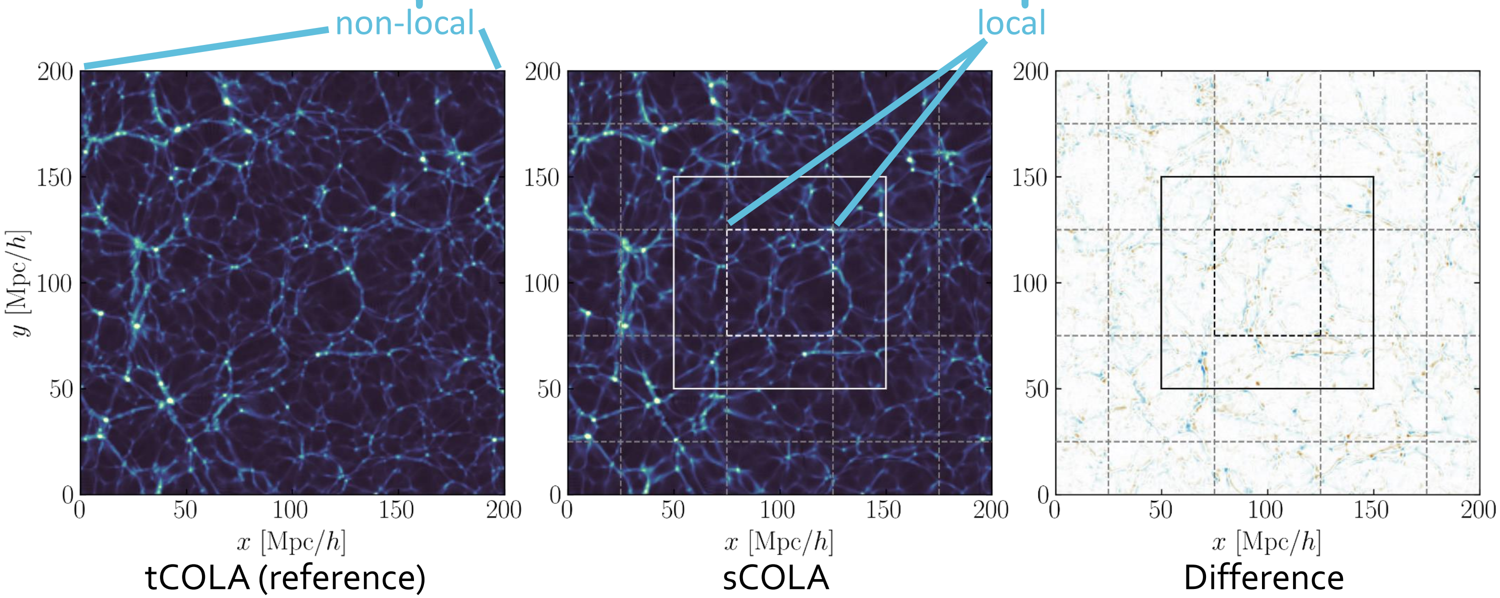

Publicly available implementation:
Bitbucket:florent-leclercq/simbelmyne/

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

The perfectly parallel algorithm and its accuracy

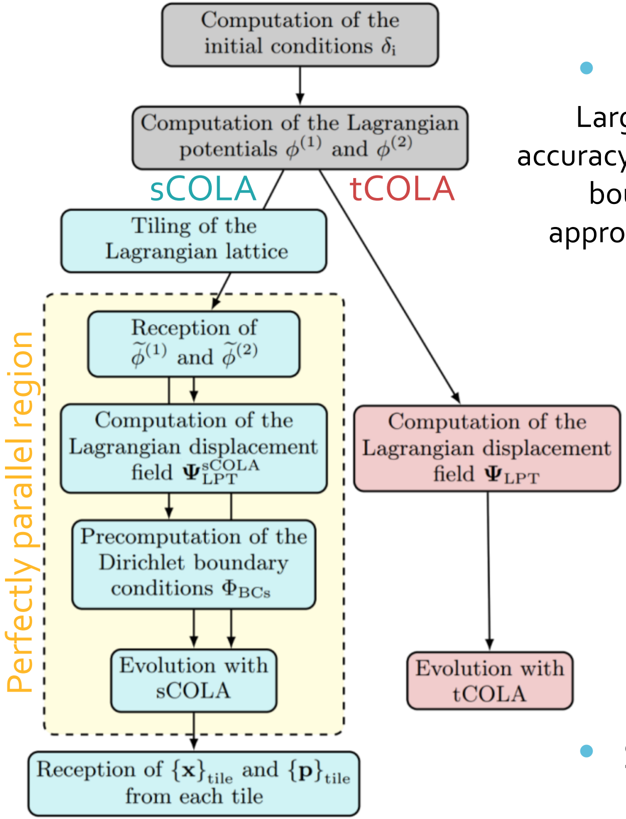

- Parameter investigation (size of tiles and buffer regions):

Large-scale accuracy is driven by boundary approximations

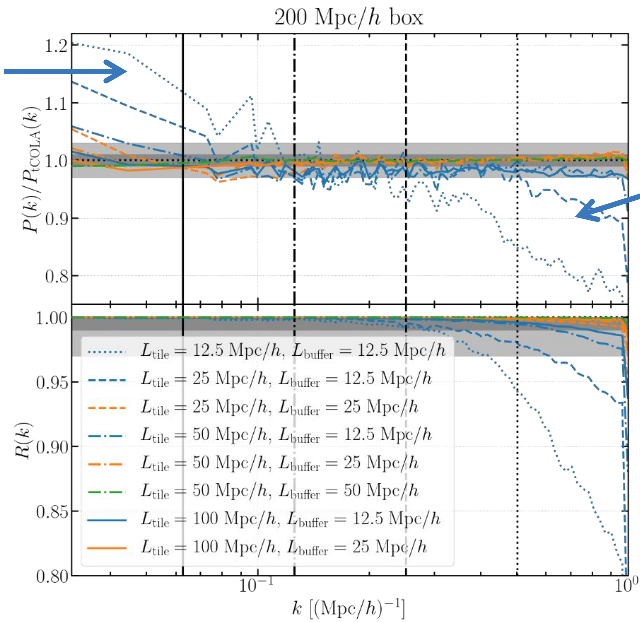

Small-scale accuracy is driven by the size of buffer regions

- Some setups reach 3 to 1% accuracy at all scales, as required for the next-generation of surveys.

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

Memory requirements, parallelisation potential & speed

- Buffer regions require to oversimulate the volume by some factor r.
- But small *N*-body simulations can be run in the L3 cache of CPUs, on GPUs or FPGAs: hardware speed-up factor of s.
- Parallelisation potential factor:

$$p = s\,\frac{N_\mathrm{tiles}}{r} = s\left(\frac{L}{L_\mathrm{sCOLA}}\right)^3$$

$$L = 1\ \mathrm{Gpc}/h$$
$$L_\mathrm{tile} = 125\ \mathrm{Mpc}/h$$
$$L_\mathrm{buffer} = 29.3\ \mathrm{Mpc}/h$$
$$r \approx 3.17$$
$$p \approx 161.59$$

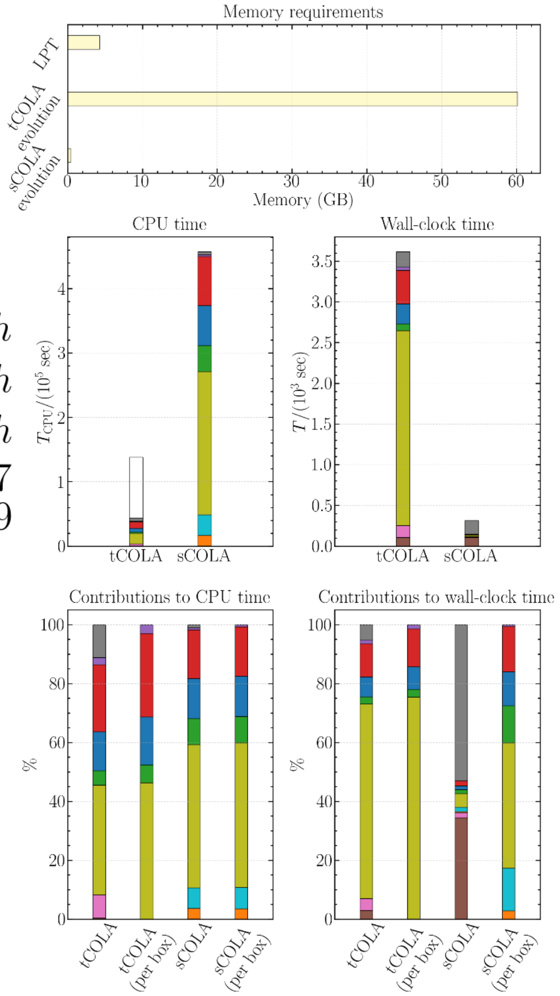

FL, Faure, Lavaux, Wandelt, Jaffe, Heavens, Percival & Noûs, 2003.04925

Additional benefits

- Light-cones and mock catalogues:
  - sCOLA boxes only need to run until they intersect the observer's past light-cone.
  - 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.

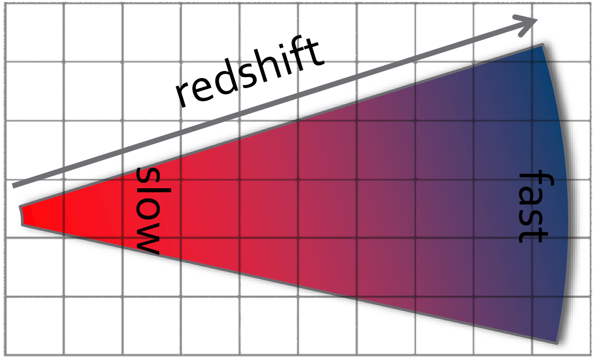

- Gravity and physics models: any gravity model (e.g. P3M, tree, or AMR) and non-gravitational physics (hydrodynamics) can be used within tiles.
- Grid computing: the algorithm is suitable for inexpensive, strongly asynchronous networks.
- Robustness to node failure.

Conclusions

- In the age of peta-/exa-scale computing, we introduced a perfectly parallel and easily applicable algorithm for cosmological simulations using sCOLA, a hybrid analytical/numerical technique.
- The approach is based on a tiling of the full simulation box, where each tile is run independently.
- Resulting larger and higher-resolution cosmological simulations can be used in the context of DESI, Euclid, LSST and upcoming extremely large-scale surveys.
- The algorithm can benefit from a variety of hardware architectures. It is suitable for participatory computing platforms such as Cosmology@Home (with potential visibility/outreach benefits).

  https://www.cosmologyathome.org

- The algorithm is implemented in the Simbelmynë code, publicly available at

  https://simbelmyne.florent-Leclercq.eu – Bitbucket:florent-leclercq/simbelmyne

Additional slides

Density field

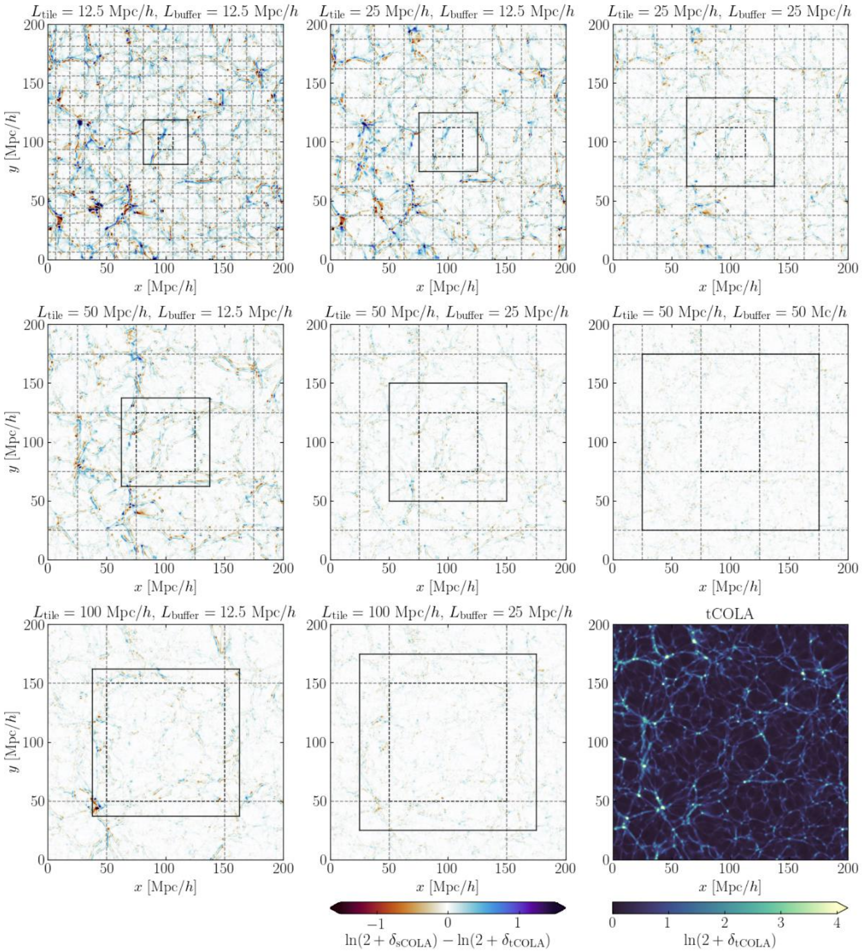

FL, Faure, Lavaux, Wandelt, Jaffe, Heavens, Percival & Noûs, 2003.04925

Velocity field

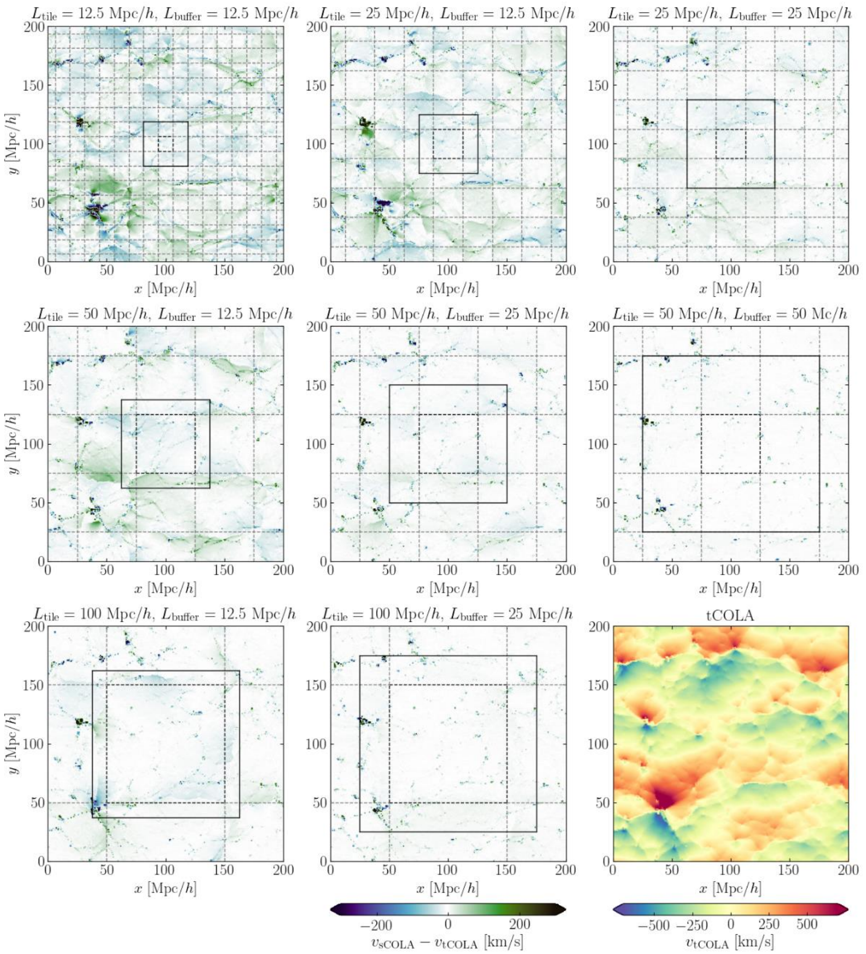

FL, Faure, Lavaux, Wandelt, Jaffe, Heavens, Percival & Noûs, 2003.04925

1 Gpc/*h* box

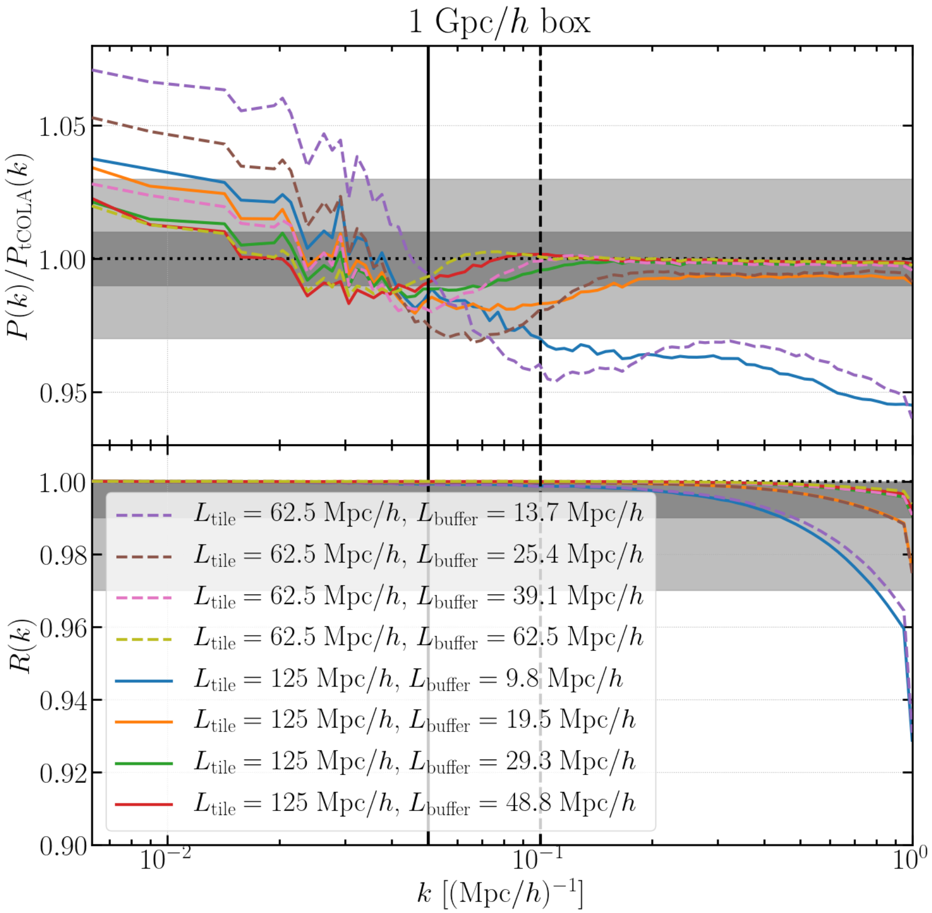

FL, Faure, Lavaux, Wandelt, Jaffe, Heavens, Percival & Noûs, 2003.04925

Test of the approximations

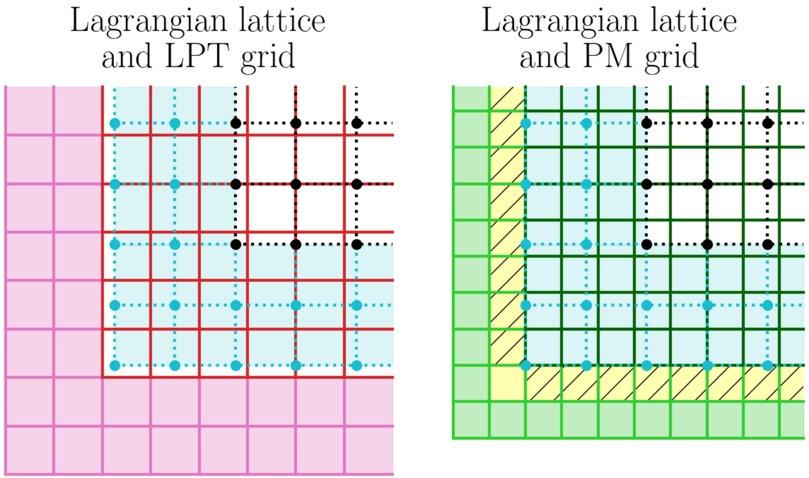

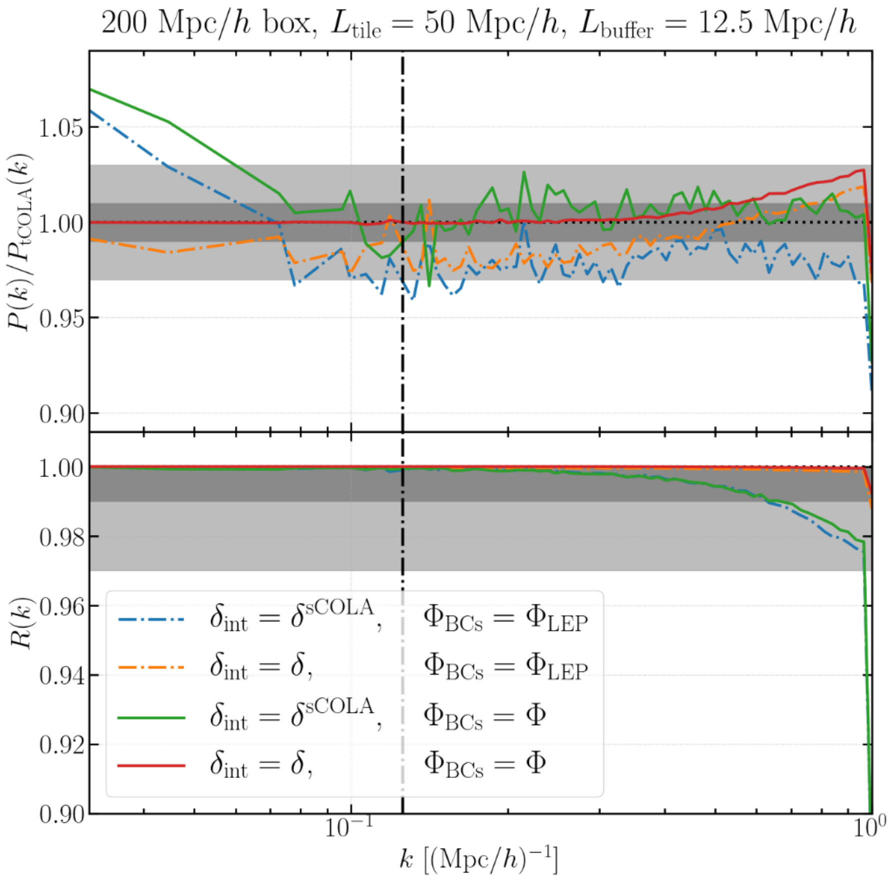

FL, Faure, Lavaux, Wandelt, Jaffe, Heavens, Percival & Noûs, 2003.04925