

SWIFT


First Cosmology Results

Josh Borrow¹
 Alexei Borrisov^{1,2}
 Richard Bower¹
 Peter Draper¹
 Pedro Gonnet³
 Matthieu Schaller^{1,4}
 James Willis¹

1: ICC, Durham University
 2: DiRAC
 3: Google Research Switzerland
 4: Leiden University

Authors listed in alphabetical order, with JB, the main text/poster author, highlighted.

More information at:
<http://www.swiftsim.com>



About SWIFT

SWIFT is an open source cosmological simulation code developed in partnership between the Institute for Computational Cosmology and computer science department at Durham University, DiRAC, and Intel through an Intel Parallel Computing Centre. SWIFT solves the equations of gravity and hydrodynamics, taking us from the big bang all the way through to the present day. To challenge such a large problem, SWIFT exploits three levels of parallelism: MPI between nodes, threads on a single node, and SIMD/AVX instructions on a single chip. SWIFT also implements task-based parallelism through the QuickSched library. These key architectural choices enable SWIFT to weak-scale nearly perfectly on highly adaptive problems, where there may only be a few active particles on the entire cluster out of billions (see Schaller et al. 2016 & Borrow et al. 2018).

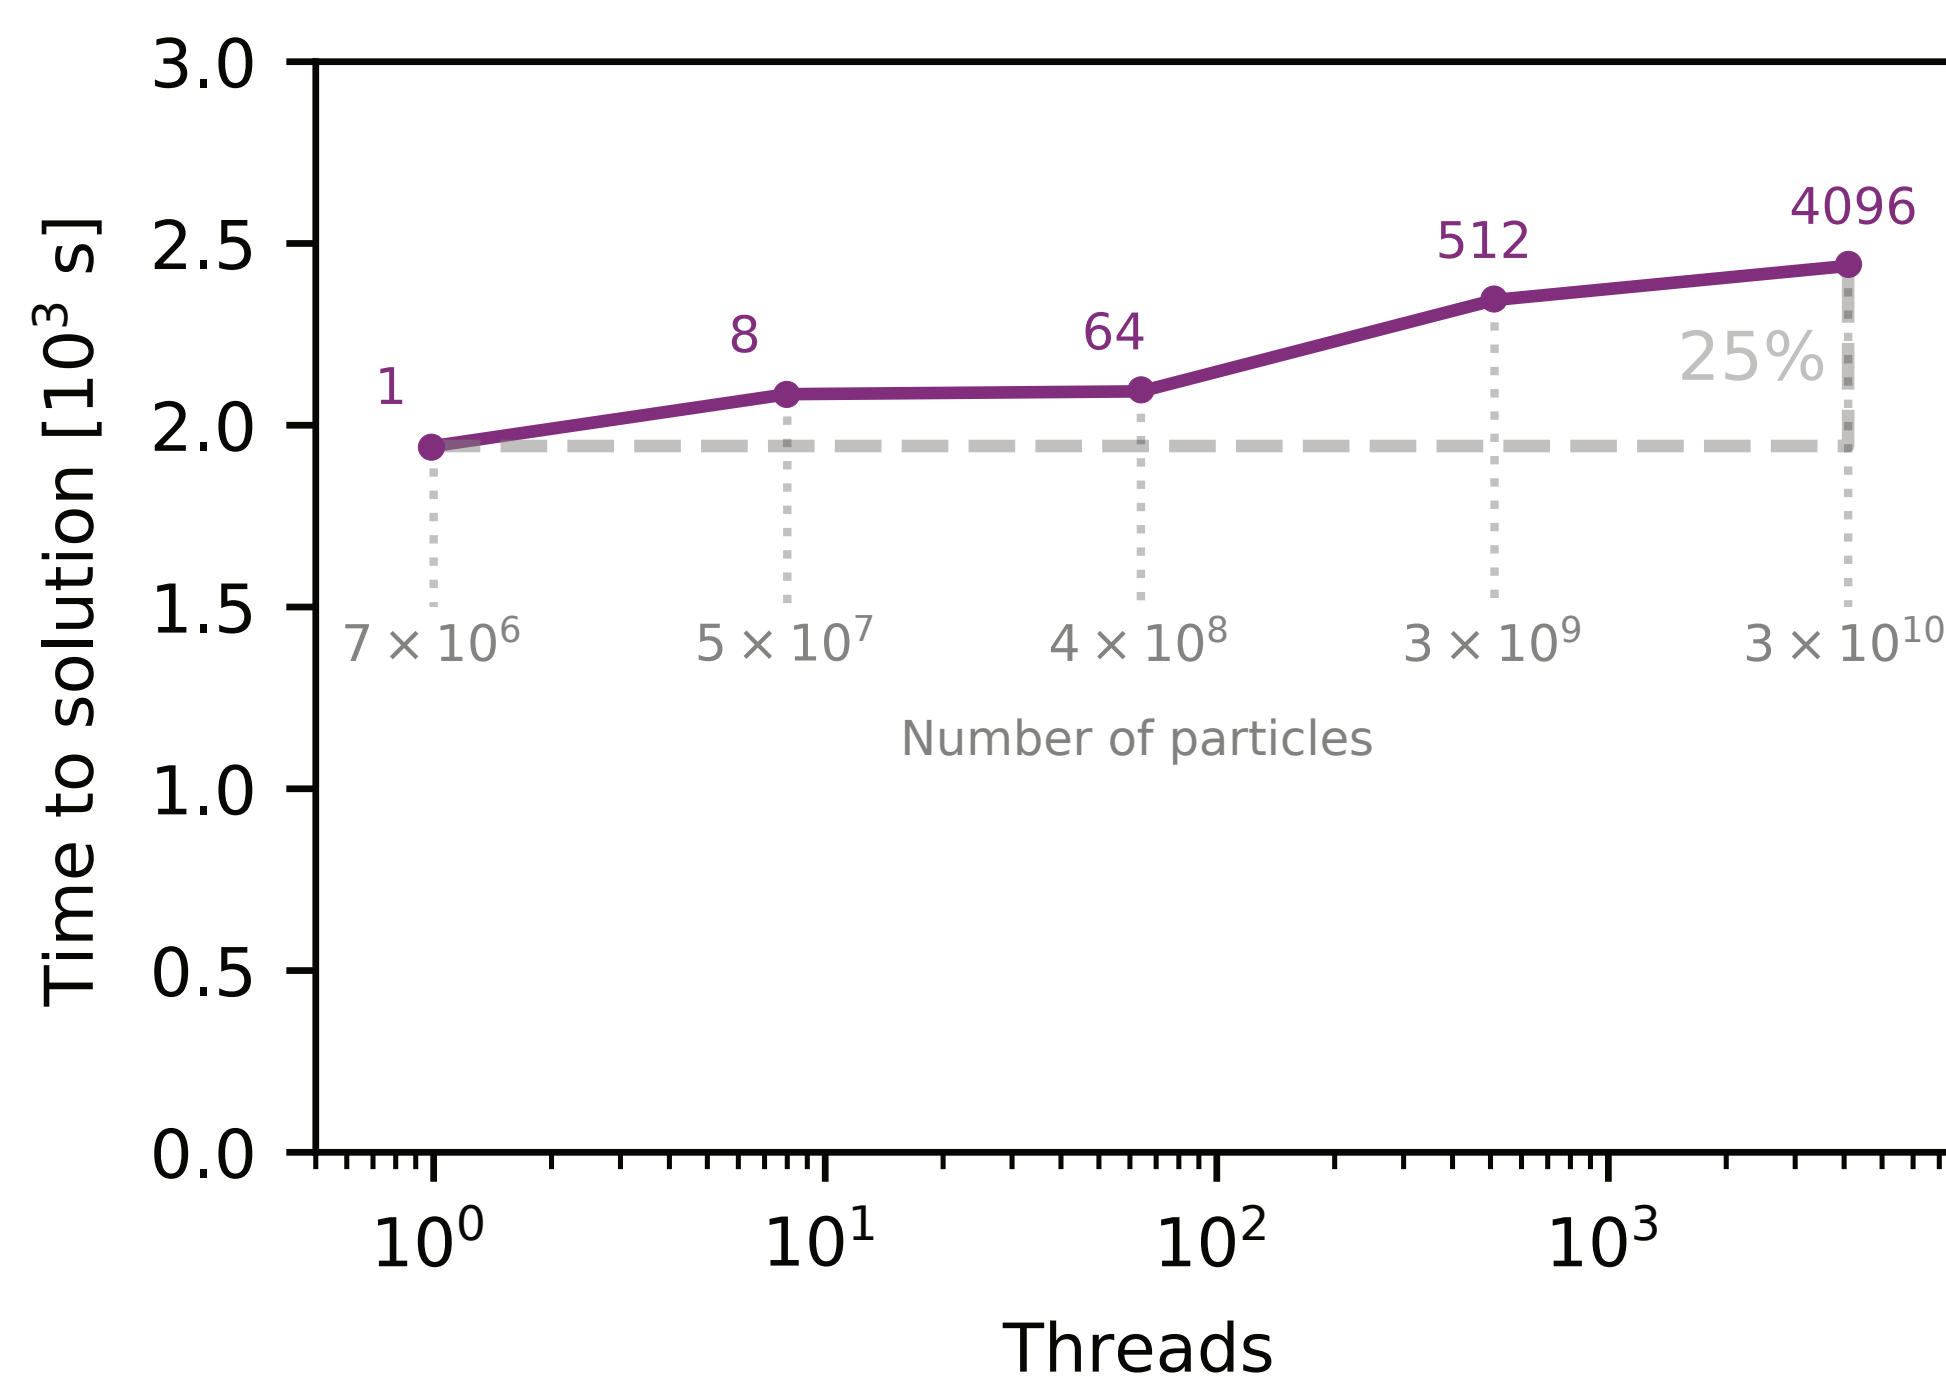


Figure 1: The weak-scaling performance of SWIFT in a worst-case cosmological problem, only increasing runtime by 25% when replicating 4000 times, using the 25 Mpc EAGLE box at $z=0.1$ (hydrodynamics only) with 6×10^6 particles per thread. Note the linear vertical axis, with logarithmic horizontal.

Simulating the Universe

SWIFT was designed to track the cosmological evolution of the Universe from the big bang to the present day. Accurately reproducing the large-scale structure is the first step; see the image below for an example. The key component in this formation, which turns an initially nearly uniform density field into gravitationally bound structures that contain clusters, groups, and galaxies, is dark matter. Dark matter makes up 80% of the mass in the Universe, and so dominates gravitational attraction. The standard model for dark matter describes it as a cold, pressureless gas, which cannot collapse easily as it has very few avenues to lose angular momentum. This gives the structure that forms, dark matter 'halos', a characteristic spherical shape and density profile.




Figure 2: A density slice through an 800 Mpc box containing 4×10^9 dark matter particles, simulated recently with SWIFT. The abundances of dark matter halos in this simulation are shown in Figure 4.

Cosmological Tests

After many years of development, SWIFT is now ready to run cosmological problems. The standard cosmological test case is the Santa Barbara Cluster (Frenk et al. 1999). This cluster of galaxies is tracked all the way from just after the big bang through to today. Below we show the gas that is present in the cluster at the end of the simulation. The equations of hydrodynamics are solved in SWIFT using Lagrangian methods, such as smoothed particle hydrodynamics (SPH, for which we provide hand-vectorised AVX-512 kernels, see Willis et al. 2018), and meshless methods which include an explicit Riemann solver.




Figure 3: Temperature projection of the Santa Barbara cluster re-simulated with SWIFT, using the same 'flavour' of SPH as is present in the de-facto standard code Gadget-2. This cluster is ubiquitous in cosmology and enables many helpful comparisons between codes, from performance to results.

Numerical Comparisons

In addition to images, it is useful to have some numerical comparisons to validate your code. One key metric in cosmology is the halo mass function (HMF); this gives the number of dark matter halos present in a volume with a given mass. To compare against the de-facto code in cosmology, Gadget-2, SWIFT was re-run on some previous problems from the P-Millennium suite (Baugh et al. 2018). The results are presented below, and show near-exact agreement between the two codes, as well as the analytical predictions.

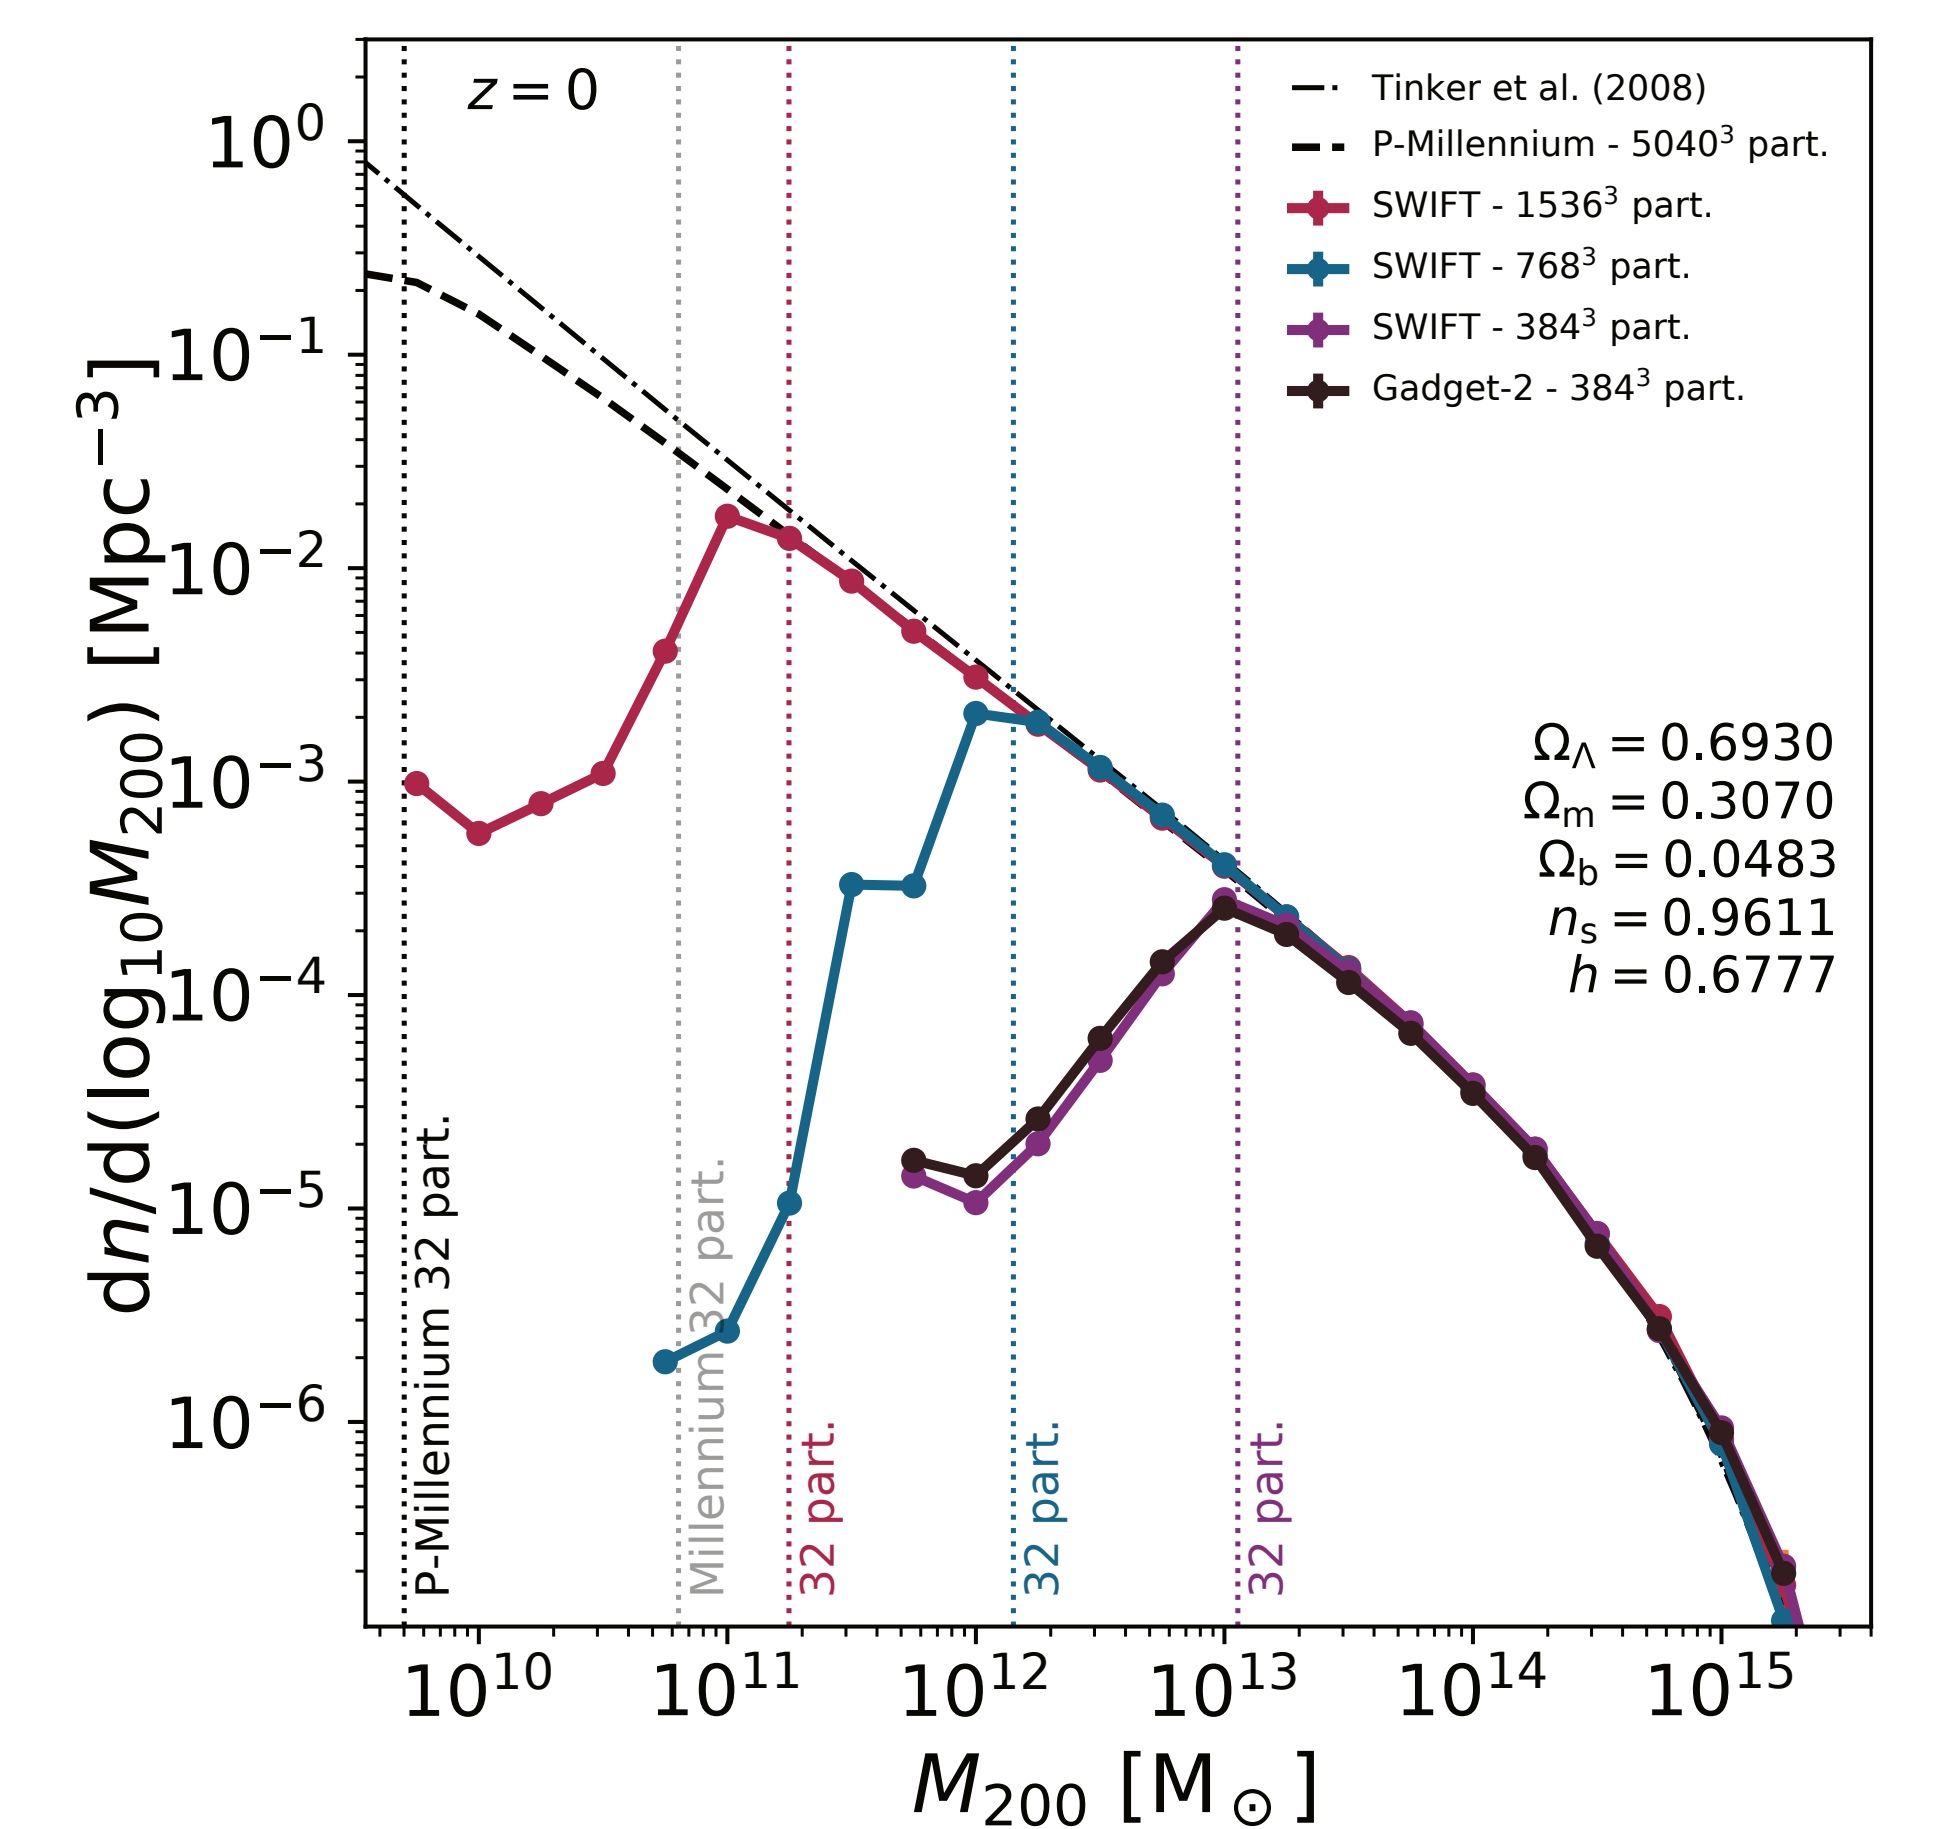


Figure 4: Convergence between SWIFT, Gadget-2, and the analytical prediction, for the HMF. The minimal number of particles required to sample a halo is at least 32. For both codes, halos were identified using the code Velociraptor, using a 6D friends-of-friends algorithm. The required run time for the 384³ particle simulation was 6.5 times higher for Gadget-2 than SWIFT; for larger boxes with more particles this gain is expected to improve due to the better scaling, load-balancing, and i/o performance of SWIFT.

References

Schaller, M., et al. (2016). SWIFT: Using task-based parallelism, fully asynchronous communication, and graph partition-based domain decomposition for strong scaling on more than 100 000 cores. In Proceedings of PASC 16 (pp. 1-10).
 Borrow et al. (2018). SWIFT: Maintaining weak-scalability with a dynamic range of 10^4 in time-step size to harness extreme adaptivity. Proceedings of the 13th SPHERIC International Workshop, Galway, Ireland, June 26-28 2018, pp. 44-51.
 Frenk, C. S. et al. (1999). The Santa Barbara Cluster Comparison Project: A Comparison of Cosmological Hydrodynamics Solutions. The Astrophysical Journal, 525(2), 554-582.
 Willis, J. S. et al. (2018). An Efficient SIMD Implementation of Pseudo-Verlet Lists for Neighbour Interactions in Particle-Based Codes. Advances in Parallel Computing, 32, 507-516.
 Baugh, C. M., et al. (2018). Galaxy formation in the Planck Millennium: the atomic hydrogen content of dark matter halos. https://arxiv.org/abs/1808.08276