Exact and approximate methods in stellar dynamics

Yohai Meiron[1,2] Long Wang[3,1], Rainer Spurzem[2,1,4]

[1] Kavli Institute for Astronomy and Astrophysics at Peking University, Yiheyuan Road 5, Haidian District, 100871 Beijing, China
[2] National Astronomical Observatories of China, Chinese Academy of Sciences, 20A Datun Road, Chaoyang District, 100012 Beijing, China
[3] Department of Astronomy, Peking University, Yiheyuan Road 5, Haidian District, 100871 Beijing, China
[4] Astronomisches Rechen-Institut, Zentrum für Astronomie, University of Heidelberg, Mönchhofstraße 12-14, D-69120 Heidelberg, Germany

ABSTRACT Stellar systems come in many shapes and sizes. We present two new GPU-accelerated N-body codes focusing on two kind of systems: dwarf spheroidal galaxies and globular clusters. *ETICS* is based on series expansion of the Poisson equation and is ideal for diffuse objects such as dwarf galaxies. Since in globular clusters close stellar encounters and binaries play very important roles in the dynamics, a much more accurate integrator is needed. *NBODY6++* is a direct-summation N-body code which can provide this kind of accuracy.

Meiron Y. et al. 2014, ApJ, 792, 98
Wang L. et al. 2015, submitted to MNRAS

BACKGROUND: COLLISIONAL AND COLLISIONLESS IN ASTROPHYSICS

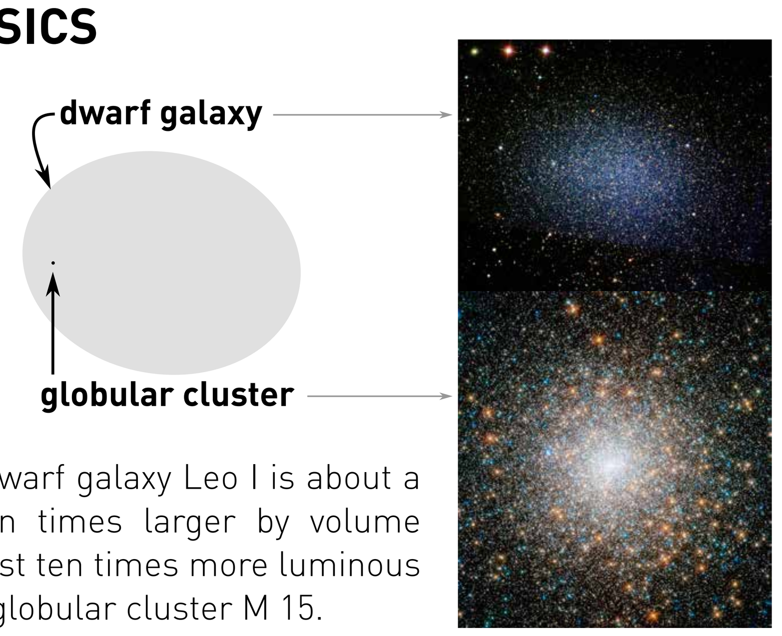

The dwarf galaxy Leo I is about a million times larger by volume but just ten times more luminous than globular cluster M 15.

Dwarf galaxies and globular clusters are two kinds of interesting stellar systems. Dwarf galaxies are often small satellite of the Milky Way galaxy, and have up to a few 10^9 stars. Globular clusters are smaller (up to 10^6 stars), and much more compact in comparison. While it seems that dwarf galaxies would pose the bigger challenge due to having more paritcles, this is not the case becasue these objects are **collisionless**, meaning that on the relevant time scales, a star is unlikely to have any strong interaction with another star. Thus, the gravitational field can be calculated in an *approximate* way that neglects these star-star interactions. For globular clusters, this is not the case: strong interactions between starts do occur during the cluster's life time and significantly affect the dynamics (these are **collisional** objects). Thus, globular clusters have to be integrated with an *exact* methods like direct summation.

ETICS – EXPANSION TECHNIQUES IN COLLISIONLESS SYSTEMS

Expansion techniques approximate the gravitational field in a way that gets rid of the small fluctuations due to finite sampling, while also reducing the code's complexity. *ETICS* is a GPU implementation of the **Multipole Expansion (MEX)** and the **Self-Consistent Field (SCF)** methods, which are two kinds of expansion techniques. MEX is a *Taylor*-like expansion of the Green's function, while SCF is a *Fourier*-like expansion of the density. In both methods the integrand below is written as a series of functions (of $\boldsymbol{r}$) with coefficients: in MEX one uses the given density to evaluate the functions, while their coefficients are known in advance; in SCF one evaluates coefficients, while the functions are known in advance.

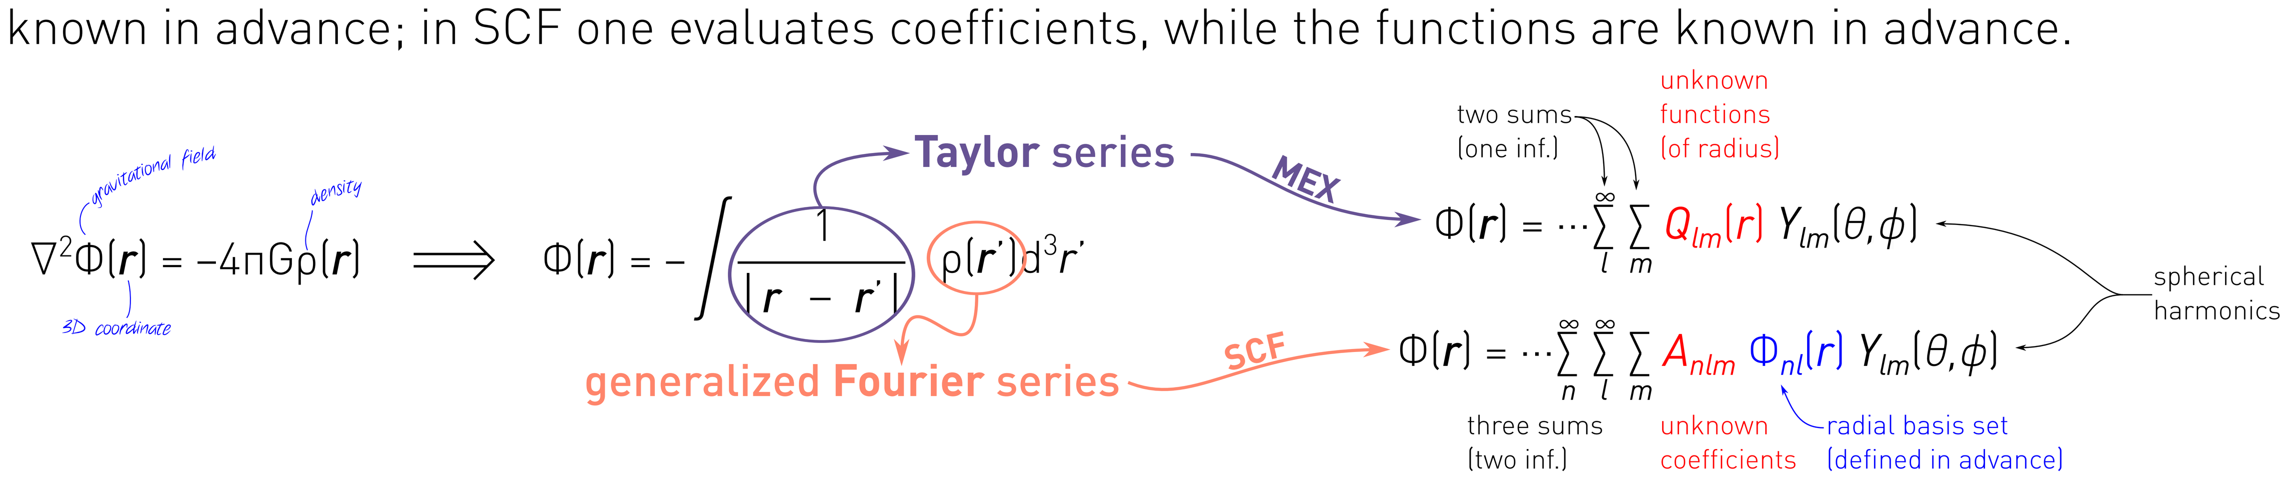

The unknown coefficients or functions are claculated from the density field by projection. The infinite sums are in practice truncated at $l_{\max}$ and $n_{\max}$ which correspond to *angular* and *radial* information, repectively.

Both algorithms are quite memory intensive, requiring many transformations of the data. To minimize access to global memory, we chose the loop order carefully so recursion relations are used when possible. MEX requires sorting and cumulative summation of the multipoles, for which we used the *Thrust* template library. In the SCF algorithm, we have to collect and sum hundreds of coefficients from every particle. There is not enough shared memory to contain these data and employ a reduction algorithm, so the process is broken up into nested loops, the innermost is executed inside a kernel.

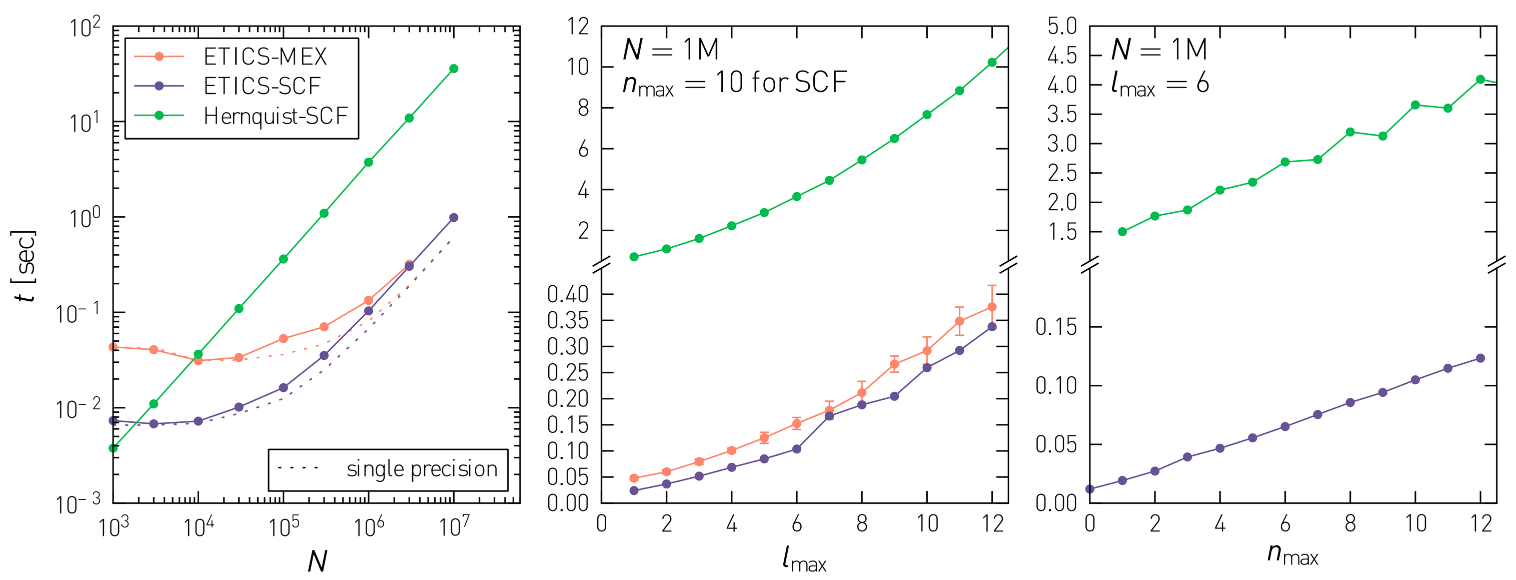

Scaling of *ETICS* with number of particles and expansion cutoff(s); t is the time for one full force calculation. Hernquist's SCF code (in green) is a CPU code and was tested on Intel Xeon E5520 (one core). *ETICS* was tested with NVIDIA Tesla K20. For the scaling with N we set $l_{\max}$ = 6, and for the SCF codes $n_{\max}$ = 10 too. The scaling is theoretically linear with N for SCF and $N \log N$ for MEX, but the theoretical behavior is only seen asymptotically for the GPU codes, since the GPU is not fully loaded at low N. Both methods scale quadratically with $l_{\max}$. SCF scales linearly with $n_{\max}$. For SCF on this CPU-GPU pair, we get acceleration factor of up to ~ 35 in double precision and ~ 60 in single precision.

NBODY6++ – DIRECT *N*-BODY CODE FOR GLOBULAR CLUSTER INTEGRATION

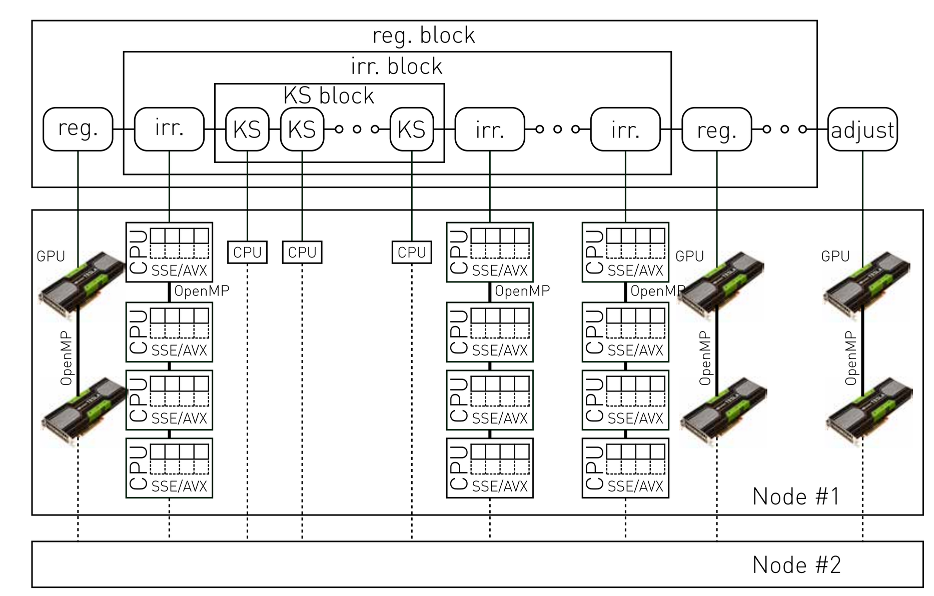

NBODY6++ is a direct integration N-body code, meaning that all intercations are considered for extremely accurate integration. This legacy code is based on Sverre Aarseth's work over the past 50 years, and uses several mathematical innovations to ease the computational demand while retaining the direct-summation nature of the code (such as the neighbor scheme, block time steps and regularization of binary stars). Additionally, the latest hardware technology is used: the bottlneck of the code (the "regular" force calculation, i.e. the force due to all stars but the nearest neighbors) has been moved to the **GPU**, while other critical parts of the code are also accelerated using GPU, **OpenMP and AVX/SSE**. *NBODY6++* is an extension of Aarseth's *NBODY6* to **supercomputers by using MPI**. The current largest *direct* N-body simulation has about $5{\times}10^5$ stars, we aim to simulate 10^6 stars, which is closer to a realistic globular cluster, to obtain better understanding of the dynamical evolution of these objects.

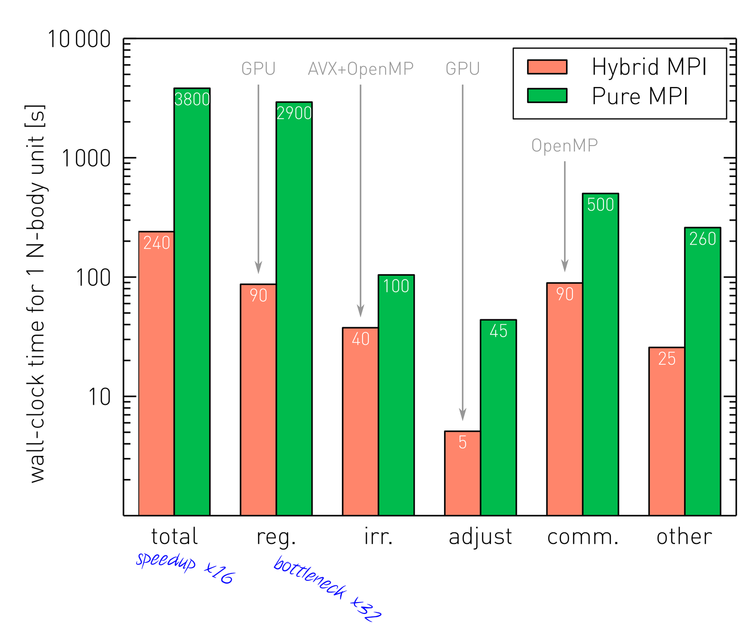

The figure on the left shows a wall-clock time comparison between the pure hybrid and pure MPI versions of *NBODY6++*. The former uses only MPI for parallelization, while the latter uses additional acceleration technilogies (GPU, OpenMP and AVX/SSE). We tested 256k particles for one N-body time unit on four nodes of the "Kepler" cluster at the ARI (University of Heidelberg), each equipped with one NVIDIA Tesla K20m GPU and two Intel Xeon E5-2650 CPUs.

CONCLUSIONS Realistic simulations of dwarf galaxies could use an approximate force calculation such as the *ETICS* code which is ×35 as fast a s CPU counterpart for double precision. For globular cluster simulations, direct methods are needed, and we have made a huge step forward by accelerating *NBODY6++* using GPUs, this will bring a total speedup of ×30 for 10^6 bodies on 16 GPUs.

the Silk Road Project 丝绸之路计划

We thank **Peter Berczik** and the **Silk Road Project team**, **Sverre Aarseth**, **Keigo Nitadori**, **Baile Li** and **Kelly Holley-Bockelmann** for close collaboration on these projects. RS thanks the Qianren Thousand Talents Plan of the People's Republic of China for support; we thank the Center of Information and Computing at National Astronomical Observatories for help and support with the "Laohu" GPU supercomputer; the team of IAS/JSC Jülich, Germany for substantial support and collaboration; the University of Heidelberg and German Science Foundation DFG for support in SFB881 "The Milky Way System" Project Z2. We thank NVIDIA/China and **John Xie** for supporting our Laohu cluster upgrade project.