arXiv:1508.02510v1 [astro-ph.IM] 11 Aug 2015

Performance analysis of parallel gravitational *N*-body codes on large GPU cluster

Siyi Huang[1], Rainer Spurzem[1,2,4], Peter Berczik[1,3,4]

[1] National Astronomical Observatories and Key Laboratory of Computational Astrophysics, Chinese Academy of Sciences, Beijing 100012, China; *huang41@nao.cas.cn*
[2] Kavli Institute for Astronomy and Astrophysics, Peking University, Beijing 100871, China
[3] Main Astronomical Observatory, Ukrainian National Academy of Sciences, 03680 Kiev, Ukraine
[4] Astronomisches Rechen-Institut, Zentrum für Astronomie, Universität Heidelberg, 69120 Heidelberg, Germany

Abstract We compare the performance of two very different parallel gravitational *N*-body codes for astrophysical simulations on large GPU clusters, both pioneer in their own fields as well as in certain mutual scales - NBODY6++ and Bonsai. We carry out the benchmark of the two codes by analyzing their performance, accuracy and efficiency through the modeling of structure decomposition and timing measurements. We find that both codes are heavily optimized to leverage the computational potential of GPUs as their performance has approached half of the maximum single precision performance of the underlying GPU cards. With such performance we predict that a speed-up of $200-300$ can be achieved when up to 1k processors and GPUs are employed simultaneously. We discuss the quantitative information about comparisons of two codes, finding that in the same cases Bonsai adopts larger time steps as well as relative energy errors than NBODY6++, typically ranging from $10-50$ times larger, depending on the chosen parameters of the codes. While the two codes are built for different astrophysical applications, in specified conditions they may overlap in performance at certain physical scale, and thus allowing the user to choose from either one with finetuned parameters accordingly.

Key words: methods: analytical — methods: data analysis — methods: numerical

1 INTRODUCTION

Algorithms for gravitational *N*-body simulations, which are widely used tools in astrophysics nowadays, have mainly evolved into two categories over the passed decades. Traditionally, by computing the pairwise force among particles, the direct summation method has been employed as the core idea of the so-called "direct *N*-body code". High accuracy can be archived by choosing smaller time steps, with higher computational costs. Some best-known examples are the NBODY series codes developed by Aarseth (1999) and the Starlab environment developed by Hut, McMillan, Makino, Portegies Zwart, etc (Hut 2003). Alternatively, the force calculation can be approximated with certain assumptions. In the late 1960s, some approximation algorithm such as tree-code or mesh-code are developed in attempt to reduce the computational complexity so that larger simulations can be scaled on limited hardware with acceptable time. One of the most prominent approximated algorithm, namely Barnes-Hut tree (Barnes & Hut 1986), also have many implementations such as GADGET developed by Springel (2005) and PEPC developed by Gibbon, Winkel and collaborators (Winkel et al. 2012) .

Since direct summation methods use an all-to-all particle force direct summation method they have a raw computational complexity $\mathcal{O}(N^2)$. Additional algorithms are developed to cut the absolute wall-clock time down in spite of the prohibitive asymptotic complexity. On the other side approximate schemes reduce complexities to $\mathcal{O}(N \log N)$ or even $\mathcal{O}(N)$ thanks to the approximate treatments in force computations and some special structures such as octree or grids. However, such intelligent approximative algorithms may not be very suitable for the simulation of certain astronomical systems such as dense star clusters, e.g. globular clusters or nuclear star clusters with or without central massive black holes. This is because in these systems two-body relaxation is important, which can only be correctly modeled by following pairwise particle interactions with high precision at large distances. They require the use of direct summation methods, which has so far great difficulties reaching even one million particles (but see Wang et al. 2015). As such, the only practical approach at the present to handle simulations with ultra-high particle numbers (e.g. cosmological structure formation) is through the employment of approximate methods, despite their lack of resolution at small scales (Shin et al. 2014; Genel et al. 2014; Vogelsberger et al. 2014).

Consequently, parallel technologies are applied for *N*-body simulations as a proper solution. With the help of parallel hardware, simulations could be accelerated multifold in accordance with the numbers of invoking processors theoretically. In practice parallel schemes were implemented, as well as performance analysis of parallel *N*-body codes on supercomputers or distributed systems, such as the work by Gualandris et al. (2007). Supercomputer clusters were used for the parallelization of *N*-body simulations, special devices were also added in order to process the hot computation sections. In the beginning, special-purpose architectures called GRAPE series (Makino et al. 2003) are designed for *N*-body simulations exclusively, which achieved speed-ups by putting the whole force calculations into the hardware that placing many pipelines on one chip, detailed performance of GRAPE measured by Harfst et al. (2007). In recent years, with the rapid development of hardware manufacturing technique the GPU (Graphics Processing Unit) as general-purpose devices are more and more used and acted the same role as GRAPE. Now architectures consisted of many processors and equipped with corresponding GPUs are prevalent solutions in *N*-body simulations.

Parallel *N*-body simulation software running partially or even entirely on GPUs were developed subsequently (Berczik et al. 2011, 2013; Spurzem et al. 2012; Bédorf et al. 2012a,b), while in practice applications the performance would not measure up to ideal speed-up because of some inevitable serial code in the code structure. The actual speed-up is limited by sequential fractions in codes and not directly proportional to the number of processor cores, the theoretical maximum value could be predicted by *Amdahl's law* (Amdahl 1967). What's more, this peak value is unapproachable on account of communication overhead between multiple processors. The effectiveness of either parallelization or GPU acceleration introduced in *N*-body software is not intuitive but interested.

In this paper, we focus on the performance analysis of two kinds of *N*-body software, direct *N*-body code NBODY6++ and tree-code Bonsai , which both can be executed in parallel and accelerated by GPUs. Section 2 describes an overview of the software and hardware we used. Section 3 describes the performance models used to analyse complicated *N*-body codes, provides the detailed measurements, performance results and reasonable predictions. Section 4 describes the performance comparisons and analysis, then makes a conclusion which gives us a better reference in the choice of opportune scheme of software type and hardware scale in *N*-body simulations.

2 SOFTWARE AND HARDWARE

2.1 Direct *N*-body implementation: `NBODY6++`

In this section we provide a brief description of NBODY6++ , which we used for the performance analysis of direct *N*-body code.

NBODY6++ developed by Rainer Spurzem is a parallel version of NBODY6 (Spurzem 1999; Khalisi et al. 2003; Spurzem et al. 2008). The standard NBODY6 is the 6^{th} generation of NBODY code initiated by Sverre Aarseth, who has a lifelong dedication to the development of the family of NBODY series

codes (Aarseth 1999). The first code NBODY1 was a basic direct *N*-body code with individual time steps. Ahmad-Cohen neighbour scheme (Ahmad & Cohen 1973) was used in NBODY2 and NBODY5 made it possible to treat larger systems. Kustaanheimo-Stiefel (1965) two-body regularization and chain regularization were applied in NBODY3 and NBODY5 to deal with close encounters. By the time NBODY6 had developed, the code included both neighbour scheme and regularizations, as well as applied with Hermite scheme integration method combined with hierarchical block time steps.

NBODY6++ is a descendant of the standard NBODY6, it kept those features of its predecessor mentioned above as well as increased the efficiency by redesigning the algorithms suitable for parallel hardware. NBODY6++ used the SPMD (Single Program Multiple Data) scheme to achieve parallelism, in this mode multiple autonomous processors simultaneously start with chunked local data and then communicate with each other through the *copy algorithm* (Makino 2002; Dorband et al. 2003), which is a parallelized algorithm assumes that each processor has a local copy of the whole system and every processor handles the subgroup of data itself then broadcasts the new data to all the other processors immediately. The parallelization scheme of NBODY6++ is implemented with the standard MPI library package.

The most major improvement of NBODY6++ is the parallelization of regular and irregular force computations, which were special concepts introduced from Ahmad-Cohen neighbour scheme that divided the full force of each particle involved by other all particles into two parts: one part called irregular force that has a frequent but short time steps for interactions with adjacent particles, another one called regular force which has a longer time steps for full interactions. By assigning the most expensive overhead sections to multiple processors NBODY6++ achieved the expected efficiency. Moreover, the heavier regular force computation component was adapted for GPU acceleration using CUDA. The performance of parallel accelerations showed in Section 3.1.

2.2 *N*-body tree-code implementation: **Bonsai**

In this section we provide a brief description of Bonsai , which we used for the performance analysis of *N*-body tree-code.

Tree-code algorithm as a widely used method nowadays for *N*-body simulations is originally introduced by Barnes & Hut (1986). This algorithm reduces the computational complexity of *N*-body simulation from $\mathcal{O}(N^2)$ to $\mathcal{O}(N \log N)$, therefore improves the simulating scale compared to the brute force methods. Here Bonsai , a parallel GPU tree-code implementation developed by Jeroen Bédorf, Evghenii Gaburov and Simon Portegies Zwart (Bédorf et al. 2012a,b), is a suitable representative of gravitational *N*-body tree-code in recent years.

Certain schemes are introduced in Bonsai to ensure the high efficiency of the code (Bédorf et al. 2012a). A sparse octree is used as the data structures, which means the structure is the three-dimension extension of a binary tree where tree-cells are not complete and equal, it is based on the underlying particle distribution. The tree is constructed layer-by-layer from top to bottom, and inverted the direction in traverse process. Tree-cell properties are updated during the steps, and the integration of the simulation are advanced. The depth of tree traversing affects both accuracy and time consumption crucially, which is determined by the multipole acceptance criterion (MAC) in the tree-code. The criterion is described as follows,

$$d > \frac{l}{\theta} + \delta \tag{1}$$

where d is the smallest distance between a group and the cell's center-of-mass, l is the length of the cell, δ is the distance between the cells center-of-geometry and the center-of-mass, θ is an opening angle parameter to control the accuracy. If the inequality is satisfied then traversing process will be interrupted while multipole moment will be used.

Like other existing *N*-body codes, Bonsai uses parallel technique to reach large scale or high resolution simulations, and applies GPUs to speed up force computations. In contrast to those GPU tree-code, Bonsai executes all parts of the algorithm on GPUs, to avoid the bottlenecks generated from CPU-GPU communications. In Section 3.2 the performance of main parts in the code are presented.

2.3 Hardware environments and initial conditions

The supercomputer we mainly used for all the tests of both software presented above is an IBM iDataPlex Cluster JUDGE named "The Milky Way System" partition provided and maintained by Jülich Supercomputing Centre in Germany, which is a dedicated GPU cluster using 2 Intel Xeon X5650 6-core processors and 2 NVIDIA Tesla M2050/M2070 GPU cards in every node, with 206 compute nodes and 239 Teraflops peak performance in total.

Our performance measurements involved two different kinds of parallel gravitational GPU-accelerated N-body codes: NBODY6++ and Bonsai . The initial conditions of all tests of both codes are consistent with each other, starting with Plummer model and running over standard 1 N-body Time Unit. The number of particles used ranges from $N = 2^{13}(8k)$ to $2^{20}(1M)$, doubled the number over the interval successively. There are additional tests using larger particle numbers up to $N = 2^{24}(16M)$ in Bonsai code runs. The number of processors we chose is the series increasing number $N_p = 1, 2, 4, 8, 16, 32$. Other parameters which are necessary but specific only in each code, such as time step factor for regular/irregular force polynomial and desired optimal neighbour number in NBODY6++ , or accuracy control parameter θ and softening value ϵ in Bonsai , will be described detailed in Section 3.

3 PERFORMANCE

In this section we evaluated the performance of these two GPU-based parallel N-body simulation codes (i.e. NBODY6++ and Bonsai) which we tested mainly on the Jülich Dedicated GPU Environment described in Section 2.3.

In spite of a vast difference between two codes derived from their own fundamental algorithms and specific details, which make it difficult to give a one-to-one comparison, there are some global values providing sufficient information. Timing variables, speed-up and hardware performance indicators like speed and bandwidth were measured below for performance analysis of codes.

3.1 Performance of NBODY6++

NBODY6++ , a parallel direct gravitational N-body code, is featured with a couple of elegant algorithms and schemes developed and maintained over the past few decades. The procedures of durative ameliorations enabled more realistic size of simulations running in achievable circumstances while increased sophistication as well. As a consequence we present a performance model for analysing the overall behaviour as well as main components of the code. Through this model we will have a better idea about the performance of a typical direct N-body code and predictions about the code behaviour further in larger scales.

3.1.1 Performance model

We measured running time directly for evaluating performance and modelling. In NBODY6++ , the total wall-clock time T_{total} required to advance the simulation for a certain integration interval can be written as

$$T_{\text{tot}} = T_{\text{force}} + T_{\text{comm}} + T_{\text{host}} \tag{2}$$

where $T_{\text{force}} = T_{\text{reg}} + T_{\text{irr}} + T_{\text{pre}}$ is time spend on both host and device involving force calculations, here T_{reg}, T_{irr} and T_{pre} are time spend on force computations of regular time step, irregular time step respectively and prediction; $T_{\text{comm}} = T_{\text{mov}} + T_{\text{mci}} + T_{\text{mcr}} + T_{\text{syn}}$ is time spend on data moving for parallel components, MPI communications after regular and irregular blocks and synchronizations of processors; T_{host} is time spend on host side which are absolutely sequential runs. All of the time variables are measured directly by standard Fortran functions ETIME in sequential mode and MPI_WTIME in parallel mode. The entire descriptions are gathered in the last glossary (Table 3).
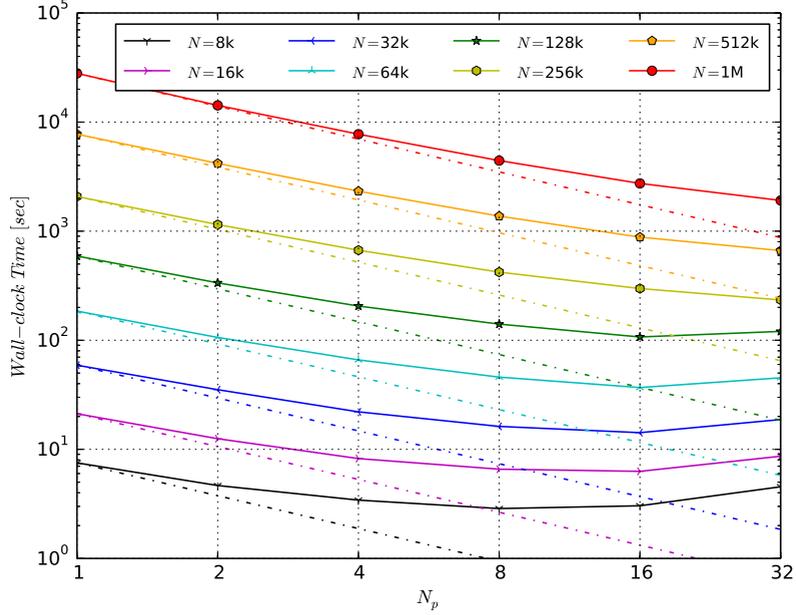
Fig. 1 Total wall-clock time (T_{tot}) of `NBODY6++` as a function of N & N_p. Solid lines are the measured values of running time, dashed lines are the ideal acceleration by increasing processor numbers. (The unit symbols in the legend have the magnitudes: $1\text{k} = 1,024$, $1\text{M} = 1\text{k}^2$ and $1\text{G} = 1\text{k}^3$, similarly hereinafter.)

Table 1 Main components of `NBODY6++`

Description	Timing variable	Expected scaling N	N_p	Fitting value [sec]
Regular force computation	T_{reg}	$\mathcal{O}(N_{\text{reg}} \cdot N)$	$\mathcal{O}(N_p^{-1})$	$(2.2 \cdot 10^{-9} \cdot N^{2.11} + 10.43) \cdot N_p^{-1}$
Irregular force computation	T_{irr}	$\mathcal{O}(N_{\text{irr}} \cdot \langle N_{nb} \rangle)$	$\mathcal{O}(N_p^{-1})$	$(3.9 \cdot 10^{-7} \cdot N^{1.76} - 16.47) \cdot N_p^{-1}$
Prediction	T_{pre}	$\mathcal{O}(N^{kn_p})$	$\mathcal{O}(N_p^{-kp_p})$	$(1.2 \cdot 10^{-6} \cdot N^{1.51} - 3.58) \cdot N_p^{-0.5}$
Data moving	T_{mov}	$\mathcal{O}(N^{kn_{m1}})$	$\mathcal{O}(1)$	$2.5 \cdot 10^{-6} \cdot N^{1.29} - 0.28$
MPI communication (regular)	T_{mcr}	$\mathcal{O}(N^{kn_{cr}})$	$\mathcal{O}(kp_{cr} \cdot \frac{N_p-1}{N_p})$	$(3.3 \cdot 10^{-6} \cdot N^{1.18} + 0.12)(1.5 \cdot \frac{N_p-1}{N_p})$
MPI communication (irregular)	T_{mci}	$\mathcal{O}(N^{kn_{ci}})$	$\mathcal{O}(kp_{ci} \cdot \frac{N_p-1}{N_p})$	$(3.6 \cdot 10^{-7} \cdot N^{1.40} + 0.56)(1.5 \cdot \frac{N_p-1}{N_p})$
Synchronization	T_{syn}	$\mathcal{O}(N^{kn_s})$	$\mathcal{O}(N_p^{kp_s})$	$(4.1 \cdot 10^{-8} \cdot N^{1.34} + 0.07) \cdot N_p$
Sequential parts on host	T_{host}	$\mathcal{O}(N^{kn_h})$	$\mathcal{O}(1)$	$4.4 \cdot 10^{-7} \cdot N^{1.49} + 1.23$

Notes: Detailed descriptions of used symbol gathered in Table 3.

According to the decomposition described above we broke down the code structure and measured these main sections which have heavy weights in code. Owing to a large amount of variables and the high complexities some insignificant components in `NBODY6++` are not counted in. For every parts to be analysed we listed the expected scaling and optimal fitting value in Table 1, which are gotten from the structure of code implementation, chronograph and fitting functions. A python function `scipy.optimize.curve_fit` are used to obtain the optimal fitting value, which based on non-linear least squares.
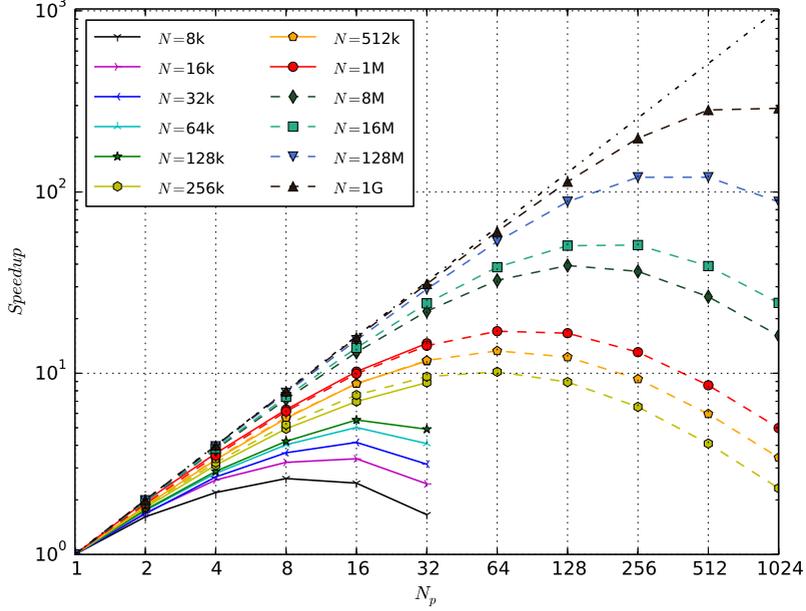
Fig. 2 The speed-up (S) of NBODY6++ as a function of N & N_p. Solid lines are the measured speed-up ratio between sequential and parallel wall-clock time, dashed lines are the predicted performance of larger scale simulations.

The conception speed-up is used for evaluating the parallelism of the code. There are a couple of definitions of speed-up with different ranges. The ideal maximum speed-up $S_i = N_p$ will never be accessible, where N_p is the number of processors used. Unreachable as well, but a more reasonable indicator to predict the theoretical maximum speed-up is so-called *Amdahl's law*, which is defined as

$$S_a(N_p) = \frac{T(1)}{T(N_p)} = \frac{1}{(1 - X) + \frac{X}{N_p}} \tag{3}$$

where X is the fraction of the algorithm that can benefit from parallelization. In practice there is another experiential speed-up to be measured through timer recording, which is given by

$$S_e(N_p) = \frac{T_{\text{tot}}(1)}{T_{\text{tot}}(N_p)} \tag{4}$$

where $T_{\text{tot}}(1)$ & $T_{\text{tot}}(N_p)$ are both the measured values of actual running time. By combining the fitting values into the speed-up formula we will have a general overall perception of the code, and by which make a prediction accordingly about the code performance in larger scale simulations.

The speed of force calculation is measured by the extent at which program reaches the peak of computing devices. Here in our tests the computing device particularly refers to the NVIDIA Tesla M2050/M2070 GPU cards, which feature up to 1,030 Gigaflops of single precision floating point performance and 515 Gigaflops of double precision floating point performance per card.

In NBODY6++ , as the total force are divided into two parts we used two speed variables P_{reg} and P_{irr} represent regular and irregular force calculating speed, which are written as

$$P_{\text{reg}} = \frac{N_{\text{reg\_tot}}}{T_{\text{reg}}} = \frac{N_{\text{reg}} \cdot N \cdot \gamma_{\text{h4}}}{T_{\text{reg}}}, \qquad P_{\text{irr}} = \frac{N_{\text{irr\_tot}}}{T_{\text{irr}}} = \frac{N_{\text{irr}} \cdot \langle N_{\text{nb}} \rangle \cdot \gamma_{\text{h4}}}{T_{\text{irr}}} \tag{5}$$
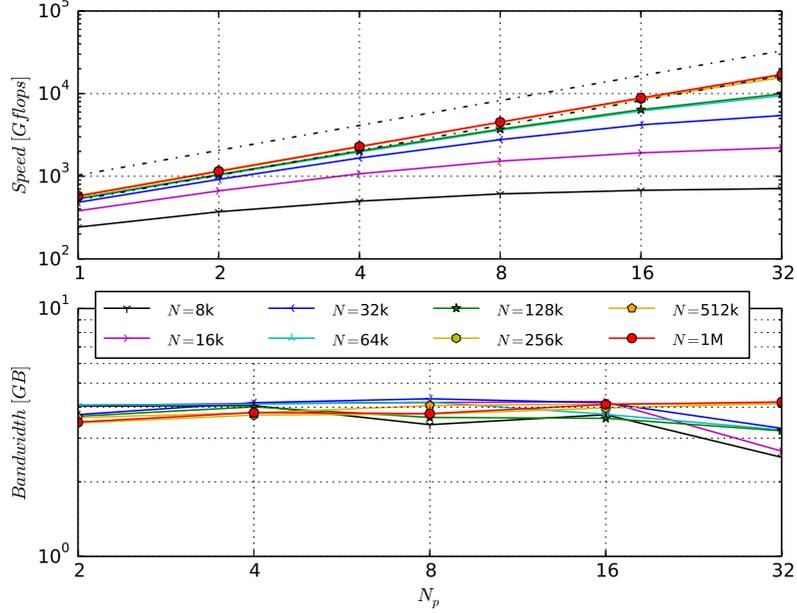
Fig. 3 Hardware performance of NBODY6++ running on "MilkyWay" GPU cluster. The upper panel corresponds to the regular force computation speed (P_{reg}), where two dashed lines refer to the peak single and double precision floating point performance. The lower panel corresponds to the bandwidth of regular part (B_{reg}).

where $N_{[\mathrm{reg}|\mathrm{irr}]\_\mathrm{tot}}$ is the total floating point operations of regular/irrgular force computations, $N_{[\mathrm{reg}|\mathrm{irr}]}$ is the cumulative number of regular/irrgular time steps, $\langle N_{\mathrm{nb}} \rangle$ is the average number of integrated "neighbour" particles, and γ_{h4} defines the floating point operations counts of 4^{th} Hermite scheme per particle per interaction per step, from the work Nitadori & Makino (2008) which is a constant value has $\gamma_{\mathrm{h4}} = 60$.

Bandwidth is measured as a part of hardware performance along with computing speed (P). In NBODY6++ , we defined bandwidth (B_{reg}, B_{irr}) as

$$B_{\mathrm{reg}} = \frac{N_{\mathrm{mcr}}}{T_{\mathrm{mcr}}} = \frac{8 \cdot (41 + l_{max}) \cdot N/N_p}{T_{\mathrm{mcr}}}, \qquad B_{\mathrm{irr}} = \frac{N_{\mathrm{mci}}}{T_{\mathrm{mci}}} = \frac{8 \cdot 19 \cdot \langle N_{\mathrm{act}} \rangle /N_p}{T_{\mathrm{mci}}} \qquad (6)$$

where $N_{[\mathrm{mcr}|\mathrm{mci}]}$ is the number of bytes transferred during MPI communication after regular/irregular blocks, constant terms derived from the size of datasets transferred, in which l_{max} is the maximum size of neighbour lists set manually. Detailed results of all these performance indicators presented in the next section.

3.1.2 Performance results

The measured total wall-clock time of NBODY6++ is shown in Figure 1.

On the whole, the result shows a good extensibility and acceleration when using more processors numbers. To be specific, we assign each part of the code with different weight. Among all of the parts the time spend for force computations always has the highest value, therefore both regular, irregular force computations and prediction have been implemented with parallel algorithm and decrease rapidly when

code running in multi-processors. Here in the heaviest part T_{force}, the regular force computation T_{reg}, which takes the highest fraction of computing time in the former code versions, has been accelerated and implemented on the specific device (GPU), as a consequence of causing a significant reduction of the whole running time costs. Other parts are currently executed on CPU side.

Table 1 shows the Main components of NBODY6++ as a function of N and N_p. As describing above for every part expected scaling is evaluated by code structure, and fitting value is based on experimental data. The fitting process includes two steps by fitting N and N_p successively but independently. Firstly we used a minimum of fixed N_p to avoid the disturbance of processor number and obtained the experimental scaling of N. The fitting values with N are obtained under circumstances which using an increasing particle numbers and fixed single processor ($N_p = 1$), while for cases of multiprocessor-related value(i.e. T_{mci}, T_{mcr}, T_{syn}, which have no numbers in single processor runs) the number of processors changed to $N_p = 2$. The fitting values with N_p are obtained by the second step. Grouping the dataset according to N, then dividing every group data by each N dependent function to get the fitting values with N_p. The fitting results of every main parts listed in the last column. Considering the expected scaling value of main parts, as T_{mov}, T_{syn} and T_{host} have no significant and direct scaling with N from code structure while T_{pre} makes up of two prediction branches that determined by N in next time step, we expect these values as a simple exponent form of N. N_{reg}, N_{irr} and $\langle N_{nb} \rangle$ which used in T_{reg} and T_{irr} are values which are completely dependent with N as $N_{\text{reg}} \propto N^{1.18}$, $N_{\text{irr}} \propto N^{1.10}$ and $\langle N_{nb} \rangle \propto N^{2/3}$, then in the fitting values they are combined together also as an exponent form of N. At last, an unified exponent form of N and N_p are used in the last column rather than other symbols used in the middle column.

By taking the fitting values into the definition of experimental speed-up, we give the prediction about the performance of NBODY6++ , which is shown in Figure 2. As a result, the optimal value of N_p needed for larger simulations of different scales showed clearly from the figure.

Figure 3 shows the hardware performance of NBODY6++ in the actual environment on a real GPU accelerated cluster. Because GPUs played as the central role in acceleration we focus on GPU relevant parts in NBODY6++ , P_{reg} and B_{reg} were drawn in the figure. For the figure of P_{reg}, two dashed lines i.e. peak single and double precision floating point performance are used as the baseline, and the computation speed of different group of N runs are increasingly closer to the peak when N doubled. In the whole NBODY6++ data structures there is two type precisions used in respective parts. Double precision type is used in main loops of the code which declared as the type "REAL*8" in global header file, while for the regular force computation part which accelerated by GPU all of the data convert to the type "REAL*4" then single precision is used in all relevant parts of CUDA routine. Therefore a mixing precision data structure is used in the regular part of NBODY6++ - double precision in data moving process and single precision in data computation process. As the computation process in GPU card dominating the regular part, we use single precision to make a comparison. As shown in the figure, the force computation speed of large N runs passed over half of the maximum single precision performance (for instance, P_{reg} for 1M particle runs have the values $530 \sim 570$ Gflops per M2050/M2070 GPU card.) This proportion concurs with the results of Berczik et al. (2011, 2013), which claimed the speed performance of another direct N-body code ϕ-GPU getting the values $\propto 550$ Gflops per C2050 GPU card and $\propto 1.48$ Tflops per K20 GPU card, both results approached half of the single precision performance peak. Considering the hardware architecture, as operations among various register and memory causes extra inevitable time consumption, the proportion is acceptable in practical environments. For the figure of B_{reg} ignoring the "dropping" points, others remained with the level of more than 80% of the maximum bandwidth performance.

3.2 Performance of **Bonsai**

3.2.1 Performance model

Similarly to Equation (2), in Bonsai the total wall-clock time T_{tot} can be written as

$$T_{\text{tot}} = T_{\text{tree}} + T_{\text{force}} + T_{\text{comm}} + T_{\text{other}} \qquad (7)$$
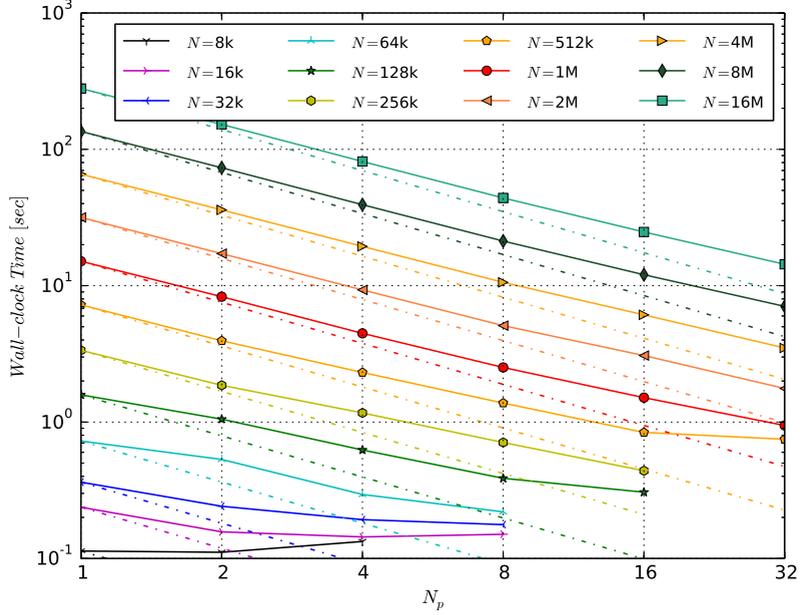
Fig. 4 Total wall-clock time (T_{tot}) of `Bonsai` as a function of N & N_p. The legend is the same as Figure 1. Opening parameters of initial condition set as $\Delta t = 0.0625$, $\epsilon = 0.01$, $\theta = 0.5$.

where $T_{\text{tree}} = T_{\text{sort}} + T_{\text{build}} + T_{\text{prop}} + T_{\text{grp}}$ is time spend on tree building, which mainly included sorting and reordering of the particles along a 1D number string mapped along *Space Filling Curve*, tree structure construction, tree-node properties computation and active groups setting for next steps; T_{force} is time spend on force computations in tree-traverse; $T_{\text{comm}} = T_{\text{dom}} + T_{\text{exch}} + T_{\text{syn}}$ is time spend on distributing or redistributing the particles interprocessor and synchronizations of processors; $T_{\text{other}} = T_{\text{pre}} + T_{\text{corr}} + T_{\text{ene}}$ is time consumptions for other mainly essential parts, like local-tree predictions before tree construction, corrections after force computations and energy check. All of the time variables are measured by CUDA C function `cuEventElapsedTime` from CUDA Event Management Driver API and counted on device (GPU) entirely. Trivial time consumptions on the host side is ignored. The entire descriptions are gathered in the last glossary (Table 3).

In hardware performance aspect, differed from Equation (5) the force calculating speed in `Bonsai` are written as

$$P_{\text{force}} = \frac{N_{\text{force\_tot}}}{T_{\text{force}}} = \frac{N_{\text{force}} \cdot \gamma_{\text{t}}}{T_{\text{force}}} \tag{8}$$

where $N_{\text{force\_tot}}$ is the total floating point operation counts; N_{force} is the cumulative number of interactions; γ_{t} is the number of operation counts for each interaction in tree-code, which we used a constant value has $\gamma_{\text{t}} = 38$ from the work Warren & Salmon (1992); Kawai et al. (1999); Hamada et al. (2009); Hamada & Nitadori (2010), and the result figure shows in Figure 6. Note that Bédorf et al. (2014) used another separated operation counts 23 & 65 for particle-particle and particle-cell interactions respectively.

Table 2 Main components of `Bonsai`

Description	Timing variable	Expected scaling N	Expected scaling N_p	Fitting value [sec]
Sorting and reordering	T_{sort}	$\mathcal{O}(N)$	$\mathcal{O}(N_p^{-1})$	$(1.5 \cdot 10^{-6} \cdot N + 2.45 \cdot 10^{-4}) \cdot N_p^{-1}$
Tree construction	T_{build}	$\mathcal{O}(N)$	$\mathcal{O}(N_p^{-1})$	$(2.8 \cdot 10^{-7} \cdot N + 2.06 \cdot 10^{-2}) \cdot N_p^{-1}$
Node properties	T_{prop}	$\mathcal{O}(N)$	$\mathcal{O}(N_p^{-1})$	$(9.1 \cdot 10^{-8} \cdot N + 5.78 \cdot 10^{-3}) \cdot N_p^{-1}$
Set active groups	T_{grp}	$\mathcal{O}(N)$	$\mathcal{O}(N_p^{-1})$	$(1.7 \cdot 10^{-9} \cdot N + 1.16 \cdot 10^{-3}) \cdot N_p^{-1}$
Force computation	T_{force}	$\mathcal{O}(N\log N)$	$\mathcal{O}(N_p^{-kp_{g1}})$	$(2.5 \cdot 10^{-6} \cdot N\log N - 0.10) \cdot N_p^{-0.88}$
Domain update	T_{dom}	$\mathcal{O}(N\log N)$	$\mathcal{O}(1)$	$5.4 \cdot 10^{-10} \cdot N\log N + 2.96 \cdot 10^{-3}$
Exchange	T_{exch}	$\mathcal{O}(N\log N)$	$\mathcal{O}(1)$	$2.1 \cdot 10^{-9} \cdot N\log N + 1.16 \cdot 10^{-2}$
Synchronization	T_{syn}	$\mathcal{O}(N^{kn_s})$	$\mathcal{O}(kp_{s1} \cdot N_p^{kp_{s2}})$	$(1.4 \cdot 10^{-4} \cdot N^{0.45} + 9.3 \cdot 10^{-4})(0.5 \cdot N_p^{0.49})$
Prediction	T_{pre}	$\mathcal{O}(N)$	$\mathcal{O}(N_p^{-1})$	$(1.5 \cdot 10^{-8} \cdot N + 1.49 \cdot 10^{-3}) \cdot N_p^{-1}$
Correction	T_{corr}	$\mathcal{O}(N)$	$\mathcal{O}(N_p^{-1})$	$(3.8 \cdot 10^{-8} \cdot N + 7.88 \cdot 10^{-4}) \cdot N_p^{-1}$
Energy check	T_{ene}	$\mathcal{O}(N)$	$\mathcal{O}(N_p^{-1})$	$(8.8 \cdot 10^{-9} \cdot N + 7.14 \cdot 10^{-4}) \cdot N_p^{-1}$

Notes: Detailed descriptions of used symbol gathered in Table 3.

3.2.2 Opening parameters in tree-code

Three opening parameters play significantly important roles in tree-code running and affect the performance consequently as a different result.

θ: It is a dimension-less parameter defined in Equation (1) that controls the accuracy. Our test results showed that a smaller θ makes the running time increasing shapely, then stopped rising on a certain range ($\theta \approx 0.01$ as an experimental value); while a bigger θ ($\theta \geqslant 0.2 \sim 0.35$ influenced by N) causes a less accuracy of simulations.

ϵ: The softening parameter ϵ does not contribute to the running time, on the other hand an optimal ϵ could lead to a best approach to the minimum error. For a too small softening the estimates of the forces will be too noisy, while for a too large softening the force estimates will be misrepresented systematically, in between there is an optimal softening. The optimal ϵ depends both on the number of particles and the size of time steps. From the work Athanassoula et al. (2000) when ϵ has different minimum values the conditions of simulations are not the same, there is a relationship between N and optimal ϵ. Through the comparison of their conclusions with groups of our test results of `Bonsai` code (using ΔE instead of *MASE*, $\theta = 0.5$; $\Delta t = 0.0625$), we conclusion that the value of ϵ what leads to a minimum ΔE is consistent with conclusions of the reference.

Δt: The value of Δt affect both running time and energy error. On the running time side, Δt has a noticeable linearly dependent with time as: $t \propto 1/\Delta t$; on the relative energy error side, the tests results are more complex. Our results shows that $\Delta E_{\min}$ varies sensitively under the chosen of combinations of Δt and ϵ, as well as the different N.

3.2.3 Performance results

The measured total wall-clock time of `Bonsai` is shown in Figure 4.

In Figure 4 & 5 we do not use any data for small particle number and large number of GPU nodes, because the performance goes down and the GPU's are not fully loaded in this regime.

We decomposed T_{tot} into main components as described in Section 3.2.1. For every components we measured the running time separately, and obtain fitting formulae as a function of N and N_p. The fitting procedures was the same as in the `NBODY6++` part described in Section 3.1.1, the results are shown in Table 2.

Figure 5 shows the experimental speed-up of `Bonsai` defined as Equation (4). Compared with the result of Figure 2, `Bonsai` has tendency to a lower peak but wider scope. Considering the weight factors of force computational part in Table 1 & 2 quantitative information of ascendant distinctness are revealed, while the different weight of communication part is the mainly determinant for descendant lines.
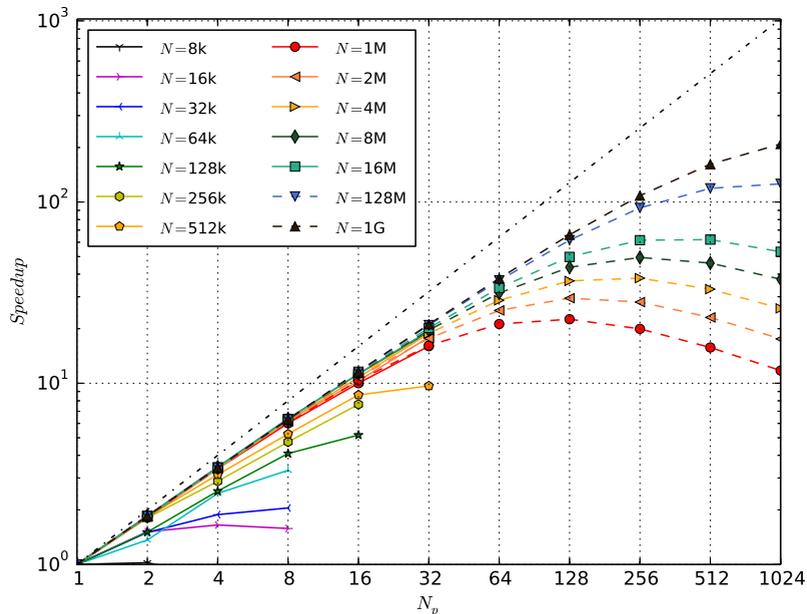
Fig. 5 The speed-up (S) of `Bonsai` as a function of particle number N & N_p. The legend is the same as Figure 2. Opening parameters of initial condition set as $\Delta t = 0.0625$, $\epsilon = 0.01$, $\theta = 0.5$.

Figure 6 shows the hardware performance of `Bonsai` in a practical environment on a real GPU accelerated cluster. The performance in floating point operations per second is given for the dominant part of the force computations only. The figure indicates that for the large enough N we get near half of the peak single precision of our M2050/M2070 GPU card and increasing steadily for large scale N_p. This proportion is similar to the result of NBODY6++ code discussed in Section 3.1. The utilization of the GPU is quite good considering the tree-code structure.

4 DISCUSSION

In this paper we analyze the performance of two very different kinds of *N*-body codes, both pioneers in their fields and both heavily optimized for GPU acceleration and parallelization - NBODY6++ and `Bonsai` . There is always the question what is the turn-even point for the codes, how do they compare with each other. Due to the very different nature of the two codes such a comparison is inevitably unfair - NBODY6++ has few-body regularizations and is aimed for high accuracy of both near and more distant gravitational forces; `Bonsai` achieves optimal performance if the opening parameter θ is relatively large, providing rather less accurate gravitational forces. But in certain ranges of parameters both codes may overlap in performance, accuracy and efficiency. It is the goal of this paper to provide a quantitative information about this.

We do this with the help of the four panels in Figure 7 - they show wall-clock time and energy accuracy as a function of the average time step; the main curves are for `Bonsai` as indicated in the caption, for two different opening parameters. However, also data for NBODY6++ are shown for comparison: wall-clock time and accuracy as a function of average time step. In addition, we show that for a fixed particle number the time step of `Bonsai` which results to the same wall-clock time as for NBODY6++ .
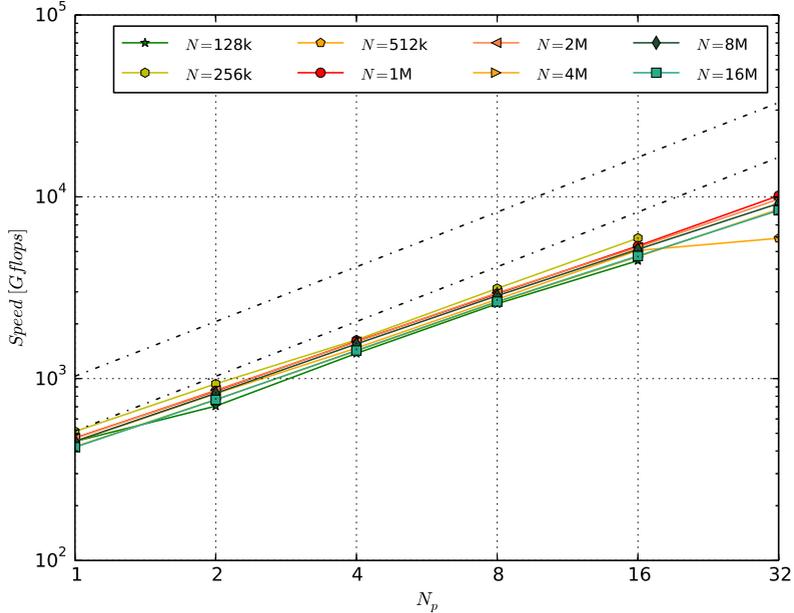
Fig. 6 Hardware performance of force computation speed (P_{force}) of `Bonsai` running on "MilkyWay" GPU cluster. Two dashed lines in the figure refer to the peak single and double precision floating point performance.

The following main conclusions can be drawn: at same wall-clock time and same particle number (and $\theta = 0.2$) `Bonsai` runs typically with time steps of a factor of $10 - 50$ larger. In other words `NBODY6++` provides a much smaller time step and a factor of 10 better accuracy (see lower panels of Figure 7). Here energy error are used as the criterion to compare the accuracy. In our case the time evolution of the energy error contain two main parts. One part comes from the machine accuracy of the potential and force calculations. This is in our cases close to the single precision machine accuracy of order $10^{-7} \sim 10^{-8}$. The other error component comes from the numerical integration process itself, which plays as the dominated role in total energy error. In the `NBODY6++` we are using the complex Hermite 4^{th} order individual block time step integration combined with the Ahmad-Cohen neighbour scheme. We have chosen the time step parameter η of the Aarseth time step criteria (for regular and irregular time step, which values set as 0.02 in initial input files) such that the energy error keeps the level of $10^{-6} \sim 10^{-7}$. How the global energy error of our integrator in `NBODY6++` behaves can be found in a comprehensive study by Makino (1991). In the case of the `Bonsai`, the code using the simple leapfrog integration scheme, which (for the reasonable computational speed to reach the 1 N-body Time Unit) have a average energy error in a level of $10^{-5} \sim 10^{-6}$. Insofar we have not discovered anything unexpected; however `Bonsai` can reach surprisingly good accuracy in total energy (like $5 \cdot 10^{-6}$) at wall-clock times comparable to `NBODY6++`. With a larger opening angle ($\theta = 0.5$) the time step and wall-clock parameters approach each other more (factor two to three, for one million bodies). In such a case still there is a quite considerable energy accuracy of order 10^{-5}. That levels of energy accuracy arguably may be sufficient even for collisional gravothermal systems, as it seems.

However, the total energy conservation is not the only criterion to judge about the use of a code and its accuracy. In `NBODY6++` close encounters and interactions of compact or hierarchical multiple systems are treated with regularization methods and zero softening, while `Bonsai` uses an artificial
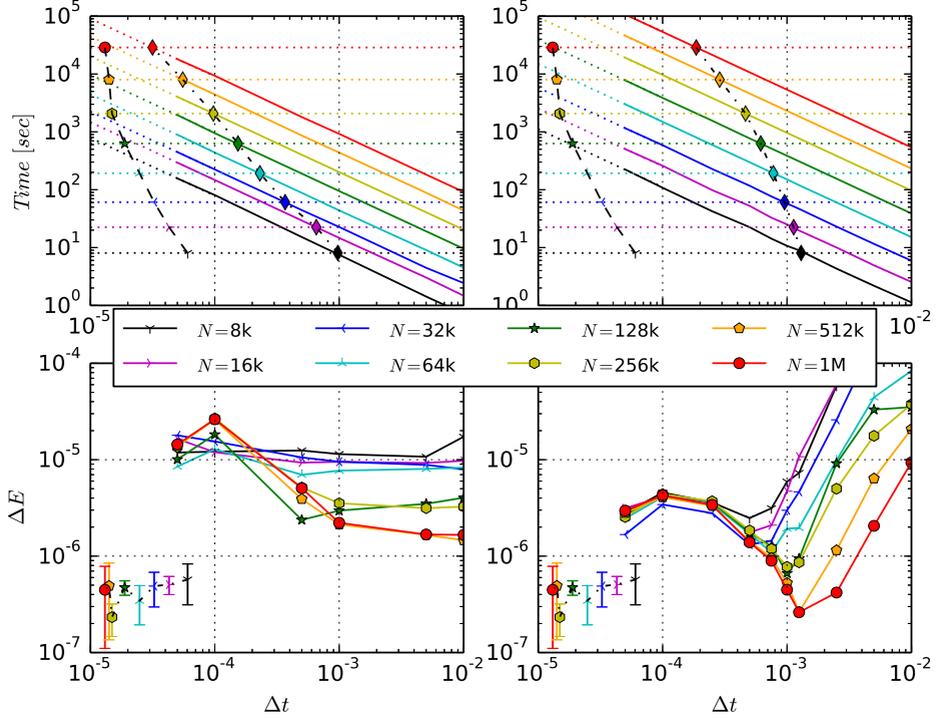
Fig. 7 Comparisons of wall-clock time and relative energy error of `NBODY6++` and `Bonsai` as a function of Δt. Opening parameters of `Bonsai` set as $\epsilon = 0.01$, $\theta = 0.5$ in the left column, and smaller value $\epsilon = 0.001$, $\theta = 0.2$ as the control group in the right column. In every panel the left dashed line corresponds `NBODY6++` benchmark data, and solid lines are `Bonsai` data. The diamond symbols indicates junctions of `Bonsai` which has the same running time as `NBODY6++` in the case of the same N.

softening of the gravitational potential at small distances. Reasonable energy conservation refers to that artificial gravitational potential including softening, which is conservative as well, but not the true few-body potential. So the additional numerical efforts necessary for `NBODY6++` goes on one hand side into the exact resolution of all kinds of close interactions below the softening length used in `Bonsai`. But also on the other side, the long-range interactions, `Bonsai` uses the standard Tree-code procedure of approximating forces from groups of particles by forces from their centers of masses and multipoles. This feature needs to be tested by simulation of core collapsing star clusters, where long range gravitational interactions determine the global evolution, which is beyond the scope of this paper.

References

Aarseth, S. J. 1999, PASP, 111, 1333 1, 3

Ahmad, A., & Cohen, L. 1973, Journal of Computational Physics, 12, 389 3

Amdahl, G. M. 1967, in Proceedings of the April 18-20, 1967, Spring Joint Computer Conference, 483–485, AFIPS '67 (Spring) (New York, NY, USA: ACM) 2

Athanassoula, E., Fady, E., Lambert, J. C., & Bosma, A. 2000, MNRAS, 314, 475 10

Barnes, J., & Hut, P. 1986, Nature, 324, 446 1, 3

Bédorf, J., Gaburov, E., Fujii, M. S., et al. 2014, ArXiv e-prints 9

Bédorf, J., Gaburov, E., & Portegies Zwart, S. 2012a, Journal of Computational Physics, 231, 2825 2, 3

Bédorf, J., Gaburov, E., & Portegies Zwart, S. 2012b, in Advances in Computational Astrophysics: Methods, Tools, and Outcome, *Astronomical Society of the Pacific Conference Series*, vol. 453, edited by R. Capuzzo-Dolcetta, M. Limongi, & A. Tornambè, 325 2, 3

Berczik, P., Nitadori, K., Zhong, S., et al. 2011, in International conference on High Performance Computing, Kyiv, Ukraine, October 8-10, 2011., p. 8-18, 8–18 2, 8

Berczik, P., Spurzem, R., & Wang, L. 2013, in Third International Conference "High Performance Computing", HPC-UA 2013, p. 52-59, 52–59 2, 8

Dorband, E. N., Hemsendorf, M., & Merritt, D. 2003, Journal of Computational Physics, 185, 484 3

Genel, S., Vogelsberger, M., Springel, V., et al. 2014, MNRAS, 445, 175 2

Gualandris, A., Portegies Zwart, S., & Tirado-Ramos, A. 2007, Parallel Computing, 33, 159 2

Hamada, T., Narumi, T., Yokota, R., et al. 2009, in Proceedings of the Conference on High Performance Computing Networking, Storage and Analysis, 62:1–62:12, SC '09 (New York, NY, USA: ACM) 9

Hamada, T., & Nitadori, K. 2010, in High Performance Computing, Networking, Storage and Analysis (SC), 2010 International Conference for, 1–9 9

Harfst, S., Gualandris, A., Merritt, D., et al. 2007, New Astron., 12, 357 2

Hut, P. 2003, in Astrophysical Supercomputing using Particle Simulations, *IAU Symposium*, vol. 208, edited by J. Makino & P. Hut, 331 1

Kawai, A., Fukushige, T., & Makino, J. 1999, in Proceedings of the 1999 ACM/IEEE Conference on Supercomputing, SC '99 (New York, NY, USA: ACM) 9

Khalisi, E., Omarov, C., Spurzem, R., Giersz, M., & Lin, D. 2003, in High Performance Computing in Science and Engineering 03, edited by E. Krause, W. J0Ł1ger, & M. Resch, 71–87 (Springer Berlin Heidelberg) 2

Makino, J. 1991, ApJ, 369, 200 12

Makino, J. 2002, New Astron., 7, 373 3

Makino, J., Fukushige, T., Koga, M., & Namura, K. 2003, PASJ, 55, 1163 2

Nitadori, K., & Makino, J. 2008, New Astron., 13, 498 7

Shin, J., Kim, J., Kim, S. S., & Park, C. 2014, Journal of Korean Astronomical Society, 47, 87 2

Springel, V. 2005, MNRAS, 364, 1105 1

Spurzem, R. 1999, Journal of Computational and Applied Mathematics, 109, 407 2

Spurzem, R., Berczik, P., Zhong, S., et al. 2012, in Advances in Computational Astrophysics: Methods, Tools, and Outcome, *Astronomical Society of the Pacific Conference Series*, vol. 453, edited by R. Capuzzo-Dolcetta, M. Limongi, & A. Tornambè, 223 2

Spurzem, R., Berentzen, I., Berczik, P., et al. 2008, in The Cambridge N-Body Lectures, *Lecture Notes in Physics, Berlin Springer Verlag*, vol. 760, edited by S. J. Aarseth, C. A. Tout, & R. A. Mardling,

Table 3 Glossary

Variable	Description
GENERAL	
N	total particle number
N_p	number of processors
S_a	theoretical maximum speed-up defined by Amdahl's law
S_i	ideal maximum speed-up equals to N_p
S_e	experiential speed-up equals to the ratio between measured time of single and multiple processor number
P	force computation speed of floating point operations per second
B	bandwidth of bytes of data transfer per second
T_{tot}	total wall-clock time
kn_x, kp_x	quantitative factors for fitting result of certain parts; $k[n|p]$ implies the factor only depends on $N|N_p$, subscript x indicates different parts
ΔE	relative energy error
Δt	time step interval of integration
NBODY6++	
$\langle N_{\mathrm{act}} \rangle$	average number of integrated active particles
$\langle N_{\mathrm{nb}} \rangle$	average neighbour number
N_{irr}	cumulative number of irregular time steps
N_{reg}	cumulative number of regular time steps
γ_{h4}	floating point operations counts per particle per interaction per step
T_{comm}	sum of communication time
T_{force}	sum of force computation time
T_{host}	time spend on the host side
T_{irr}	neighbour (irregular) force computation time
T_{mci}	MPI communication after irregular blocks
T_{mcr}	MPI communication after regular blocks
T_{mov}	time spend on data moving for parallel runs
T_{pre}	particle prediction time
T_{reg}	full (regular) force computation time
T_{syn}	interprocessor synchronization time
Bonsai	
N_{force}	cumulative number of interactions
γ_{t}	floating point operations counts per particle per interaction per step
T_{build}	time spend on tree structure building
T_{comm}	sum of communication time
T_{corr}	particle correction time
T_{dom}	time spend on update of particle domain
T_{ene}	energy check time
T_{exch}	time spend on particle exchange
T_{force}	force computation time
T_{grp}	time spend on setting active groups
T_{pre}	local tree prediction time
T_{prop}	node properties computation time
T_{sort}	sorting and data-reordering time
T_{syn}	interprocessor synchronization time
T_{tree}	sum of the whole tree construction time
ϵ	softening to diminish the effect of graininess
θ	opening angle to control the accuracy

377 2

Vogelsberger, M., Genel, S., Springel, V., et al. 2014, MNRAS, 444, 1518 2

Wang, L., Spurzem, R., Aarseth, S., et al. 2015, ArXiv e-prints 2

Warren, M. S., & Salmon, J. K. 1992, in Proceedings of the 1992 ACM/IEEE Conference on Supercomputing, 570–576, Supercomputing '92 (Los Alamitos, CA, USA: IEEE Computer Society Press) 9

Winkel, M., Speck, R., Hübner, H., et al. 2012, Computer Physics Communications, 183, 880 1