

RESEARCH

Open Access



GPU-enabled particle-particle particle-tree scheme for simulating dense stellar cluster system

Masaki Iwasawa^{1,2*}, Simon Portegies Zwart² and Junichiro Makino^{1,3}

Abstract

We describe the implementation and performance of the P^3T (Particle-Particle Particle-Tree) scheme for simulating dense stellar systems. In P^3T , the force experienced by a particle is split into short-range and long-range contributions. Short-range forces are evaluated by direct summation and integrated with the fourth order Hermite predictor-corrector method with the block timesteps. For long-range forces, we use a combination of the Barnes-Hut tree code and the leapfrog integrator. The tree part of our simulation environment is accelerated using graphical processing units (GPU), whereas the direct summation is carried out on the host CPU. Our code gives excellent performance and accuracy for star cluster simulations with a large number of particles even when the core size of the star cluster is small.

PACS Codes: 95.10.Ce; 98.10.+z

Keywords: methods: N -body simulations

1 Background

Direct N -body simulation has been the most useful tool for the study of the evolution of collisional stellar systems such as star clusters and the center of the galaxy (Aarseth 1963). The force calculations, of which the cost is $O(N^2)$, are the most compute-intensive part of direct N -body simulations. Barnes and Hut (1986) developed a scheme which reduces the calculation cost to $O(N \log N)$ by constructing the tree structure and evaluating the multipole expansions. Dehnen (2002, 2014) developed a scheme to reduce the calculation cost to $O(N)$ by combining the fast multipole method (Greengard and Rokhlin 1987) and the tree code. Recently, the graphical processing units (GPU), which is a device originally developed for rendering the graphical image, started to be used for scientific simulations. The tree code is also implemented on GPUs and it is much faster than it is on CPUs (Gaburov et al. 2010; Bédorf et al. 2012). Bédorf et al. (2014) parallelized the tree

code on GPUs and showed good scalability up to 18,600 GPUs. They also simulated the Milky Way Galaxy with N of up to 242 billion and reported that the average calculation time per iteration on 18,600 GPUs was 4.8 seconds.

The tree schemes are widely used for collisionless system simulations. However, for collisional system simulations, the use of the tree code has been very limited. One reason might be that a collisional stellar system spans a wide range in timescales. Thus it is essential that each particle has its own integration timestep. This scheme is called the individual timestep or the block timestep (McMillan 1986). However, when we use the tree code and the block timestep together, the tree structure is reconstructed at every block timestep, because the positions of integrated particle are updated. The cost of the usual complete reconstruction of the tree is $O(N \log N)$ and not negligible.

To reduce the cost of the reconstruction of the tree, McMillan and Aarseth (1993) introduced local reconstruction of tree. They demonstrated a good performance, but there seems to be no obvious way to parallelize their scheme.

* Correspondence: masaki.iwasawa@riken.jp

¹RIKEN Advanced Institute for Computational Science, Minatojima-minamimachi, Chuo-ku, Kobe, Japan

²Sterrewacht Leiden, P.O. Box 9513, Leiden, 2300 RA, The Netherlands

Recently, Oshino et al. (2011) introduced another approach to combine the tree code and the block timesteps which they called the P³T scheme. This scheme is based on the idea of Hamiltonian splitting (Kinoshita et al. 1991; Wisdom and Holman 1991; Duncan et al. 1998; Chambers 1999; Brunini and Viturro 2003; Fujii et al. 2007; Moore and Quillen 2011). In the P³T scheme, the Hamiltonian of the system is split into short-range and long-range parts and they are integrated with different integrators. The long-range part is evaluated with the tree code and is integrated using the leapfrog scheme with a shared timestep. The short range part is evaluated with direct summation and integrated using the fourth-order Hermite scheme (Makino and Aarseth 1992) with the block timesteps. They investigated the accuracy and the performance of the P³T scheme for planetary formation simulations and showed that the P³T scheme achieves high performance.

In this paper, we present the implementation of the P³T scheme on GPUs and report its accuracy and performance for star cluster simulations. We found that the P³T scheme demonstrates a very good performance for star cluster simulations, even when the core of the cluster becomes small.

The structure of this paper is as follows. In Section 2, we briefly describe the P³T scheme. In Section 3, we report the accuracy and performance of the P³T scheme. We summarize these results in Section 4.

2 Methods

2.1 Formulation

In this section, we describe the P³T scheme. The Hamiltonian H of a gravitational N -body system is given by

$$H = \sum_i^N \frac{|\mathbf{p}_i|^2}{2m_i} - \sum_i^N \sum_{i<j}^N \frac{Gm_i m_j}{s_{ij}}, \quad (1)$$

$$s_{ij} = \sqrt{|\mathbf{q}_{ij}|^2 + \epsilon^2}, \quad (2)$$

$$\mathbf{q}_{ij} = \mathbf{q}_i - \mathbf{q}_j, \quad (3)$$

where \mathbf{p}_i , m_i and \mathbf{q}_i are momentum, mass and position of the particle i , respectively. To avoid the singularity of the $1/r$ potential, we use the Plummer softening ϵ (Aarseth 1963). With the P³T scheme, H is split into H_{hard} and H_{soft} as follows (Oshino et al. 2011):

$$H = H_{\text{hard}} + H_{\text{soft}}, \quad (4)$$

$$H_{\text{hard}} = \sum_i^N \frac{|\mathbf{p}_i|^2}{2m_i} - \sum_i^N \sum_{i<j}^N \frac{m_i m_j}{s_{ij}} [1 - W(s_{ij})], \quad (5)$$

$$H_{\text{soft}} = - \sum_i^N \sum_{i<j}^N \frac{m_i m_j}{s_{ij}} W(s_{ij}). \quad (6)$$

Here $W(s_{ij})$ is a smooth transition function. A suitable form of $W(s_{ij})$ should be zero when a distance between two particles is smaller than the inner cutoff radius r_{in} and should be unity if the distance is larger than the outer cutoff radius r_{cut} . This splitting is introduced by Chambers (1999) to avoid undesirable energy error from close encounters between particles. Similar splitting has been used with P³M (Particle-Particle Particle-Mesh) scheme, in which the long-range part of the interaction is evaluated by using FFT (Hockney and Eastwood 1981).

Forces derived from H_{hard} and H_{soft} are given by

$$\mathbf{F}_{\text{hard},i} = - \frac{\partial H_{\text{hard}}}{\partial \mathbf{q}_i} = - \sum_{j \neq i}^N \frac{m_i m_j}{s_{ij}^3} (1 - K(s_{ij})) \mathbf{q}_{ij}, \quad (7)$$

$$\mathbf{F}_{\text{soft},i} = - \frac{\partial H_{\text{soft}}}{\partial \mathbf{q}_i} = - \sum_{j \neq i}^N \frac{m_i m_j}{s_{ij}^3} K(s_{ij}) \mathbf{q}_{ij}, \quad (8)$$

$$K(s_{ij}) = W(s_{ij}) - s_{ij} \frac{dW(s_{ij})}{ds_{ij}}. \quad (9)$$

We call $K(s_{ij})$ the cutoff function.

The tree algorithm is used for the evaluation of $\mathbf{F}_{\text{soft},i}$ to reduce the calculation cost.

The formal solution of the equation of motion for the phase space coordinate $\mathbf{w} = (\mathbf{q}, \mathbf{p})$ at time $t + \delta t$ for the given Hamiltonian H is

$$\mathbf{w}(t + \delta t) = e^{\delta t \{H\}} \mathbf{w}(t) = e^{\delta t \{H_{\text{soft}} + H_{\text{hard}}\}} \mathbf{w}(t). \quad (10)$$

Here the braces $\{, \}$ stand for the Poisson bracket. In the P³T scheme, we use the second order approximation;

$$\mathbf{w}(t + \delta t) = e^{\delta t/2 \{H_{\text{soft}}\}} e^{\delta t \{H_{\text{hard}}\}} e^{\delta t/2 \{H_{\text{soft}}\}} \mathbf{w}(t) + O(\delta t^3). \quad (11)$$

Here, the formal solution for the H_{soft} term is the simple velocity kick, since H_{soft} contains the potential only. We numerically integrate the H_{hard} term, since it cannot be solved analytically. We use the fourth-order Hermite scheme with the block timestep (Makino and Aarseth 1992). The fourth-order integrator requires $K(s_{ij})$ to be three-times differentiable with respect to position. We use the following formula:

$$K(x) = \begin{cases} 0 & (x < 0), \\ -20x^7 + 70x^6 - 84x^5 + 35x^4 & (0 \leq x < 1), \\ 1 & (1 \leq x), \end{cases} \quad (12)$$

$$x = \frac{y - \gamma}{1 - \gamma}, \quad (13)$$

$$y = \frac{s_{ij}}{r_{\text{cut}}}, \quad (14)$$

$$\gamma = \frac{r_{\text{in}}}{r_{\text{cut}}}. \quad (15)$$

This $K(x)$ is the lowest-order polynomial which satisfies the requirement that derivatives up to the third order is zero for $x = 0$ and 1 (i.e. the highest-order term of the lowest-order polynomial is the seventh, because there are eight boundary conditions at $x = 0$ and $x = 1$).

In Figure 1, we plot $K(y)$ (top panel) and forces (bottom panel) with $\gamma = 0.1$. According to Oshino et al. (2011), Chambers (1999), $K(y)$ with $\gamma = 0.1$, is smooth enough to be integrated. Thus, for all calculations, we use $\gamma = 0.1$. The functional form of $W(y; \gamma)$ is given by

$$W(y; \gamma) = \begin{cases} \frac{7(\gamma^6 - 9\gamma^5 + 45\gamma^4 - 60\gamma^3 \log \gamma - 45\gamma^2 + 9\gamma - 1)}{3(\gamma - 1)^7} y & (y < \gamma), \\ G(y; \gamma) + (1 - G(1; \gamma))y & (\gamma \leq y < 1), \\ 1 & (1 \leq y), \end{cases} \quad (16)$$

$$G(y; \gamma) = (-10/3y^7 + 14(\gamma + 1)y^6 - 21(\gamma^2 + 3\gamma + 1)y^5 + (35(\gamma^3 + 9\gamma^2 + 9\gamma + 1)/3)y^4 - 70(\gamma^3 + 3\gamma^2 + \gamma)y^3 + 210(\gamma^3 + \gamma^2)y^2 - 140\gamma^3 y \log(y) + (\gamma^7 - 7\gamma^6 + 21\gamma^5 - 35\gamma^4))/(1 - \gamma)^7. \quad (17)$$

With the P³T scheme, the time integration proceeds as follows:

- (1) At time t , by using the tree code, calculate the acceleration due to H_{soft} , $\mathbf{a}_{\text{soft},i}$, and construct a list of all particles which come within r_{cut} from particle i for Δt_{soft} . Here, Δt_{soft} is the timestep for the soft Hamiltonian.
- (2) Update the velocities of all particles with $\mathbf{v}_{\text{new},i} = \mathbf{v}_{\text{old},i} + (1/2)\Delta t_{\text{soft}}\mathbf{a}_{\text{soft},i}$.
- (3) Integrate all particles to time $t + \Delta t_{\text{soft}}$ under H_{hard} , using the neighbour list and the fourth order Hermite integrator with the block timesteps.
- (4) Calculate the acceleration due to H_{soft} at new time $t + \Delta t_{\text{soft}}$ and update the velocity
- (5) Go back to step 2.

For the timestep criterion for the block timestep, we use the following form (Oshino et al. 2011).

$$\Delta t_i = \min\left(\eta \sqrt{\frac{\sqrt{|\mathbf{a}_i^{(0)}|^2 + a_0^2} |\mathbf{a}_i^{(2)}| + |\mathbf{a}_i^{(1)}|^2}{|\mathbf{a}_i^{(0)}| |\mathbf{a}_i^{(3)}| + |\mathbf{a}_i^{(2)}|^2}}, \Delta t_{\text{max}}\right), \quad (18)$$

$$a_0 = \alpha \frac{m}{r_{\text{cut}}^2}. \quad (19)$$




Figure 1 The cutoff function $K(y)$ (top) and the forces (bottom) as functions of $y = s_{ij}/r_{\text{cut}}$.

Here η is the accuracy parameter of the timestep and its typical value is 0.1. Δt_{max} is the maximum timestep which should be smaller than Δt_{soft} , $\mathbf{a}_i^{(n)}$ is the n th time derivative of the acceleration of particle i , a_0 is a constant introduced to prevent Δt_i from becoming too small when the distance to the nearest neighbor is close to r_{cut} and α is a parameter to control a_0 . In this case, the acceleration from H_{hard} becomes very small and there is no need to use very small Δt_i . According to Oshino et al. (2011), when we choose $\alpha \leq 1$, α hardly affects the energy error. Thus we set $\alpha = 0.1$ for all simulations.

In our Hermite implementation, $\mathbf{a}_i^{(2)}$ and $\mathbf{a}_i^{(3)}$ are derived using interpolation of $\mathbf{a}_i^{(0)}$ and $\mathbf{a}_i^{(1)}$, and as a consequence we cannot use equation (18) for the first step. We use:

$$\Delta t_i = \min\left(\eta_s \sqrt{\frac{|\mathbf{a}_i^{(0)}|^2 + a_0^2}{|\mathbf{a}_i^{(1)}|^2}}, \Delta t_{\text{max}}\right). \quad (20)$$

This criterion does not contain the 2nd and 3rd time derivatives of the acceleration. To prevent the timestep derived by equation (20) from becoming too large, we set η_s to be the one-tenth of η for all simulation in this paper.

We summarize all accuracy parameters in Table 1.

Table 1 Symbols and definitions for the accuracy parameters of the P³T scheme

α	timestep softening. For all runs $\alpha = 0.1$
γ	ratio of inner and outer cutoff radius ($r_{\text{in}}/r_{\text{cut}}$). For all runs $\gamma = 0.1$
Δr_{buff}	width of the buffer shell. $\Delta r_{\text{buff}} = 3\sigma \Delta t_{\text{soft}}$, as a standard value
Δt_{soft}	timestep of the soft part. $\Delta t_{\text{soft}} = (1/256)(N/16K)^{-1/3}$, as a standard value
Δt_{max}	maximum timestep of the hard part. $\Delta t_{\text{max}} = \Delta t_{\text{soft}}/4$, as a standard value
ϵ	plummer softening length. $\epsilon = (4/N)$, as a standard value
η	accuracy parameter for timestep criterion. $\eta = 0.1$, as a standard value
r_{cut}	outer cutoff radius of smooth transition functions W and K . $r_{\text{cut}} = 4\Delta t_{\text{soft}}$, as a standard value
r_{in}	inner cutoff radius of smooth transition functions W and K ($r_{\text{in}} = \gamma r_{\text{cut}}$)
θ	opening criterion for tree. $\theta = 0.4$, as a standard value

2.2 Implementation on GPUs

Even with the Barnes-Hut tree algorithm, obtaining $\mathbf{F}_{\text{soft},i}$ is still costly and dominates the total calculation time (Os-hino et al. 2011). To accelerate this part, we use GPUs, by modifying the *sequoia* library (Bédorf, Gaburov and Portegies Zwart, submitted to ComAC), on which the high-performance tree code for parallel GPUs *Bonsai* (Bédorf et al. 2012) is based. Our library calculates the long range forces on all particles, $\mathbf{F}_{\text{soft},i}$ by the Barnes-Hut tree algorithm (up to the quadrupole moment). On the other hand, we calculate $\mathbf{F}_{\text{hard},i}$ on the host computer. The library also returns, for each particle, the list of particles within the distance h from it. We use this list of neighbors to calculate $\mathbf{F}_{\text{hard},i}$. The value of h should be sufficiently larger than r_{cut} to guarantee that the particles which are not on the list of the neighbors of particle i do not enter the sphere of the radius r_{cut} around particle i during the time interval Δt_{soft} .

We call the sphere with a radius of r_{cut} the neighbor sphere and the shell between the sphere with a radius of h and the neighbor sphere the buffer shell. The particles of which the nearest neighbor is outside the sphere with radius h are considered isolated and the particles on the list of neighbors are considered neighbor particles. We denote the width of the buffer shell as Δr_{buff} (i.e. $h = r_{\text{cut}} + \Delta r_{\text{buff}}$).

The compute procedures of our implementation of the P³T scheme on GPU is as follows:

- (1) Evaluate long range forces on all particles $\mathbf{F}_{\text{soft},i}$ using GPU.
- (2) Particles are divided into two groups; isolated and non-isolated, by using the neighbour list made on GPU.
- (3) For non-isolated particles, $\mathbf{F}_{\text{hard},i}$ are calculated on the host computer.
- (4) All particles receive a velocity kick through $\mathbf{F}_{\text{soft},i}$ for $\Delta t_{\text{soft}}/2$.
- (5) Isolated particles are drifted by $\mathbf{r}_i \leftarrow \mathbf{r}_i + \Delta t_{\text{soft}} \mathbf{v}_i$.
- (6) Non-isolated particles are integrated with the fourth-order Hermite scheme for Δt_{soft} .
- (7) Evaluate $\mathbf{F}_{\text{soft},i}$ and make the neighbour list in the same way as in step 1-2.

- (8) All particles obtain the velocity kick again for $\Delta t_{\text{soft}}/2$.
- (9) go back to step 3.

3 Results

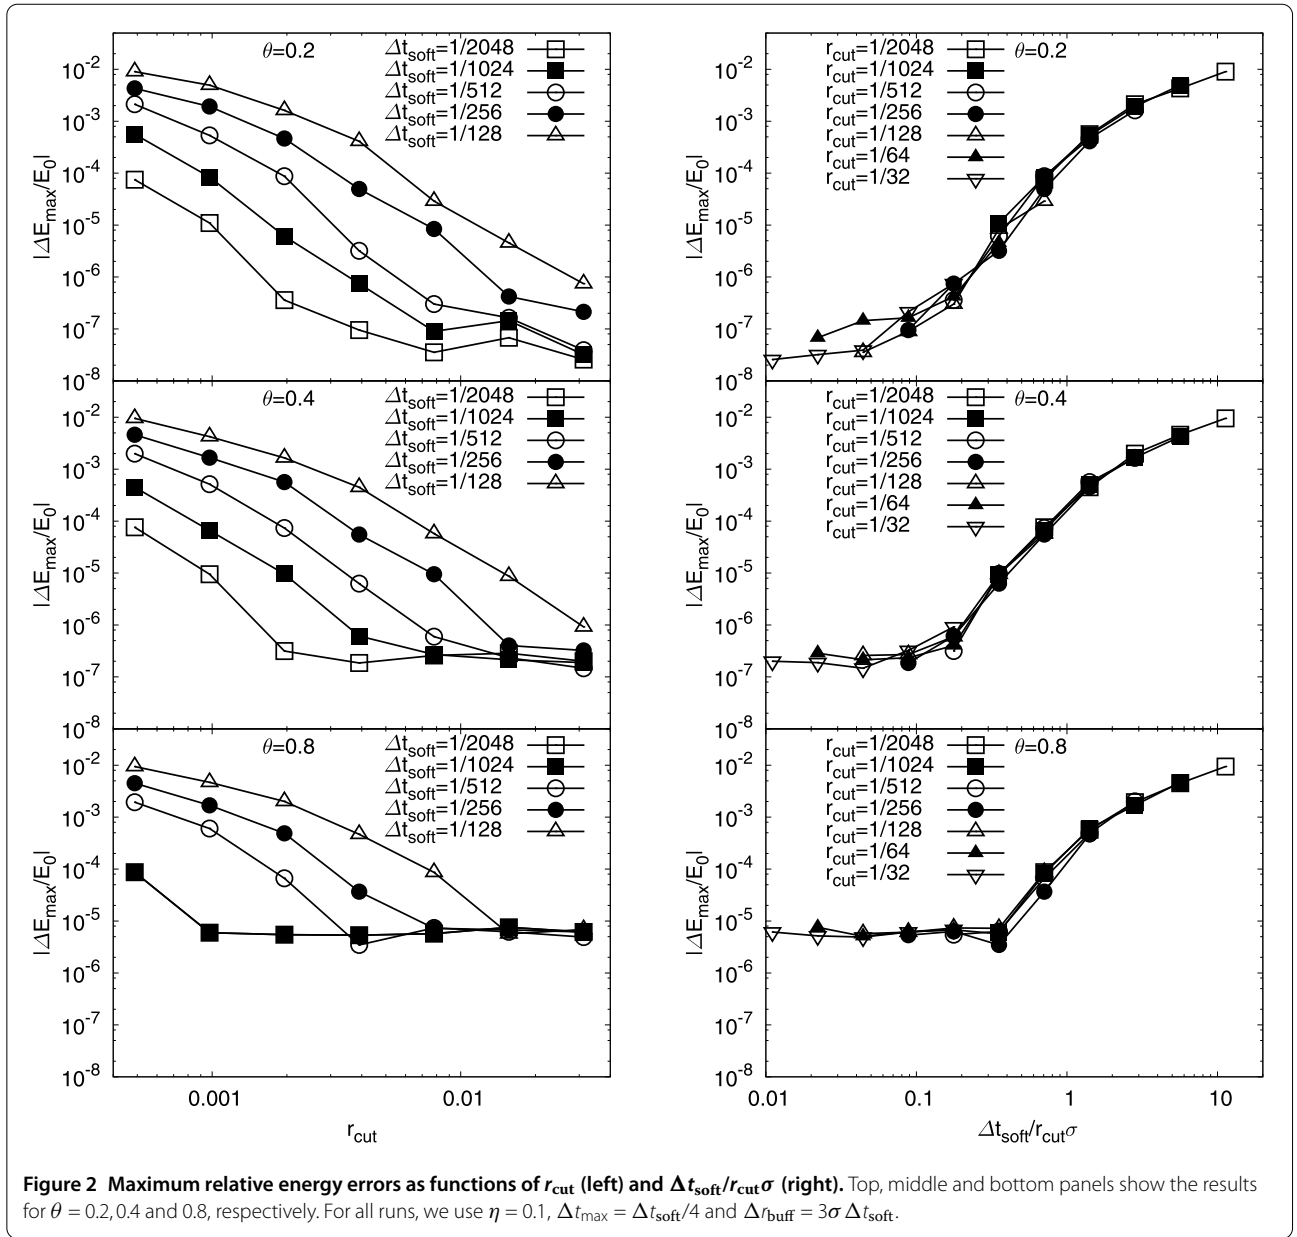
3.1 Accuracy and performance

We performed a number of test calculations using the P³T scheme on GPUs, to study its accuracy and performance. In this section, we describe the result of these tests. For most of them we adopted a Plummer model (Plummer 1911) with 128K (hereafter $K = 2^{10}$) equal-mass particles as the initial condition. We use the so-called N -body unit or Heggie unit, in which total mass $M = 1$, the gravitational constant $G = 1$ and total energy $E = -1/4$ (Heggie and Mathieu 1986). To avoid the singularity of the gravitational potential, we use the Plummer softening and set $\epsilon = 4/N$. Since this value is a typical separation of a hard binary in the N -body unit, we can follow the evolution of the system up to the moment of the core collapse.

Note, in this paper, we use the energy errors as an indicator of the accuracy of the scheme. However, energy conservation dose not guarantee accuracy of simulations (though it is necessary). Thus we will perform realistic simulations in Section 3.2 and check the statistical character of stellar systems by comparing the results with the Hermite scheme, which is widely used in collisional stellar system simulations. As we will see later, for simulations of the core collapse of the star cluster, when the relative energy error is $\lesssim 10^{-3}$ at the moment of the core collapse, the behavior of the core collapse with the P³T scheme agrees with that with the Hermite scheme very well.

3.1.1 Accuracy

With the P³T scheme, we have six accuracy parameters. First, we discuss how each parameter controls the accuracy of the P³T scheme. Finally, we describe the accumulation of the energy error in a long-term integration. To measure energy errors accurately, we calculate potential energies by the direct summation instead of the tree code for all runs in this paper.

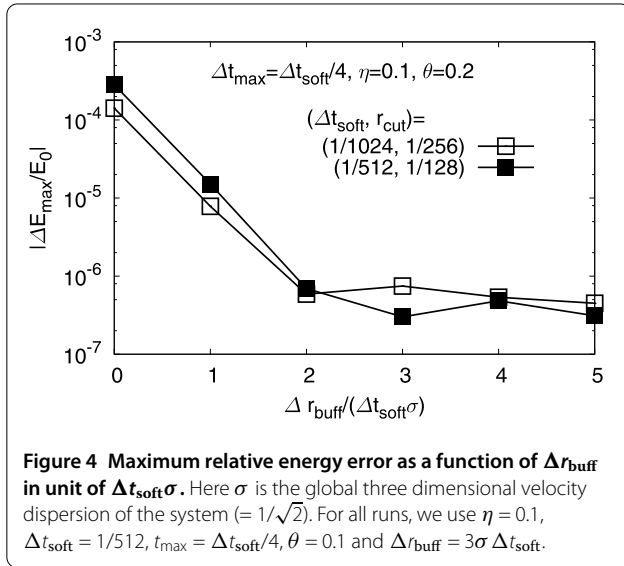
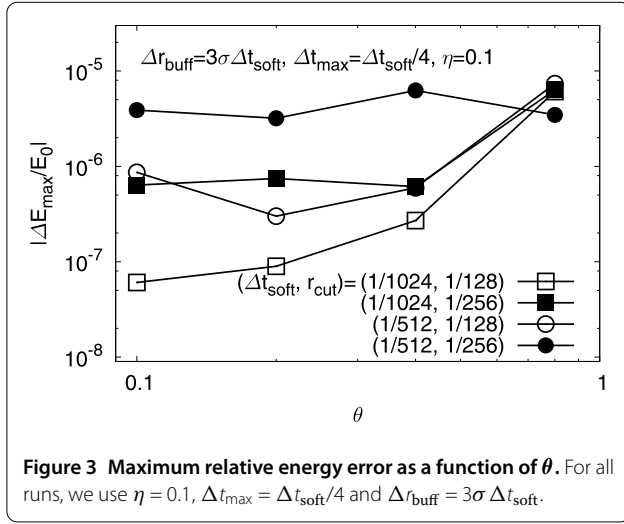


Effect of r_{cut} , Δt_{soft} and θ In Figure 2, we present the maximum relative energy error $|\Delta E_{\text{max}}/E_0|$ over 10 N -body time units as a function of r_{cut} and Δt_{soft} for several different values of the opening criterion of the tree, θ . Here ΔE_{max} is the maximum energy error and E_0 is the initial energy. We choose $\eta = 0.1$, $\Delta t_{\text{max}} = \Delta t_{\text{soft}}/4$ and $\Delta r_{\text{buff}} = 3\sigma \Delta t_{\text{soft}}$, where σ is the global three dimensional velocity dispersion and we adopt $\sigma = 1/\sqrt{2}$.

We can see that the error is smaller for smaller θ , smaller Δt_{soft} , or larger r_{cut} . Roughly speaking, the error depends on two terms, $\Delta t_{\text{soft}}/r_{\text{cut}}\sigma$ and θ . If $\Delta t_{\text{soft}}/r_{\text{cut}}\sigma$ is large, it determines the error. In this regime, the error is dominated by the truncation error of the leapfrog integrator. If it is

small enough, θ determines the error, in other words, the tree force error dominates the total error. Even for a very small value of θ like 0.2, the tree force error dominates if $\Delta t_{\text{soft}}/r_{\text{cut}}\sigma \lesssim 0.05$.


In Figure 3, we plot the maximum energy error as a function of θ . We use the same η , Δt_{max} and Δr_{buff} as in Figure 2. For the runs with $r_{\text{cut}} = 1/256$ and $\Delta t_{\text{soft}} = 1/512$, the energy error does not drop below 10^{-6} because the error of the leapfrog integrator is larger than the tree force error. In an chaotic system like the model used in our simulations such energy error is sufficient to warrant a scientifically reliable result (Portegies Zwart and Boekholt 2014). On the other hand, for the run with $r_{\text{cut}} = 1/128$ and



$\Delta t_{\text{soft}} = 1/1,024$, integration error is smaller than the tree force error.

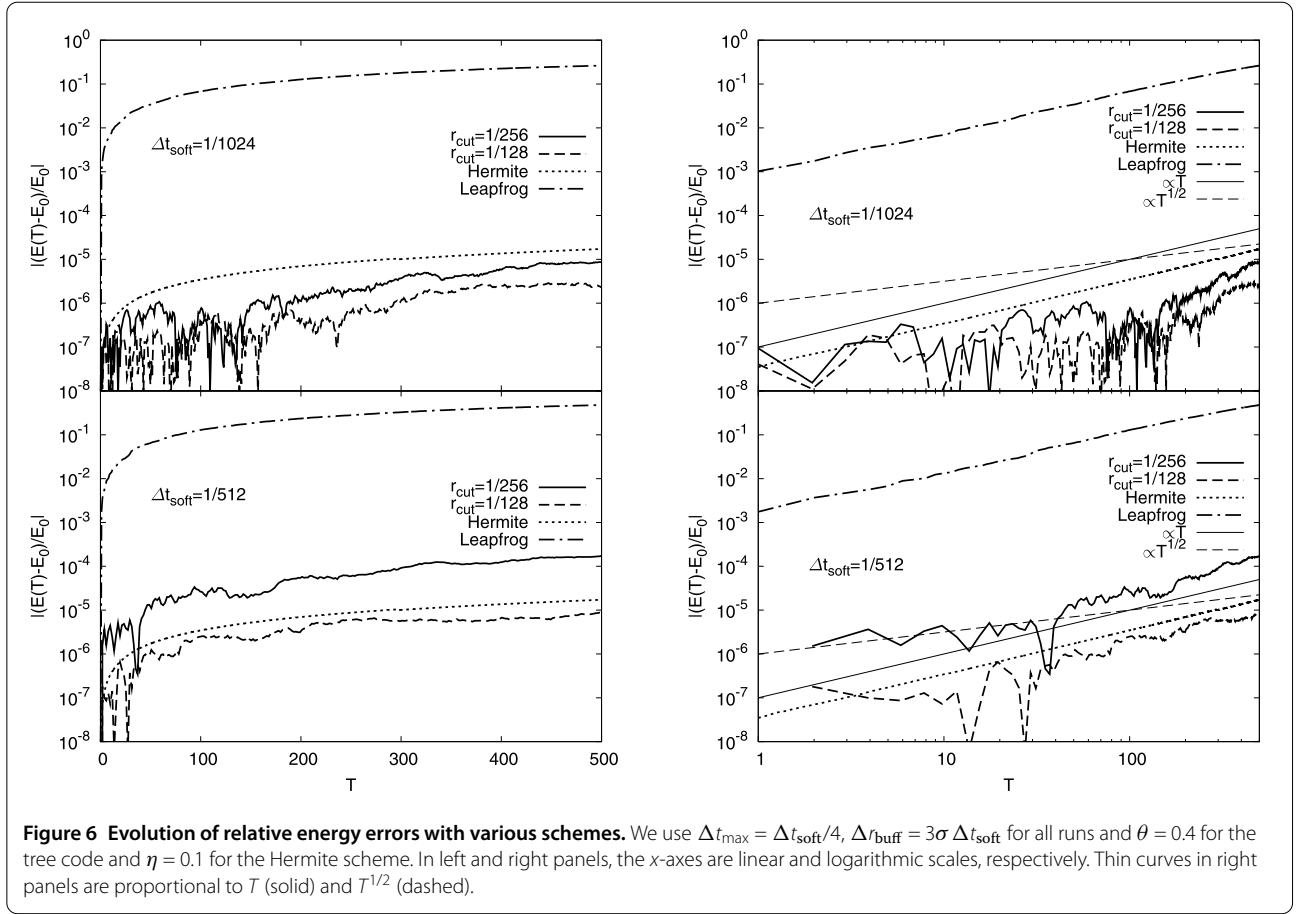
Effect of Δr_{buff} In Figure 4, we show the maximum relative energy error as a function of Δr_{buff} for the runs with $\Delta t_{\max} = \Delta t_{\text{soft}}/4$, $\eta = 0.1$, $\theta = 0.2$, for $(\Delta t_{\text{soft}}, r_{\text{cut}}) = (1/512, 1/128)$ and $(1/1,024, 1/256)$. The energy error is almost constant for $\Delta r_{\text{buff}} \gtrsim 2\Delta t_{\text{soft}}\sigma$, which indicates that the energy error for $\Delta r_{\text{buff}} < 2\Delta t_{\text{soft}}\sigma$ is caused by particles that are initially outside the buffer shell (with radius $r_{\text{cut}} + \Delta r_{\text{buff}}$) and plunge into the neighbour sphere (with radius r_{cut}) during the timestep Δt_{soft} . We can prevent this by adopting $\Delta r_{\text{buff}} \gtrsim 2\Delta t_{\text{soft}}\sigma$.

Effect of Δt_{\max} and η The maximum relative energy errors over 10 N -body time units are shown in the top panel of Figure 5 as a function of η and the number of steps for



the Hermite part (per particle per unit time, N_{step}) are presented in the bottom panel. The energy errors go down as η decrease until $\eta \sim 0.2$. For $\eta \lesssim 0.2$, the errors hardly depend on Δt_{\max} .

Long term integration In Figure 6, we show the time evolution of the relative energy error until $T = 500$. We compare the accuracy of our P³T scheme with two other schemes, the direct fourth-order Hermite scheme and the leapfrog scheme with the Barnes-Hut tree code. The calculations with the direct Hermite scheme are performed by using the *Sapporo* library on GPU (Gaburov and Harfst 2009), and the calculations with the leapfrog scheme are performed by using the *Bonsai* library on GPU (Bédorf et al. 2012). The energy error of the P³T scheme behaves like a random walk whereas that of the leapfrog and the Hermite schemes grow monotonically. In the right-hand panels of Figure 6, we show the same evolution of the error as in the left panels, but time is plotted with a logarithmic scale. This allows us to realize that the error growth of Hermite and tree schemes are linear, whereas the error in the P³T scheme grows as $\propto T^{1/2}$. This latter proportionality is caused by the short-term error of the P³T scheme,



which is dominated by the randomly changing tree-force error. For long-term integration the P³T scheme conserves energy better than the Hermite or leapfrog schemes.

3.1.2 Calculation cost

In this section, we discuss the calculation cost of the P³T scheme and its dependence on the number of particles N , required accuracy, and other parameters.

First, we construct a simple theoretical model of the dependence of the calculation cost on parameters of the integration scheme such as N , Δt_{soft} , θ and r_{cut} . Finally, we derive the optimal set of parameters from the model and compare this model with the result of the numerical tests. We found that the calculation cost per unit time is proportional to $N^{4/3}$.

Theoretical model The calculation cost for the force evaluations in P³T is split into the tree part and the Hermite part. For the tree part, the calculation cost of evaluating forces for all particles per tree step is proportional to $O(\theta^{-3}N \log N)$. Since we use constant timestep for the tree part, the calculation costs of the integration of particles per unit time for the tree part is proportional to $O(\theta^{-3}N \log N / \Delta t_{\text{soft}})$.

For the Hermite part, since each particle has its own neighbour particles and timesteps, the number of interactions for all particles per unit timestep is given by

$$N_{\text{int,hard}} = \sum_i^N N_{\text{ngh},i} N_{\text{step},i} \quad (21)$$

$$\sim \sum_i^N 4\pi/3 (r_{\text{cut}} + \Delta r_{\text{buff}})^3 n_i \langle \Delta t_i \rangle^{-1} \quad (22)$$

$$\propto N^2 (r_{\text{cut}} + \Delta r_{\text{buff}})^3 \langle \langle \Delta t \rangle \rangle^{-1}. \quad (23)$$

Here $N_{\text{ngh},i}$ is the number of the neighbour particles around particle i , $N_{\text{step},i}$ is the number of timesteps required to integrate particle i for one unit time, n_i is the local density around particle i , $\langle \Delta t_i \rangle$ is the average timestep of particle i over one unit time and $\langle \langle \Delta t \rangle \rangle$ is the average of $\langle \Delta t_i \rangle$ over all particles. Here we assume n_i is constant within the radius of $r_{\text{cut}} + \Delta r_{\text{buff}}$ around particle i .

Next we express the $\langle \langle \Delta t \rangle \rangle$ as a function of N and r_{cut} . To simplify the discussion, we define the timestep of the particle through the relative position and velocity from its nearest neighbour particle; $\langle \langle \Delta t \rangle \rangle \propto r_{\text{NN}}/v_{\text{NN}}$, where r_{NN}

and v_{NN} are the relative position and the velocity of the nearest neighbour particle. We can replace v_{NN} to the velocity dispersion σ . Thus average timestep is given by

$$\langle\langle\Delta t\rangle\rangle \propto r_{\text{NN}}/v_{\text{NN}} \sim r_{\text{NN}}/\sigma. \quad (24)$$

To further simplify the derivation we assume that the number density of particles in the system is uniform. If r_{cut} is larger than the mean inter-particle distance $\langle r \rangle$ (i.e. if most particles have neighbour particles), the average timestep is roughly given by

$$\langle\langle\Delta t\rangle\rangle \sim \min\left(\eta \frac{R}{\sigma} N^{-1/3}, \Delta t_{\text{max}}\right), \quad (25)$$

where R is the typical size of the system. In this case, the average timestep depend only on N (does not depend on r_{cut}).

If r_{cut} is small compared to $\langle r \rangle$, most particles are isolated and most of the non-isolated particles have only one neighbour particle. In this case, $\langle\langle\Delta t\rangle\rangle$ is given by

$$\langle\langle\Delta t\rangle\rangle \sim \min\left(\eta \frac{r_{\text{cut}}}{\sigma}, \Delta t_{\text{max}}\right). \quad (26)$$

In Figure 7 we show the number of steps per particle per unit time N_{step} for a plummer sphere as a function of N (top panel) and as a function of r_{cut} (bottom panel). In the top panel, we can see that N_{step} is roughly proportional to $N^{1/3}$ for large N (i.e. $\langle r \rangle$ is small). On the other hand when N is small N_{step} is almost constant because $\langle r \rangle$ is large (see equation (26)).

The bottom panel of Figure 7 shows that all curves eventually approach to constant values for both of large and small r_{cut} . For large r_{cut} , the timesteps of the non-isolated particles are determined by N , not by r_{cut} (see equation (25)), whereas for small values of r_{cut} the non-isolated particles have a timesteps Δt_{max} . This is because most neighbouring particles are in the buffer shell and not in the neighbour sphere. For runs with $\Delta t_{\text{soft}} = 1/2,048, 1/1,024$ and $1/512$, we can see bumps of N_{step} at $r_{\text{cut}} \sim 1/512$ due to the dependence on r_{cut} shown in equation (26).

Using above discussions, the number of interactions for all particles per unit time of the Hermite part $N_{\text{int,hard}}$ and the tree part $N_{\text{int,soft}}$ are given by

$$N_{\text{int,hard}} \propto \begin{cases} N^{7/3}(r_{\text{cut}} + \Delta r_{\text{buff}})^3 & (\text{for } r_{\text{cut}} \gg \langle r \rangle), \\ N^2(r_{\text{cut}} + \Delta r_{\text{buff}})^3 & (\text{for } r_{\text{cut}} \ll \langle r \rangle), \end{cases} \quad (27)$$

$$N_{\text{int,soft}} \propto \theta^{-3} N \log N / \Delta t_{\text{soft}}, \quad (28)$$

Optimal set of accuracy parameters In this section, we derive the optimal values of r_{cut} and Δt_{soft} from the point of view of the balance of the calculation costs between the

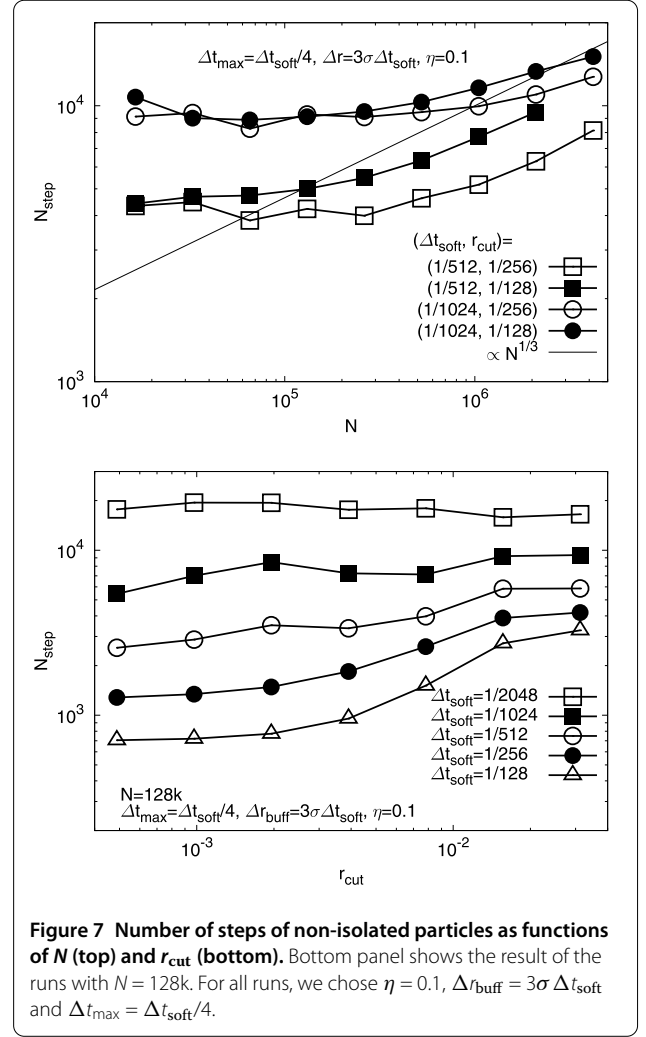


Figure 7 Number of steps of non-isolated particles as functions of N (top) and r_{cut} (bottom). Bottom panel shows the result of the runs with $N = 128\text{k}$. For all runs, we chose $\eta = 0.1$, $\Delta r_{\text{buff}} = 3\sigma \Delta t_{\text{soft}}$ and $\Delta t_{\text{max}} = \Delta t_{\text{soft}}/4$.

tree and the Hermite parts, in other words we express r_{cut} and Δt_{soft} as functions of N such that $N_{\text{int,hard}}/N_{\text{int,soft}}$ is independent of N . Following the discussion in Section 3.1.1 and because the energy errors can be controlled through $\Delta t_{\text{soft}}/r_{\text{cut}}$ and $\Delta t_{\text{soft}}/\Delta r_{\text{buff}}$, r_{cut} and Δr_{buff} should be proportional to Δt_{soft} .

The requirements are met for $N_{\text{int,hard}} \propto N^{7/3}(r_{\text{cut}} + \Delta r_{\text{buff}})^3$ (or $N^2(r_{\text{cut}} + \Delta r_{\text{buff}})^3$), $\Delta t_{\text{soft}} \propto N^{-1/3}$ and $r_{\text{cut}} \propto N^{-1/4}$ and both $N_{\text{int,hard}}$ and $N_{\text{int,soft}}$ are proportional to $N^{4/3}$ (or $N^{5/4}$). Here we have neglected the $\log N$ dependence in the tree part.

This is illustrated in Figure 8, where we plot $N_{\text{int,hard}}$ for a plummer sphere as a function of N . Following above discussions, we use the N -dependent tree timestep: $\Delta t_{\text{soft}} = (1/256)(N/16\text{K})^{-1/3}$ and $N_{\text{int,hard}}$ as well as $N_{\text{int,soft}}$ are proportional to $N^{4/3}$.

In Figures 9 and 10, we plot the wall-clock time of execution T_{cal} and the maximum relative energy errors $|\Delta E_{\text{max}}/E_0|$ for the time integration for 10 N -body units against N . Figure 9 shows the results of the runs with

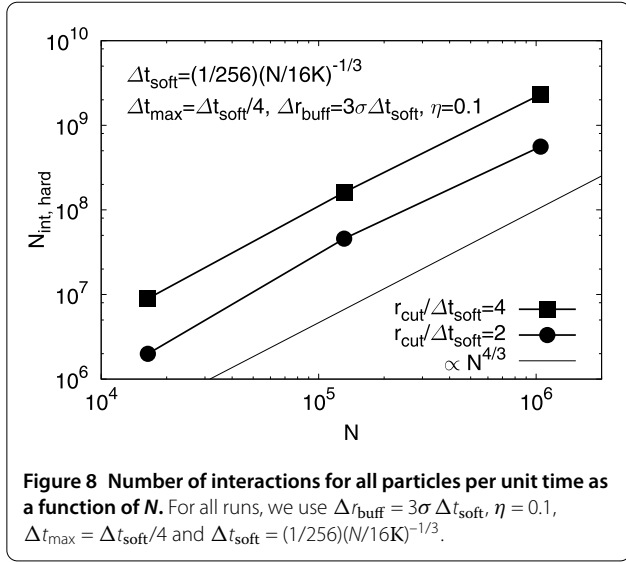


Figure 8 Number of interactions for all particles per unit time as a function of N . For all runs, we use $\Delta r_{\text{buff}} = 3\sigma \Delta t_{\text{soft}}$, $\eta = 0.1$, $\Delta t_{\text{max}} = \Delta t_{\text{soft}}/4$ and $\Delta t_{\text{soft}} = (1/256)(N/16K)^{-1/3}$.

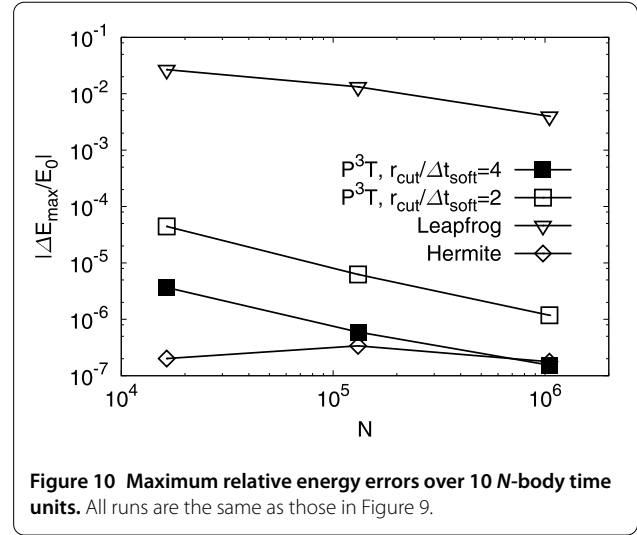


Figure 10 Maximum relative energy errors over 10 N -body time units. All runs are the same as those in Figure 9.




Figure 9 Wall-clock time of execution as a function of N . Top (bottom) panel shows the results of the runs with $r_{\text{cut}}/\Delta t_{\text{soft}} = 2(4)$. We use $\theta = 0.4$, $\eta = 0.1$, $\Delta r_{\text{buff}} = 3\sigma \Delta t_{\text{soft}}$ and $\Delta t_{\text{soft}} = (1/256)(N/16K)^{-1/3}$.

$r_{\text{cut}}/\Delta t_{\text{soft}} = 2$ (top panel) and 4 (bottom panel). All runs in these figures are carried out on NVIDIA GeForce GTX680^a GPU and Intel Core i7-3770K CPU. For each run, we use one CPU core and one GPU card.

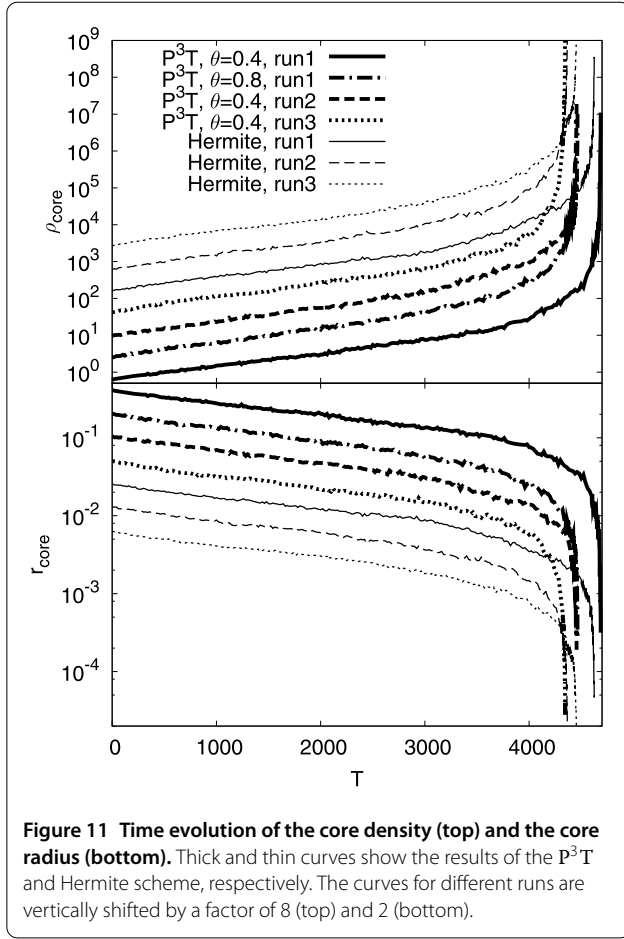
We also perform the simulations using the direct Hermite integrator with the same η and the standard tree code with the same θ and Δt_{soft} . These calculations are performed with the Sapporo GPU library (Gaburov and Harfst 2009) and a standard tree code with the same θ and Δt_{soft} using the Bonsai GPU library (Bédorf et al. 2012). The calculation time for our P³T implementation is also proportional to $N^{4/3}$, as we presented above, while for the Hermite integrator it is proportional to $N^{7/3}$. The P³T scheme is faster than the direct Hermite integrator for $N > 16K$ and when $N = 1M$ ($M = 2^{20}$), the P³T scheme is about 50 times faster than the direct Hermite scheme. The pure tree code is slightly faster than the P³T scheme, but the integration errors are worse by several orders of magnitude (see Figures 6 and 10).

3.2 Examples of practical applications

In Sections 3.1.1 and 3.1.2, we presented a detailed discussion on the accuracy and performance of our P³T scheme. However, we performed simple simulations, where the stellar systems are in the dynamical equilibrium. In this section, we study the performance of our P³T scheme when applied to more realistic, or more difficult, simulations by comparing the results of the Hermite scheme. In Section 3.2.1, we discuss the case of the simulation of star clusters up to core collapse. In Section 3.2.2, we discuss the case of a galaxy model with massive central black hole binary.

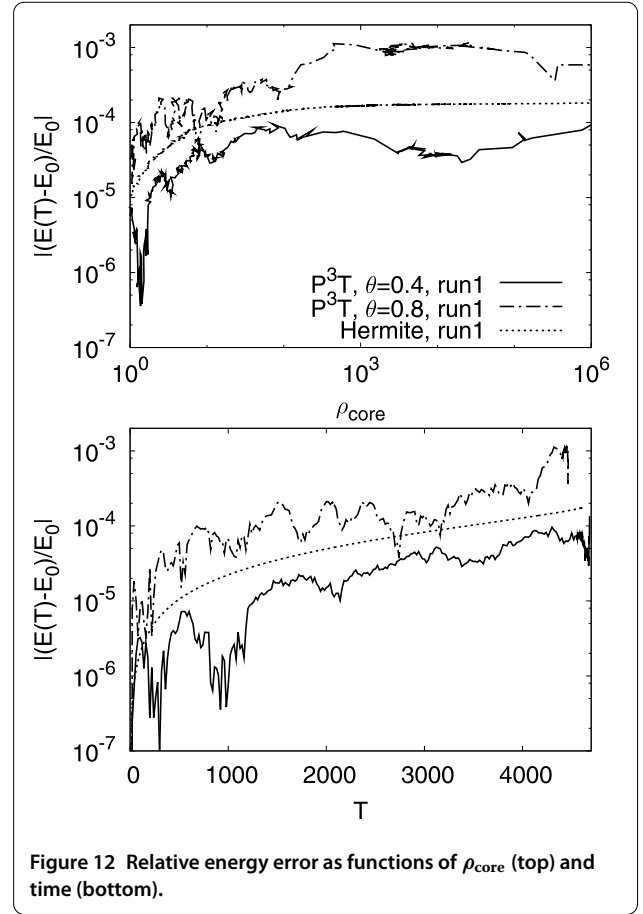
3.2.1 Star cluster down to core collapse

In this section, we discuss the performance of our P³T scheme for the simulation of the core collapse of a star cluster. First, we describe the initial condition and parameters of the integration scheme. Next, we compare the calculation results obtained by the P³T and Hermite schemes, and finally, the calculation speed.



Initial conditions We apply the P^3T scheme to the evolution of a star cluster consisting of 16K stars to the moment of the core collapse (Lynden-Bell and Eggleton 1980). We use an equal-mass plummer model as an initial density profile and we adopt $\eta = 0.1$. We apply the Plummer softening $\epsilon = 4/N = 1/4,096$. The simulations are terminated when the core number-density exceeds 10^6 , at which point the mean interparticle distance in the core is comparable to ϵ . Next, we set θ . We must choose θ so that the tree force error is smaller than the force due to the two-body relaxation. Hernquist et al. (1993) pointed out that, for $\theta = 0.5$ with monopole and quadrupole, the tree-force error is much smaller than the force due to the two-body relaxation. Thus we choose $\theta = 0.4$ with quadrupole as a standard model. For comparison, we also perform a run with $\theta = 0.8$.


To resolve the motions of the particles in the core, we impose Δt_{soft} to be smaller than $1/128$ of the dynamical time of the core ($\sim \sqrt{3\pi/16\rho_{\text{core}}}$, where ρ_{core} is the core density). To reduce the calculation cost for the Hermite part we require $r_{\text{cut}} \propto \rho_{\text{core}}^{-1/3}$ and set the initial value of $r_{\text{cut}} = 1/64$. We also change $\Delta r_{\text{buff}} = 3\sigma_{\text{core}}\Delta t_{\text{soft}}$, where σ_{core} is the velocity dispersion in the core, and $\Delta t_{\text{max}} = \Delta t_{\text{soft}}/4$, as Δt_{soft} and



σ_{core} are changing. Here, to calculate ρ_{core} and σ_{core} , we use the formula proposed by Casertano and Hut (1985). The same simulation is repeated using the fourth-order Hermite scheme with the block timesteps with the same value of $\eta = 0.1$.

Results In Figure 11 we present the evolution of the core densities ρ_{core} (top panel) and the core radii r_{core} (bottom panel) for P^3T and Hermite schemes. For each scheme, we perform three runs, changing the initial random seed for generating the initial conditions of the Plummer model. The behaviors of the cores for all runs are similar. The differences between two schemes are smaller than run-to-run variations.

Figure 12 shows the relative energy errors of the runs with the same initial seed as functions of the core density (top panel) and the time (bottom panel). The energy errors of the runs with P^3T scheme change randomly, whereas those of the Hermite code grow monotonically. As a result, the P^3T scheme with $\theta = 0.4$ conserves energy better than the Hermite scheme in the long run. The errors for the P^3T scheme with $\theta = 0.8$ is slightly worse than that of the Hermite scheme, but the behavior of the core are sim-



ilar with other runs. Thus the choice of $\theta = 0.4$ is enough to follow the core collapse simulations.

Calculation speed Figure 13 shows the calculation time of the P^3T scheme ($\theta = 0.4$) and Hermite scheme on GPU. As shown in this figure, the calculation time of the P^3T scheme is dominated by the tree (soft) part calculation.

Initially the P^3T scheme is much faster than the Hermite scheme, but after the time when $\rho_{\text{core}} \sim 10^4$, the P^3T scheme is slightly slower than the Hermite scheme because in the P^3T scheme, Δt_{soft} is proportional to $\rho_{\text{core}}^{-1/2}$. However, even for the P^3T scheme, the CPU time spent after ρ_{core} reaches 10^4 is small. As a result, the calculation time to the moment of the core collapse of the P^3T scheme is smaller than that of the Hermite scheme by a factor of two.

3.2.2 Orbital evolution of SMBH binary

In this section, we also discuss the performance of the P^3T scheme applied to simulations of a galaxy with a supermassive black hole (SMBH) binary. First, we describe the initial conditions and parameters of the integration scheme. Next, we compare the calculation results obtained by the P^3T and Hermite schemes, and finally, the calculation speed.

Initial conditions and methods We use the Plummer model with $N = 16K, 128K$ and $256K$ as the initial galaxy model. Two SMBH particles with a mass of 1% of that of the galaxy are placed at the positions $(\pm 0.5, 0.0, 0.0)$ with the velocities $(0.0, \pm 0.5, 0.0)$. We use three values for the cut off radius with respect to three different kinds of interactions. For the interaction between field stars (FSs), we set $r_{\text{cut,FS-FS}} = 1/256$. For the interaction between SMBHs, the force is not split and $F_{\text{soft}} = 0$. In other words, the force between SMBHs is integrated with the pure Hermite scheme. We set the cut off radius between SMBH and FS $r_{\text{cut,BH-FS}} = 1/32$ which is large enough that Δt_{soft} is smaller than the Kepler time of a particle in orbit around the SMBH binary at a distance of $r_{\text{cut,BH-FS}}$. We use the Plummer softening $\epsilon = 10^{-4}$ for the interactions between FS-FS and FS-SMBH. For the SMBH-SMBH interaction, we do not use the softening. The accuracy parameter of timestep criterion for FS η_{FS} is 0.1, and for SMBH η_{BH} is 0.03. We adopt $\Delta r_{\text{buff}} = 3\sigma \Delta t_{\text{soft}}$, $\Delta t_{\text{max}} = \Delta t_{\text{soft}}/4$ and $\theta = 0.4$.


We use $\Delta t_{\text{soft}} = 1/1,024$ at $T = 0$, and as the binary becomes harder, we decrease Δt_{soft} to suppress the aliasing error of the binary. As a standard model, we set Δt_{soft} to be less than half of the Kepler time of the SMBH binary t_{kep} . Only for $N = 128K$, we also perform two other runs, where $\Delta t_{\text{soft}} < t_{\text{kep}}/4$ and t_{kep} .

We also perform the same simulations by the Hermite scheme with the same η_{FS} and η_{BH} .


Results Figure 14 shows the evolution of the semi-major axis (top panel) and eccentricity (middle panel) of the SMBH binary and the relative energy error (bottom panel) as functions of time for our standard models ($\Delta t_{\text{soft}} < t_{\text{kep}}/2$). The behaviors of the semi-major axis of the SMBH binary for the runs with the same N agree very well. The hardening rate of the binary depends on N because of the loss-cone refilling through the two-body relaxation (Begelman et al. 1980; Makino and Funato 2004; Berczik et al. 2005). The evolution of the eccentricity has large variation, because this evolution is sensitive to small N fluctuation (Merritt et al. 2007). In the cases of $N = 16K$ with the Hermite scheme, the relative energy error increases dramatically after $T = 150$ because the binding energy and the eccentricity of the binary are very high.

Figure 15 is the same as Figure 14 but for several different values of Δt_{soft} . Thick solid, dashed and dotted curves indicate the results for $\Delta t_{\text{soft}} < t_{\text{kep}}/4$, $t_{\text{kep}}/2$ and t_{kep} , respectively. The orbital parameters show similar behaviors for all runs. The absolute value of the energy errors of P^3T runs ($\sim 10^{-5}$) are small compared with the binding energy of SMBH binary, which is roughly 0.05.

Calculation speed Figure 16 shows the calculation time for runs for several different values of N with $\Delta t_{\text{soft}} < t_{\text{kep}}/2$.



Initially, the P^3T scheme is much faster than the Hermite scheme. As the SMBH binary becomes harder, the P^3T scheme slows down more significantly than the direct Hermite scheme does. We can see that T_{cal} of the Hermite scheme is roughly proportional to a^{-1} for $a^{-1} > 300$, whereas that of the P^3T scheme is roughly proportional to $a^{-5/2}$, because Δt_{soft} is proportional to the Kepler time of the binary ($\propto a^{3/2}$). However, the calculation time for all runs with the P^3T scheme is shorter than that with the Hermite scheme by $a = 1/800$. We can also confirm that as we use more N , the ratio of the calculation time of the P^3T scheme to the Hermite scheme become larger. The reason why the P^3T scheme becomes slower for large a^{-1} is simply that we force the timestep of all particles to be smaller than the orbital period of the SMBH binary. For the Hermite scheme, we do not put such constraint. Thus, in the Hermite scheme, particles far away from the SMBH have the timestep much larger than the orbital period of the SMBH



binary. This large timestep can cause accuracy problem (Nitadori and Makino 2008). With P^3T , it is possible to apply the perturbation approximation to F_{soft} between the SMBH binary and other particles. Such a treatment should improve the accuracy and speed of the P^3T scheme when the SMBH binary becomes very hard.

In Figure 17, we plot the calculation time of the hard and soft parts for the standard model with $N = 128k$. We can see that the soft parts dominate the calculation time.

In Figure 18, we compare the calculation time for the runs with various Δt_{soft} ($< t_{\text{kep}}, t_{\text{kep}}/2, t_{\text{kep}}/4$). Since the most of the calculation time is spent after the binary becomes hard, the calculation time strongly depends on the criterion of the Δt_{soft} . From Figure 15, the evolution of the orbital parameters for all runs with the P^3T scheme are




Figure 16 Wall-clock times as a function of $1/a$ (top and middle) and the system time of the simulations (bottom) for several different values of N . In top and middle panels, the x- and y-axis are logarithmic and linear scales, respectively. In the bottom panel, the x-axis is scaled by $N/16K$.

similar for various Δt_{soft} criteria. Thus we could choose larger $\Delta t_{\text{soft}} \gtrsim t_{\text{kep}}$ after the binary formation.

4 Conclusions

We have described the implementation and performance of the P^3T scheme for simulating dense stellar systems. In our implementation, the tree part is accelerated using

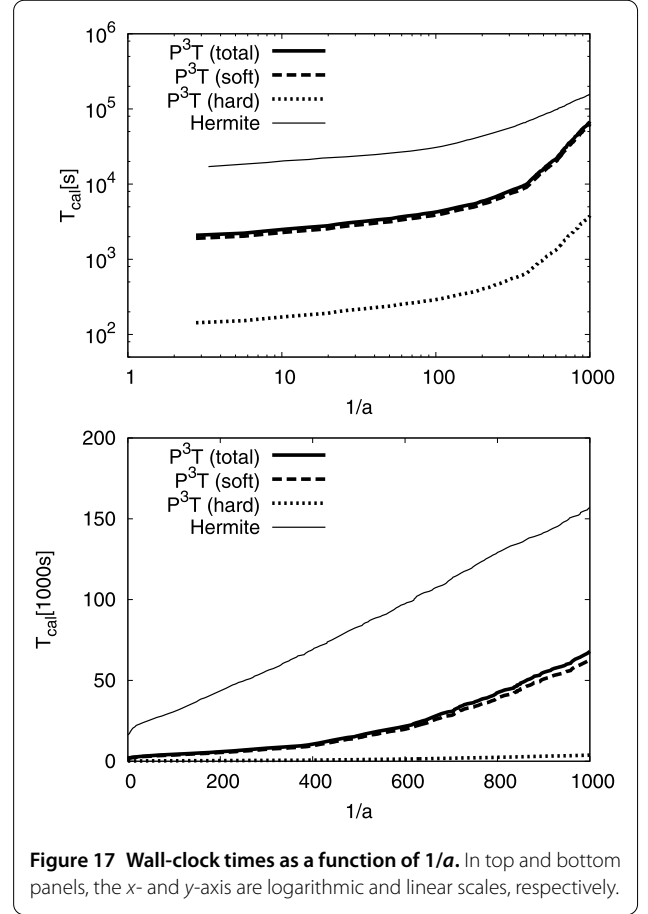


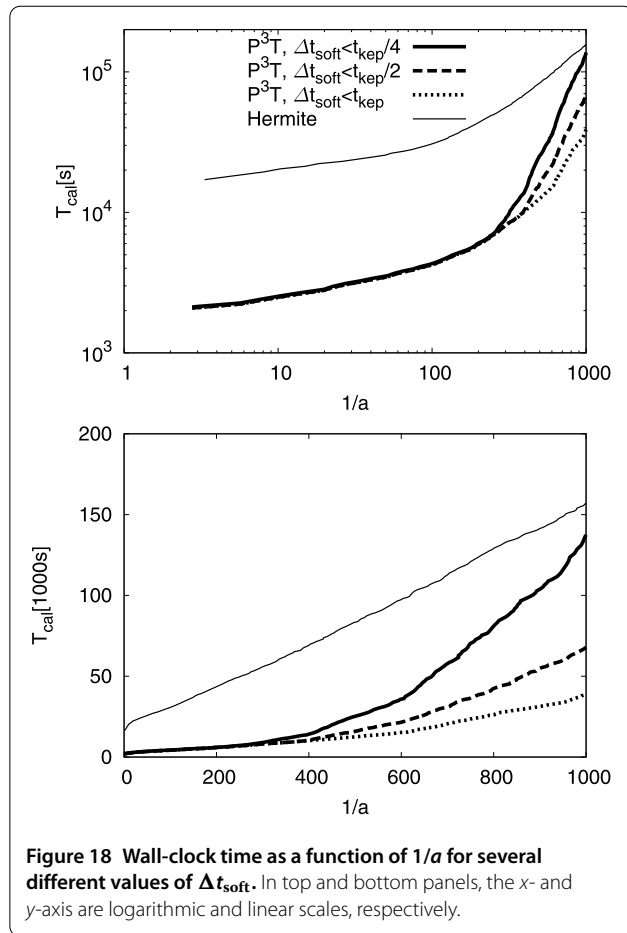
Figure 17 Wall-clock times as a function of $1/a$. In top and bottom panels, the x- and y-axis are logarithmic and linear scales, respectively.

GPU. The accuracy and performance of the P^3T scheme can be controlled through six parameters: Δr_{cut} , Δr_{buff} , Δt_{soft} , Δt_{max} , η and θ . We find that $\Delta r_{\text{buff}} \gtrsim 2\sigma \Delta t_{\text{soft}}$ is a good choice to prevent non-neighbour particles from entering the neighbour sphere. The integration errors can be controlled through $\Delta t_{\text{soft}}/\Delta r_{\text{cut}}\sigma$. For $\theta = 0.2$, if we set Δt_{soft} to be less than $0.05\Delta r_{\text{cut}}/\sigma$, the integration error is smaller than the tree force error. For the Hermite part, if we choose $\eta \lesssim 0.2$, the errors hardly depend on Δt_{max} .

From the point of view of the balance of the calculation costs between the tree and Hermite parts, we derive the optimal set of accuracy parameters, and find that the calculation cost is proportional to $N^{4/3}$.

The P^3T scheme is suitable for simulating large N stellar clusters with a high density contrast, such as star clusters or galactic nuclei. We demonstrate the efficiency of the code and show that it is able to integrate N -body systems to the moment of the core collapse. We also performed the simulations of the galaxy with the SMBH binary and found that the P^3T scheme can be applied to these simulations.

Finally, we discuss the possibilities of implementing two important effects on star cluster evolution in P^3T . The first is the effect of a tidal field which dramatically changes the



collapse time and the evaporation time of a star cluster. The tidal field effect can be included in the soft part.

The other is an effect of stellar-mass binaries. A stellar-mass binary plays an important role in halting the core collapse. In this paper, we introduce the Plummer softening and neglect these binary effect. However, we could treat these effects by integrating stellar-mass binaries in the hard part.

Our P^3T code is incorporated in the AMUSE frameworks and free for use (Portegies Zwart et al. 2013; Pelupessy et al. 2013).

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

All authors, MI, SPZ and JM conceived of the study. MI developed the code, performed all simulations and drafted the manuscript. SPZ and JM helped to draft the manuscript. All authors read and approved the final manuscript.

Author details

¹RIKEN Advanced Institute for Computational Science, Minatojima-minamimachi, Chuo-ku, Kobe, Japan. ²Sterrewacht Leiden, P.O. Box 9513, Leiden, 2300 RA, The Netherlands. ³Earth-Life Science Institute, Tokyo Institute of Technology, Ookayama, Meguro-ku, Tokyo, Japan.

Acknowledgements

We are grateful to Jeroen Bédorf, for preparations of the GPU cluster and GPU library. We also thank to Shoichi Oshino, Daniel Caputo and Keigo Nitadori for stimulating discussion, and to Edwin van der Helm for carefully reading the manuscript. This work was supported by NWO (grants VICI [#639.073.803], AMUSE [#614.061.608] and LGM [# 612.071.503]), NOVA and the LKBF.

Endnote

- ^a GTX680 does not have ECC (Error Check and Correct) memories. However, as we will see later, we do not observe any large energy error in any of our runs, which means the hardware error does not affect our result. Betz et al. (2014) performed Molecular Dynamics simulations, in order to investigate the rate of bit-flip error events. They observed a single bit-flip error event in about 4,700 GPU*hours without ECC and conclude that the bit-flip error is exceedingly rare.

Received: 19 February 2014 Accepted: 15 June 2015

Published online: 03 July 2015

References

- Aarseth, SJ: Dynamical evolution of clusters of galaxies. *I. Mon. Not. R. Astron. Soc.* **126**, 223-255 (1963)
- Barnes, J, Hut, P: A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature* **324**, 446-449 (1986)
- Dehnen, W: A hierarchical $O(N)$ force calculation algorithm. *J. Comput. Phys.* **179**, 27-42 (2002)
- Dehnen, W: A fast multipole method for stellar dynamics. *Comput. Astrophys. Cosmol.* **1**, 1 (2014)
- Greengard, L, Rokhlin, V: A fast algorithm for particle simulations. *J. Comput. Phys.* **73**, 325-348 (1987)
- Gaburov, E, Bédorf, J, Portegies Zwart, S: Gravitational tree-code on graphics processing units: implementation in CUDA. *Proc. Comput. Sci.* **1**, 1119-1127 (2010)
- Bédorf, J, Gaburov, E, Portegies Zwart, S: A sparse octree gravitational N -body code that runs entirely on the GPU processor. *J. Comput. Phys.* **231**, 2825-2839 (2012)
- Bédorf, J, Gaburov, E, Fujii, MS, Nitadori, K, Ishiyama, T, Portegies Zwart, S: 24.77 Pflops on a gravitational tree-code to simulate the Milky Way Galaxy with 18600 GPUs. In: *SC'14 Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 54-65. IEEE Press, Piscataway (2014). doi:10.1109/SC.2014.10
- McMillan, SLW: The vectorization of small- n integrators. In: Hut, P, McMillan, SLW (eds.) *The Use of Supercomputers in Stellar Dynamics*. Lecture Notes in Physics, vol. 267, pp. 156-161. Springer, Berlin (1986)
- McMillan, SLW, Aarseth, SJ: An $O(N \log N)$ integration scheme for collisional stellar systems. *Astrophys. J.* **414**, 200-212 (1993)
- Oshino, S, Funato, Y, Makino, J: Particle-particle particle-tree: a direct-tree hybrid scheme for collisional N -body simulations. *Publ. Astron. Soc. Jpn.* **63**, 881-892 (2011)
- Kinoshita, H, Yoshida, H, Nakai, H: Symplectic integrators and their application to dynamical astronomy. *Celest. Mech. Dyn. Astron.* **50**, 59-71 (1991)
- Wisdom, J, Holman, M: Symplectic maps for the n -body problem. *Astron. J.* **102**, 1528-1538 (1991)
- Duncan, MJ, Levison, HF, Lee, MH: A multiple time step symplectic algorithm for integrating close encounters. *Astron. J.* **116**, 2067-2077 (1998)
- Chambers, JE: A hybrid symplectic integrator that permits close encounters between massive bodies. *Mon. Not. R. Astron. Soc.* **304**, 793-799 (1999)
- Brunini, A, Viturro, HR: A tree code for planetesimal dynamics: comparison with a hybrid direct code. *Mon. Not. R. Astron. Soc.* **346**, 924-932 (2003)
- Fujii, M, Iwasawa, M, Funato, Y, Makino, J: BRIDGE: a direct-tree hybrid N -body algorithm for fully self-consistent simulations of star clusters and their parent galaxies. *Publ. Astron. Soc. Jpn.* **59**, 1095-1106 (2007)
- Moore, A, Quillen, AC: QYMSYM: a GPU-accelerated hybrid symplectic integrator that permits close encounters. *New Astron.* **16**, 445-455 (2011)
- Makino, J, Aarseth, SJ: On a Hermite integrator with Ahmad-Cohen scheme for gravitational many-body problems. *Publ. Astron. Soc. Jpn.* **44**, 141-151 (1992)
- Hockney, RW, Eastwood, JW: *Computer Simulation Using Particles* (1981)
- Plummer, HC: On the problem of distribution in globular star clusters. *Mon. Not. R. Astron. Soc.* **71**, 460-470 (1911)
- Heggie, DC, Mathieu, RD: Standardised units and time scales. In: Hut, P, McMillan, SLW (eds.) *The Use of Supercomputers in Stellar Dynamics*. Lecture Notes in Physics, vol. 267, pp. 233-235. Springer, Berlin (1986)

- Portegies Zwart, S, Boekholt, T: On the minimal accuracy required for simulating self-gravitating systems by means of direct N -body methods. *Astrophys. J. Lett.* **785**, L3 (2014)
- Gaburov, E, Harfst, S, Portegies Zwart, S: SAPPORO: a way to turn your graphics cards into a GRAPE-6. *New Astron.* **14**, 630-637 (2009)
- Betz, RM, DeBardeleben, NA, Walker, RC: An investigation of the effects of hard and soft errors on graphics processing unit-accelerated molecular dynamics simulations. *Concurr. Comput., Pract. Exp.* **26**, 2134-2140 (2014)
- Lynden-Bell, D, Eggleton, PP: On the consequences of the gravothermal catastrophe. *Mon. Not. R. Astron. Soc.* **191**, 483-498 (1980)
- Hernquist, L, Hut, P, Makino, J: Discreteness noise versus force errors in N -body simulations. *Astrophys. J. Lett.* **402**, L85 (1993)
- Casertano, S, Hut, P: Core radius and density measurements in N -body experiments connections with theoretical and observational definitions. *Astrophys. J.* **298**, 80-94 (1985)
- Begelman, MC, Blandford, RD, Rees, MJ: Massive black hole binaries in active galactic nuclei. *Nature* **287**, 307-309 (1980)
- Makino, J, Funato, Y: Evolution of massive black hole binaries. *Astrophys. J.* **602**, 93-102 (2004)
- Berczik, P, Merritt, D, Spurzem, R: Long-term evolution of massive black hole binaries. II. Binary evolution in low-density galaxies. *Astrophys. J.* **633**, 680-687 (2005)
- Merritt, D, Mikkola, S, Szell, A: Long-term evolution of massive black hole binaries. III. Binary evolution in collisional nuclei. *Astrophys. J.* **671**, 53-72 (2007)
- Nitadori, K, Makino, J: Sixth- and eighth-order Hermite integrator for N -body simulations. *New Astron.* **13**, 498-507 (2008)
- Portegies Zwart, S, McMillan, SLW, van Elteren, E, Pelupessy, I, de Vries, N: Multi-physics simulations using a hierarchical interchangeable software interface. *Comput. Phys. Commun.* **183**, 456-468 (2013)
- Pelupessy, FI, van Elteren, A, de Vries, N, McMillan, SLW, Drost, N, Portegies Zwart, SF: The astrophysical multipurpose software environment. *Astron. Astrophys.* **557**, A84 (2013)

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Immediate publication on acceptance
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► springeropen.com
