

RESEARCH

Open Access

A fast multipole method for stellar dynamics

Walter Dehnen*

Abstract

The approximate computation of all gravitational forces between N interacting particles via the fast multipole method (FMM) can be made as accurate as direct summation, but requires less than $\mathcal{O}(N)$ operations. FMM groups particles into spatially bounded cells and uses cell-cell interactions to approximate the force at *any* position within the sink cell by a Taylor expansion obtained from the multipole expansion of the source cell. By employing a novel estimate for the errors incurred in this process, I minimise the computational effort required for a given accuracy and obtain a well-behaved distribution of force errors. For relative force errors of $\sim 10^{-7}$, the computational costs exhibit an empirical scaling of $\propto N^{0.87}$. My implementation (running on a 16 core node) out-performs a GPU-based direct summation with comparable force errors for $N \gtrsim 10^5$.

Keywords: methods: numerical; methods: N -body simulations

1 Background

The computation of the mutual gravitational forces at every time step dominates the computational costs of all N -body simulations. When simulating collisionless stellar dynamics, the N -body model is merely a Monte-Carlo representation of a smooth phase-space distribution and the N -body force is only ever an *estimate* for the smooth force field of the continuous system modelled (see also Dehnen and Read 2011). In particular, the N -body force unavoidably carries an *estimation error*. This motivates the use of *approximate* methods for computing the N -body force, such as the Barnes and Hut (1986) tree code, as long as the approximation errors are small compared to the estimation errors.

N -body simulations of collisional stellar dynamics are of a completely different nature. Here, the particles simulate individual stars and the N -body force carries no estimation error. Consequently, the (negative) gravitational potential

$$\Psi(\mathbf{x}_b) = \sum_{a \neq b} G\mu_a \psi(\mathbf{x}_b - \mathbf{x}_a) \quad \text{with } \psi(\mathbf{r}) = |\mathbf{r}|^{-1} \quad (1)$$

and its derivative, the acceleration, must be calculated with high accuracy. This is typically achieved by *direct summation*, when equation (1) is translated into computer code

and the only errors are owed to finite computational precision.

This computation incurs a cost of $\mathcal{O}(N)$ for a single particle and thus $\mathcal{O}(N^2)$ per unit time for running a full simulation. As a consequence, realistic simulations with $N \sim 10^{6-7}$ for globular clusters and galactic centres are still very challenging and large parameter studies impossible. Measures employed to ameliorate this situation include the usage of powerful special-purpose hardware devices (Makino and Taiji 1998) or graphical processing units (GPUs, Gaburov et al. 2009), as well as separating the highly fluctuating forces due to close neighbours, in order to reduce the frequency of expensive far-field force computations (Ahmad and Cohen 1973).

While these measures substantially reduce the effective costs, the complexity of N^2 remains. The alternative of using approximate methods also for collisional stellar dynamics is so far untested. The requirements for such a method differ from that in collisionless N -body methods in two important aspects: (i) there is no gravitational softening and (ii) to preserve the validity of the N -body model, the approximation errors must be much smaller than what is common in collisionless N -body simulations.

A straightforward approach is to use the tree code with a small opening angle and/or high expansion order, resulting in a scheme with $\mathcal{O}(N \ln N)$ costs. A more efficient approach is to use the *fast multipole method* (FMM; Greengard and Rokhlin 1987; Cheng et al. 1999) which has costs of only $\mathcal{O}(N)$. An initial attempt by Capuzzo-Dolcetta and

* Correspondence: walter.dehnen@le.ac.uk

Department for Physics & Astronomy, University of Leicester, Leicester, LE1 7RH, United Kingdom

Miocchi (1998) to port this technique from its original realm of molecular dynamics to astrophysics failed to obtain better practical efficiency than the tree code. However, when adapting the FMM to the inhomogeneity of stellar systems and the low force accuracy required in collisionless dynamics (by using a hierarchical tree data structure and a flexible opening angle), it is substantially faster than the tree code (Dehnen 2000, 2002).

The critical question here is whether FMM can be tuned to be more efficient than direct summation at force accuracies and particle numbers required by collisional N -body techniques. The goal of this study is to address this question by tuning FMM for the application to collisional N -body simulations, investigating the resulting dependence of computational costs and numerical accuracy on the various numerical parameters, and assessing its practical efficiency.

This paper is organised as follows. In Section 2 and Appendix A, the mathematical (and algorithmic) foundations of FMM are derived and laid down. Section 3 (and Appendix B) introduces and motivates my approach for quantifying the resulting acceleration errors; Section 4 provides useful estimates for the errors of individual FMM interactions; Section 5 deals with optimising the multipole-acceptance criterion; and in Section 6 the method is tuned to obtain a force accuracy target with minimal computational effort. Finally, in Section 7 possible extensions and applications are discussed, and Section 8 concludes.

2 FMM basics

The tree code approximates the sum (1) by first dividing source particles a into groups bounded by geometric cells, each of which is well-separated from the sink position \mathbf{x}_b , and then computing the forces of each source cell from their multipole moments. This corresponds to Taylor expanding the Greens function $\psi(\mathbf{x}_b - \mathbf{x}_a)$ about the distance to an appropriate centre \mathbf{z} of each source cell.

The essence of the fast multipole method is to Taylor expand the Greens function not only at the source positions \mathbf{x}_a , but also at the sink positions \mathbf{x}_b . This latter amounts to approximating (a contribution to) the gravitational field within each sink cell by its local Taylor expansion about some appropriate potential expansion centre \mathbf{s} . Obviously, this approach is beneficial only if the forces for a large fraction of the sinks within a cell are to be computed simultaneously.

2.1 Mathematical background

The FMM relations are most easily derived using Cartesian coordinates. However, for Newtonian gravity, $\psi = |\mathbf{r}|^{-1}$, the resulting relations are inefficient. Instead, exploiting that this Greens function satisfies $\nabla^2 \psi = 0$ for $\mathbf{r} \neq 0$ naturally leads to spherical harmonics. Cheng et al. (1999) have already given (without derivation) the corresponding

FMM relations, but in a form ill-suited for computer code. In Appendix A, I derive equivalent but much more compact and computationally convenient relations. These are summarised here.

Let $\mathbf{r} = (x, y, z)$ with spherical polar coordinates r, θ, ϕ , then

$$\Theta_n^m(\mathbf{r}) = (-1)^m \frac{(n-m)!}{r^{n+1}} P_n^m(\cos \theta) e^{im\phi}, \quad (2a)$$

$$\Upsilon_n^m(\mathbf{r}) = (-1)^m \frac{r^n}{(n+m)!} P_n^m(\cos \theta) e^{im\phi} \quad (2b)$$

with integer indices $0 \leq |m| \leq n$ are (complex-valued) harmonic functions, i.e. $\nabla^2 \Upsilon_n^m = 0$ for all \mathbf{r} and $\nabla^2 \Theta_n^m = 0$ for all $\mathbf{r} \neq 0$. The Υ_n^m are homogeneous polynomials of total degree n in x, y , and z (they are defined in Appendix A.3 without reference to polar coordinates; see also Table 3). With these definitions, the FMM relations for the computation of the potential due to all particles within source cell A and at any position \mathbf{x}_b within sink cell B are

$$\Psi_{A \rightarrow B}(\mathbf{x}_b) = \sum_{n=0}^p \sum_{m=-n}^n \Upsilon_n^{m*}(\mathbf{s}_B - \mathbf{x}_b) \mathcal{F}_n^m(\mathbf{s}_B) + \delta \Psi_{A \rightarrow B}, \quad (3a)$$

$$\mathcal{F}_n^m(\mathbf{s}_B) = \sum_{k=0}^{p-n} \sum_{l=-k}^k \mathcal{M}_k^{l*}(\mathbf{z}_A) \Theta_{n+k}^{m+l}(\mathbf{s}_B - \mathbf{z}_A), \quad (3b)$$

$$\mathcal{M}_n^m(\mathbf{z}_A) = \sum_{a \in A} \mu_a \Upsilon_n^m(\mathbf{x}_a - \mathbf{z}_A). \quad (3c)$$

Here, p is the expansion order and $\delta \Psi_{A \rightarrow B}$ the error of the approximated potential. This expansion converges with increasing p if $\max_{a \in A} \{|\mathbf{x}_b - \mathbf{x}_a - \mathbf{r}|\} < |\mathbf{r}|$ with $\mathbf{r} \equiv \mathbf{s}_B - \mathbf{z}_A$.

Other important relations are those for the *multipoles* \mathcal{M}_n^m with respect to another expansion centre

$$\mathcal{M}_n^m(\mathbf{z}') = \sum_{k=0}^n \sum_{l=-k}^k \Upsilon_k^l(\mathbf{z} - \mathbf{z}') \mathcal{M}_{n-k}^{m-l}(\mathbf{z}), \quad (3d)$$

and for the *field tensors* \mathcal{F}_n^m of the local expansion (3a) with respect to another expansion centre

$$\mathcal{F}_n^m(\mathbf{s}') = \sum_{k=0}^{p-n} \sum_{l=-k}^k \Upsilon_k^{l*}(\mathbf{s} - \mathbf{s}') \mathcal{F}_{n+k}^{m+l}(\mathbf{s}). \quad (3e)$$

Moreover, the computation of the acceleration \mathbf{a} from the local expansion (3a) requires

$$\Psi_n^m(\mathbf{x}_b) = \sum_{k=0}^{p-n} \sum_{l=-k}^k \Upsilon_k^{l*}(\mathbf{s}_B - \mathbf{x}_b) \mathcal{F}_{n+k}^{m+l}(\mathbf{s}_B) + \delta \Psi_n^m, \quad (3f)$$




Figure 1 Timings for the M2L kernel. The time taken by the M2L kernel (using double precision without vectorisation) as function of expansion order p . The rotation accelerated version of Appendix A.6.2 is faster than the direct implementation of equation (3b) for $p \geq 5$.

when $\mathbf{a} = \nabla \Psi_0^0 = -(\Re\{\Psi_1^1\}, \Im\{\Psi_1^1\}, \Psi_1^0)$. Finally, the gravity generated from a source distribution with given multipoles is given by

$$\Psi_n^m(\mathbf{x}_b) = \sum_{k=0}^{p-n} \sum_{l=-k}^k \mathcal{M}_k^{l*}(\mathbf{z}_A) \Theta_{n+k}^{m+l}(\mathbf{x}_b - \mathbf{z}_A) + \delta \Psi_n^m. \quad (3g)$$

Relations (3b), (3d), and (3e) are equivalent to the much more complicated equations (17), (13), and (21) of Cheng et al. (1999), given without derivation.^a

There are $(p+1)^2$ independent real-valued numbers \mathcal{F}_n^m (as well as \mathcal{M}_n^m , see also Appendix A.5.2), and their computation via equations (3b), (3d), and (3e) requires $\mathcal{O}(p^4)$ operations.^b These operation counts can be reduced to $\mathcal{O}(p^3)$ by rotating \mathbf{r} into the z direction (see Appendix A.6). Figure 1 plots the time required per interaction computation as function of expansion order p , showing an effective $p^{2.3}$ scaling of the computational costs at $p \leq 10$, shallower than the $\mathcal{O}(p^3)$ asymptote.

2.2 Algorithmic approach

2.2.1 The tree code: walking the tree

Let us first consider the tree code, which also uses the multipole expansion but is algorithmically simpler than FMM. The basic data structure is a hierarchical tree of spatial cells, which are either cubic with eight daughter cells (oct-tree) or cuboidal with two daughters (binary tree). In a first step, the multipoles \mathcal{M}_n^m have to be computed for each cell from those of their daughter cells, using the M2M kernel (equation (3d), see also Table 1), or (in case of final cells) of their particles, using the P2M kernel (equation (3c)).

Next, the force for each sink position is computed using a separate *tree walk* starting with the root cell. The force

Table 1 The FMM kernels

Name	Meaning	Equation
P2P	particle to particle	(1)
P2M	particle to multipole	(3c)
M2M	multipole to multipole	(3d)
M2P	multipole to particle	(3g)
M2L	multipole to local expansion	(3b)
P2L	particle to local expansion	see table legend
L2L	local expansion to local expansion	(3e)
L2P	local expansion to particle	(3f)

The tree code replaces direct summation (P2P) with P2M-M2M-M2P, while FMM uses P2M-M2M-M2L-L2L-L2P, see also Figure 2. The P2L kernel corresponds to equation (3b) with $\mathbf{z}_A \rightarrow \mathbf{x}_A$ and $\mathcal{M}_k^l \rightarrow m_A \delta_{k0}$, i.e. $\mathcal{F}_n^m(\mathbf{s}_g) = m_A \Theta_n^m(\mathbf{s}_g - \mathbf{x}_A)$.

generated by a cell C is computed via its multipole expansion, using the M2P kernel (equation (3g)), if a *multipole-acceptance criterion* is met, i.e. if the cell is considered to be *well-separated* from the sink position. Otherwise, the cell is opened: the force is computed as the sum of the forces generated by the daughter cells (recursing if necessary). Thus, the tree code replaces direct summation's P2P kernel with the P2M, M2M, and M2P kernels, see the left panel of Figure 2 for a schematic view.

2.2.2 FMM: the dual tree walk

An adaptive FMM algorithm also uses a hierarchical tree data structure. As with the tree code, the cell multipoles \mathcal{M}_n^m have to be precomputed for every cell in a first step.

Next, the forces for *all* sink positions and generated by *all* source particles are approximated using a single *dual tree walk* (Dehnen 2002). This algorithm considers cell \rightarrow cell interactions and starts with the root \rightarrow root interaction. If the interacting cells are well separated, the interaction is approximated using the M2L kernel (equation (3b)), which computes and accumulates the local field tensors $\mathcal{F}_n^m(\mathbf{s}_B)$ for the expansion of gravity within the sink cell B and due to all sources within the source cell A (in a mutual version of the algorithm, the interactions $A \rightarrow B$ and $B \rightarrow A$ are considered simultaneously). Otherwise, the interaction is split, typically into those between the daughters of the larger of the two interacting cells with the smaller.

Finally, the local field tensors $\mathcal{F}_n^m(\mathbf{s})$ are passed down the tree using the L2L kernel, and the local expansions are evaluated at the sink positions using the L2P kernel. Thus, the FMM replaces the M2P kernel of the tree code with the M2L, L2L and L2P kernels, see also Figure 2.

Of course, in both tree code and FMM, direct summation (P2P kernel) is used whenever computationally preferable, i.e. for interactions involving only a few sources and sinks.

3 Quantifying the approximation accuracy

Before the method can be optimised for accuracy, a sensible quantitative measure for this accuracy is needed as well as an acceptable value for this measure.

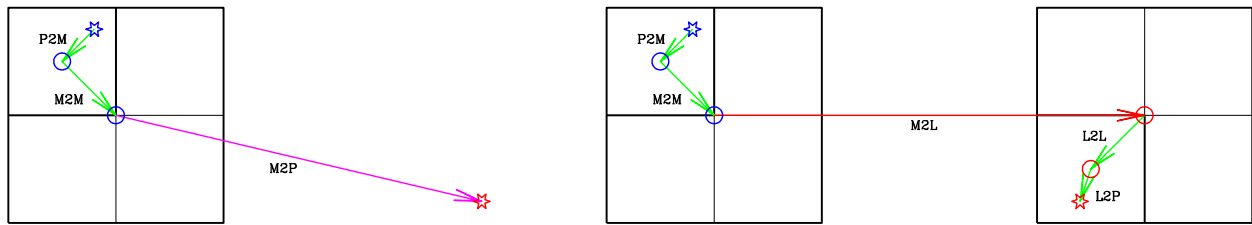


Figure 2 Schematic view of the tree code (left) and FMM (right). The tree code approximates the force from a source particle (blue star) at a sink position (red star) using the P2M and M2M kernels (green arrows) to compute the multipoles at the cell centres (blue circles) followed by the M2P kernel (pink arrow). The P2M and M2M kernels are called once per source particle and cell, respectively, while the M2P kernel is called many times for each sink position. FMM replaces many calls to the M2P kernel by a single call to the M2L kernel (red arrow) followed by the L2L and L2P kernels (green arrows). Again, the L2L and L2P kernels are called once per sink cell and particle, respectively, but a single M2L kernel replaces many M2P kernels of the tree code, because it accounts for *all* sink positions within the sink cell.

With direct-summation, the accuracy is limited only by the finite precision of computer arithmetic (round-off error). If double (64-bit) precision is not used throughout, it is customary to use the conservation of the total energy for quality control (e.g. Gaburov et al. 2009). However, as shown in Appendix B, the relative energy error is much smaller than the typical relative force error, simply because it is an average over many force errors. Even worse, the computation of the total energy, required for measuring its error, typically incurs a larger error. Thus, any measured non-conservation of the total energy is dominated by measurement error rather than true non-conservation due to acceleration errors.

With the tree code and FMM, the situation is subtly different, as discussed in Appendix B.3. Here, the measured non-conservation of energy actually reflects the amplitude of the acceleration errors in an average sense. However, an average measure for the effect of approximation errors cannot reflect their effect on the correctness of the simulation. For example, a single large force error has hardly any effect on the energy conservation but may seriously affect the validity of the simulation. While this latter goal is difficult to quantify, it is certainly better to consider the whole distribution of acceleration errors and pay particular attention to large-error outliers, than merely monitor an average.

3.1 Scaling acceleration errors

Obviously, the absolute errors $\delta a = |\mathbf{a}_{\text{computed}} - \mathbf{a}_{\text{true}}|$ are not very useful by themselves and must be normalised to be meaningful. One option is to divide δa by some mean field strength \bar{a} . While this makes sense for the average particle, it fails for those in the outskirts of the stellar system, where the field strength diminishes well below its mean.

To overcome such issues, a natural choice is the relative error $\delta a/a$. However, this is still problematic in the centre of a stellar system, where forces from the outward lying

parts largely cancel. In such a situation, a can be small and hence the relative error large, even if each individual pair-wise force has been computed with high accuracy. One option for avoiding this problem is - in analogy to the error estimate of numerical quadrature in case of an integrand oscillating around zero - to normalise δa with the sum

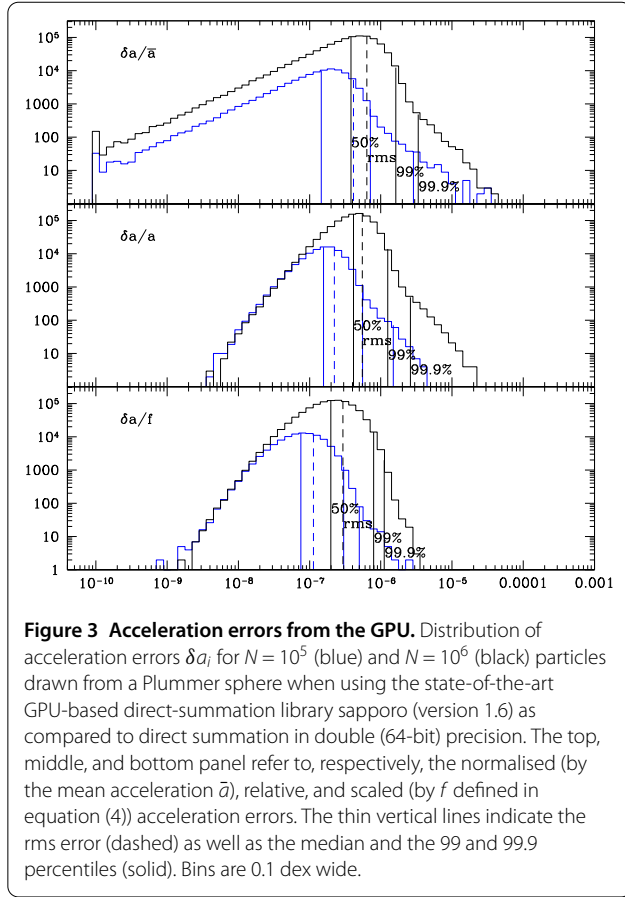
$$f_b \equiv \sum_{a \neq b} \frac{G\mu_a}{|\mathbf{x}_a - \mathbf{x}_b|^2} \quad (4)$$

of the absolute values of all pair-wise accelerations. In general $f_b \geq a_b$, while in the outskirts of a stellar system $f \rightarrow a \approx GM/r^2$ such that the *scaled error* $\delta a/f$ approaches the relative error $\delta a/a$ as desired. Conversely, in the centre $f \gg a$ (for a Plummer sphere, for example, $f \rightarrow 2GM/r_s^2$ as $r \rightarrow 0$ in the continuum limit) and $\delta a/f$ behaves sensibly if $a \rightarrow 0$.

3.2 The acceleration errors of direct summation

In order to assess the errors currently tolerated in collisional N -body simulations, the GPU-based direct-summation library *sapporo* (Gaburov et al. 2009) was applied to two sets of, respectively, $N = 10^5$ and $N = 10^6$ equal-mass particles, drawn randomly from a (Plummer 1911) sphere (without any outer truncation). Figure 3 shows the resulting distributions of acceleration errors as compared to direct summation in double (64-bit) precision. As expected, the typical relative (or scaled) error is $\sim 10^{-7}$, comparable to the relative round-off error of single-precision floating-point arithmetic. However, there is a clear tail of large relative errors (middle panel). This is due to particles at small radii, whose acceleration is small, because the pair-wise forces with other particles mostly cancel out, while the (round-off) errors accumulate.

There is a significant increase in the error amplitude with particle number N : the errors for $N = 10^6$ are on average $\sim \sqrt{10}$ larger than for $N = 10^5$. This worrying property suggests that the fidelity of simulations using *sapporo*



diminishes with N , implying that using this library with $N \gtrsim 10^7$ is not advisable.

From this exercise I conclude that in practice relative (or scaled) acceleration errors with an rms value of a few 10^{-7} and maximum ~ 10 times larger are accepted in N -body simulations of collisional stellar dynamics.

4 Assessing the approximation errors

In order to optimise any implementation of FMM for high accuracy and low computational costs, a good understanding of and accurate estimates for the errors incurred by each individual FMM interaction are required. To this end, I now perform some numerical experiments.

I create a Plummer sphere of $N = 10^6$ particles and build an oct-tree. For each cell, the centre \mathbf{z}_{ses} of the smallest enclosing sphere for all its particles is found (see Section 5.1.1). I use $\mathbf{z} = \mathbf{s} = \mathbf{z}_{\text{ses}}$ for each cell and pre-compute the cells' multipole moments $\mathcal{M}_n^m(\mathbf{z})$. Finally, the dual tree walk is performed using the multipole-acceptance criterion

$$\theta < \theta_{\text{crit}} \quad (5)$$

with the *opening angle*

$$\theta \equiv (\rho_{\mathbf{z},A} + \rho_{\mathbf{s},B})/r, \quad r \equiv |\mathbf{s}_B - \mathbf{z}_A|. \quad (6)$$

Here, for each cell C

$$\rho_{\mathbf{z},C} \equiv \max_{\text{sources } a \in C} \{|\mathbf{x}_a - \mathbf{z}_C|\} \quad \text{and} \quad (7)$$

$$\rho_{\mathbf{s},C} \equiv \max_{\text{sinks } b \in C} \{|\mathbf{x}_b - \mathbf{s}_C|\} \quad (8)$$

are (approximations for) the radii of the smallest spheres centred on \mathbf{z} and \mathbf{s} and containing all sources and sinks, respectively. In the experiments of this section $\rho_{\mathbf{z}} = \rho_{\mathbf{s}}$ for each cell, because $\mathbf{z} = \mathbf{s}$ and because all particles are source and sink simultaneously, but in general $\rho_{\mathbf{z}}$ and $\rho_{\mathbf{s}}$ may differ.

With the simple criterion (5) the multipole expansion is guaranteed to converge and have bounded errors.^c Cell \rightarrow cell interactions with $N_A N_B < p^3$, cell \rightarrow particle interactions with $N_C < 4p^2$, and particle \rightarrow cell interactions with $N_C < p^2$ are ignored, because direct summation is faster than FMM and will be preferred in a practical application. For the remaining well-separated interactions, the accelerations of all particles within the sink cell and due to all particles within the source cell are calculated in 64-bit precision using both FMM and direct summation. I then evaluate for each sink particle the acceleration error

$$\delta a \equiv |\mathbf{a}_{\text{fmm}} - \mathbf{a}_{\text{true}}| \quad (9)$$

with \mathbf{a}_{true} obtained by direct summation.

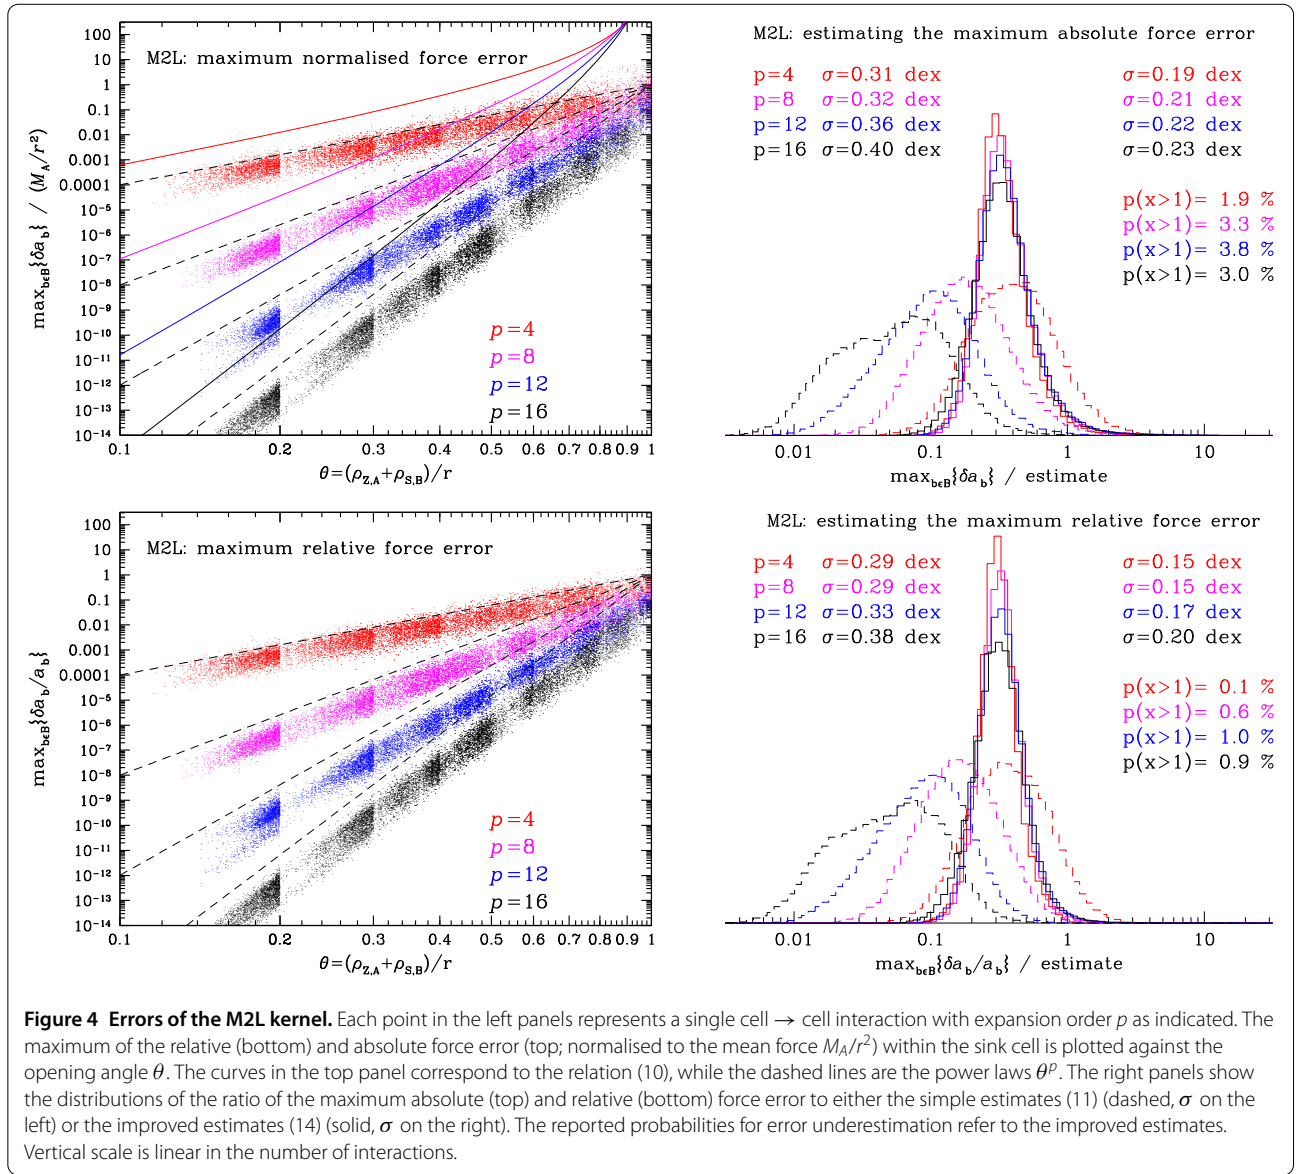
4.1 Cell-cell interactions

Cell-cell interactions involve the M2L kernel of the P2M + [M2M] + M2L + [L2L] + L2P chain of kernels. They are by far the most common and most important of all interactions encountered in the dual tree walk. For a random subset of cell-cell interactions generated by my experiments, the top panel of Figure 4 plots the maximum (over all particles within the sink cell) of δa normalised by the average acceleration M_A/r^2 against θ , while the bottom panel plots the maximum relative force error $\delta a/a$. As expected, the errors decrease with smaller θ and increasing p , though there is substantial scatter at any given θ and p . At $\theta \sim 1$, the expansion order has little effect on the errors, implying that $\theta \ll 1$ is required for small errors.

4.1.1 Comparing with simple error estimates

The approximation error from a single FMM interaction with $\theta < 1$ has the theoretical strict upper bound (Dehnen 2002)

$$\delta a \leq \frac{M_A}{r^2} \frac{(p+1)\theta^p}{(1-\theta)^2}, \quad (10)$$



which is plotted as thin curves in the top panel of Figure 4. Obviously, this upper bound is satisfied, but typically it is 10-100 times larger than the actual largest error.

Moreover, equation (10) predicts diverging errors for $\theta \rightarrow 1$, while the actual errors behave much nicer. This is presumably because diverging errors only occur for rare sink positions combined with extreme source distributions (such as all particles concentrated near one point at the edge of the source sphere), which are not realised in these experiments.

Figure 4 also shows as dashed lines the simple power laws θ^p , which give closer, though not strict, bounds

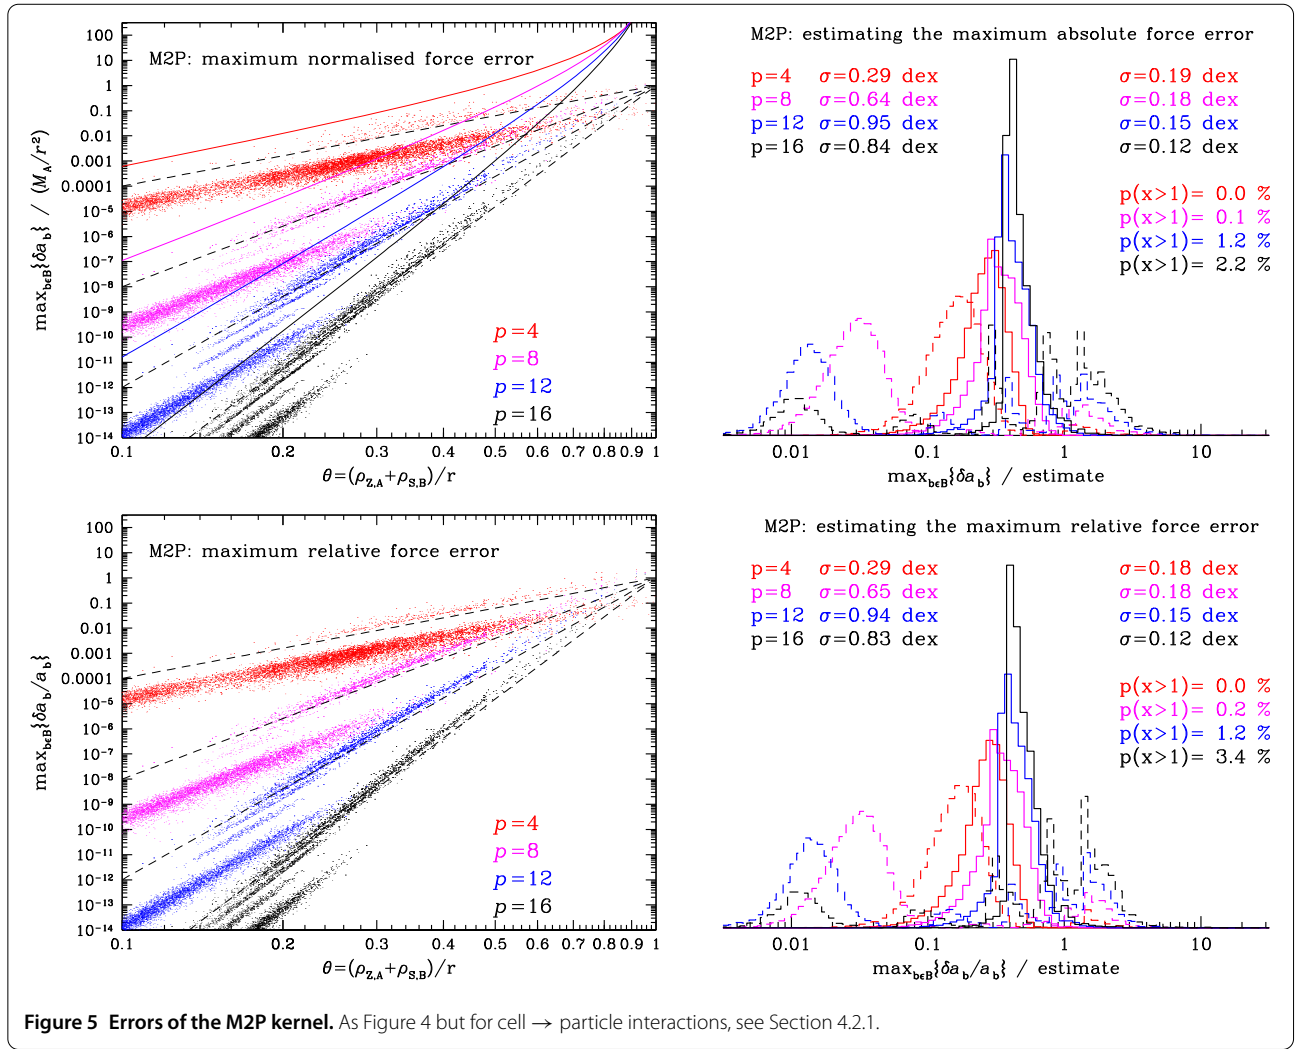
$$\delta a \lesssim \theta^p M_A / r^2 \quad \text{and} \quad \delta a / a \lesssim \theta^p \quad (11)$$

to the actual errors.

4.1.2 Better error estimates

The simple error estimates (11) are still quite inaccurate: the maximal error is often much smaller (see also the dashed histograms in the right panels of Figure 4). The offsets of θ^p from the actual errors increase with p (see left panels of Figure 4). This effect vanishes if the same limit for $N_A N_B$ is used for all p , suggesting that it is caused by smoother distributions for larger numbers N_A of sources. Indeed, if I simply divide the estimates (11) by $\sqrt{N_A}$ the scatter of the residuals is much reduced, but a systematic trend with p remains.

However, there is more information about the distribution of sources than merely their number: their multipole moments \mathcal{M}_n^m for $n \leq p$. In order to incorporate this information into an error estimate, I first compute for each



cell the multipole power

$$\mathcal{P}_n^2 \equiv \sum_{m=-n}^n (n-m)!(n+m)! |\mathcal{M}_n^m(\mathbf{z})|^2. \quad (12)$$

By design these (i) satisfy $\mathcal{P}_{n,A} \leq M_A \rho_{z,A}^n$ for any distribution of sources; (ii) are invariant under rotation (of the coordinate system) and hence independent of the interaction direction; and (iii) provide an upper bound for the amplitude of the multipole: $|\mathcal{M}_n^m(\mathbf{z})| \leq \mathcal{P}_n/n!$. Having computed \mathcal{P}_n for each source cell, one can evaluate

$$E_{A \rightarrow B} \equiv \frac{1}{M_A} \sum_{k=0}^p \binom{p}{k} \frac{\mathcal{P}_{k,A} \rho_{s,B}^{p-k}}{r^p} \quad (13)$$

with $\mathcal{O}(p)$ operations. Note that $E_{A \rightarrow B} \leq \theta^p$ with equality only for $\mathcal{P}_{n,A} = M_A \rho_{z,A}^n$. The new error estimates are

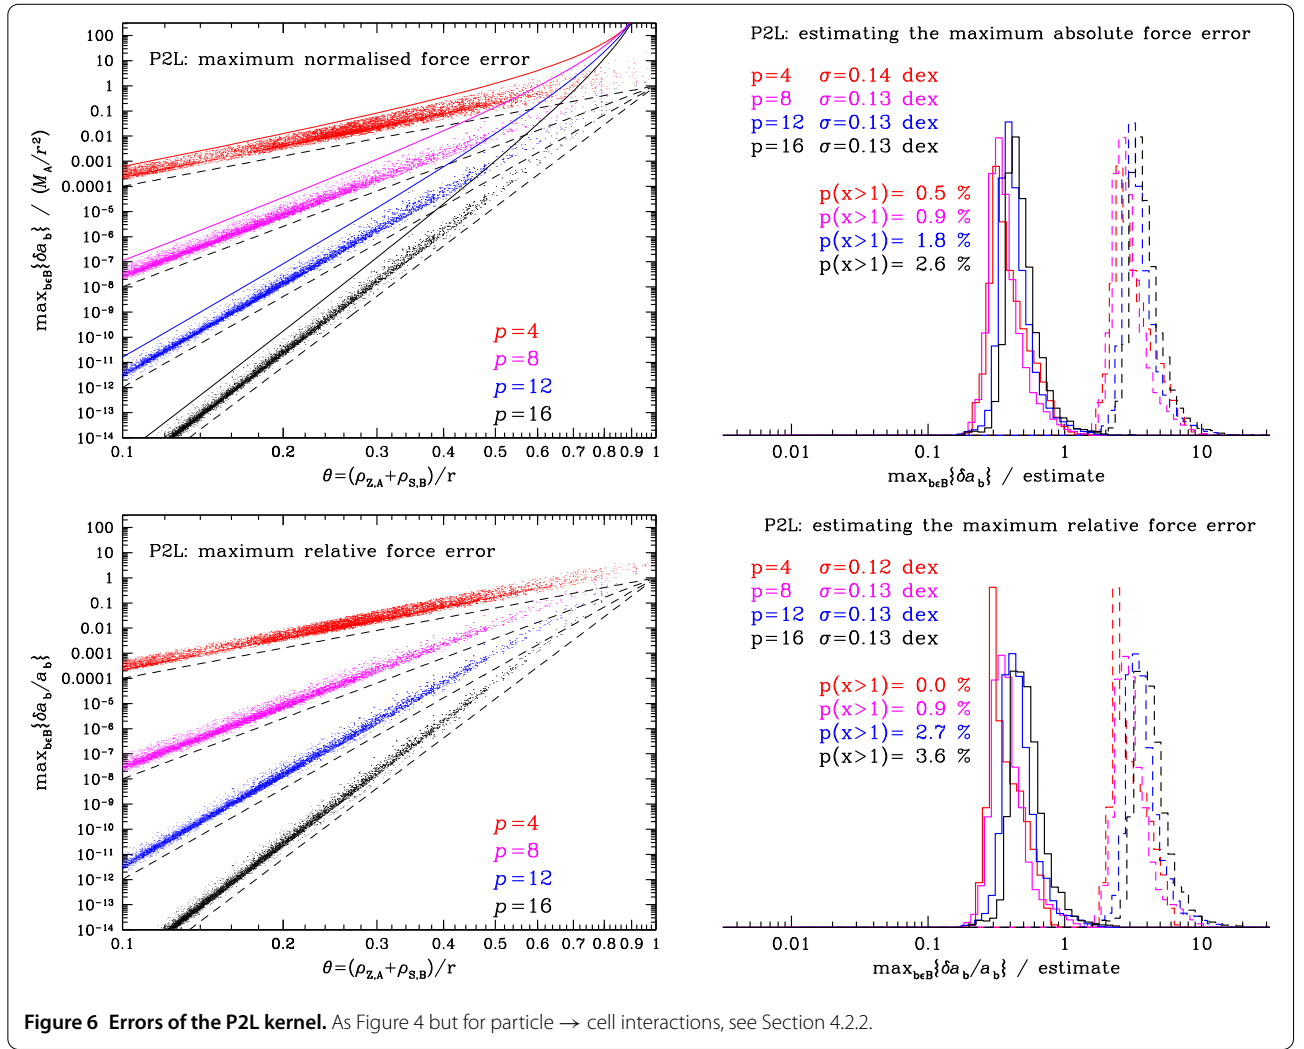
then

$$\delta a \lesssim \tilde{E}_{A \rightarrow B} \frac{M_A}{r^2} \quad \text{and} \quad \frac{\delta a}{a} \lesssim \tilde{E}_{A \rightarrow B} \quad (14)$$

with

$$\tilde{E}_{A \rightarrow B} = \frac{8 \max\{\rho_{z,A}, \rho_{s,B}\}}{\rho_{z,A} + \rho_{s,B}} E_{A \rightarrow B}. \quad (15)$$

In the right panels of Figure 4, these new error estimates are compared with the simple estimates (11) of the last subsection by displaying the distributions of the ratio of the actual maximum error to these estimates. The main difference between the two sets of estimators is their accuracy: there is much less scatter for the new (solid histograms) than for the old estimators (dashed). Consequently, there are hardly any interactions for which the force error is overestimated by more than a factor ten, while the simple estimators (11) overestimated the force error by more



than that for many interactions, in particular at large p . Another remarkable property of the new error estimator is its consistency with respect to expansion order: there is no systematic drift with expansion order.

The number of underestimated force errors (abscissa > 1 in the right panels of Figure 4) is small but there is a clear tail of underestimated absolute errors (top panel). As this is not present for the relative errors, it must be caused by the deviation of the acceleration from the mean M_A/r^2 . Indeed, the maximum error is expected to occur on the side of the sink towards the source, where the acceleration is larger, about $M_A/(r - \rho_{s,B})^2$. When accounting for this by simply replacing r in (14) with $r - \rho_{s,B}$, the tail of underestimated force errors is diminished, but the overall distributions widens and a tail of overestimated errors appears.

4.2 Particle-cell interactions

Just occasionally, the dual tree walk algorithm encounters particle-cell interactions. Most of them will be computed

using direct summation, leaving only the few with populous cells for the FMM approximation.

For particle \rightarrow cell and cell \rightarrow particle interactions the FMM approximation uses the P2L and M2P kernels, respectively. Because these kernels correspond to the M2L kernel in the limits of $\rho_{z,A} \rightarrow 0$ and $\rho_{s,B} \rightarrow 0$, respectively, all the algebra developed in the previous sub-section still applies.

4.2.1 Cell \rightarrow particle interactions

Figure 5 is equivalent to Figure 4 for cell \rightarrow particle interactions (which dominate in the tree code). The most notable difference to Figure 4 is the streaky nature of the relations in the left panels, implying a multi-modal distribution of errors at any given θ and p , as also evident from the dashed histograms in the right panels. The cause for this is simply that in an oct-tree cell size is quantised. In fact, the improved error estimates (14) account for this ef-

fect resulting in narrow mono-modal distributions of error offsets.

4.2.2 Particle \rightarrow cell interactions

Figure 6 is equivalent to Figure 4 for particle \rightarrow cell interactions. Clearly, at any given θ and p , the errors are larger than for any other type of interactions and are in fact approaching the theoretical limit (solid curves in the top left panel). What is more, not much can be done about this in terms of error estimates: since the source is just a particle without inner structure, the improved estimates (14) are simply a rescaling by a factor 8 from the simple power laws (a simple shift between the dashed and solid histograms in the right panels). They are nonetheless equally accurate as for the cell \rightarrow cell interactions and suffer from a similar level of force underestimation (for a few percent of interactions and by less than a factor two).

5 Optimising the multipole-acceptance criterion

With the improved error estimates in hand, the practical implementation of FMM for high accuracy can finally be considered. The main questions arising in this context are:

- What to pick for the expansion centres \mathbf{z} and \mathbf{s} ?
- When to consider two cells well-separated?
- What expansion order p to use?

The possible answers to these questions affect both the computational cost and the approximation accuracy. Hence, for a given accuracy target, there exists an optimal choice for all these parameters, in the sense of minimal CPU time (and memory) consumption. This section deals with the algorithmic aspects of this problem, i.e. the choice for \mathbf{z} and \mathbf{s} and the functional form of the multipole-acceptance criterion. The tuning of the parameters (of the multipole-acceptance criterion as well as the expansion order) with the aim of minimal computational effort for a given accuracy is the subject of the next section.

Astonishingly, this issue of optimal choice for \mathbf{z} and \mathbf{s} and the multipole-acceptance criterion has not been much investigated. Instead, implementations of multipole methods often employ either of two simple strategies. The tree code generally uses a fixed order p and an expansion centred on the cells' centres of mass, while two cells are considered well-separated if the simple geometric multipole-acceptance criterion (5) is satisfied, such that θ_{crit} controls the accuracy.

With traditional FMM, on the other hand, the expansion centres \mathbf{z} and \mathbf{s} are both taken to be the geometric cell centres and two cells are deemed well-separated as soon as the expansion converges, corresponding to $\theta_{\text{crit}} = 1$. When using hierarchical cubic grids (instead of an adaptive tree), this is implemented by interacting only between non-neighbouring cells on the same grid level whose parent

cells are neighbours (e.g. Cheng et al. 1999). The accuracy is then only controlled by the expansion order p .

5.1 Choice of expansion centres \mathbf{z} and \mathbf{s}

As far as I am aware, all existing FMM implementations use the same position for the multipole and potential expansion centres, i.e. $\mathbf{z} = \mathbf{s}$, for each cell. For traditional FMM, these are equal to the geometric cell centres. This has the benefit of a finite number of possible interaction directions $\hat{\mathbf{r}}$, in particular when $\theta_{\text{crit}} = 1$, for which the coefficients $\Theta_n^m(\hat{\mathbf{r}})$ could be pre-computed. However, the computation of these coefficients on the fly is often faster than a table look-up. Moreover, in view of Figure 4 $\theta_{\text{crit}} = 1$ appears ill-suited for high accuracy.

In fact, the restriction $\mathbf{z} = \mathbf{s}$ reduces the freedom and hence the potential for optimising the method. Nonetheless, when aiming for low accuracy, choosing $\mathbf{z} = \mathbf{s} = \mathbf{z}_{\text{com}}$, the cells' centres of mass, has some advantages. First, the dipoles vanish and the low-order multipoles tend to be near-minimal. Second, if using a mutual version of the algorithm (when the interactions $A \rightarrow B$ and $B \rightarrow A$ are done simultaneously), the computational costs are reduced and the approximated forces satisfy Newton's third law exactly, i.e. $\mathbf{F}_{ab} + \mathbf{F}_{ba} = 0$ (Dehnen 2002).

However, in practice there is no benefit from such an exact obedience of Newton's law, as the total momentum is not exactly conserved, because of integration errors arising from the fact that the particles have individual time steps. Moreover, the degree of deviation from exact momentum conservation in such a case does not reflect the true accumulated force errors. In a more general method, the approximated forces will deviate from the ideal $\mathbf{F}_{ab} + \mathbf{F}_{ba} = 0$ by an amount comparable to their actual force errors and the non-conservation of total momentum is somewhat indicative of the accumulated effect of the force errors (see also Appendix B.3).

5.1.1 Choice of the potential expansion centre \mathbf{s}

The results of Section 4, in particular the functional form of $E_{A \rightarrow B}$ in equation (13), suggest to choose the potential expansion centres \mathbf{s} such that the resulting sink radii ρ_s , and hence the estimated interaction errors, are minimal. Thus, $\mathbf{s} = \mathbf{z}_{\text{ses}}$, the centre of the *smallest enclosing sphere*. Finding the smallest enclosing sphere for a set of n points has complexity $\mathcal{O}(n)$. Doing this for every sink cell would incur a total cost of $\mathcal{O}(N \ln N)$ and be prohibitively expensive.

Instead, I use an accurate approximation by finding for each cell the smallest sphere enclosing the spheres of its grand-daughter cells. This incurs a total cost of $\mathcal{O}(N)$ and is implemented via the Computational Geometry Algorithms Library (www.cgal.org, Fischer et al. 2013), using an algorithm of Matoušek et al. (1996).

5.1.2 Choice of the multipole expansion centre \mathbf{z}

As already mentioned above, setting $\mathbf{z} = \mathbf{z}_{\text{com}}$ has some virtue for low expansion orders p . However, for high expansion orders, the high-order multipoles become ever more important, suggesting that $\mathbf{z} = \mathbf{z}_{\text{ses}}$ may be a better choice. In order to assess the relative merits of these methods, I repeated the experiments of Section 4 for both methods and compared the resulting maximum absolute and relative force errors incurred for the *same* cell \rightarrow cell interactions (for which the two methods give different θ).

I found that the errors for the two methods are very similar with an rms deviation of ~ 0.15 dex, but a very small mean deviation. At $p \lesssim 8$ there is a trend of more accurate forces for $\mathbf{z} = \mathbf{z}_{\text{com}}$, while at $p \gtrsim 8$ smaller errors are obtained with $\mathbf{z} = \mathbf{z}_{\text{ses}}$. This trend is simply a consequence of \mathcal{P}_k being smaller for $\mathbf{z} = \mathbf{z}_{\text{com}}$ than for $\mathbf{z} = \mathbf{z}_{\text{ses}}$ at low k and larger at high k . This together with the improved error estimates (14) also explains that (for an interaction $A \rightarrow B$) $\mathbf{z} = \mathbf{z}_{\text{com}}$ tends to give more accurate forces if $\rho_{\mathbf{z},A} < \rho_{\mathbf{z},B}$, while $\mathbf{z} = \mathbf{z}_{\text{ses}}$ tends to be more accurate if $\rho_{\mathbf{z},A} > \rho_{\mathbf{z},B}$.

5.2 A simple FMM implementation

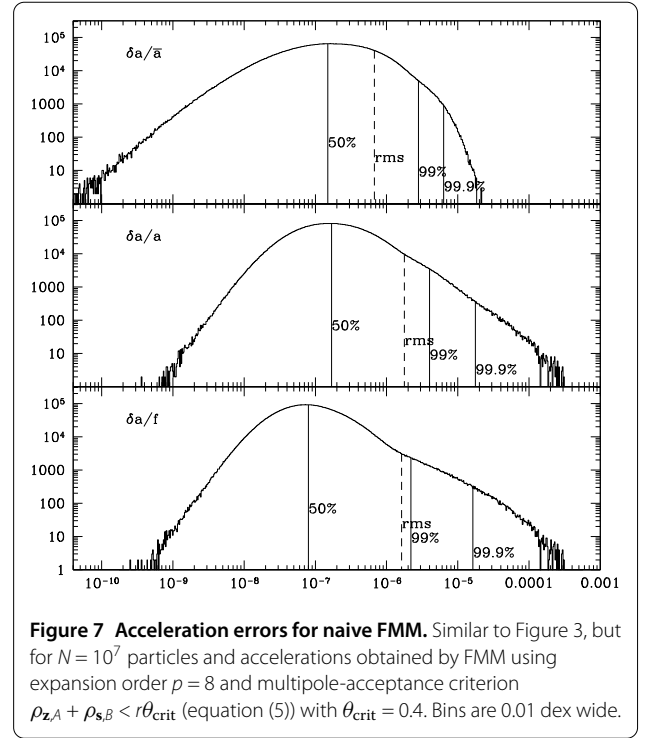
Let us first experiment with an implementation that uses the simple multipole-acceptance criterion (5) and a fixed expansion order p . This is the standard choice for the tree code and as such implemented in many gravity solvers used in astrophysics. The computational costs of such an implementation roughly scale as $p^\alpha / \theta_{\text{crit}}^3$ with $\alpha \sim 2.3$, because the number of interactions increases as $\theta_{\text{crit}}^{-3}$ for large N , while the cost for one is $\propto p^{2.3}$. Together with the simple error estimate (11), this means that *if* one aims each FMM interaction to satisfy $\delta a/a < \epsilon$, then the minimum cost for fixed ϵ occurs for

$$\theta_{\text{crit}} = e^{-\alpha/3} \approx 0.46.$$

Thus, the optimal opening angle is independent of p . The accuracy is then controlled by the expansion order, requiring $p \gtrsim 16$ for $\epsilon < 10^{-8}$ (according to Figure 4). The computational costs rise roughly like $|\ln \epsilon|^\alpha$ with decreasing ϵ .

I applied the FMM method with $\mathbf{z} = \mathbf{z}_{\text{ses}}$, expansion order $p = 8$, and $\theta_{\text{crit}} = 0.4$ to $N = 10^7$ equal-mass particles drawn from a Plummer sphere. Figure 7 plots the resulting distributions of absolute (top), relative (middle), and scaled (bottom) acceleration errors. All three distributions are mono-modal, but very wide, much wider than those obtain from GPU-based direct summation (Figure 3). In particular, there are extended tails towards very large relative or scaled errors, containing only $\lesssim 1\%$ of the particles but reaching up to 1000 times the median error. These tails are due to particles at large radii and, for the relative errors only, also at small radii (see the discussion in Section 3.1).

There are two main effects responsible for these properties of the error distributions. First, errors from a single FMM interaction follow a distribution with variance of



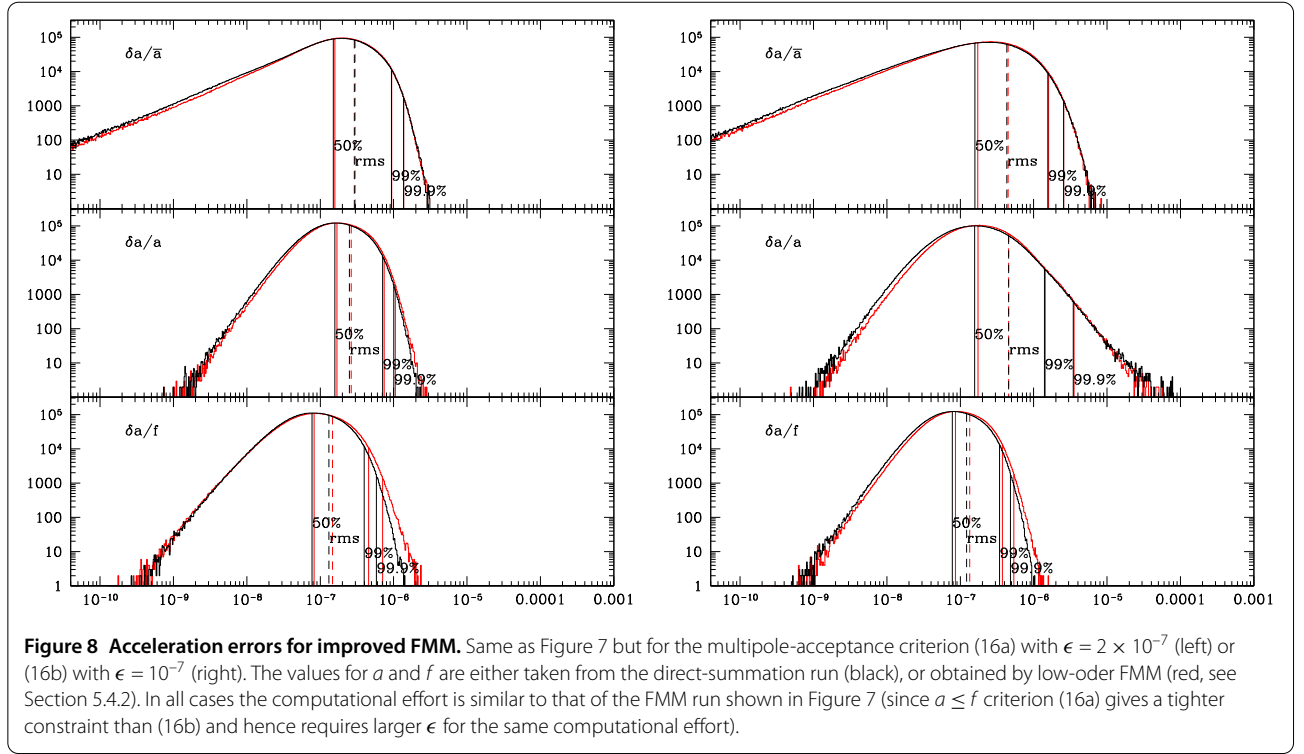
1-2 dex. The maximum errors reported in Section 4 only occur for particles near the edges and corners of the sink cell, while most have smaller errors. Moreover, the force errors due to FMM interactions of the same sink cell with source cells in opposing directions tend to partially cancel rather than add up. Both explain why the median errors reported in Figure 7 are much smaller than the maximum relative error incurred by a single cell \rightarrow cell interaction, which according to Figure 4 is $\sim 10^{-4}$.

More important is a second effect: the final force errors are not the sum of the relative errors of individual FMM interactions, which are controlled by the simple multipole-acceptance criterion, but of their *absolute* errors δa . Since, according to equation (11), $\delta a \sim \theta^p M_A / r^2 \sim \theta^{p+2} M_A / 4 \rho_{\mathbf{z},A}^2$, the FMM interactions with cells of large surface density $M/\rho_{\mathbf{z}}^2$ dominate the error budget. In fact, the particles at very large radii have $\delta a/a \approx \delta a/f \sim 10^{-4}$, exactly as expected from a few FMM interactions with near maximal errors.

5.3 Towards better multipole-acceptance criteria

This discussion suggests that multipole-acceptance criteria which balance the absolute force errors of individual FMM interactions are preferable. When working with the simple estimators (11) or the error bound (10), this leads to critical opening angles which depend on the properties of the interacting cells, such as their mass or surface density.

Such an approach can indeed be made to work (Dehnen 2002), but the aim here is to go beyond that and use the im-



proved error estimates (14). This results in the multipole-acceptance criteria

$$\theta < 1 \wedge \tilde{E}_{A \rightarrow B} \frac{M_A}{r^2} < \epsilon \min_{b \in B} \{a_b\}, \quad (16a)$$

$$\theta < 1 \wedge \tilde{E}_{A \rightarrow B} \frac{M_A}{r^2} < \epsilon \min_{b \in B} \{f_b\} \quad (16b)$$

with the aim to obtain $\delta a/a \lesssim \epsilon$ and $\delta a/f \lesssim \epsilon$, respectively.

The black histograms in Figure 8 show the error distributions resulting from these criteria, when the values for a_b and f_b used in equation (16) have been taken from the direct-summation comparison run. The distributions for $\delta a/a$ in the left and $\delta a/f$ in both panels are remarkably narrow with a median error $\sim \epsilon$ as targeted, a steep truncation towards large errors, and a maximum error $\sim 10\epsilon$. The tail of large $\delta a/a$ in the right panel is due to particles at small radii, for which $a \ll f$ such that criterion (16b) allows large $\delta a/a$.

The difference between these error distributions and those shown in Figure 7 and resulting from the simple geometric multipole-acceptance criterion (5) is remarkable. While the median errors are comparable, the criteria (16) do not produce extended tails of large errors of the quantity controlled ($\delta a/a$ in left and $\delta a/f$ in the right panels of Figure 8), and the maximum errors are more than 2 orders of magnitude smaller. What is more, the tails towards small errors have also been somewhat reduced, indicating

that the improved criterion avoids overly accurate individual FMM interactions.

This improvement has been achieved without increasing the overall computational effort, but by carefully considering the error contribution from each approximated interaction.

5.4 Practical multipole-acceptance criteria

In a real application one has, of course, no a priori knowledge of a_b or f_b for any particle and must instead use something else in the multipole-acceptance criteria (16). In some situations, a suitable scale can be gleaned from the properties of the system modelled. For example, if simulating a star cluster of known mass profile $M(r)$ and centre \mathbf{x}_0 , one may simply use $a_b \sim GM(r_b)r_b^{-2}$ with $r_b = |\mathbf{x}_b - \mathbf{x}_0|$. I now consider other options.

5.4.1 Using accelerations from the previous time step

Employing the accelerations \mathbf{a}_b from the previous time step in equation (16a) requires no extra computations. However, it means that the gravity solver is not self-contained, but requires some starter to get the initial accelerations.

Also, using information from the previous time step subtly introduces an artificial arrow of time into the simulation, because $\delta a_{\text{new}} < \epsilon a_{\text{old}}$ implies $\delta a_{\text{new}}/a_{\text{new}} < \epsilon a_{\text{old}}/a_{\text{new}}$. Hence, a particle moving in a direction of increasing acceleration has, on average, smaller $\delta a/a$ than when moving in the opposite direction, or in reversed time. However,

the time integration methods currently employed almost exclusively in N -body simulations of collisional stellar dynamics are irreversible and introduce their own arrow of time. This suggests, that the additional breach of time symmetry by the magnitude (not the direction) of the force error may not be a serious problem in practice.^d

5.4.2 Estimating a_b or f_b using low-order FMM

As Section 4 has shown, the error estimate $\tilde{E}_{A \rightarrow B}$ used in the multipole-acceptance criteria (16) still has significant uncertainty, and using highly accurate values for a_b or f_b in equation (16) is unnecessary. Instead, rough estimates should suffice. Such estimates can be obtained via a low-order FMM. This amounts to running the FMM twice: once with a simple multipole-acceptance criterion to obtain rough estimates for a_b or f_b , and then again using the sophisticated criteria (16) employing the results of the first run.

The acceleration scale f (defined in equation (4)) is similar to the gravitational potential (1), except that its Greens function is $|\mathbf{r}|^{-2}$. This implies that it too can be estimated using FMM, albeit not using an explicitly harmonic formulation.

I implemented both options, estimating \mathbf{a} or f via FMM, using the lowest possible order ($p = 0$ for f and $p = 1$ for gravity - recall that $\mathbf{a} = \nabla \Psi$ is approximated at one order lower than the potential Ψ) and multipole-acceptance criterion $\theta < 1$. To this end, I use $\mathbf{s} = \mathbf{z} = \mathbf{z}_{\text{com}}$ and a mutual version of the dual tree walk. The resulting estimates for f or $a = |\mathbf{a}|$ have rms relative errors of $\sim 15\%$. The additional computational effort is still much smaller than that of the high-accuracy approximation of gravity itself, though estimating f is faster because it is a scalar rather than a vector and because no square-root needs to be calculated.

The distributions of acceleration errors resulting from using these estimates in equation (16) are shown in red in Figure 8. They are only very slightly worse than those in black, which have been obtained using the exact values of a_b and f_b in equation (16).

6 Optimising adaptive FMM

The previous section provided answers to the first two questions asked at its beginning, but not to the one after the optimal expansion order p . To answer this question I now report on some experiments, which also provide the actual computational costs for a given required force accuracy.

All experiments are run on a single compute node with 16 Intel Xeon E5-2670 CPUs, which support the AVX instruction set (see below), and using code generated by the gcc compiler (version 4.8.2).

6.1 Implementation details

The FMM relations of Section 2 and Appendix A (using the rotation-accelerated M2L kernel of Appendix A.6

when faster) have been implemented in computer code. The code employs a one-sided version of the dual tree walk, which considers the interactions $A \rightarrow B$ and $B \rightarrow A$ independently. The code is written in the C++ programming language and has been tested using various compilers and hardware. The implementation employs vectorisation and shared-memory parallelism as outlined below.

6.1.1 Vectorisation

Most current CPUs support vector sizes of 16 (SSE), 32 (AVX), or 64 (MIC) bytes, allowing $K = 2, 4$, or 8 identical simultaneous double-precision floating-point operations (or twice as many in single precision). Because the FMM kernels do not (usually) relate adjacent elements, their efficient vectorisation is not straightforward (and well beyond compiler optimisation). I explicitly implement a method computing K M2L kernels simultaneously. To this end, the multipole moments of the K source cells are loaded into a properly aligned buffer (similar to transposing a matrix) before, and afterwards the K field tensors are added from their vector-buffer to the sink cells' field tensors. Unfortunately, this loading and storing (which cannot be vectorised) reduces the speed-up obtained by the simultaneous kernel computations.

Conversely, direct summation is perfectly suitable for vectorisation and a speed-up of a factor K is achievable. The code prefers direct summation whenever this is deemed to be faster, based on a threshold for the number of particle-particle interactions 'caught' in a given cell-cell interaction.

6.1.2 Multi-threading

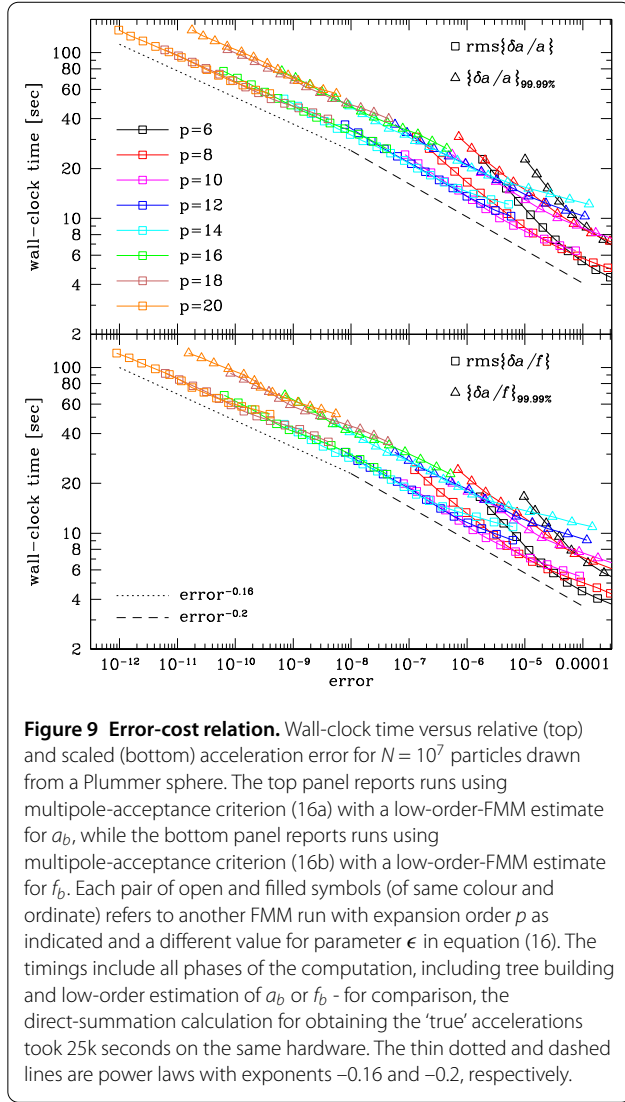
All parts of the implementation use multi-threading and benefit from multi-core architectures. This is done via hierarchical task-based parallelism implemented via *threading building blocks* (tbb, Reinders 2007), an open source task parallel library with a work-stealing scheduler. The algorithms for multi-threaded tree building and dual tree walk are quite similar to those described by Taura et al. (2012) and I refrain from giving details here.

6.1.3 Precision and expansion order

This study reports only on one particular implementation aimed at high accuracy. It uses double precision (64 bits) floating-point arithmetic throughout, $\mathbf{z} = \mathbf{z}_{\text{ses}}$, and expansion orders $p \leq 20$.

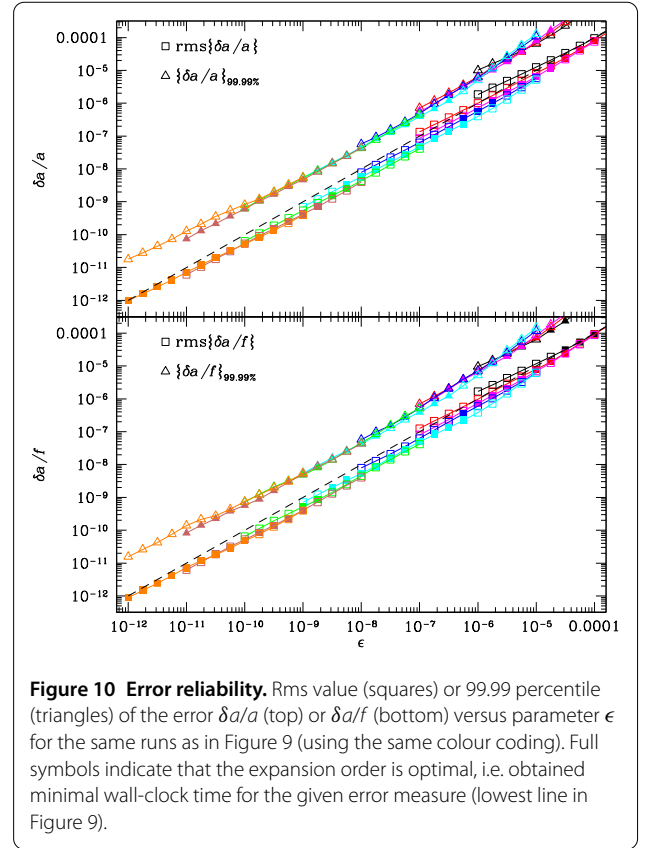
6.2 Wall-clock time versus accuracy

I applied my implementation with criteria (16a) and (16b) to $N = 10^7$ particles drawn from a Plummer sphere, and using low-order estimates for a_b and f_b in equation (16). I varied the expansion order p and the accuracy parameter ϵ and for each run plot in Figure 9 the total wall-clock time against the rms and the 99.99 percentile acceleration errors.



The rms error is always ten times smaller than the 99.99 percentile,^e implying the absence of extended large-error tails. For any fixed expansion order p , the relation between time and error can be approximated by a constant plus a power law that becomes flatter for larger p . At any given error, there is an optimal expansion order p in the sense of providing the fastest approximation. When using this optimal expansion order, the fastest FMM computation for a given error scales very nearly like a power law with exponent ~ -0.18 . Thus when reducing the error by a factor ten, the computational costs rise only by a factor ~ 1.5 .

Constraining the relative error (top panel of Figure 9) is slightly more costly than constraining the scaled error (bottom panel). This is largely because $f > a$ as discussed in the caption to Figure 8, but also because estimating f is easier and faster than estimating a . Of course, the estimation of a can be easily avoided in practice by using the accelerations from the previous time step.



6.3 Accuracy versus parameter ϵ

In any practical application there is, of course, no possibility to check on the actual error, so it is important to test how well it is reflected by the parameter ϵ . As can be seen from Figure 10, the rms value for the respective error ($\delta a/a$ if using criterion (16a) and $\delta a/f$ if using criterion (16b)) is typically slightly less than ϵ for the optimal expansion order p . At intermediate values ($\epsilon \sim 10^{-8}$) the error is actually a factor ~ 2 smaller. The 99.99 percentile of the errors is typically a factor ten larger.

6.4 Complexity: scaling with the number N of particles

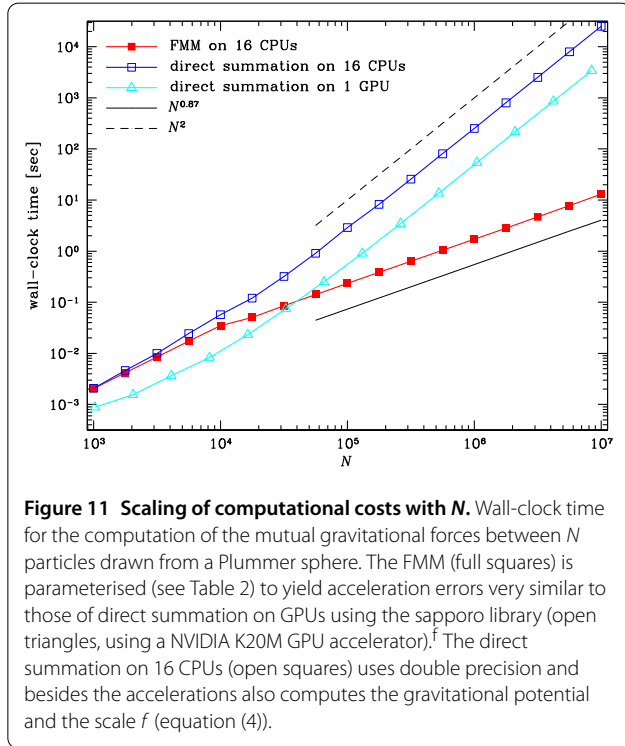
The overall cost of my high-accuracy FMM implementation is dominated by the computation of all node-node interactions during the dual tree walk. All other phases (establishing the hierarchical tree structure, computing \mathbf{z} , \mathbf{s} , and \mathcal{M}_n^m for each cell; passing down \mathcal{F}_n^m and evaluating gravity for each sink position) contribute much less (see Table 2). When using a simple geometric multipole-acceptance criterion, such as equation (5), the FMM is well known to have complexity $\mathcal{O}(N)$ (e.g. Cheng et al. 1999). This is because distant interactions contribute less than $\mathcal{O}(N)$, so that the overall costs are dominated by the local interactions only (Dehnen 2002).

I am not aware of theoretical estimates for the complexity for the case of more sophisticated multipole-acceptance

Table 2 Timings and errors for FMM runs for different N

N	t_{tree}	t_{fmm}	t_{up}	t_{walk}	t_{down}	$\{\delta a/f\}_{\text{rms}}$	$\{\delta a/f\}_{99.99\%}$
10^4	0.0004	0.0011	0.0012	0.0317	0.0005	2.81×10^{-7}	1.32×10^{-6}
10^5	0.0019	0.0042	0.0097	0.2131	0.0044	3.61×10^{-7}	2.51×10^{-6}
10^6	0.0290	0.0245	0.0957	1.5366	0.0431	3.85×10^{-7}	3.32×10^{-6}
10^7	0.2559	0.1897	0.8870	11.358	0.3999	4.05×10^{-7}	2.20×10^{-6}


The runs used $p = 10$ and $\epsilon = 10^{-6.25}$. The timings are given in seconds and refer to, respectively, the tree building; the estimation of f via low-order FMM; the passing up of $z, s, \min(f)$, and \mathcal{M}_n^m ; the dual tree walk; and the passing down of \mathcal{F}_n^m and evaluation of gravity. See also Figure 11.



criteria, but Dehnen (2002) reports an empirical scaling proportional to $N^{0.93}$ for his approach of a mass-dependent opening angle. Table 2 and Figure 11 present the timings obtained with my implementation using $p = 10$, $\epsilon = 10^{-6.25}$, and low-order FMM estimates of f_b in equation (16b). With these settings, the acceleration errors are comparable to those generated via the sapporo library on a GPU (the current state-of-the-art force solver for collisional N -body simulations), as reported in Section 3.2.

From Table 2, it can be seen that the costs for tree building grow faster than linearly with N ($N \ln N$ is expected), those for the upward and downward passes roughly linearly with N (as expected), but those for the FMM estimation of f and the dual tree walk less than linearly. As a result, the total computational costs are very well fit by the power law $N^{0.87}$ for $N > 10^4$, see Figure 11.

Figure 11 also shows the timings for a (double-precision) direct-summation on the same hardware (yielding much more accurate accelerations) and for a mixed-precision



direct-summation on a GPU using the sapporo library (yielding comparably accurate accelerations). At large (but realistic) N FMM out-performs direct summation, even if accelerated using a GPU.

6.5 Scaling with the number of CPUs

Figure 12 plots the strong scaling factor t_1/nt_n for my multi-threaded implementation. The scaling drops to 80% for 16 cores, which is not untypical for multi-threaded programs. This drop is presumably caused by imbalances at synchronisation points, of which the implementation has many. Most of these are not algorithmically required, but allow for a much easier implementation. Clearly, any massively parallel implementation needs to address this issue to retain good scaling for large numbers of processors.

7 Beyond simple gravity approximation

So far, I have considered the approximate computation of the unsoftened gravitational potential and acceleration at all particle positions with equal relative (or scaled) accuracy. However, the fast multipole method can be easily modified or extended beyond that.

For example, one may want to have individual accuracy parameters ϵ_b instead of a global one. This is easily accommodated by replacing $\epsilon \min_{b \in B} \{a_b\}$ in criterion (16a) with $\min_{b \in B} \{\epsilon_b a_b\}$ and analogously for criterion (16b).

When using individual ϵ_b , but also in general, it may be beneficial to adapt the expansion order p to the accuracy actually required for a given cell \rightarrow cell interaction. This could be implemented by using the lowest $p \leq p_{\max}$ for which the multipole-acceptance criterion is satisfied.

7.1 Force computation for a subset of particles

Most N -body codes employ adaptive individual time steps for each particle. The standard technique is Makino's (1991) block-step scheme, where the forces of all *active* particles are computed synchronously. Active are those particles with time step smaller than some threshold (which varies from one force computation to the next).

When using FMM in such a situation, only interactions with sink cells contain at least one active particle must be considered. If the fraction of active particles in such cells is small (but non-zero), FMM becomes much less efficient per force computation. Fortunately, however, active particles are typically spatially correlated (because the time steps of adjacent particles are similar), such that the fraction of active particles is either zero or large.

I performed some practical tests, where only particles within some distance from the origin of the system were considered active. Figure 13 plots the wall-clock time vs. the number N_a of active particles for $N = 10^7$. As expected the costs for preparation phase (tree building and upward pass) are largely independent of N_a (the slight increase of the red curve at large N_a is because \mathbf{s} and ρ_s are computed as part of the upward pass, but only for cells with active particles).




Figure 13 Computational costs for a subset of sinks. Wall-clock time for the computation of gravity for the innermost N_a of $N = 10^7$ particles (of a Plummer sphere), using the same parameters (p and ϵ) as in Figure 11.

The costs for the interaction and downward pass, on the other hand, decrease roughly like $N_a^{0.87}$ for $N_a \gtrsim 10^4$. The net effect is that for $N_a/N \lesssim 0.01$, the costs are almost completely dominated by the preparation phase, and hence independent of N_a . The precise point of this transition depends on N and the FMM parameters. For smaller N and/or more accurate forces, the relative contribution of the tree walk phase increases and the transition occurs at smaller N_a .

There is certainly some room for improvement by, e.g. using a smaller expansion order p than is optimal for $N_a = N$ and/or re-cycling the tree structure from the previous time step. Both measures reduce the costs of the preparation phase and increase that of the interaction phase (at given ϵ), but shall reduce the overall costs if $N_a \ll N$.

7.2 Softened gravity or far-field force

Gravitational softening amounts to replacing the Newtonian Greens function $\psi = |\mathbf{r}|^{-1}$ by Dehnen (2001)

$$\psi(\mathbf{r}) = h^{-1} \varphi(|\mathbf{r}|/h) \quad (17)$$

with softening length h and *softening kernel* $\varphi(q) \rightarrow q^{-1}$ as $q \rightarrow \infty$. This corresponds to replacing each source point by a smooth mass distribution with density $\mu_b \varrho(\mathbf{x} - \mathbf{x}_b)$, where

$$\varrho(\mathbf{r}) = -(4\pi)^{-1} \nabla^2 \psi(\mathbf{r}). \quad (18)$$

This Greens function (17) is no longer harmonic and harmonic FMM cannot be used. One obvious option is to use the more general Cartesian FMM of Appendix A.1 (Dehnen 2002). The computational costs of this approach grow faster with expansion order p , such that small approximation errors (requiring high p) become significantly more expensive. However, small approximation errors are hardly required in situations where gravitational softening is employed. Alternatively, if softening is restricted to a finite region, i.e. if $\varrho(\mathbf{r}) = 0$ for $|\mathbf{r}| \geq h$, harmonic FMM can still be used to compute gravity from all sources at distances $|\mathbf{r}| \geq h$, while direct summation could be used for *neighbours*, sources at $|\mathbf{r}| < h$. This approach is sensible only if the number of neighbours is sufficiently bounded (so that the cost incurred by the direction summation remains small). This is the case, in particular, if the number of neighbours is kept (nearly) constant by adapting the individual softening lengths h_i in order to adapt the numerical resolution (Price and Monaghan 2007).

In practice, this requires to carry with each cell the radius $h_z > \rho_z$ of the smallest sphere centred on \mathbf{z} which contains all softening spheres of its sources, and allow a FMM interaction $A \rightarrow B$ only if $|\mathbf{z}_A - \mathbf{s}_B| > h_{z,A} + \rho_{s,B}$.

The same technique can be used to restrict the FMM approximation to the far field for each particle, i.e. the force generated by all sources outside of a sphere of known radius h_b around \mathbf{x}_b .

7.3 Jerk, snap, crackle, and pop

The jerk is the total time derivative of the acceleration

$$\mathbf{j} = \dot{\mathbf{a}} = \frac{d}{dt} \frac{\partial \Psi}{\partial \mathbf{x}}. \quad (19)$$

The simplest way to estimate this using FMM, is to not allow the expansion centres to have any velocity ($\dot{\mathbf{z}} = \dot{\mathbf{s}} = 0$), such that differentiating the FMM relations (3) w.r.t. time gives

$$\begin{aligned} \dot{\Psi}_n^m(\mathbf{x}_b) = & \sum_{k=0}^{p-n} \sum_{l=-k}^k [\gamma_k^{l*}(\mathbf{s}_B - \mathbf{x}_b) \dot{\mathcal{F}}_{n+k}^{m+l}(\mathbf{s}_B) \\ & + \dot{\gamma}_k^{l*}(\mathbf{s}_B - \mathbf{x}_b, -\dot{\mathbf{x}}_b) \mathcal{F}_{n+k}^{m+l}(\mathbf{s}_B)] + \delta \dot{\Psi}_n^m, \end{aligned} \quad (20a)$$

$$\dot{\mathcal{F}}_n^m(\mathbf{s}_B) = \sum_{k=0}^{p-n} \sum_{l=-k}^k \mathcal{M}_k^{l*}(\mathbf{z}_A) \Theta_{n+k}^{m+l}(\mathbf{s}_B - \mathbf{z}_A), \quad (20b)$$

$$\mathcal{M}_n^m(\mathbf{z}_A) = \sum_{a \in A} \mu_a \dot{\gamma}_n^m(\mathbf{x}_a - \mathbf{z}_A, \dot{\mathbf{x}}_a), \quad (20c)$$

and the jerk follows from $\mathbf{j} = -(\Re\{\dot{\Psi}_1^1\}, \Im\{\dot{\Psi}_1^1\}, \dot{\Psi}_1^0)$. Since $\dot{\mathbf{z}} = \dot{\mathbf{s}} = 0$, the M2M and L2L kernels (equations (3d) and (3e)) work also for the time derivatives $\dot{\mathcal{M}}_n^m$ and $\dot{\mathcal{F}}_n^m$ of the multipoles and field tensors, respectively. The relations for the next order, the snap $\mathbf{s} = \ddot{\mathbf{a}}$, can be derived by differentiating yet again.

With each additional order (jerk, snap, crackle, pop, ...), the computational cost of the combined M2L kernels is not more than the corresponding multiple of the ordinary M2L kernel (i.e. acceleration plus jerk are twice as costly as just acceleration). This is a direct consequence of not allowing cell-centre velocities hence preventing the terms depending on \mathbf{z} or \mathbf{s} in equation (3) to carry any time dependence. In contrast, the computational costs of the P2M and L2P kernels grows quadratically with the order of time derivative. This is not really a problem, since those kernels are only needed once per particle, while the M2L kernel is typically used $\gtrsim 100$ times more often.

7.4 The tidal field

FMM can also be used to approximate the Hessian

$$\mathbf{T}_b = \nabla \nabla \Psi(\mathbf{x}_b) \quad (21)$$

of the potential, which is given by the components of Ψ_2^m via

$$\begin{aligned} T_{xx} &= -\frac{1}{2} \Psi_2^0 + \frac{1}{2} \Re\{\Psi_2^2\}, \\ T_{xy} &= -\frac{1}{2} \Im\{\Psi_2^2\}, \\ T_{xz} &= \Re\{\Psi_2^1\}, \\ T_{yy} &= -\frac{1}{2} \Psi_2^0 - \frac{1}{2} \Re\{\Psi_2^2\}, \\ T_{yz} &= \Im\{\Psi_2^1\}, \quad \text{and} \\ T_{zz} &= \Psi_2^0 \end{aligned} \quad (22)$$

(in particular $\text{tr}(\mathbf{T}) = 0$ as expected). Note, however, that the accuracy of this approximation is lower than that for the acceleration. \mathbf{T} is of particular interest in collisionless N -body modelling, when

$$\tau_b \leq \eta \|\mathbf{T}_b\|^{-1/4} \quad (23)$$

with dimensionless parameter $\eta \ll 1$ has been suggested as criterion for individual particle time steps (Dehnen and Read 2011). The matrix norm of \mathbf{T} may be directly computed from Ψ_2^m as

$$\|\mathbf{T}_b\| = \frac{3}{2} |\Psi_2^0|^2 + 2 |\Psi_2^1|^2 + \frac{1}{2} |\Psi_2^2|^2. \quad (24)$$

8 Discussion and conclusions

The fast multipole method (FMM) approximates the computation of the mutual forces between N particles. I have derived the relevant mathematical background, giving much simpler formulae than the existing literature, for the case of unsoftened gravity, when the harmony of the Greens function allows significant reduction of the computational complexity.

Like the tree code, my FMM implementation uses a hierarchical tree of spatial cells. Unlike the tree code, FMM uses cell \rightarrow cell interactions, which account for all interactions between sources in the first cell and sinks in the second. Almost all distant particle \rightarrow particle interactions are ‘caught’ by fewer than $\mathcal{O}(N)$ cell \rightarrow cell interactions, such that local interactions, requiring $\mathcal{O}(N)$ computations, dominate the overall workload (Dehnen 2002). With the tree code, the situation is reversed: the distant interactions require $\mathcal{O}(N \ln N)$ computations and dominate the overall work. This implies that FMM has the best complexity of all known force solvers. What is more, the predominance of local as opposed to distant interactions makes FMM ideally suited for applications on super-computers, where communications (required by distant interactions) are increasingly more costly than computations. However, FMM is inherently difficult to parallelise and this study considered only a multi-threaded implementation with a task-parallel dual tree walk (the core of FMM).

Most previous implementations of FMM considered simple choices for the cell's multipole- and force-expansion centres and the multipole-acceptance criterion which decides whether a given cell \rightarrow cell interaction shall be processed via the multipole expansion or be split into daughter interactions. Traditionally, a simple opening-angle based multipole-acceptance criterion has been used and cell centres equal to either the cell's geometric centre or its centre of mass. These choices, which presumably were based on computational convenience and intuition, inevitably result in a wide distribution of individual relative force errors with extended tails reaching ~ 1000 times the median.

The main goal of this study was avoid such extended tails of large force errors and to minimise the computational effort at a given force accuracy. The key for achieving this goal is a reasonably accurate estimate, based on the multipole power of the source cell and the size of the sink cell, for the actual force error incurred by individual cell \rightarrow cell interactions. Based on the insight from this estimate, I set the cell's force-expansion centres to (an approximation of) the centre of the smallest sphere enclosing all its particles, when the cell size and hence the error estimates are minimal. I also use the new estimates in the multipole-acceptance criterion, such that each cell \rightarrow cell interaction is considered on the merit of the error it likely incurs. This results in very well behaved distributions of the relative force errors, provided an initial estimate for the forces is at hand. This can either be taken from the previous time step or obtained via low-accuracy FMM.

After these improvements, the method has only two free parameters: the expansion order p and a parameter ϵ for the relative force error. Experiments showed that the actual rms relative force error is typically somewhat less than ϵ , while for any given ϵ there is an optimum p at which the computational cost are minimal. For $\epsilon = 10^{-6.25}$, for example, $p = 10$ is optimal and the accelerations errors are comparable to those of direct summation on a GPU (the current state-of-the-art method for collisional N -body simulations). With these parameter settings, the computational costs scale like $N^{0.87}$ for large N and the method out-performs any direct-summation implementation for $N \gtrsim 10^5$. When computing only the forces for $N_a < N$ of N particles, the costs are roughly proportional to $N_a^{0.87}$ for $N_a/N \gtrsim 0.01$, but become independent of N_a below that (where the costs for tree building dominate). For large N , this is still significantly faster than direct summation.

An implementation of the FMM on a GPU accelerator should yield a further significant speed-up compared to my CPU-based implementation, though this is certainly a challenging task, given that FMM is algorithmically more complex than direct summation or a tree code (both of which have been successfully ported to the GPU). Presumably a somewhat lesser challenge is a massively parallel im-

plementation of the method, which can be run on a super computer.

A practical application of FMM in an actual collisional N -body simulation would be very interesting. Since the force between close neighbours is always computed directly (in double precision) as explained earlier, close encounters can be treated essentially in the same fashion as with existing techniques. However, an unfortunate hindrance to an application of the presented techniques originates from the long marriage of existing collisional N -body techniques with direct summation. Methods, such as the Ahmad-Cohen neighbour scheme, to reduce the need for the costly far-field force summations are not necessary with FMM, and the existing N -body tools are not well suited for an immediate application of FMM.

Appendix A: Derivation of the FMM relations

Here, the FMM relations given in Section 2 are derived and motivated. Differently from the main text, the multipole and force expansion centres, \mathbf{z} and \mathbf{s} , are not explicitly distinguished and instead \mathbf{z} is used for either. The general case $\mathbf{z} \neq \mathbf{s}$ is a trivial generalisation.

A.1 Cartesian FMM

The distance vector $\mathbf{x}_b - \mathbf{x}_a$ between two particles residing in two well-separated cells A and B , respectively, can be decomposed into three components (see also Figure 14)

$$\mathbf{x}_b - \mathbf{x}_a = \mathbf{r} + \mathbf{r}_b - \mathbf{r}_a \quad (25)$$

with $\mathbf{r}_a \equiv \mathbf{x}_a - \mathbf{z}_A$, $\mathbf{r}_b \equiv \mathbf{x}_b - \mathbf{z}_B$, and $\mathbf{r} \equiv \mathbf{z}_B - \mathbf{z}_A$. The Taylor expansion of the general Greens function $\psi(\mathbf{x}_b - \mathbf{x}_a)$ in \mathbf{r}_a and \mathbf{r}_b up to order p then reads using multi-index notation^g

$$\psi(\mathbf{x}_b - \mathbf{x}_a) = \sum_{|\mathbf{n}|=0}^p \sum_{|\mathbf{m}|=0}^{p-|\mathbf{n}|} \frac{(-1)^{|\mathbf{m}|}}{\mathbf{n}! \mathbf{m}!} \mathbf{r}_b^{\mathbf{n}} \mathbf{r}_a^{\mathbf{m}} \nabla^{\mathbf{n}+\mathbf{m}} \psi(\mathbf{r}) + \mathcal{R}_p. \quad (26)$$

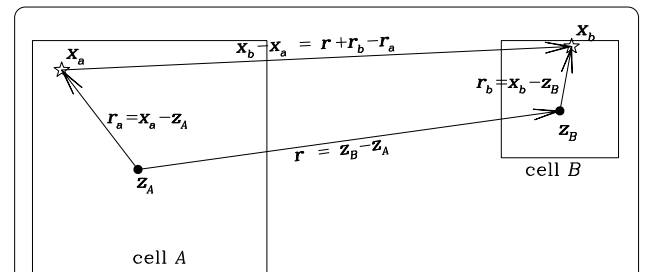


Figure 14 Cell-cell interaction. The Greens function for the interaction between particles a and b at positions \mathbf{x}_a and \mathbf{x}_b and residing in cells A and B , is approximated by Taylor expanding the Greens function about \mathbf{r} , the distance vector between the cell centres \mathbf{z}_A and \mathbf{z}_B .

This series converges (the remainder $\mathcal{R}_p \rightarrow 0$) as $p \rightarrow \infty$, if $|\mathbf{r}_a + \mathbf{r}_b| < |\mathbf{r}|$. Inserting (26) into the expression

$$\Psi_{A \rightarrow B}(\mathbf{x}_b) = \sum_{a \in A} \mu_a \psi(\mathbf{x}_b - \mathbf{x}_a) \quad (27)$$

for the (negative) potential due to all source points in cell A and for any sink position \mathbf{x}_b in cell B , one obtains after re-arranging

$$\Psi_{A \rightarrow B}(\mathbf{x}_b) = \sum_{|\mathbf{n}|=0}^p \frac{1}{\mathbf{n}!} (\mathbf{x}_b - \mathbf{z}_B)^{\mathbf{n}} F_{\mathbf{n}}(\mathbf{z}_B) + \delta \Psi_{A \rightarrow B}, \quad (28a)$$

$$F_{\mathbf{n}}(\mathbf{z}_B) \equiv \sum_{|\mathbf{m}|=0}^{p-|\mathbf{n}|} M_{\mathbf{m}}(\mathbf{z}_A) D_{\mathbf{n}+\mathbf{m}}(\mathbf{z}_B - \mathbf{z}_A), \quad (28b)$$

$$M_{\mathbf{m}}(\mathbf{z}_A) \equiv \sum_{a \in A} \mu_a \frac{(-1)^{\mathbf{m}}}{\mathbf{m}!} (\mathbf{x}_a - \mathbf{z}_A)^{\mathbf{m}} \quad (28c)$$

with the derivatives $D_{\mathbf{n}}(\mathbf{r}) \equiv \nabla^{\mathbf{n}} \psi(\mathbf{r})$. The FMM algorithm essentially works these equations backwards: in a first step, the *multipoles* $M_{\mathbf{m}}(\mathbf{z})$ are computed for each cell via (28c) and by utilising those of daughter cells via the shifting formula

$$M_{\mathbf{m}}(\mathbf{z} + \mathbf{x}) = \sum_{\mathbf{n}=0}^{\mathbf{m}} \frac{\mathbf{x}^{\mathbf{n}}}{\mathbf{n}!} M_{\mathbf{m}-\mathbf{n}}(\mathbf{z}). \quad (28d)$$

Second, for each cell the *field tensors* $F_{\mathbf{n}}(\mathbf{z})$ of all its interactions are computed via (28b) and added up. Finally, the field tensors are passed down the tree, utilising the shifting formula

$$F_{\mathbf{n}}(\mathbf{z} + \mathbf{x}) = \sum_{|\mathbf{m}|=0}^{p-|\mathbf{n}|} \frac{\mathbf{x}^{\mathbf{m}}}{\mathbf{m}!} F_{\mathbf{n}+\mathbf{m}}(\mathbf{z}), \quad (28e)$$

and the potential (and its derivative, the acceleration) is evaluated via (28a) at each sink position. Equations (28) are the basis of Cartesian FMM, such as implemented in Dehnen's (2000, 2002) falcON algorithm.

At each order $n = |\mathbf{n}|$, there are $\binom{n+2}{2}$ coefficients $F_{\mathbf{n}}$ (as well as $M_{\mathbf{n}}$ and $D_{\mathbf{n}}$), and the total number of coefficients up to order p is $\binom{p+3}{3}$. The computational effort of the resulting algorithm is dominated by their computation in (28b), which requires about $\binom{p+6}{6}$ multiplications. Thus at large p a straightforward application of this method approaches an operation count of $\mathcal{O}(p^6)$. The computation (28b) of the field tensors is essentially a convolution in index space and hence can be accelerated using a fast Fourier technique with costs $\mathcal{O}(p^3 \ln p)$ (but see endnote b).

A.2 Harmonic tensors

For the important case $\psi = |\mathbf{r}|^{-1}$ corresponding to gravitational and electrostatic forces, the above method can be

improved by exploiting that this Greens function is *harmonic*, i.e. $\nabla^2 \psi = 0$ for $|\mathbf{r}| > 0$. As a consequence, the $D_{\mathbf{n}} = \nabla^{\mathbf{n}} \psi$ are harmonic too and satisfy

$$\nabla^2 D_{\mathbf{k}} = D_{\mathbf{k}+(2,0,0)} + D_{\mathbf{k}+(0,2,0)} + D_{\mathbf{k}+(0,0,2)} = 0. \quad (29)$$

In other words: $D_{\mathbf{n}}$ is *traceless*. At given degree $n = k + 2$, equation (29) gives $\binom{n}{2}$ constraints such that of the $\binom{n+2}{2}$ terms only $2n + 1$ are truly independent. In inner products, a traceless tensor only 'sees' the traceless part of its co-operand:

$$\sum_{|\mathbf{n}|=n} \bar{A}_{\mathbf{n}} B_{\mathbf{n}} = \sum_{|\mathbf{n}|=n} \bar{A}_{\mathbf{n}} \bar{B}_{\mathbf{n}} = \sum_{|\mathbf{n}|=n} A_{\mathbf{n}} \bar{B}_{\mathbf{n}}, \quad (30)$$

where the 'reduced' tensor $\bar{A}_{\mathbf{n}}$ denotes the traceless part of $A_{\mathbf{n}}$. Furthermore, $\bar{\mathbf{r}}^{\mathbf{n}}$ is related to $D_{\mathbf{n}}$ via

$$D_{\mathbf{n}}(\mathbf{r}) = (-1)^n (2n - 1)!! \frac{\bar{\mathbf{r}}^{\mathbf{n}}}{r^{2n+1}}. \quad (31)$$

With these relations, the Taylor series of the harmonic Greens function becomes, for $r > x$

$$\frac{1}{|\mathbf{r} - \mathbf{x}|} = \sum_{n=0}^{\infty} \frac{x^n}{r^{n+1}} \sum_{|\mathbf{n}|=n} \frac{(2n - 1)!!}{\mathbf{n}!} \bar{\mathbf{x}}^{\mathbf{n}} \bar{\mathbf{r}}^{\mathbf{n}}, \quad (32)$$

which is the Cartesian equivalent to the spherical harmonic expansion

$$\frac{1}{|\mathbf{r} - \mathbf{x}|} = \sum_{n=0}^{\infty} \frac{x^n}{r^{n+1}} \sum_{m=-n}^n Y_n^m(\hat{\mathbf{x}}) Y_n^{m*}(\hat{\mathbf{r}}) \quad (33)$$

(see equation (42) for a definition of Y_n^m). While at each order n there are only $2n + 1$ truly independent terms, the expansion (32) still carries all $\binom{n+2}{2}$ terms, amounting to a total of $\binom{p+3}{3}$ terms in an expansion up to order p . The equivalent spherical harmonic expansion (33) only carries $2n + 1$ terms per order^h amounting to a total of $(p + 1)^2$, i.e. at large p is much preferable.

The number of terms actually used can be reduced to $(2n + 1)$ per order, for example, by omitting all terms with $n_z > 1$ and recover their contribution via recursive application of

$$D_{\mathbf{k}+(0,0,2)} = -D_{\mathbf{k}+(2,0,0)} - D_{\mathbf{k}+(0,2,0)} \quad (34)$$

(Applequist 1989; Hinsen and Felderhof 1992). However, the resulting algebraic challenges are considerable, though the overall computational effort could well be reduced to $\mathcal{O}(p^3)$ operations (Joachim Stadel, private communication), but I am not aware of a systematic demonstration.

A.3 Spherical harmonics

The algebraic complications with obtaining an efficient Cartesian FMM stem from the fact that the Laplace operator involves three terms, such that the resulting recovery relation (34) has two terms instead of one on the right-hand side. This problem can be avoided by Taylor expanding in other than Cartesian coordinates where the Laplace operator involves only two instead of three terms.

The simplest possibility is a linear combination of Cartesian coordinates with complex coefficients. The standard FMM relations emerge from replacing x and y with

$$\xi \equiv \frac{1}{2}(x + iy) \quad \text{and} \quad \eta \equiv -\frac{1}{2}(x - iy) = -\xi^*, \quad (35)$$

while keeping z . Then $\partial_\xi = \partial_x - i\partial_y$ and $\partial_\eta = -\partial_x - i\partial_y$, such that $\partial_x^2 + \partial_y^2 = -\partial_\xi \partial_\eta$ and hence for harmonic functions

$$\partial_\xi \partial_\eta = \partial_z^2, \quad (36)$$

or $D_{\mathbf{k}+(0,0,2)} = D_{\mathbf{k}+(1,1,0)}$ in place of equation (34). With this relation one can eliminate all mixed ξ - η derivatives in favour of z derivatives. This in turn allows a reduction in the number of indices from three to two by using the total number n of derivatives and the number $|m|$ of ξ (for $m < 0$) or η derivatives (for $m > 0$).

Somewhat surprisingly, the relations required for FMM are hardly covered by the rich literature on spherical harmonics (and FMM). To derive the relevant formulae, I follow the ideas of Maxwell (1892, see also James 1969) and define the differential operator

$$\Delta_n^m \equiv \begin{cases} (-\partial_z)^{n-|m|} \partial_\xi^{|m|}, & -n \leq m < 0, \\ (-\partial_z)^{n-|m|} \partial_\eta^{|m|}, & 0 \leq m \leq n, \\ 0, & |m| > n. \end{cases} \quad (37)$$

When applied to harmonic functions, this operator satisfies

$$\Delta_n^m \Delta_k^l = \Delta_{n+k}^{l+m}, \quad (38)$$

which can be shown via equation (36) and is inevitably linked to

$$\Delta_n^{-m} = (-1)^m \Delta_n^{m*}. \quad (39)$$

Since $\psi = |\mathbf{r}|^{-1}$ is harmonic, its derivatives

$$\Theta_n^m(\mathbf{r}) \equiv \Delta_n^m |\mathbf{r}|^{-1} \quad (40)$$

are harmonic too. Moreover, the functions $\Theta_n^m(\mathbf{r})$ are homogeneous of degree $-(n+1)$, i.e. $\Theta_n^m(\alpha\mathbf{r}) = \alpha^{-(n+1)}\Theta_n^m(\mathbf{r})$. I also define the *solid spherical harmonic* of degree n as

$$\Upsilon_n^m(\mathbf{r}) \equiv \frac{r^{-1}\Theta_n^m(\mathbf{r}/r^2)}{(n-m)!(n+m)!} = \frac{r^{2n+1}\Theta_n^m(\mathbf{r})}{(n-m)!(n+m)!}. \quad (41)$$

That Υ_n^m is harmonic follows from the fact that if $f(\mathbf{r})$ is harmonic, then so is $r^{-1}f(\mathbf{r}/r^2)$ (Hobson 1931) (try this with your undergraduate students). Note that $\Upsilon_n^m(\mathbf{r})$ is just a homogeneous polynomial of total degree n in x , y and z . These harmonics are related to the usual normalised *surface spherical harmonic*

$$Y_n^m(\hat{\mathbf{r}}) \equiv (-1)^m \sqrt{\frac{(n-m)!}{(n+m)!}} P_n^m(\cos\theta) e^{im\phi} \quad (42)$$

via

$$\Theta_n^m(\mathbf{r}) = [(n-m)!(n+m)!]^{1/2} r^{-n-1} Y_n^m(\hat{\mathbf{r}}), \quad (43a)$$

$$\Upsilon_n^m(\mathbf{r}) = [(n-m)!(n+m)!]^{-1/2} r^n Y_n^m(\hat{\mathbf{r}}). \quad (43b)$$

Table 3 gives the first few harmonics in terms of x , y , z .

Table 3 The real-valued functions $(n-m)!(n+m)!U_n^m(\mathbf{r}) = r^{2n+1}\Upsilon_n^m(\mathbf{r})$ for $n \leq 6$

$m \backslash n$	0	1	2	3	4	5	6
6							$10395(x^6 - 15x^4y^2 + 15x^2y^4 - y^6)$
5						$945x(x^4 - 10x^2y^2 + 5y^4)$	$10395xz(x^4 - 10x^2y^2 + 5y^4)$
4				$105(x^4 - 6x^2y^2 + y^4)$	$945z(x^4 - 6x^2y^2 + y^4)$	$945(x^4 - 6x^2y^2 + y^4)(11z^2 - r^2)$	
3			$15x(x^2 - 3y^2)$	$105xz(x^2 - 3y^2)$	$105x(x^2 - 3y^2)(9z^2 - r^2)$	$945zx(x^2 - 3y^2)(11z^2 - 3r^2)$	
2		$3(x^2 - y^2)$	$15z(x^2 - y^2)$	$15(x^2 - y^2)(7z^2 - r^2)$	$315z(x^2 - y^2)(3z^2 - r^2)$	$315(x^2 - y^2)(33z^4 - 18z^2r^2 + r^4)$	
1	x	$3xz$	$3x(5z^2 - r^2)$	$15xz(7z^2 - 3r^2)$	$45x(21z^4 - 14z^2r^2 + r^4)$	$315xz(33z^4 - 30z^2r^2 + 5r^4)$	
0	1	z	$3z^2 - r^2$	$3z(5z^2 - 3r^2)$	$3(35z^4 - 30z^2r^2 + 3r^4)$	$15z(63z^4 - 70z^2r^2 + 15r^4)$	$45(231z^6 - 315z^4r^2 + 105z^2r^4 - 5r^6)$
-1		y	$3yz$	$3y(5z^2 - r^2)$	$15yz(7z^2 - 3r^2)$	$45y(21z^4 - 14z^2r^2 + r^4)$	$315yz(33z^4 - 30z^2r^2 + 5r^4)$
-2			$6xy$	$30xyz$	$30xy(7z^2 - r^2)$	$630xyz(3z^2 - r^2)$	$630xy(33z^4 - 18z^2r^2 + r^4)$
-3				$15y(3x^2 - y^2)$	$105yz(3x^2 - y^2)$	$105y(3x^2 - y^2)(9z^2 - r^2)$	$945zy(3x^2 - y^2)(11z^2 - 3r^2)$
-4					$420xy(x^2 - y^2)$	$3780xyz(x^2 - y^2)$	$3780xy(x^2 - y^2)(11z^2 - r^2)$
-5						$945y(5x^4 - 10x^2y^2 + y^4)$	$10395yz(5x^4 - 10x^2y^2 + y^4)$
-6							$20790xy(3x^2 - y^2)(x^2 - 3y^2)$

See equations (58) for the relations to $\Upsilon_n^m(\mathbf{r})$ and $\Theta_n^m(\mathbf{r})$.

A.4 Spherical-harmonic FMM

In order to derive the relations for spherical-harmonic FMM, one must obtain the equivalent to the Cartesian Taylor expansion (26) and shift operations (28d), (28e). Via induction one can show that when applied to harmonic functions

$$\begin{aligned} \frac{1}{n!}(\mathbf{x} \cdot \nabla)^n &= \sum_{m=-n}^n \gamma_n^{m*}(-\mathbf{x}) \Delta_n^m \\ &= \sum_{m=-n}^n \gamma_n^m(-\mathbf{x}) \Delta_n^{m*}, \end{aligned} \quad (44)$$

which gives the translation operator for harmonic functions

$$\exp(\mathbf{x} \cdot \nabla) = \sum_{n=0}^{\infty} \sum_{m=-n}^n \gamma_n^m(-\mathbf{x}) \Delta_n^{m*}. \quad (45)$$

When applying this to the harmonic Greens function, one gets

$$\frac{1}{|\mathbf{r} - \mathbf{x}|} = \sum_{n=0}^{\infty} \sum_{m=-n}^n \gamma_n^m(\mathbf{x}) \Theta_n^{m*}(\mathbf{r}), \quad (46)$$

which, because of equations (43), is equivalent to the standard form (33) and converges for $r > x$. Translating once again and employing (38) yields

$$\frac{1}{|\mathbf{r} - \mathbf{x} - \mathbf{y}|} = \sum_{k=0}^{\infty} \sum_{l=-k}^k \sum_{n=0}^{\infty} \sum_{m=-n}^n \gamma_n^m(\mathbf{x}) \gamma_k^l(\mathbf{y}) \Theta_{n+k}^{m+l*}(\mathbf{r}), \quad (47)$$

which converges for $r > |\mathbf{x} + \mathbf{y}|$. Comparing (46) and (47) one finds immediately the translation formula

$$\gamma_n^m(\mathbf{x} + \mathbf{y}) = \sum_{k=0}^n \sum_{l=-k}^k \gamma_k^l(\mathbf{x}) \gamma_{n-k}^{m-l}(\mathbf{y}). \quad (48)$$

When applying the translation operator (45) to Θ_n^m , one gets

$$\Theta_n^m(\mathbf{x} + \mathbf{y}) = \sum_{k=0}^{\infty} \sum_{l=-k}^k \gamma_k^{l*}(-\mathbf{y}) \Theta_{n+k}^{m+l}(\mathbf{x}). \quad (49)$$

As the Cartesian FMM relations (28) were based on equation (26), the spherical harmonic FMM relations (3) are based on equation (47), which for $\psi = |\mathbf{r}|^{-1}$ is completely equivalent but computationally more efficient.

A.5 Implementation details

A.5.1 Recursive evaluation of spherical harmonics

One may also obtain the relations

$$\Delta_k^l \Theta_n^m(\mathbf{x}) = \Theta_{n+k}^{m+l}(\mathbf{x}), \quad (50)$$

$$\Delta_k^l \gamma_n^m(\mathbf{x}) = (-1)^{k+l} \gamma_{n-k}^{m+l}(\mathbf{x}). \quad (51)$$

The first one follows immediately from equations (38) and (40), while the second can be deduced by equating (48) to $\gamma_n^m(\mathbf{x} + \mathbf{y})$ obtained by applying the translation operator (45). From these two relations combined with the operator relation (38) and the definitions (40) and (41), one can obtain numerous recurrence relations. For example (omitting the arguments for brevity),

$$r^2 \Theta_n^m = (2n-1)z \Theta_{n-1}^m - ([n-1]^2 - m^2) \Theta_{n-2}^m, \quad (52)$$

$$(n^2 - m^2) \gamma_n^m = (2n-1)z \gamma_{n-1}^m - r^2 \gamma_{n-2}^m, \quad (53)$$

which are equivalent to the recurrence relation (Gradshteyn and Ryzhik 1994, Eq. 8.733.2) for associated Legendre functions and, together with

$$\Theta_n^m = (2n-1) \frac{x+iy}{r^2} \Theta_{n-1}^{m-1} \quad \text{and} \quad \gamma_n^m = \frac{x+iy}{2n} \gamma_{n-1}^{m-1} \quad (54)$$

as well as their counterparts for $m = -n$, allow for an efficient and stable evaluation of $\Theta_n^m(\mathbf{r})$ and $\gamma_n^m(\mathbf{r})$.

Differentiating these relations with respect to time, one obtains recursion relations for the time derivatives of the harmonic functions. For example,

$$\begin{aligned} (n^2 - m^2) \dot{\gamma}_n^m &= (2n-1) [\dot{z} \gamma_{n-1}^m + z \dot{\gamma}_{n-1}^m] \\ &\quad - 2\mathbf{r} \cdot \dot{\mathbf{r}} \gamma_{n-2}^m - r^2 \dot{\gamma}_{n-2}^m. \end{aligned} \quad (55)$$

Alternatively, from equations (50) and (51) one may also directly derive

$$\dot{\Theta}_n^m = \dot{\xi} \Theta_{n+1}^{m+1} - \dot{z} \Theta_{n+1}^m + \dot{\eta} \Theta_{n+1}^{m-1} \quad \text{and} \quad (56)$$

$$\dot{\gamma}_n^m = \dot{\xi} \gamma_{n+1}^{m+1} + \dot{z} \gamma_{n+1}^m + \dot{\eta} \gamma_{n+1}^{m-1}. \quad (57)$$

A.5.2 Real-valued spherical harmonics

Because of the anti-symmetry relation (39), the complex spherical harmonics defined above are redundant: there are only $2n+1$ independent (real) harmonics per order, in agreement with the counting in Section A.1. Hence, for any practical application one needs an appropriately reduced set of $2n+1$ real-valued independent spherical harmonics per order. The simplest option is to consider real and imaginary parts of the complex-valued harmonics with $m \geq 0$:

$$U_n^m \equiv \begin{cases} \Im(\gamma_n^{|m|}) = \frac{1}{2i}(\gamma_n^{-m} - \gamma_n^{-m*}), & m < 0, \\ \Re(\gamma_n^{|m|}) = \frac{1}{2}(\gamma_n^m + \gamma_n^{m*}), & m \geq 0; \end{cases} \quad (58a)$$

and

$$T_n^m \equiv \begin{cases} \Im(\Theta_n^{|m|}) = \frac{1}{2i}(\Theta_n^{-m} - \Theta_n^{-m*}), & m < 0, \\ \Re(\Theta_n^{|m|}) = \frac{1}{2}(\Theta_n^m + \Theta_n^{m*}), & m \geq 0. \end{cases} \quad (58b)$$

The relevant relations for these real-valued spherical harmonics are best directly transcribed from the corresponding complex relations.

A.6 Accelerating FMM relations

The FMM kernels M2L, M2M, and L2L (equations (3b), (3d), (3e)) all require $\mathcal{O}(p^4)$ operations. However, if the interactions or translations are along the z -axis, the costs are only $\mathcal{O}(p^3)$ because $\Upsilon_n^m(\hat{\mathbf{z}}) = \delta_{m0}/n!$.

One method to exploit this is to first translate along the z -axis and then perpendicular to the z -axis. For a vector \mathbf{r}_\perp perpendicular to the z -axis, $\Upsilon_n^m(\mathbf{r}_\perp)$ vanishes whenever $n + m$ is even. This implies that a translation along \mathbf{r}_\perp can be done faster than a general translation (in the limit of $p \rightarrow \infty$, twice as fast).

This splitting method cannot be applied to the M2L kernel (3b) (because it is not a translation), which occurs many more times in the FMM algorithm than the M2M and L2L kernels. To accelerate the M2L kernel, one can exploit that a rotation only costs $\mathcal{O}(p^3)$ operations, too. Thus, if one first rotates into a frame in which the interaction is along the z -axis, applies the M2L kernel in the rotated frame, and finally rotates back into the original frame, the total costs are still $\mathcal{O}(p^3)$.

A.6.1 Fast rotations

Since the spherical harmonics are homogeneous, a rotation (as opposed to a translation) does not mix between different orders n , and consequently the operation count is $\mathcal{O}(p^3)$. Thus, a general rotation is of the form

$$Y_n^m(\tilde{\mathbf{r}}) = \sum_{l=-n}^n \Gamma_n^{ml} Y_n^l(\mathbf{r}), \quad (59)$$

where $\tilde{\mathbf{r}}$ denotes the vector \mathbf{r} in the rotated frame. Unfortunately, the matrices Γ_n , also known as Wigner functions, are generally dense and non-trivial functions of the Euler angles. However, a rotation by angle α around the z -axis is simple:

$$Y_n^m(\tilde{\mathbf{r}}) = e^{-im\alpha} Y_n^m(\mathbf{r}) \quad (60)$$

with an operation count of only $\mathcal{O}(p^2)$. With this one can build a general rotation by first rotating around the z -axis, then swapping z and x , rotating again about the z -axis (the x -axis of the original frame), swapping z and x again, and

performing a final rotation around the z -axis. Like rotations, swapping coordinate axes does not mix between different orders n and can be represented as

$$\Theta_n^m(\tilde{\mathbf{r}}) = \sum_{l=-n}^n B_n^{ml} \Theta_n^l(\mathbf{r}), \quad (61)$$

where now $\tilde{\mathbf{r}}$ denotes the vector \mathbf{r} in the frame obtained by swapping two Cartesian coordinates. The important difference between equations (59) and (61) is that the matrices B_n are *constants*. Recursive relations for these swap matrices can be derived via the operator algebra of Section A.3. For example, for swapping x and z , one finds

$$2\tilde{\Delta}_1^0 = \Delta_1^1 - \Delta_1^{-1} \quad \text{and} \quad 2\tilde{\Delta}_1^{\pm 1} = \Delta_1^1 + \Delta_1^{-1} \pm 2\Delta_1^0 \quad (62)$$

with which one can derive the recurrence relations

$$\begin{aligned} 2B_{n+1}^{ml} &= B_n^{ml-1} - B_n^{ml+1}, \\ 2B_{n+1}^{m\pm 1l} &= B_n^{ml-1} + B_n^{ml+1} \pm 2B_n^{ml}, \end{aligned} \quad (63)$$

where it is understood that $B_n^{ml} = 0$ for $|l| > n$. A similar exercise for swapping y and z reveals that the swap matrices are given by $i^{m-l} B_n^{ml}$, while the corresponding swap matrices for Υ_n^m are given by the transpose (because these matrices are orthonormal and the product (46) is invariant under coordinate swapping). Whereas the matrices B_n are dense, the corresponding matrices for the real-valued harmonics (equations (58)) are not (Pinchon and Hoggan 2007). For example, the matrices for swapping x and z for Θ_4^m and T_4^m are (omitting zero entries)

$$\frac{1}{16} \times \begin{bmatrix} 1 & -8 & 28 & -56 & 70 & -56 & 28 & -8 & 1 \\ -1 & 6 & -14 & 14 & & -14 & 14 & -6 & 1 \\ 1 & -4 & 4 & 4 & -10 & 4 & 4 & -4 & 1 \\ -1 & 2 & 2 & -6 & & 6 & -2 & -2 & 1 \\ 1 & & -4 & & 6 & & -4 & & 1 \\ -1 & -2 & 2 & 6 & & -6 & -2 & 2 & 1 \\ 1 & 4 & 4 & -4 & -10 & -4 & 4 & 4 & 1 \\ -1 & -6 & -14 & -14 & & 14 & 14 & 6 & 1 \\ 1 & 8 & 28 & 56 & 70 & 56 & 28 & 8 & 1 \end{bmatrix} \quad (64)$$

and

$$\frac{1}{8} \times \begin{bmatrix} & 8 & & 56 & & & & & \\ 1 & & 14 & & & & & & \\ & 4 & & -4 & & & & & \\ 1 & & -2 & & & & & & \\ & & & & 3 & & -4 & & 1 \\ & & & & & -6 & & 2 & \\ & & & & -5 & & 4 & & 1 \\ & & & & & 14 & & 6 & \\ & & & & & & 35 & 28 & 1 \end{bmatrix}, \quad (65)$$

respectively. Thus, this method of achieving a general rotation not only avoids the (recursive) computation of the Wigner functions Γ_n (which itself costs $\mathcal{O}(p^3)$ operations), but also benefits from the facts that the swap matrices B_n have ≈ 4 times fewer non-zero entries than the Γ_n and are known a priori, such that they can be ‘hard-wired’ into computer code.

A.6.2 A fast M2L kernel

With these preliminaries, one can finally put together an accelerated $\mathcal{O}(p^3)$ version for performing the M2L kernel (3b). Let $(x, y, z) = \mathbf{r}$, then one first rotates the multipoles \mathcal{M}_k^l (around the z -axis) by angle $\alpha_z = \arctan(y/x)$, swaps x and z , rotates by $\alpha_x = \arctan(\sqrt{x^2 + y^2}/z)$, and swaps x and z back. The obtained $\tilde{\mathcal{M}}_k^l$ has z -axis aligned with the interaction direction, and the M2L kernel can be performed via

$$\tilde{\mathcal{F}}_n^m(\mathbf{z}_B) = \sum_{k=|m|}^{p-n} (-1)^m \tilde{\mathcal{M}}_k^m(\mathbf{z}_A) \frac{(n+k)!}{r^{n+k+1}}. \quad (66)$$

Finally, one must rotate $\tilde{\mathcal{F}}_n^m$ back to the original frame by first swapping x and z , rotating by $-\alpha_x$, swapping x and z again, followed by a final rotation by $-\alpha_z$.

These rotations and swaps can be accelerated further by exploiting that in (66) only multipoles $\tilde{\mathcal{M}}_n^m$ with $|m| \leq \min\{n, p-n\}$ are needed and, similarly, that $\tilde{\mathcal{F}}_n^m = 0$ for $|m| > \min\{n, p-n\}$. As Figure 1 demonstrates, the overhead due to the rotations pays off already for $p = 5$.

Appendix B: The energy error of a simulation

The gravitational forces (and potentials) used in N -body simulations always carry some error. When using direct summation, this is solely due to round-off errors, while for approximate methods the approximation error should dominate round-off. Here, I investigate the consequences of these errors for the non-conservation of the total energy.

B.1 The energy error due to force errors

Consider, the energy error generated by acceleration errors $\delta \mathbf{a}_b$ after one time step τ ,

$$\delta E_{\text{tot}} = \tau \sum_b \mu_b \dot{\mathbf{x}}_b \cdot \delta \mathbf{a}_b. \quad (67)$$

Because the $\delta \mathbf{a}_b$ are not correlated with the velocities $\dot{\mathbf{x}}_b$, their dot products largely cancel and δE_{tot} will be small. In order to estimate its amplitude, let us assume $\tau = \eta \sigma / \bar{a}$ with $\eta \ll 1$, velocity dispersion σ , and typical acceleration \bar{a} . If further assuming virial equilibrium and a relative acceleration error ε ,

$$|\delta E_{\text{tot}}|/|E_{\text{tot}}| \sim \eta \varepsilon / \sqrt{N}. \quad (68)$$

Over time this accumulates in the fashion of a random walk and after one dynamical time or $\sim \eta^{-1}$ time steps

$$|\delta E_{\text{tot}}|/|E_{\text{tot}}| \sim \varepsilon \sqrt{\eta/N}. \quad (69)$$

Thus, the relative energy error resulting from the force errors alone is much smaller than ε , simply because it is some average over many force errors.

B.2 The measurement error

In order to measure the total energy, one must also calculate the individual particle potentials Ψ_b (which are otherwise not required for the simulation). Assuming that the Ψ_b are computed with relative error ε , the resulting error for the total energy is

$$|\delta E_{\text{tot}}|/|E_{\text{tot}}| \approx 2\varepsilon/\sqrt{N}. \quad (70)$$

If the same precision ε is used for computing the particle potentials and accelerations, this is much larger than the energy error (69) due to force errors.

B.3 Approximate gravity solvers

The situation is different for approximative methods, such as the tree code, FMM, and mesh-based techniques. All of these approximate the true potential, but use the exact derivatives of the approximated potential for the accelerations. Therefore, the total *approximated* energy should be conserved (modulo round-off errors), even if the approximation is poor.

For the FMM and the tree code the situation is actually different, because the approximated potential is not globally continuous but only piece-wise. This is because the concrete form of the approximation used for a given particle depends on its position (which determines how FMM approximates each pair-wise force). A particle crossing a boundary between such continuous regions suffers a jump in the (approximated) potential, and hence energy, while the corresponding kick in velocity (to conserve energy) is ignored. These discontinuities are part of the approximation error and their amplitudes proportional. The implication is that for the tree code and FMM energy is not conserved (even for accurate time integration) and the degree of non-conservation actually reflects the amplitude of the approximation errors in an average sense.

Competing interests

The author declares that he has no competing interests.

Acknowledgements

The author thanks Joachim Stadel for many helpful discussions and the suggestion to allow $\mathbf{z} \neq \mathbf{s}$, Alessia Gualandris for running sapporo to provide the data for Figure 3, and Simon Portegies Zwart and Jeroen Bédorf for providing the timings for sapporo 2 in Figure 11. This work was supported by STFC consolidated grant ST/K001000/1.

Endnotes

- ^a Cheng et al.'s (1999) expressions are quite cumbersome because they are given in terms of the surface spherical harmonics Y_n^m in polar coordinates and because they contain phase-factors like i^{m-l-m} owing to their unconventional definition for the Y_n^m which implies $Y_n^{-m} = Y_n^{m*}$ instead of $Y_n^{-m} = (-1)^m Y_n^{m*}$.
- ^b Expressions like $\mathcal{O}(p^n)$ for the operation count relate to the asymptotic behaviour at large expansion orders p . While this is straightforward to specify, it is not necessarily very relevant, since in the range up to $p \sim 10$, as required in practice, the actual costs usually grow more slowly than implied by the asymptotic behaviour (see Figure 1 for a typical example) and because the numerical implementation may be data-dominated rather than computation dominated.
- ^c The original definition used in the tree code of Barnes and Hut (1986) did not ensure bounded errors, causing the infamous 'exploding galaxies' bug first reported by Salmon and Warren (1994).
- ^d The situation is different for N -body simulations of collisionless stellar dynamics, where reversible integrators are used and the accepted force errors, and thus their time asymmetries, are significantly larger.
- ^e The increase of this ratio to ≈ 20 towards small errors may well be caused by inaccuracies of the direct summation used for calculating the errors.
- ^f The timings for the sapporo library also include additional computations (nearest neighbour finding and neighbour listing). These contribute negligibly at large N , but at small N they are, together with latency on the GPU, responsible for the deviation of the observed complexity from N^2 .
- ^g Using multi-index notation $\mathbf{n} \equiv (\mathbf{n}_x, \mathbf{n}_y, \mathbf{n}_z)$ with $n \equiv |\mathbf{n}| \equiv \mathbf{n}_x + \mathbf{n}_y + \mathbf{n}_z$, such that the first sum in (26) is over non-negative integer triples \mathbf{n} with $\mathbf{n}_x + \mathbf{n}_y + \mathbf{n}_z \leq p$. Furthermore $\mathbf{n}! \equiv \mathbf{n}_x! \mathbf{n}_y! \mathbf{n}_z!$ and $\mathbf{r}^{\mathbf{n}} \equiv r_x^{\mathbf{n}_x} r_y^{\mathbf{n}_y} r_z^{\mathbf{n}_z}$.
- ^h In equation (33), the Y_l^m are complex-valued for $m \neq 0$, but because of their symmetry $Y_l^{m*} = (-1)^m Y_l^{-m}$ there are only $2n + 1$ independent real-valued components per order n .

- Hinsen, K, Felderhof, BU: Reduced description of electric multipole potential in Cartesian coordinates. *J. Math. Phys.* **33**, 3731-3735 (1992)
- Hobson, EW: The Theory of Spherical and Ellipsoidal Harmonics. Cambridge University Press, Cambridge (1931)
- James, RW: Transformation of spherical harmonics under change of reference frame. *Geophys. J. Int.* **17**, 305-316 (1969)
- Makino, J: Optimal order and time-step criterion for Aarseth-type N -body integrators. *Astrophys. J.* **369**, 200-212 (1991). doi:10.1086/169751
- Makino, J, Taiji, M: Scientific Simulations with Special-Purpose Computers: The GRAPE Systems. Wiley, New York (1998)
- Matoušek, J, Sharir, M, Welzl, E: A subexponential bound for linear programming. *Algorithmica* **16**, 498-516 (1996)
- Maxwell, JC: Treatise on Electricity and Magnetism. Oxford University Press, Oxford (1892)
- Pinchon, D, Hoggan, PE: Rotation matrices for real spherical harmonics: general rotations of atomic orbitals in space-fixed axes. *J. Phys. A, Math. Theor.* **40**, 1597-1610 (2007). doi:10.1088/1751-8113/40/7/011
- Plummer, HC: On the problem of distribution in globular star clusters. *Mon. Not. R. Astron. Soc.* **71**, 460-470 (1911)
- Price, DJ, Monaghan, JJ: An energy-conserving formalism for adaptive gravitational force softening in smoothed particle hydrodynamics and N -body codes. *Mon. Not. R. Astron. Soc.* **374**, 1347-1358 (2007). doi:10.1111/j.1365-2966.2006.11241.x
- Reinders, J: Intel Threading Building Blocks. O'Reilly Media, Sebastopol (2007)
- Salmon, JK, Warren, MS: Skeletons from the treecode closet. *J. Comput. Phys.* **111**, 136-155 (1994). doi:10.1006/jcph.1994.1050
- Taura, K, Nakashima, J, Yokota, R, Maruyama, N: A task parallel implementation of fast multipole methods. In: Proceedings of the 2012 SC Companion: High Performance Computing, Networking Storage and Analysis, pp. 617-625 (2012). doi:10.1109/SC.Companion.2012.86

doi:10.1186/s40668-014-0001-7

Cite this article as: Dehnen: A fast multipole method for stellar dynamics. *Computational Astrophysics and Cosmology* 2014 **1**:1.

Received: 30 January 2014 Accepted: 30 April 2014

Published online: 11 September 2014

References

- Ahmad, A, Cohen, L: A numerical integration scheme for the N -body gravitational problem. *J. Comput. Phys.* **12**, 389-402 (1973). doi:10.1016/0021-9991(73)90160-5
- Applequist, J: Traceless Cartesian tensor forms for spherical harmonic functions: new theorems and applications to electrostatics of dielectric media. *J. Phys. A, Math. Gen.* **22**, 4303-4330 (1989). doi:10.1088/0305-4470/22/20/011
- Barnes, J, Hut, P: A hierarchical $O(N \log N)$ force-calculation algorithm. *Nature* **324**, 446-449 (1986)
- Capuzzo-Dolcetta, R, Miocchi, P: A comparison between the fast multipole algorithm and the tree-code to evaluate gravitational forces in 3-D. *J. Comput. Phys.* **143**, 29-48 (1998). doi:10.1006/jcph.1998.5949
- Cheng, H, Greengard, L, Rokhlin, V: A fast adaptive multipole algorithm in three dimensions. *J. Comput. Phys.* **155**, 468-498 (1999). doi:10.1006/jcph.1999.6355
- Dehnen, W: A very fast and momentum-conserving tree code. *Astrophys. J.* **536**, L39-L42 (2000)
- Dehnen, W: Towards optimal softening in three-dimensional N -body codes - I. Minimizing the force error. *Mon. Not. R. Astron. Soc.* **324**, 273-291 (2001)
- Dehnen, W: A hierarchical $O(N)$ force calculation algorithm. *J. Comput. Phys.* **179**, 27-42 (2002)
- Dehnen, W, Read, JI: N -Body simulations of gravitational dynamics. *Eur. Phys. J. Plus* **126**, 55 (2011). doi:10.1140/epjp/i2011-11055-3
- Fischer, K, Gärtner, B, Herrmann, T, Hoffmann, M, Schönherr, S: Bounding volumes. In: CGAL User and Reference Manual, 4.2 edn., CGAL Editorial Board (2013). http://www.cgal.org/Manual/4.2
- Gaburov, E, Harfst, S, Portegies Zwart, S: SAPPORO: a way to turn your graphics cards into a GRAPE-6. *New Astron.* **14**(7), 630-637 (2009). doi:10.1016/j.newast.2009.03.002
- Gradshteyn, IS, Ryzhik, I: Table of Integrals, Series, and Products, 5th edn. Academic Press, London (1994)
- Greengard, L, Rokhlin, V: A fast algorithm for particle simulations. *J. Comput. Phys.* **73**, 325-348 (1987). doi:10.1016/0021-9991(87)90140-9

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