

A FAST POTENTIAL AND SELF-GRAVITY SOLVER FOR NONAXISYMMETRIC DISKS

SHENGTAI LI¹, MATTHEW J. BUONI², AND HUI LI³

¹ Theoretical Division, MS B284, Los Alamos National Laboratory, NM 87545, USA; sli@lanl.gov

² Department of Mechanical Engineering, University of California, Santa Barbara, CA 93107, USA; buoni@engineering.ucsb.edu

³ Theoretical Division, MS B227, Los Alamos National Laboratory, NM 87545, USA; hli@lanl.gov

Received 2008 July 29; accepted 2008 October 29; published 2009 March 5

ABSTRACT

Disk self-gravity could play an important role in the dynamic evolution of interaction between disks and embedded protoplanets. We have developed a fast and accurate solver to calculate the disk potential and disk self-gravity forces for disk systems on a uniform polar grid. Our method closely follows the method given by Chan et al., in which a fast Fourier transform in the azimuthal direction is performed and a direct integral approach in the frequency domain in the radial direction is implemented on a uniform polar grid. This method can be very effective for disks with vertical structures that depend only on the disk radius, achieving the same computational efficiency as for zero-thickness disks. We describe how to parallelize the solver efficiently on distributed parallel computers. We propose a mode-cutoff procedure to reduce the parallel communication cost and achieve nearly linear scalability for a large number of processors. For comparison, we have also developed a particle-based fast tree code to calculate the self-gravity of the disk system with a vertical structure. The numerical results show that our direct integral method is at least two orders of magnitude faster than our optimized tree-code approach.

Key words: accretion, accretion discs – methods: numerical – planetary systems: protoplanetary discs

1. INTRODUCTION

Well before extrasolar planets were discovered, Goldreich & Tremaine (1979, 1980) and Lin & Papaloizou (1986a, 1986b) speculated that tidal interactions between disks and embedded protoplanets would lead to planet migration. Ward (1997) suggested that two different types of migration could occur. Nelson & Benz (2003a, 2003b) numerically studied the effects of disk self-gravity in two-dimensional simulations of planet-disk interactions. Pierens & Huré (2005a) reported by an analytical derivation that the disk gravity accelerates the planetary migration. Recently, simulations by Baruteau & Masset (2008) confirmed that the self-gravity indeed accelerates the type I migration. They implemented a two-dimensional self-gravity solver in their code using the fast Fourier transform (FFT) method of Binney & Tremaine (1987), which required a logarithmic radial spacing ($\log(r)$).

In Newtonian gravity, we can define the gravitational potential Ψ associated with the mass density, ρ , by the volume integral

$$\Psi(\mathbf{x}) = -G \iiint \frac{\rho(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d^3x' \quad (1)$$

over all space, where G is the gravitational constant. Rewriting Equation (1) in a differential form, we obtain Poisson's equation

$$\nabla^2 \Psi = 4\pi G\rho, \quad (2)$$

with Ψ satisfying the boundary condition $\Psi(\infty) = 0$ at all times. The numerical “Poisson solvers” to Equation (2) can be classified into two categories: difference methods and integral methods. The difference methods solve Equation (2) directly by either finite-difference or finite-element methods with necessary boundary conditions. The known boundary conditions at infinity are usually not very useful if a finite domain is considered. User-specified boundary conditions or Dirichlet boundary conditions obtained by direct summation are often required. A key advantage of difference methods is that,

generally, they are relatively fast once initialized. However, they have low or limited accuracy, and they often rely on the integral method to provide the boundary conditions.

The integral methods are to integrate Equation (1) directly. They have the advantage that the summation stops naturally at the domain boundaries. However, an integral method has two difficulties in a practical implementation. One is that the integral has a point mass singularity (i.e., when $\mathbf{x}' \rightarrow \mathbf{x}$ in Equation (1)), which is often circumvented by introducing softening. The other difficulty is that it is computationally prohibitive for a large system. The computational cost can be significantly lowered if the FFT method can be used.

The gravitational field, $g = -\nabla\Psi$, is more convenient in many applications. It can be calculated via numerical derivatives, which generally have poor precision. To obtain high accuracy, we can integrate the field directly by

$$g(\mathbf{x}) = G \iiint \frac{\rho(\mathbf{x}')(\mathbf{x}' - \mathbf{x})}{|\mathbf{x} - \mathbf{x}'|^3} d^3x', \quad (3)$$

which shares the same two difficulties as calculating Equation (1).

In this paper, we present a method for computing the disk self-gravity for quasi-two-dimensional disk models. We consider a disk with the cylindrical grid (r, ϕ, z) . We assume that the scale height of the disk (semi-thickness) is only radius dependent, that is, $H = H(r)$, and it is quite small, $H(r)/r \ll 1$ (namely, geometrically thin disks). We also assume that the vertical structure of the density can be described by some function $Z(r, z)$ that is independent of ϕ and time t , that is,

$$\rho(t, r, \phi, z) = \Sigma(t, r, \phi)Z(r, z).$$

In this paper, we are particularly interested in the gravitational potential and field force at the $z = 0$ plane. Although our method is valid for any function of $Z(r, z)$, we assume that the density vertically has a Gaussian distribution. Under this assumption,

$$Z(r, z) = \frac{1}{\sqrt{2\pi}H(r)^2} \exp\left(-\frac{z^2}{2H(r)^2}\right). \quad (4)$$

With the presence of the vertical structure (4), the migration rate of the protoplanet is reduced up to 50% (see Koller 2004). We assume that the disk has a constant sound speed c_s . Then, from the scale height

$$\frac{H}{r} = \frac{c_s}{v_K},$$

where v_K is the Keplerian velocity $v_K(r) = \sqrt{GM_\star/r}$, we obtain

$$H(r) = \frac{c_s r^{3/2}}{\sqrt{GM_\star}}.$$

Other profiles for the scale height are also easy to handle. For example, if the aspect ratio, $h = H/r$, is constant, then $H(r) = rh$. The numerical verification in Section 3 actually uses a constant H in the whole domain.

Since we are interested in the potential and field in the $z = 0$ plane, the easiest method to obtain Ψ is to directly integrate the equation

$$\begin{aligned} \Psi(t, r, \phi) = & \int_{r_{\min}}^{r_{\max}} \int_0^{2\pi} \Sigma(t, r', \phi') r' dr' d\phi' \\ & \times \int_{-\infty}^{\infty} -\frac{GZ(r', z')}{\sqrt{r^2 + r'^2 - 2rr' \cos(\phi - \phi') + z'^2}} dz', \end{aligned} \quad (5)$$

where r_{\min} and r_{\max} are the inner and outer radial boundaries, respectively, and Σ is the vertically integrated density. As in Huré (2005), we have employed two different coordinate systems in the radial direction: r as the field grid and r' as the source grid. For simplicity, we set $G = 1$ and $M_\star = 1$.

The rest of the paper is organized as follows. In Section 2, we describe the method given by Chan et al. (2006) on how to solve Equation (5) using a direct integral via the Green function method. Our method, though essentially follows the one given by Chan et al. (2006), differs in that we use a direct summation method on a uniform grid in the radial direction, which is the same grid used in our hydro code. In addition, we present two approaches to circumvent the singularity in Equation (5) and two methods to calculate the force field. For comparison, we have also implemented a two-dimensional tree code with a simplified three-dimensional treatment to calculate the self-gravity for disks with vertical structures. In Section 3, we present an efficient parallel implementation scheme on distributed memory computers. We describe a new algorithm to calculate the gravity force at an arbitrary point. We also present numerical test results to compare different approaches. A few concluding remarks are given in Section 4.

2. GREEN FUNCTION METHOD

We present here a numerical method to compute integral (5). For a zero-thickness nonaxisymmetric disk, the classical FFT method (e.g., Binney & Tremaine 1987) based on a polar grid, which has been implemented by Baruteau & Masset (2008), is often used. However, as pointed out by Huré & Pierens (2005), the FFT method has few drawbacks. First, it requires a grid with logarithmic radial spacing, which could be inconvenient to most hydro-solvers using uniform spacing. Second, to avoid the well-known alias issue, the FFT method requires to double the number of cells along the radial direction. The FFT calculation is thus done on a grid that has twice the extent of the hydrodynamic grid along the radial direction,

which induces some complications in the calculation of the convolution kernels, as well as many communications between both grids. Furthermore, with the presence of a vertical structure, it is impossible to apply the FFT method in three-dimensional cylindrical coordinates because neither the potential nor the gravitational field can be represented as convolution products in z . There is no such coordinate transformation as described in Binney & Tremaine (1987) to make the integral (5) become a convolution product in all coordinates.

The Green function method given by Chan et al. (2006) avoids the FFT in the radial direction. Instead, a pseudospectral method on a scaled cosine radial grid is used to achieve the high-order accuracy. Moreover, a known, time-independent vertical structure is easy to incorporate. In the following, we describe modifications to their method so that it can be directly applied to a uniform radial grid.

2.1. Modified Method Given by Chan et al. (2006)

Chan et al. (2006) proposed a direct integral method to solve Equation (5). For the sake of completeness, we recap the key steps in their method here. First, we introduce a softening ε to Equation (5) and denote

$$\begin{aligned} \mathcal{G}(r, r', \phi - \phi') \\ = \int_{-\infty}^{\infty} -\frac{Z(r', z')}{\sqrt{r^2 + r'^2 - 2rr' \cos(\phi - \phi') + \varepsilon^2 + z'^2}} dz', \end{aligned} \quad (6)$$

where ε can be zero or a radius-dependent parameter that will be discussed later in Section 2.2. Function $\mathcal{G}(r, r', \phi - \phi')$ can be computed either analytically or by numerical quadrature. For the special case of $Z(r, z)$ defined by (4),

$$\mathcal{G}(r, r', \phi - \phi') = -\frac{e^{R^2/4} K_0(R^2/4)}{\sqrt{2\pi} H(r')},$$

where $R^2 = (r^2 + r'^2 - 2rr' \cos(\phi - \phi') + \varepsilon^2)/H^2(r')$, and K_0 denotes the modified Bessel function of the second kind. Let

$$I(r, r', \phi - \phi') = 2\pi r' \mathcal{G}(r, r', \phi - \phi'). \quad (7)$$

Then, Equation (5) becomes

$$\Psi(t, r, \phi) = \int_{r_{\min}}^{r_{\max}} \frac{1}{2\pi} \int_0^{2\pi} \Sigma(t, r', \phi') I(r, r', \phi - \phi') dr' d\phi'. \quad (8)$$

Note that both Σ and I are periodic functions with period 2π , and they are in a natural convolution representation in Equation (8). Applying the Fourier transform to (8) with respect to (w.r.t.) ϕ , and using the convolution theorem, we obtain

$$\hat{\Psi}_m(t, r) = \int_{r_{\min}}^{r_{\max}} \hat{\Sigma}_m(t, r') \hat{I}_m(r, r') dr', \quad m \in [-\infty, +\infty], \quad (9)$$

where \hat{f}_m represents the coefficients in the Fourier series expansion of f , which is given as

$$\hat{f}_m = \frac{1}{2\pi} \int_0^{2\pi} f(\phi) e^{-im\phi} d\phi.$$

Note that our Equation (9) is slightly different from Equation (39) in Chan et al. (2006), since we have put all the relevant coefficients in the representation of I in Equation (7).

The integral over r in Equation (9) was evaluated using the Chebyshev spectral method on a ‘‘Chebyshev-roots grid’’ by Chan et al. (2006). However, for most hydro-codes that use a uniform grid, it is desirable to use the same uniform grid for both hydro and self-gravity solutions to avoid interpolation between the hydro solver and self-gravity solver (which is inconvenient and introduces interpolation error).

We propose to integrate (9) directly with numerical quadrature using the available discrete values of $\hat{\Sigma}_m$ and \hat{I}_m on a uniform grid. Note that $I(r, r', \phi - \phi')$ has an analytic expression and is not changed with time as long as the vertical structure remains the same. As proposed by Chan et al. (2006), we can precompute its Fourier transform, $\hat{I}_m(r, r')$, to speed up the algorithm. The whole algorithm can be summarized as follows:

1. For each pair (r, r') , we precompute $\hat{I}_m(r, r')$. Take the Fourier transform of $I(r, r', \phi - \phi')$, which is defined in Equation (7), w.r.t. $\phi - \phi'$,

$$\hat{I}_m(r, r') = \text{FFT}(I(r, r', \phi - \phi')), \quad m = 1, 2, \dots, N_\phi,$$

where N_ϕ is the number of cells in the ϕ -direction. Since $I(r, r, \phi - \phi')$ is a real and even function w.r.t. $\phi - \phi'$, only the discrete cosine transform is needed and $N_\phi/2$ modes need to be stored.

2. Take the Fourier transform of $\Sigma(r', \phi')$ w.r.t. ϕ' to obtain $\hat{\Sigma}_m(r')$, $m = 1, 2, \dots, N_\phi$.
3. Calculate Equation (9) by numerical quadrature in the radial direction to obtain $\hat{\Psi}_m(r)$. Either the midpoint or the trapezoidal rule can be used, depending on the grid point distribution of the source grid.
4. Take the inverse Fourier transform of $\hat{\Psi}(r)$ w.r.t. ϕ to obtain $\Psi(r, \phi)$.

We should emphasize that the precomputing of $\hat{I}_m(r, r')$ plays a major role in the efficiency of our self-gravity solver. This step can be calculated once for all during a long time simulation of the disk–planet interaction system. We find that it costs at least 50 times more than the rest of steps. One cannot afford to calculate it in every time step.

The overall computational cost is $O(N_r N_\phi \log N_\phi + N_\phi N_r^2)$, where N_r is the number of cells in the radial direction. Again, these steps are essentially the same as in Chan et al. (2006), except that the integration is done on a uniform polar grid using direct summation.

2.2. Treatment of Singularity

When $\varepsilon = 0$, the integrand in Equation (6) contains a singularity. Note that the density splitting method of Pierens & Huré (2005b) cannot be used in our case because the density profile is unknown during the precomputing of $\hat{I}_m(r, r')$.

We have tested two approaches to solve the singularity issue. One approach is to use a nonzero softening, $\varepsilon(r')$. We propose a function that scales roughly linearly with Δr , that is, $\varepsilon(r') = \alpha(r')\Delta r$. The idea is that the mass distribution from the small spheres at each cell center should approximate a smooth mass distribution instead of a sum of δ -functions. We optimize $\alpha(r')$ by minimizing the relative error in $\Psi(r, \phi)$ for sharply peaked Gaussian mass distributions located at r' ranging from 0.6 to 1.8 (scaled value for the disk–planet problem) with

$c_s = 0.05$. We find a piecewise linear function for $\alpha(r)$:

$$\alpha(r) = \begin{cases} 0.17 + \frac{0.06}{0.6}(r - r_{\min}), & r < 1.0 \\ 0.23 + \frac{0.03}{0.2}(r - 1.0), & 1.0 \leq r < 1.2 \\ 0.26 + \frac{0.04}{0.3}(r - 1.2), & 1.2 \leq r < 1.5 \\ 0.30 + \frac{0.03}{0.2}(r - 1.5), & 1.5 \leq r < 1.7 \\ 0.33 + \frac{0.04}{0.3}(r - 1.7), & 1.7 \leq r \end{cases} \quad (10)$$

Our experiments show that $\alpha(r)$ is fairly insensitive to different resolutions in the radial direction. Baruteau & Masset (2008), however, have argued that the softening must scale with r . This might be because the logarithmic grid in the radial direction is used in their FFT implementation. For the uniform grid along the radial direction, we find that even a constant α , for example $\alpha(r) = 0.23$, yields good results for the density peak near $r = 1$.

We remark that our choice of the softening should not be mixed with the softening that is used for calculating the gravity from the planet. It has been pointed out by Nelson & Benz (2003a) that the softening used between the disk and planet should be at least 0.75 times the physical size of the grid zone. Baruteau & Masset (2008) suggested using $\varepsilon(r) = 0.3H(r)$ for both disk self-gravity and the gravity between disk cells and planet. We find that this choice is too big and produces large error for disk self-gravity in our implementation. Therefore, we use two different softenings in our simulations: for the disk self-gravity $\varepsilon = \alpha(r)\Delta r$ and for the gravity between the disk and planet $\varepsilon = 0.1H(r_p)$ (r_p is the planet position in the radial direction). The coefficient 0.1, which is still much smaller than the values generally used in the literature, is chosen partially due to the presence of a vertical structure, and it very well matches with the Tanaka et al. (2002) theory on the torque evaluation.

A question arises that the different softening choices may introduce inconsistency between the disk and planet. This concern can be alleviated by the fact that the gravity force for disk cells near the planet are dominated by the planet gravity. Therefore, this inconsistency does not have much impact on the dynamics of the disk.

To eliminate the dependence on the softening, we have implemented the other approach, which is to use different grids for r and r' so that r'_i is never equal to r_j . This approach has been used by Chan et al. (2006) for their spectral method. In our case, the integral of Equation (9) is performed on a fixed uniform grid of r' . Because r' is the cell-centered grid, we define r as the node-centered grid, that is,

$$r_0 = r'_1 - 0.5\Delta r, \quad r_i = r'_i + 0.5\Delta r, \quad i = 1, 2, \dots, N. \quad (11)$$

Note that the r grid has one more point than the r' grid. After calculating $\mathcal{G}(r, r', \phi - \phi')$ for r at the node center, we can obtain $\mathcal{G}(r, r', \phi - \phi')$ for r at the cell center by interpolation, and then calculate the potential $\Psi(r, \phi)$ at the cell center. We can also directly calculate the potential $\Psi(r, \phi)$ with r at the node center for use of the field calculation.

2.3. Field Calculation

We have implemented two approaches to calculate the field components. The first approach is the difference method: we calculate the potential first and then use the finite-difference method to calculate the field.

If the potential is calculated at the cell center (r_i, ϕ_j) , we use the central difference to obtain the field:

$$\Psi_r(r_i, \phi_j) = \frac{\Psi(r_{i+1}, \phi_j) - \Psi(r_{i-1}, \phi_j)}{2\Delta r} \quad (12)$$

$$\Psi_\phi(r_i, \phi_j) = \frac{\Psi(r_i, \phi_{j+1}) - \Psi(r_i, \phi_{j-1})}{2r_i \Delta\phi}, \quad (13)$$

where the notation $(\cdot)_\phi = \partial(\cdot)/(r\partial\phi)$. Note that we need values of $\Psi(r_0, \phi)$ and $\Psi(r_{N+1}, \phi)$ for the Ψ_r component at the boundary cells located at r_1 and r_N , respectively; otherwise, one-side finite difference, which is not very accurate, should be used.

If the potential is calculated at the edge center $(r_{i+\frac{1}{2}}, \phi_j)$, the field can be calculated as

$$\begin{aligned} \Psi_r(r_i, \phi_j) &= \frac{\Psi(r_{i+\frac{1}{2}}, \phi_j) - \Psi(r_{i-\frac{1}{2}}, \phi_j)}{\Delta r} \\ \Psi_\phi(r_i, \phi_j) &= \frac{\Psi(r_{i+\frac{1}{2}}, \phi_{j+1}) + \Psi(r_{i-\frac{1}{2}}, \phi_{j+1}) - \Psi(r_{i+\frac{1}{2}}, \phi_{j-1}) - \Psi(r_{i-\frac{1}{2}}, \phi_{j-1})}{4r_i \Delta\phi}. \end{aligned} \quad (14)$$

High-order methods that use wide stencils are also possible. In fact, we can obtain the derivative Ψ_ϕ with only one extra FFT in the above algorithm of calculating the potential. After obtaining $\hat{\Psi}_m(r)$, we take the inverse transformation of $-im\hat{\Psi}_m(r)$ to get Ψ_ϕ .

The second approach for the field calculation is the integral method: we directly convolute the field Green function with the mass distribution. The differentiation in Ψ_r and Ψ_ϕ can be directly applied to the Green function, which yields

$$\Psi_r(r, \phi) = \int_{r_{\min}}^{r_{\max}} \int_0^{2\pi} \Sigma(t, r', \phi') r' \mathcal{G}_r(r, r', \phi - \phi') dr' d\phi' \quad (16)$$

$$\Psi_\phi(r, \phi) = \int_{r_{\min}}^{r_{\max}} \int_0^{2\pi} \Sigma(t, r', \phi') r' \mathcal{G}_\phi(r, r', \phi - \phi') dr' d\phi'. \quad (17)$$

Since we have two field components, Ψ_r and Ψ_ϕ , the integral method takes twice the computation time of the difference method. But we expect that the integral method might have higher accuracy than the finite-difference approximation. However, if the mode cutoff approach, which will be described in Section 3.2, is used, it is possible that the integral method may lose this advantage. In fact, we observed in our numerical tests that the integral method of (16) and (17) with the mode cutoff gave larger errors and more than double the computation and communication times of the finite-difference method.

2.4. Tree-Code Force Calculation

We explored yet another approach to solve for the field involved using a hierarchical tree code. In this approach, we treat each cell as a particle, and hence an effective force law between two points, (r, ϕ) and (r', ϕ') , in the disk that accounts for the vertical structure is required. Just as before, we assume $\rho(\mathbf{x}) = \Sigma(r, \phi)Z(r, z)$. And since symmetry arguments require $Z(r, z)$ to be an even function of z , the z -component of the force must be zero at the $z = 0$ plane. Therefore, the force is directed to the plane of the disk and it has a magnitude factor given by

$$F(d) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{Z(r, z)Z(r', z')d}{(d^2 + (z - z')^2)^{3/2}} dz dz', \quad (18)$$

where $d = (r^2 + r'^2 - 2rr' \cos(\phi - \phi'))^{1/2}$.

If we further assume that $Z(r, z)$ is given by Equation (4), then we can rewrite Equation (19) as

$$F(d; H, H') = \frac{C_{3D}(d; H, H')}{d^2}, \quad (19)$$

where

$$C_{3D}(d; H, H') = \frac{d^3}{2\pi H H'} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{e^{-z^2/2H^2} e^{-z'^2/2H'^2}}{(d^2 + (z - z')^2)^{3/2}} dz dz' \quad (20)$$

is called the three-dimensional correction factor with respect to the two-dimensional force without the vertical structure. As $H \rightarrow 0$, (20) approaches the three-dimensional factor defined by Koller (2004).

We note that this force law is a line–line force (giving the net interaction between the mass distributed along two vertical lines at (r, ϕ) and (r', ϕ') , respectively). This is different from what is done in the Green function method, which uses a point–line interaction. By setting $H = 0$ in Equation (19), we can recover the interaction force law of the Green function method. However, if we wished to include line–line interactions in the Green function method, we would have to abandon the more efficient potential calculation of Equation (6) and instead use Equations (16) and (17) with \mathcal{G}_r and \mathcal{G}_ϕ defined appropriately.

Rescaling the integration in Equation (20) and introducing the nondimensional variables $d' = \frac{d}{\sqrt{H^2 + H'^2}}$ and $\alpha = \frac{H'}{H}$ give

$$C_{3D}(d'; \alpha) = \frac{d'^3}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{e^{-z^2/2} e^{-z'^2/2}}{\left(d'^2 + \frac{(z - z'\alpha)^2}{1 + \alpha^2}\right)^{3/2}} dz dz'. \quad (21)$$

Finally, we perform an orthogonal coordinate transformation using the variables

$$\begin{aligned} \zeta &= \frac{1}{\sqrt{1 + \alpha^2}}(z - z'\alpha) \\ \zeta' &= \frac{1}{\sqrt{1 + \alpha^2}}(\alpha z + z'). \end{aligned}$$

This allows us to simplify Equation (21) to

$$C_{3D}(d') = \frac{d'^3}{2\pi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{e^{-\zeta^2/2} e^{-\zeta'^2/2}}{(d'^2 + \zeta^2)^{3/2}} d\zeta d\zeta', \quad (22)$$

revealing that it is actually independent of α . In fact, similar to Equation (6), we can express Equation (22) as

$$C_{3D}(d') = -\frac{d'^3}{2\sqrt{2\pi}} e^{d'^2/4} (K_0(d'^2/4) + K'_0(d'^2/4)), \quad (23)$$

where K_0 and K'_0 denote the modified Bessel function of the second kind and its derivative, respectively.

The tree code requires the calculation of $C_{3D}(d')$ and its first and second derivatives over a wide range of d' (ranging from 10^{-2} to 10^3). Calculating Equation (23) on the fly is prohibitively expensive and should be avoided. Another option is to precompute Equation (23) and its derivatives at a discrete set of d' values. Then when $C_{3D}(d')$ is needed, we interpolate the precomputed data. This approach can provide a significant speedup, but because the data are stored in main memory, it is still too slow for our purposes.




Figure 1. Three-dimensional correction factor (Equation (23)) and relative error plotted vs. nondimensional distance (d') using model (24).




Figure 2. Three-dimensional force factor (Equation (19)) and relative error plotted vs. nondimensional distance (d') using model (24).

Another approach we explored involves approximating Equation (23) with a model function that is accurate yet inexpensive enough to compute on the fly. Plotting $C_{3D}(d')$ reveals its form (Figure 1) and suggests a model function of the form

$$C_{3D}^{(\text{model})}(d') = \frac{(d'/d'_{1/2})^p + (d'/d'_{1/2})^{2p}}{2 + (d'/d'_{1/2})^p + (d'/d'_{1/2})^{2p}}. \quad (24)$$

We find optimal parameter values to be $d'_{1/2} = 0.8252$ and $p = 0.957$. The relative error introduced by using this model is plotted with C_{3D} and F in Figures 1 and 2, respectively. Since generally $d' = d/H(r) > \varepsilon/H(r)$ with $H' = 0$, the model is acceptable for distance or softening larger than $0.005H(r)$. This model requires the calculation of one power function, a few additions and multiplications, and one division. For our application, we found this to be the most efficient way to include the three-dimensional structure using a tree code.

Our tree code is implemented in a distributed memory parallel architecture using MPI, just like our Green function FFT algorithm. There are many papers detailing parallel tree codes (Dubinski 1996; Mocchi & Capuzzo-Dolcetta 2002). So, we simply highlight the important points of our algorithm here.

1. We begin by precomputing the tree structure of our grid, which does not change with time. Each processor retains a copy of this *quad-tree* to be used throughout the simulation.

2. All processors compute partial sums (using local particles) of the moments for all nodes in the tree.
3. These partial sums are added and broadcast back to all the processors.
4. A complete copy of the density profile is distributed to all the processors. This is required for load balancing and computing the direct summation part of the force (Step 5(c)).
5. Forces are computed at every grid cell. The disk is partitioned azimuthally (the symmetry providing ideal load balancing), and processors are assigned equal-sized sectors of particles (grid cells). Each particle traverses the *quad-tree* to include interactions with all other particles.
 - (a) If a particle–node interaction meets the accuracy criterion, then the particle interacts with that node via the moments computed in Step 3.
 - (b) If the accuracy criterion is not met (the node is too large or the particle–node separation too small), then we drop one level in the tree and check all the rejected node’s subnodes for particle–node interactions.
 - (c) Last, if we reach a leaf node (a node with no subnodes) and the accuracy criterion is not met, we perform direct summation over all the leaf-node’s particles.

3. NUMERICAL IMPLEMENTATION AND EXPERIMENTS

3.1. Description of the Test Problem

As in Chan et al. (2006), we consider the density function to be a Gaussian sphere given by

$$\rho(r, \phi, z) = \frac{M}{(2\pi\sigma^2)^{3/2}} \exp\left(-\frac{r^2 + z^2}{2\sigma^2}\right),$$

where M is the normalized total mass, σ is a parameter that controls the width of the mass distribution, and it has a similar role as the scale height in Equation (4). The potential on the $z = 0$ plane has been given by Chan et al. (2006):

$$\psi(r, \phi) = -\frac{1}{r} \operatorname{erf}\left(\frac{r}{\sqrt{2}\sigma}\right), \quad (25)$$

where $\operatorname{erf}(x)$ is the error function

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt.$$

For a collection of Gaussian spheres centered on the $z = 0$ plane with the same σ , the z -dependent vertical structure can be factored out, that is,

$$\rho(r, \phi, z) = \sum_i \rho_{(r_i, \phi_i)}(r, \phi, z) = Z(r, z) \sum_i \Sigma_{(r_i, \phi_i)}(r, \phi),$$

where (r_i, ϕ_i) is the center of each sphere, $Z(r, z)$ gives the vertical structure

$$Z(r, z) = \frac{1}{\sqrt{2\pi}\sigma^2} \exp\left(-\frac{z^2}{2\sigma^2}\right), \quad (26)$$

$\Sigma_{(r_i, \phi_i)}$ is the surface density

$$\Sigma_{(r_i, \phi_i)} = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{R_i^2}{2\sigma^2}\right),$$

and $R_i(r, \phi) = \sqrt{r^2 + r_i^2 - 2rr' \cos(\phi - \phi_i)}$ is the distance between (r, ϕ) and (r_i, ϕ_i) . Notice that Equation (26) is the same as Equation (4) with a constant scale height $H(r) = \sigma$.

We use three Gaussian spheres located at $(r, \phi) = (1, 0)$, $(0.9, 3\pi/4)$, and $(1, -\pi/2)$ with a total mass of 2, 1/2, and 1, respectively, that is,

$$\Sigma_{\text{ana}}(r, \phi) = 2\Sigma_{(1,0)}(r, \phi) + \frac{1}{2}\Sigma_{(0.9,3\pi/4)}(r, \phi) + \Sigma_{(1,-\pi/2)}(r, \phi), \quad (27)$$

where $(*)_{\text{ana}}$ represents the analytic solution. The exact potential will be

$$\Psi_{\text{ana}}(r, \phi) = 2\psi(R_{(1,0)}) + \frac{1}{2}\psi(R_{(0.9,3\pi/4)}) + \psi(R_{(1,-\pi/2)}), \quad (28)$$

where ψ is defined by Equation (25). The exact fields Ψ_r and Ψ_ϕ can also be calculated accordingly.

The tests are done in the computational domain $[0.4, 2.0] \times [0, 2\pi]$ with grid $N_r \times N_\phi$. The normalized total mass for the disk is $M = 0.002M_*$. We choose $N_\phi = 4N_r$, so that the cells near $r = 1$ are nearly square. Without specification, we always use $N_r = 800$, which reaches the convergence in both the torque calculation on the planet and the azimuthal-averaged

potential vorticity distribution on the disk in our disk–planet interaction simulations (Li et al. 2005). Since the surface density (27) approaches zero outside the computational domain, the exact potential (28) is still valid for our truncated domain. For convenience, we use the following notations for comparison: N_p denotes the number of processors, E_{max} denotes the maximum error over the whole domain, RE_{max} denotes the maximum relative error, and p is the convergence order in E_{max} . Since the force is a vector that has two components, we use the following formula to calculate E_{max} :

$$E_{\text{max}} = \sqrt{|f^r - f_{\text{ana}}^r|^2 + |f^\phi - f_{\text{ana}}^\phi|^2},$$

where f^r and f^ϕ are the force components in r - and ϕ -directions, respectively. The local relative error for a variable u is defined as $\text{RE}_{\text{max}}(u) = E_{\text{max}}/|u_{\text{ana}}|$. For the force, $|u_{\text{ana}}| = \sqrt{(f_{\text{ana}}^r)^2 + (f_{\text{ana}}^\phi)^2}$. The global relative error for a variable u is defined as

$$\text{RE}(u) = \frac{\sum_{j=1}^{N_\phi} \sum_{i=1}^{N_r} |u_{i,j} - u_{i,j,\text{ana}}|}{\sum_{j=1}^{N_\phi} \sum_{i=1}^{N_r} |u_{i,j,\text{ana}}|}. \quad (29)$$

All of the computation are performed on a parallel Linux cluster at the Los Alamos National Lab. Each node of the cluster is a dual-core AMD Opteron™ processor with 2.8 GHz and 2 GB local memory.

3.2. Parallel Implementation and Comparison

Our hydro simulation for the interaction of disk and protoplanet problem is performed on a high-resolution grid, for example, $N_r \times N_\phi = 800 \times 3200$ grid. The whole domain is split into annular regions for parallel computation. Each annular region has the same number of cells in the radial direction.

The Fourier transform and its inverse can be parallelized without much difficulty because they require no communication between different processors. However, the numerical quadrature of Equation (9) is done in the radial direction, where each processor holds only part of the information about the density distribution in the whole domain.

We have tested two approaches to parallelize the numerical quadrature. In the first approach, we start by computing a partial quadrature including source terms only from the local density distribution, that is, each processor computes a partial sum of Equation (9) for all r and m . We then apply a global communication to obtain the complete summation.

In the second approach, we begin by performing a global communication over the whole domain so that every processor has a complete copy of the density profile for $r' \in [r_{\text{min}}, r_{\text{max}}]$. Then, the complete numerical quadrature (9) can be performed to obtain $\hat{\Psi}_m(r)$ for the local portion of the r grid.

We remark that in the precomputing stage, no matter which approach is used, each processor can compute its own portion of the Fourier transform $\hat{I}_m(r, r')$ independently and store it for later use. It does not need to have a copy of the whole profile for all r and r' .

In the following, we will first propose a strategy to reduce the parallel communication cost and then compare the above two approaches in calculating potential $\Psi(r, \phi)$. For simplicity, we denote the first approach as approach I and the second as approach II.

3.2.1. Mode Cutoff Approach to Reduce the Communication Cost

The communication cost in both approaches is proportional to both the number of processors and the amount of data being communicated (total grid size). Consequently, we expect that the communication cost will dominate the computation cost as the number of processors is increased, eventually becoming a bottleneck for a large number of processors. Note that the communication is done in the frequency domain, where the high modes decay exponentially for a smooth function. Thus, we can truncate the Fourier modes above a cutoff parameter (M_{cut}) without a significant loss of accuracy. If M_{cut} is far smaller than N_ϕ (the total number of modes), the communication cost can be reduced to a small fraction of the cost before the cutoff.

For both approaches, we can calculate M_{cut} based on the Fourier transformation of the density profile. For a specific radius r_i , we define the energy contained in $\hat{\Sigma}(r_i)$ as $\sum_{m=0}^{N_\phi} |\hat{\Sigma}_m(r_i)|^2$. Each processor then computes a $m_{\text{cut}}(r)$ based on satisfying some preset fraction ϵ_{cut} (e.g., $\epsilon_{\text{cut}} = 10^{-4}$) of the total energy for modes above $m_{\text{cut}}(r)$. Next, each processor finds its maximum m_{cut} over its range of r . Last, all the processors compare their m_{cut} values to obtain the global maximum M_{cut} that is used to truncate the Fourier series. Note that M_{cut} calculated in this way does not vary with the number of processors.

For approach I, $m_{\text{cut}}(r)$ can also be evaluated based on $\hat{\Psi}(r)$ instead of $\hat{\Sigma}(r)$, because the communication is done after the partial summation. To be more accurate in the field calculation, we can instead evaluate $m_{\text{cut}}(r)$ based on the energy defined by $\sum_{m=0}^{N_\phi} |m\hat{\Psi}_m(r)|^2$. The extra factor of m is included so that we obtain the energy of the force, which is the derivative of Ψ . Notice, however, that if we compute m_{cut} and M_{cut} as described above, M_{cut} will vary with the number of processors. This is because each processor computes only a partial sum of $\hat{\Psi}_m(r)$. To produce a relatively constant M_{cut} , after each processor finds its maximum m_{cut} , all the processors take an energy-weighted average of their m_{cut} to obtain the global M_{cut} . It is observed that the energy-weighting average gives an M_{cut} that remains nearly constant while the number of processors varies between 1 and N_p . For verification reason, we can require that M_{cut} be independent of the number of processors. However, this requires a costly global communication between different processors. Nonetheless, we find that M_{cut} will only increase slightly with the number of processors, which means that increasing the number of processors will not degrade the numerical accuracy. We define a preset fraction ϵ_{cut} similar to the previous density-mode cutoff approach. After extensive experiments, we observe that $\epsilon_{\text{cut}} \approx 10^{-7}$ in $\hat{\Psi}(r)$ cutoff produces results similar to $\epsilon_{\text{cut}} = 10^{-4}$ in $\hat{\Sigma}(r)$ cutoff.

For energy distributions in disk–planet simulations, the energy of the zero mode ($m = 0$) often dominates that of all other modes combined. Thus, when calculating the energy using $\hat{\Sigma}$ or $\hat{\Psi}$, it is convenient to exclude the zero mode from the total energy. Notice that if the energy is defined as $|m\hat{\Psi}_m(r)|^2$, the zero mode is naturally excluded. To remove the numerical noise, we multiply the energy of the zero mode by a small number (e.g., ϵ) and add it to the total energy. In our disk–planet simulations, the initial density distribution is axisymmetric, and, hence, the zero mode is enough. With the time evolution of the surface density, M_{cut} begins to increase quickly until it reaches a certain number, which is insensitive to different resolutions and different initial power-law density profiles. However, M_{cut} does vary a lot with

the planet mass and sound speed. As an example, we obtain the following data for simulations of an isothermal disk with different planet mass (M_p) and sound speed (c_s):

$$M_{\text{cut}} = \begin{cases} 190, & \text{for } M_p = 3 \times 10^{-5} M_\star, c_s = 0.05 \\ 297, & \text{for } M_p = 3 \times 10^{-4} M_\star, c_s = 0.05 \\ 298, & \text{for } M_p = 3 \times 10^{-5} M_\star, c_s = 0.035 \end{cases}.$$

Note that M_{cut} is much smaller than the total number of modes in the ϕ -direction, $N_\phi = 3200$.

3.2.2. Impact of the Cutoff Modes and Cutoff Thresholds

Here, we study how the number of cutoff modes, M_{cut} , impacts the computational efficiency and accuracy. We test the problem with parallel computation using 100 processors. We will show the results only for approach II. Similar results have also been obtained for approach I.

Figure 3 shows the variation in the computation and communication time with different numbers of cutoff modes. The total cost displayed in terms of the CPU time includes the computation of one FFT of the density field, the global communication of the density spectrum up to the cutoff mode M_{cut} right after the FFT, and a quadrature calculation in the radial direction. Figure 3 shows that the communication time increases at a slower rate with M_{cut} than the computation time.

The energy cutoff approach depends on the cutoff threshold ϵ_{cut} . We also vary the size of ϵ_{cut} to see how it impacts the cutoff mode. Figure 4 shows the simulation results for different ϵ_{cut} . Figures 3 and 4 show that our cutoff threshold, $\epsilon_{\text{cut}} = 10^{-4}$, which corresponds to $M_{\text{cut}} = 153$, appears to be the optimal number to balance accuracy and efficiency for this problem. This energy cutoff threshold is insensitive to grids with different resolutions.

3.2.3. Parallelization via Domain Decomposition of the Source Grid: Approach I

Approach I is implemented via domain decomposition of the source grid r' . In terms of the communication cost, it involves a global reduction of $N_r M_{\text{cut}}$ data from every processor and a distribution of $N_r M_{\text{cut}} / N_p$ data to each processor, where N_p is the number of processors. Here, we use a fixed mode cutoff number $M_{\text{cut}} = 153$.

Figure 5 shows how the communication and computation cost vary with the number of processors. Note that the communication cost remains relatively constant no matter how many processors we use. This is a surprise, because we expect that the communication cost is proportional to the length of the communicated data, which is proportional to the number of processors. We remark that one might obtain a different performance for a different MPI implementation or a different parallel cluster. In year 2007, we observed using the same code that the communication cost linearly increases with the number of processors. In year 2008, our parallel cluster has been upgraded with new parallel software and a new interconnection. The data shown in Figure 5 are calculated on the new cluster.

3.2.4. Parallelization via Domain Decomposition of the Field Grid: Approach II

Approach II described above is implemented via the domain decomposition of the field grid r . It involves a global gathering of $N_r M_{\text{cut}} / N_p$ data from every processor and a broadcast of $N_r M_{\text{cut}}$




Figure 3. The computation and communication cost for evaluating $\hat{\Psi}(r, \phi)$ (left) and the error variation (right) as a function of M_{cut} .




Figure 4. The computation and communication cost for evaluating $\hat{\Psi}(r, \phi)$ (left) and the error variation (right) as a function of energy cutoff thresholds ϵ_{cut} . The cutoff is based on the energy of $\hat{\Sigma}$.




Figure 5. The communication and total cost for evaluating $\hat{\Psi}(r, \phi)$ varying with the number of processors in MPI. $M_{\text{cut}} = 153$.

data to each processor. While the amount of communicated data is the same for both approaches, approach I also involves a global summation operation resulting in a total communication cost of about 10% more than approach II with the same value M_{cut} . Again, we show some results using only a fixed cutoff mode number $M_{\text{cut}} = 153$.

Figure 5 (approach II) shows the performance comparison with the first approach. Similar to the results of approach I described in the last subsection, the communication cost remains nearly constant. It is clear that approach II (this approach) is faster than approach I. Therefore, we will use approach II as our choice of the methods in the tests hereafter.

Figure 6 shows the parallel efficiency for different numbers of processors ($N_p = 2^n, 0 \leq n \leq 7$) and grids with different resolutions. The parallel efficiency is defined by

$$E(N_p) = \frac{T_{\text{seq}}(1)}{N_p T(N_p)},$$

where $T(N_p)$ is the runtime of the parallel algorithm and $T_{\text{seq}}(1)$ is the runtime of the sequential algorithm using one processor. For the fine resolution grid 1024×4096 , we replace $T_{\text{seq}}(1)$ with $4T(4)$ due to the memory limitation of the processors. It is clear that for a smaller number of processors and a larger grid, the parallel efficiency is better. Figure 6 shows that the parallel efficiency can be larger than 1 at certain stages. The reason is not clear to us.

3.3. Comparison of Different Approaches for Potential Calculation

We have tested two approaches for computing the potential and field: one is with the softening Equation (10) and the other




Figure 6. The parallel efficiency for grids with different resolutions and different number of processors. $M_{\text{cut}} = 153$ is used.

Table 1

Numerical Results for the Potential Calculation With and Without Softening Equation (10)

Grid N_r	Force Error			Potential Error		
	E_{max} (1)	RE (10^{-2})	p	E_{max} (10^{-2})	RE (10^{-5})	p
128	2.6514	3.1388		3.2631	7.2648	
256	0.7922	0.9205	1.74	1.7595	3.0652	0.89
512	0.2317	0.2983	1.77	0.6925	1.4737	1.35
1024	0.0725	0.1298	1.68	0.2485	0.9461	1.49
128	3.3427	2.4634		5.5060	5.1819	
256	0.8539	0.6580	1.97	1.4317	1.8465	1.94
512	0.2106	0.2105	2.02	0.3775	0.9995	1.92
1024	0.0592	0.1036	1.83	0.1115	0.7909	1.76

Note. The top half is for a shifted grid with softening and the bottom half is for a shifted grid without softening.

is without softening but with a shifted grid defined by Equation (11). Note that the field is calculated differently using different central-difference methods. The field force with softening is calculated using central-difference Equations (12) and (13). The field force without softening is calculated using central-difference Equations (14) and (15).

To minimize the impact of the cutoff mode, we use a fixed cutoff number $M_{\text{cut}} = 153$, which produces nearly the same results as without a cutoff (see Section 3.2.2). Table 1 shows the numerical errors for both approaches. The convergence order is calculated based on the maximum absolute error. The results show nearly second-order convergence as the mesh is refined. This is in agreement with the accuracy of our method, because both the quadrature rule to calculate the potential and the finite-difference method to calculate the force are of second-order accuracy.

Based on global relative errors (RE in Table 1) for both potential and field, we see that the potential method using the shifted grid without softening is more accurate than with softening. Yet the difference is relatively small.

3.4. Tests of Point Force Calculation

In the previous subsection, we have tested different approaches for disk self-gravity computed at the cell centers. In

this subsection, we describe and compare methods to calculate the gravitational force of the disk on the planet at any point (r_p, ϕ_p) . We assume a unit planet mass at that point.

We investigated the force calculations in three different ways. In the first approach, we performed a direct summation of \mathcal{G}_r and \mathcal{G}_ϕ from the point (r_p, ϕ_p) to each of the cell centers (Equations (16) and (17) with $(r, \phi) = (r_p, \phi_p)$). This is similar to a particle method and is widely used in calculating the disk force exerted on the planet in the disk–planet simulations. In the second approach, we applied bilinear interpolation to the force computed at the cell centers. In the third approach, we applied a finite difference formula directly to Ψ at the grid cells closest to the point location. This approach is very similar to bilinear interpolation of the force; only it uses a smaller grid stencil, which gives better accuracy.

To calculate the force at a specific point and account for the fact that the point can be at an arbitrary location within a cell, we take a cell that contains point $(r_p, \phi_p) = (0.99, 0.0)$, which is close to the center of one of the first Gaussian spheres, and compute the force for 11×11 equidistributed points within the cell. Since the third approach is similar to the second approach and verified to be more accurate, we only compare the first and third approaches.

We use the softening defined by (10) in the direct summation approach to avoid singularities. Figure 7 shows the absolute error and relative error throughout the grid cell. We find that the maximum error and maximum relative error are 6.12 and 16.6%, respectively. The top two plots of Figure 7 show that, although softening (10) works very well for the disk self-gravity calculated at the cell center and point force at the node and edge centers, it gives large error for the point force at other locations. The relative error range is [0.0152%, 16.6%], which means that the direct summation with this softening is very sensitive to the location of the point. We have tried two other different softenings. The middle two plots of Figure 7 show the results with softening $\varepsilon = \min(\Delta r, r\Delta\phi)$, which approximates one grid spacing. Both the maximum error and the maximum relative error are much reduced. Also, the range of the relative error becomes [0.63%, 0.88%], which means that the direction summation with this softening is relatively insensitive to the location of the point. The bottom two plots of Figure 7 show the results with softening $\varepsilon = 0.3H(r)$, which has been adopted by Baruteau & Masset (2008) and corresponds to seven grid spacings in our 800×3200 grid layout. Both the maximum error and the relative error become large everywhere. The error range is [6.04, 7.40] and the relative error range is [16.5%, 16.6%].

The third approach using finite difference on Ψ achieves much high accuracy, with maximum error and maximum relative error of 0.0315% and 0.078%, respectively. Figure 8 shows the absolute error and relative error throughout the grid cell. The relative error range is [0.0264%, 0.0778%], which means that this approach is relatively insensitive to the location of the point.

We remark that in the actual disk–planet interaction simulations, the accuracy of the force on the planet also depends on how accurately the disk density around the planet is resolved. Since the disk within one Roche lobe is usually not well resolved, the third approach in the point force calculation, though by itself is very accurate, may not always give the most “accurate” force calculation if the disk density is poorly determined. In fact, we find that the direct summation approach with a relatively large softening, for example $\varepsilon = 0.1H(r)$, gives results that are in

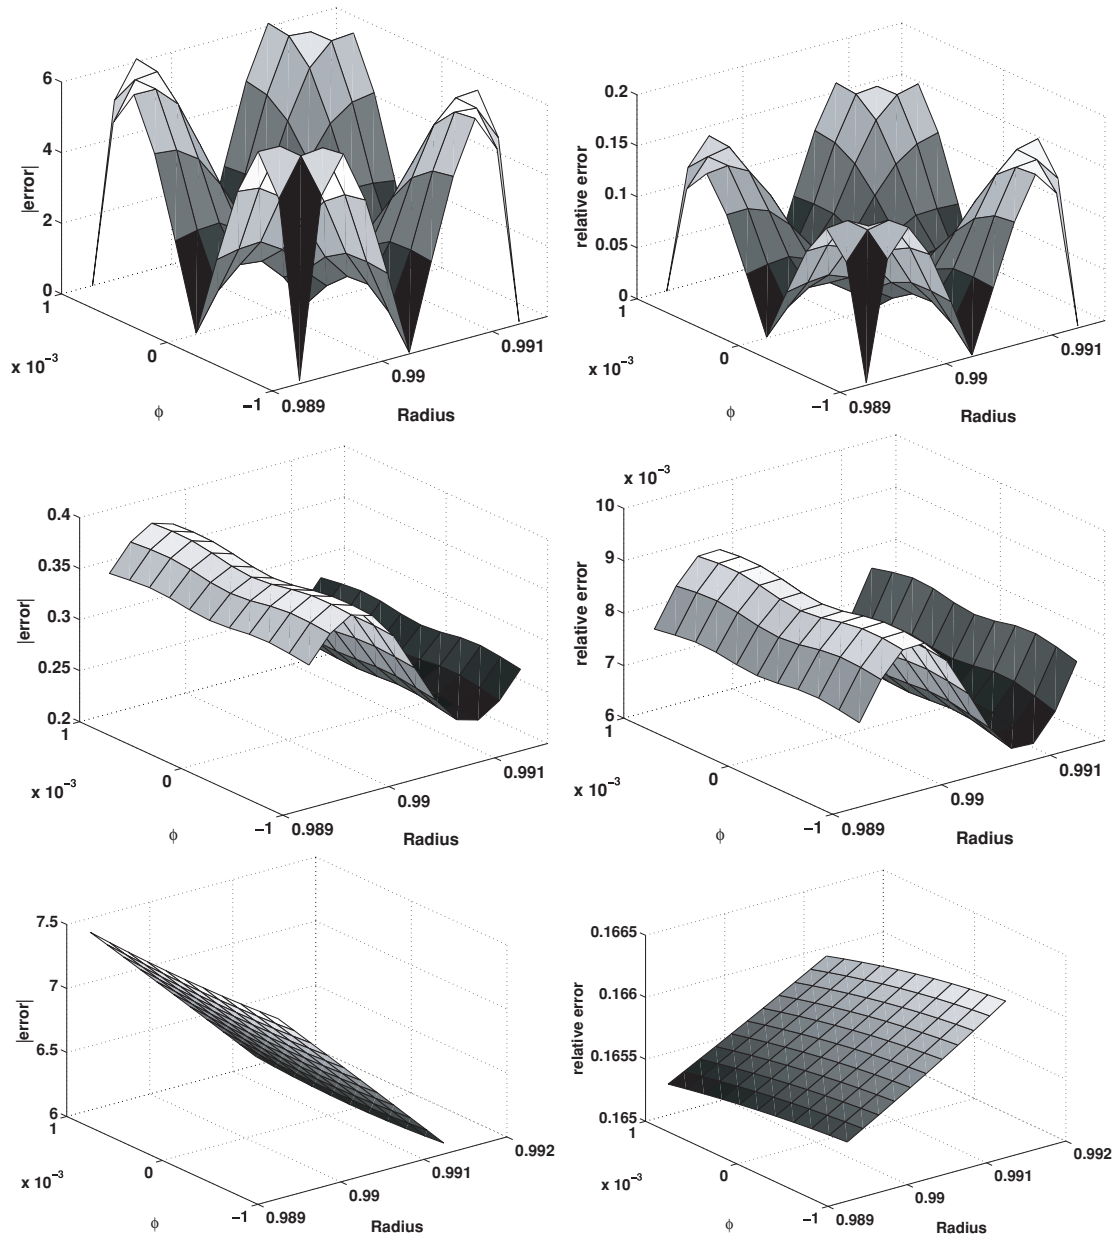


Figure 7. The accuracy of the point force calculation by the direct summation approach with different softenings: softening (10) (top two), $\varepsilon = \min(\Delta r, r\Delta\phi)$ (middle two), and $\varepsilon = 0.3H(r)$ (bottom two). The left plots show the magnitude of the errors and the right plots show the relative errors, both throughout one grid cell.

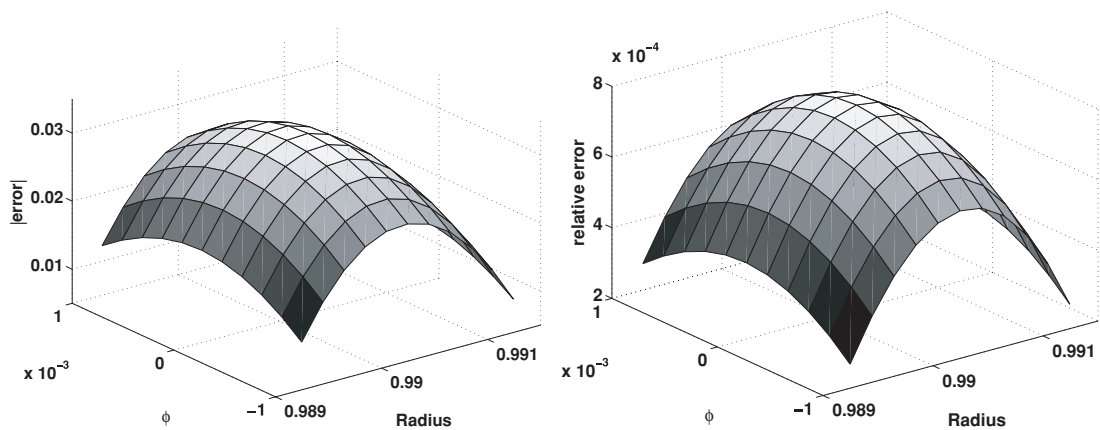


Figure 8. The accuracy of the point force calculation finite difference of the potential. The left plot show the magnitude of the errors and the right plot show the relative errors, both throughout one grid cell.

Table 2
Accuracy and Efficiency Data for Tree-Code Force Calculation on $N_p=100$ Processors

Grid	F_{\max}	E_{\max}	$RE(10^{-2})$	ρ	Total CPU Time (s)
100 × 400	171.68	5.617	1.5216		0.1078
200 × 800	171.72	2.324	0.8452	1.27	0.6844
400 × 1600	171.73	1.390	0.6291	0.74	3.7167
800 × 3200	171.74	1.114	0.5463	0.32	16.543
800 × 3200 ^a	171.74	0.08912	0.0126		0.02093

Note. ^a The last line is for the Green function method, where the CPU time does not include the precomputing time, which takes 1.001 s.

good agreement with the linear theory results by Tanaka et al. (2002).

3.5. Tests of the Tree-Code Force Calculation

Here, we test the performance of our tree-code method described in Section 2.4. First, we test the accuracy with the same error and convergence measures used in Section 3.3. For simplicity, we use a fixed softening $\varepsilon = \alpha(1)\Delta r = 0.23\Delta r$ throughout the grid. The approximate model (24) is used to calculate a three-dimensional correction factor. To obtain the field in the $z = 0$ plane, we set $H = 0$ in Equation (20) resulting in an interaction force consistent with the Green function method. Also, we use the same density function with a known analytical force field described in Section 3. Second, we test the parallel efficiency by varying the number of processors for a highly resolved grid (800 × 3200).

From the results given in Table 2, we draw the following conclusions. First, this method does not appear to converge to the exact solution. This is actually expected since the model force function (Equation (24)) with our chosen softening does not become more accurate as the grid is refined. We observe that the maximum error asymptotically approaches 1.0, which is much larger than the error using the Green function method for highly resolved grids. Next, we see that the time complexity scales approximately with $O(N_r N_\phi \log(N_r N_\phi))$, which is consistent with the theoretical prediction. Finally, we should remark that the softening has a large impact on the accuracy of the solutions for the tree code. If $\varepsilon = 0.3H(r)$ is used, the maximum error and global relative error become 20.57% and 3.045%, respectively.

Parallel efficiency results can be seen in Table 3. Although we find that our tree code scales well with the number of processors, the actual value of the CPU time is more than two orders of magnitude greater than our Green function method! We have experimented with tuning tree-code parameters and other memory optimizations and have found that while it may be possible to gain a factor of 2–4 in speedup over the results of Table 3, our tree-code method always performs *much* slower than our Green function method.

4. CONCLUSION

In this paper, we have presented a fast and accurate solver to calculate the potential and self-gravity forces for the disk systems. This method is implemented on a polar grid, and the FFT can be used in the azimuthal direction. The precalculation of the Green function and its FFT play a major role in the algorithm to reduce the computational cost. We think it could be the main reason why the Green function method is much faster than the

Table 3
Parallel Efficiency Data for Tree-Code Force Calculation on a 800 × 3200 Fixed Grid

N_p	CPU Time		Parallel Efficiency
	Total Time (s)	Comp. Time (s)	
10	182.32	180.97	0.99
20	92.78	91.83	0.98
40	46.92	46.06	0.96
50	37.71	36.85	0.96
80	24.07	23.05	0.94
100	20.07	18.94	0.90
160	13.80	12.25	0.82
200	11.28	9.49	0.80

particle tree-code method. We also presented an efficient method in implementing the solver on parallel computers. We find the computational cost for the self-gravity solver to be comparable with that of the hydro-solver for large, highly resolved grids run on a distributed memory parallel architecture. We also developed a two-dimensional tree-code solver, which uses a relatively inexpensive model force to accurately account for the vertical structure of the disk.

Finally, we notice that if the disk vertical structure varies with time, our precalculation must be done in every time step making our solver inefficient.

We have applied our self-gravity solver to simulations of the disk–planet interaction system. Compared with the simulations without self-gravity, the total computation time is increased by only 30% for a parallel computation with 100 processors. We have also confirmed that the two-dimensional self-gravity indeed accelerates the planet migration. These results will be reported elsewhere.

We would like to thank Dr. C.-K. Chan for helpful discussion during his stay at Los Alamos. We also thank the referee for many useful comments. This research was performed under the auspices of the Department of Energy. It was supported by the Laboratory Directed Research and Development (LDRD) Program at Los Alamos. It is also available as Los Alamos National Laboratory Report, Los Alamos Report LA-UR-07-5882.

REFERENCES

- Baruteau, C., & Masset, F. 2008, *ApJ*, **678**, 483
 Binney, J., & Tremaine, S. 1987, in *Galactic Dynamics*, (Princeton, NJ: Princeton Univ. Press)
 Chan, C. K., Psaltis, D., & Özel, F. 2006, *ApJ*, **645**, 506
 Dubinski, J. 1996, *New Astron.*, **1**, 133
 Goldreich, P., & Tremaine, S. 1979, *ApJ*, **233**, 857
 Goldreich, P., & Tremaine, S. 1980, *ApJ*, **241**, 425
 Huré, J. M. 2005, *A&A*, **434**, 17
 Huré, J. M., & Pierens, A. 2005, *ApJ*, **624**, 289
 Koller, J. 2004, Ph.D. Thesis, Rice Univ.
 Li, H., Li, S., Wendroff, B., Koller, J., & Liska, R. 2005, *ApJ*, **624**, 1003
 Lin, D. N. C., & Papaloizou, J. 1986a, *ApJ*, **307**, 395
 Lin, D. N. C., & Papaloizou, J. 1986b, *ApJ*, **309**, 846
 Mocchi, P., & Capuzzo-Dolcetta, R. 2002, *A&A*, **382**, 758
 Nelson, A. F., & Benz, W. 2003a, *ApJ*, **589**, 556
 Nelson, A. F., & Benz, W. 2003b, *ApJ*, **589**, 578
 Pierens, A., & Huré, J. M. 2005a, *A&A*, **433**, L37
 Pierens, A., & Huré, J. M. 2005b, *A&A*, **434**, 1
 Tanaka, H., Takeuchi, T., & Ward, W. R. 2002, *ApJ*, **565**, 1257
 Ward, W. R. 1997, *ApJ Lett.*, **482**, L211