Hierarchical N-Body Calculations

STEPHEN ROBERTS

Mathematics Department, The Faculties

Australian National University

Abstract. In this paper we describe some of the hierarchical methods that have been used to produce computationally efficient algorithms for the calculation of the mutual interactions of a collection of n particles. The work required for such calculations grows quadratically as the number of particles. In this paper we survey the so called "tree codes" or "hierarchical" methods which provide accurate approximations of the required interactions, but for which the computational work is $O(n \log n)$ for "simple" hierarchical methods, and $O(n)$ for the more sophisticated Fast Multipole method. This survey paper follows closely the article by Greengard ["Numerical Solution of the N-Body Problem", Computers in Physics, 1990].

1. Introduction

It is a very common problem to calculate sums of the following form

$$\begin{aligned} \phi_i &= \sum_{i \neq j} \log |x_i - x_j| q_j \\ E_i &= \sum_{i \neq j} \frac{1}{|x_i - x_j|^p} q_j \end{aligned} \tag{1.1}$$

where the points $x_i \in \mathbb{R}^d$, $d = 2, 3$, $i = 1, .., n$ can be considered as the positions of a collection of n particles. These sums may be interpreted in many ways. For instance, if q_i denotes the mass or charge of the i-th particle, in a collection of n massed or charged particles, then the force on the i-th particle, due to the other n particles will be given by E_i. Using Newton's equations, we can use the forces thus calculated to evolve large collections of particles which simulate large self-interacting systems, such as colliding galaxies (see [2]) or plasmas (see [6]). If q_i denotes the vorticity of a so-called "vortex" particle, then the velocity field due to a collection of n point vortex particles is of the

This work was supported by the Australian Research Council through grant A68830148, "A Theoretical and Numerical Study of Models for the Evolution of Interfaces".

above form. This representation of a fluid has been used by Chorin [5] (and others, see [12]) in a numerical method called the vortex method which is used to simulate turbulent incompressible fluid flow. Finally, if Laplace's equation is formulated as a boundary integral equation and the resulting equation discretized, then sums of the same form once again appear. See the survey article by Greengard [10] for more details on these and other applications.

It is evident that an efficient method for calculating the sums of the form (1.1) is needed, especially as applications routinely need to use on the order of 10^5 or more particles. A direct calculation of the sum (1.1) at each of n particle positions requires $O(n^2)$ floating point operations. Even with the fastest supercomputer available (say a machine capable of 10^9 floating point operations per second) we would still need 100 seconds to complete one calculation of the self interaction of 10^5 particles.

By a clever combination of analysis and algorithm design it is possible to obtain accurate approximations of the mutual interactions of n particles in which only $O(n \log n)$ operations are required in the case of tree codes [1], [4] and $O(n)$ operations for the more sophisticated Fast Multipole Method (FMM) as originated by Roklin [14] and further developed by Greengard [9]. In both cases, a hierarchy of boxes is formed and the influence of the particles in each box is combined into a multipole expansion with a finite number of terms p [15, pp. 305ff]. Consider the influence of particles in a box A, at a point x outside the box. Normally we would use the direct calculation to obtain the field at the point x due to all the individual particles in the box, but we can also use the multipole expansion associated with the box to obtain an approximation of the field at x. It can be shown that the error in using the multipole expansion instead of the exact calculation is of $O(\delta^p)$ where

$$\delta = \frac{s}{d}$$

where s is the size of the box and d is the distance the point x is from the box A (see figure 1). The reduction in computing complexity is obtained by observing that particles further from the point of interest can be combined into larger boxes and still maintain the required accuracy.

For clarity of exposition we will describe the method in the two dimensional case where we can use the notation of complex analysis to simplify the discussion. So we

will consider sums of the form

$$\begin{aligned} \phi(z) &= \sum_{j=1}^{n} \log(z - z_j) q_j \\ E(z) &= \sum_{j=1}^{n} \frac{1}{(z - z_j)} q_j \end{aligned} \tag{1.2}$$

where $z,\ z_j \in \mathbb{C}$. We will first discuss a simple $O(n \log n)$ method and then in section 3 give a quick overview of the $O(n)$ method, the Fast Multipole method.

It should be noted that all the results that follow can be obtained via simple Taylor series arguments. In the three dimensional case the discussion can be simplified by using expansions in terms of spherical harmonics (see [8]).

2. An $O(n \log n)$ Hierarchical Method

First we will consider the situation depicted in figure 1 in which n particles situated at positions $z_j, j = 1, .., n$ are distributed in two boxes, with $n/2$ particles in box A and $n/2$ particles in box B. The boxes are separated by a distance d and the boxes have size s. Consider an arbitrary point $z \in B$. For each $z_j \in A$ the potential is given by

$$\log(z - z_j) q_j = \log(z - z_A) q_j - \sum_{k=1}^{p-1} \frac{1}{k} \left(\frac{z_j - z_A}{z - z_A} \right)^k q_j + O\left(\frac{z_j - z_A}{z - z_A} \right)^p q_j \tag{2.1}$$

and the force is given by

$$\frac{1}{z - z_j} q_j = \frac{1}{z - z_A} \sum_{k=0}^{p-1} \left(\frac{z_j - z_A}{z - z_A} \right)^k q_j + O\left(\frac{z_j - z_A}{z - z_A} \right)^p q_j \tag{2.2}$$

where q_j is the charge of the j-th particle. It is important to note that

$$\left| \frac{z_j - z_A}{z - z_A} \right| \leq \frac{s}{d}.$$

Hence accuracy of the approximation can be improved if s is made smaller or d is made larger.

Let

$$a_k = \sum_{z_j \in A} (z_j - z_A)^k q_j$$

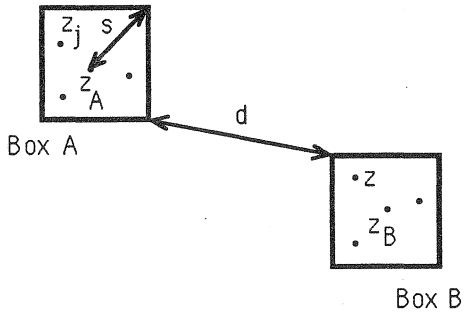

Figure 1
Two well separated boxes A and B.

denote the k-th moment of the particles in box A with respect to the centre of box A. In particular a_0 is the total charge (or mass) of the particles in A, a_1 is the corresponding dipole moment and a_2 is the quadrapole moment. Now the potential $\phi_A(z)$ and force $E_A(z)$ at the point $z \in B$ due to all the particles in box A is given by summing expressions (2.1) and (2.2) over all the particles in box A to obtain

$$\begin{aligned}\phi_A(z) &= \sum_{z_j \in A} \log(z - z_j) q_j \\ &= \log(z - z_A) a_0 - \sum_{k=1}^{p-1} \frac{a_k}{k} \frac{1}{(z - z_A)^k} + a_0 O\left(\frac{s}{d}\right)^p\end{aligned}$$

and

$$\begin{aligned}E_A(z) &= \sum_{z_j \in A} \frac{1}{z - z_j} q_j \\ &= \frac{1}{z - z_A} \sum_{k=0}^{p-1} \frac{a_k}{(z - z_A)^k} + \frac{a_0}{z - z_A} O\left(\frac{s}{d}\right)^p .\end{aligned}$$

These expansions are known as truncated multipole expansions.

The work required to calculate the coefficients for the multipole expansion in box A is $O(n)$. It takes only $O(1)$ work to calculate the influence of the $n/2$ particles in box A on any particle in box B if the multipole expansion is used. Hence the total work to calculate the sums due to the particles in box A at all the $n/2$ particle positions in box B is $O(n)$ (as opposed to $O(n^2)$ for direct calculation).

Of course we are still left with the problem of calculating the mutual interactions of the particles in each of the individual boxes in an efficient manner. The new idea is to

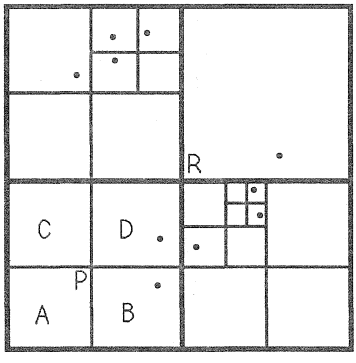

Figure 2
Barnes Hut Tree Structure
The Box R contains all the particles. Boxes are recursively divided until one or zero particles are left in a box. A generic parent box P contains more than one particle and so is divided into four child boxes A, B, C and D

use a divide and conquer algorithm. Suppose all the particles are contained in a box R (see figure 2). Here R denotes the root box. Using the Barnes-Hut [4] algorithm we recursively define a tree structure of boxes in which each parent box P containing more than 1 particle is divided into four sub-boxes A, B, C, D obtained by dividing the parent as in figure 2 where a parent box P with two particles has been divided into four boxes with boxes B and D containing particles and boxes A and C empty. To each box we associate a size corresponding to half the length of the diagonal of the box. The size of a child box is half that of its parent.

Alternative methods for dividing parent boxes exist, for instance a binary tree can be obtained by dividing each parent in a two stage sequence, first the parent is divided by a vertical line which leaves (as far as possible) the same number of particles in each child, and then each such child is divided by a horizontal line which produces two new children with approximately the same number of particles (see figure 3).

Now for a given accuracy parameter δ, the field at a point z due to a collection of particles contained in a box far from the point can be approximated to order δ^p via a multipole expansion of order p provided the box satisfies

$$\frac{s}{d} \leq \delta$$

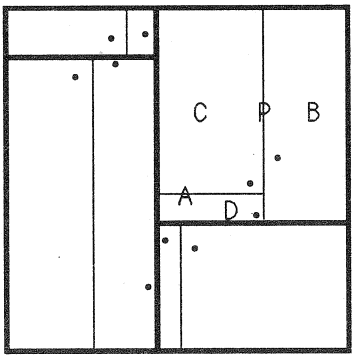

Figure 3
Alternative Tree structure
A generic parent box P is first divided by a vertical line to produce boxes A and B. Box B has only one particle and so is not divided. Box A is divided by an horizontal line to produce boxes C and D.

where s is the size of a box and d is the distance of the point z from the box. If this condition holds we say that the point z is well separated from the box.

The interactions of particles which are far from z can be combined into large well separated boxes and a multipole expansion used to provide an accurate approximation of the interactions. There will only be $O(\log n)$ well separated boxes that need to be used in this approximation of the field at a point z. This will leave a fixed small number of boxes neighbouring the point z in which direct calculation of the interaction will be necessary. So if we have a full set of p moments for each of the boxes making up the tree structure, then the computational complexity of calculating the interaction at any point z is $O(\log n)$. If we have to carry out this calculation at n points z_i then the complexity becomes $O(n \log n)$.

To complete the argument we must calculate the computational complexity of calculating a full set of p moments for the boxes in the tree structure. Suppose we have calculated the moments for the children of a parent box. Then the moments for the parent box can be calculated in $O(1)$ operations by translating the child moments to the centre of the parent box and then accumulating to obtain the moments for the parent. Since we can expect the total number of boxes to be proportional to the number of particles, we see that the calculation of the moments for all the boxes takes $O(n)$

operations. Hence the full calculation involves $O(n \log n)$ operations.

The accuracy of the method can be increased in essentially two ways: (1) Decrease the accuracy parameter δ; (2) Increase the number of terms in the truncated multipole expansion.

The first method is easy to implement but will significantly degrade performance if less than approximately 1 percent error is required. In the second case if we choose $\delta = 1/2$, then we will obtain a relative accuracy ϵ if we use an expansion of order $-\log_2 \epsilon$. Hence to obtain an accuracy comparable to single precision machine accuracy (10^{-7}) we will need approximately 20 terms in the two dimensional case and 200 terms in the three dimensional case.

It is interesting to note that this tree algorithm, in the two dimensional case, with 20 terms and $\delta = 1/2$ is actually slightly more accurate than the direct $O(n^2)$ calculation. This is due to the substantial reduction in the overall number of floating point operations that are needed for the calculation.

3. The Fast Multipole Method

Again consider the case depicted in figure 1. Using the method just considered we combine the influence of all the particles in box A into a multipole expansion about the centre of box A. For each of the particles in box B we then use this expansion to calculate the effect due to the particles in A. But if the size of box B is small relative to the distance between the boxes (that is if box B is well separated from box A), then the effect on a particle in B can be well approximated by a power series expansion centred about z_B. Let us derive the form of this power series. First note that

$$\log(z - z_A) = \log(z_B - z_A) - \sum_{n=1}^{p-1} \frac{1}{n} \left(\frac{z - z_B}{z_A - z_B} \right)^n + O\left(\frac{s}{d}\right)^p$$

and

$$\frac{1}{(z - z_A)^k} = \frac{1}{(z_B - z_A)^k} \sum_{n=0}^{p-1} \binom{k+n-1}{n} \left(\frac{z - z_B}{z_A - z_B} \right)^n + O\left(\frac{s}{d}\right)^p$$

Consequently the potential due to the particles in A can be approximated as follows:

$$\phi_A(z) = \log(z - z_A)a_0 - \sum_{k=1}^{p-1} \frac{a_k}{k} \frac{1}{(z - z_A)^k} + O\left(\frac{s}{d}\right)^p$$

$$= \log(z_B - z_A)a_0 - \sum_{n=1}^{p-1} \frac{a_0}{n} \left(\frac{z - z_B}{z_A - z_B}\right)^n$$
$$- \sum_{k=1}^{p-1} \frac{a_k}{k} \frac{1}{(z_B - z_A)^k} \left[\sum_{n=0}^{p-1} \binom{k+n-1}{n} \left(\frac{z - z_B}{z_A - z_B}\right)^n\right]$$
$$+ O\left(\frac{s}{d}\right)^p$$

$$= \log(z_B - z_A)a_0 - \sum_{k=1}^{p-1} \frac{a_k}{(z_B - z_A)^k}$$
$$- \sum_{n=1}^{p-1} \left[\frac{a_0}{n} + \sum_{k=1}^{p-1} \binom{k+n-1}{n} \frac{a_k}{k} \frac{1}{(z_B - z_A)^k}\right] \left(\frac{z - z_B}{z_A - z_B}\right)^n$$
$$+ O\left(\frac{s}{d}\right)^p$$

$$= \sum_{n=0}^{p-1} b_n (z - z_B)^n + O\left(\frac{s}{d}\right)^p$$

where b_n denotes the appropriate coefficient of the power series for $\phi_A(z)$ about z_B. Similarly, the force $E_A(z)$ satisfies

$$E_A(z) = \frac{1}{z - z_A} \sum_{k=0}^{p-1} \frac{a_k}{(z - z_A)^k} + O\left(\frac{s}{d}\right)^p$$

$$= \frac{1}{z_B - z_A} \sum_{n=0}^{p-1} \left[\sum_{k=1}^{p-1} \binom{k+n}{n} \frac{a_k}{(z_B - z_A)^k}\right] \left(\frac{z - z_B}{z_A - z_B}\right)^n$$
$$+ O\left(\frac{s}{d}\right)^p$$

$$= \sum_{n=0}^{p-1} c_n (z - z_B)^n + O\left(\frac{s}{d}\right)^p$$

where c_n denotes the appropriate coefficient of the power series for $E_A(z)$ about z_B.

So the functions ϕ_A and E_A can be well approximated by truncated power series

about the centre of box B provided the boxes A and B are mutually well separated. These truncated power series will be called dual expansions.

It is easy to see how to use the dual expansions in the situation depicted in figure 1. First the multipole expansion associated with box A is formed. This takes $O(n/2)$ operations. Then the dual expansions about the point z_B are calculated. This takes $O(1)$ operations. Then the field at points z_j in B are calculated from the dual expansions. This takes $O(n/2)$ operations. Hence the number of operations to calculate the interactions of particles in A with particles in B is $O(n)$.

Of course this is the same computational complexity as the method described in section 2 when applied to the situation in figure 1. The formation of the dual expansion has not helped us in this situation, but it will help us in the situation depicted in figure 4. Here we assume that $n/2$ particles are in box B, but the other $n/2$ particles are evenly distributed in m boxes A_i, $i = 1, ..., m$ which are all mutually well separated from B. Now let us compare the complexity of the method from section 2 with a method that uses dual expansions.

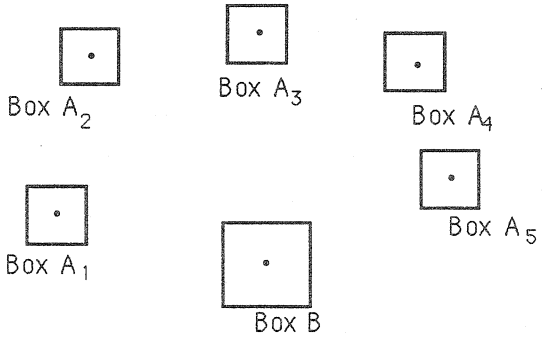

Figure 4
A boxes mutually well separated from box B

Using the method from section 2 we have that the complexity of calculating the influence of the particles in the A boxes on the particles in box B is the complexity of calculating the moments for each of the A boxes, which is $O(m \cdot (n/m))$, plus the complexity of calculating the mutual interaction of m multipoles (for each A box) with

$n/2$ particles in box B, which is $O(m \cdot n/2)$ operations. In total the complexity is $O(n+m \cdot n/2)$. Since we expect there to be $m = O(\log n)$ boxes mutually well separated from B, we see that the complexity is $O(n \log n)$ as expected.

Now suppose that we use dual expansions. Once again it takes $O(n)$ operations to calculate the moments for all the A boxes. It takes $O(m)$ operations to transfer the m multipole expansions to dual expansions about z_B, and then it takes $O(n/2)$ operations to calculate the influence at each of the $n/2$ particle positions in B, due to the accumulated dual expansions. In total the algorithm using dual expansions takes $O(n+m)$ operations. If $m = O(\log n)$ then we see that the overall complexity is $O(n)$. So in this case the use of dual expansions has helped us.

We are once again left with the problem of the mutual interactions within a box. As before, this is dealt with by using a divide and conquer strategy. The use of both multipole and dual expansions together with a divide and conquer strategy leads us to the so called Fast Multipole Method (FMM) originated by Roklin [14] and further refined by Greengard [9].

To see how the pieces fit for the FMM, we will first make some observations.

(1) Moments of parent can be computed via a merging process. The moments of the child boxes are translated to the centre of the parent box and these are then accumulated to produce the moments of the parent. For instance, the charge and dipole moment of a parent box P is obtained from the charge and dipole moments of its children A_i, $i = 1, 2, 3, 4$ as follows:

$$a_{0P} = \sum_{i=1}^{4} a_{0i}$$

$$a_{1P} = \sum_{i=1}^{4} a_{1i} + a_{0i}(z_{A_i} - z_P)$$

where a_{0P} and a_{1P} are the charge and dipole moment of the parent and a_{0i} and a_{1i} are the charge and dipole moment of the ith child box.

(2) The dual expansion of the child boxes can be obtained from their parents via a simple distribution process. If the dual expansion for the potential due to boxes well

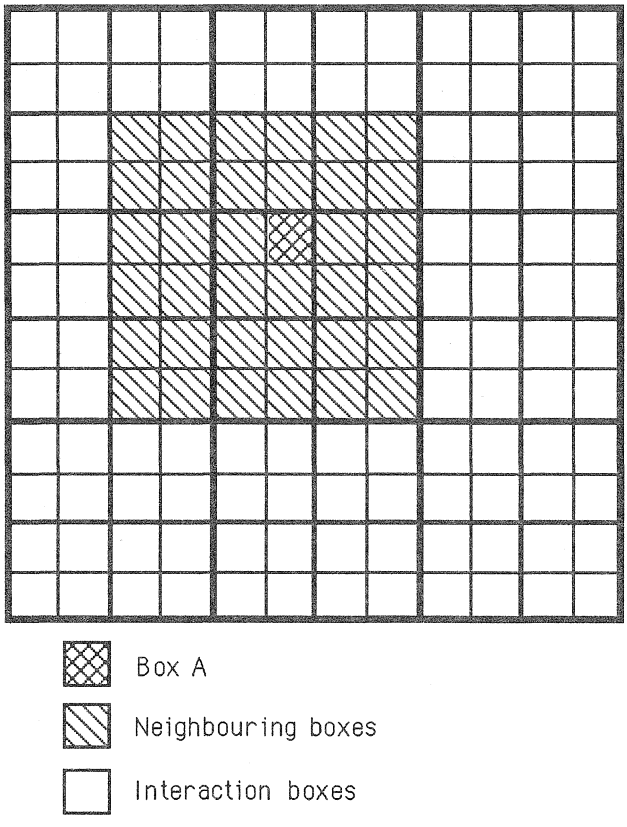

Figure 5
Neighbouring and Interacting boxes

separated from a parent box P about the centre of a parent box is given by

$$\phi_1(z) = \sum_n^{p-1} b_n (z - z_P)^n$$

then the dual expansions about the centre of any child B of P is given by

$$\begin{aligned}
\phi_1(z) &= \sum_n^{p-1} b_n (z - z_P)^n \\
&= \sum_n^{p-1} b_n (z - z_B + z_B - z_P)^n \\
&= \sum_{k=0}^{p-1} \left[\sum_{n=k}^{p-1} \binom{n}{k} (z_B - z_P)^{n-k} \right] (z - z_B)^k
\end{aligned}$$

(3) Let $\delta > 0$ be given. For each box A we define two collections of boxes which we call "neighbouring" boxes and "interacting" boxes. We define the neighbouring boxes as those boxes of the same size as A which have siblings which are less than $1/\delta$ distant from the centre of A. To simplify the structure of the neighbouring boxes we measure distance using the maximum norm, so that the neighbouring boxes will form $2n$ by $2n$ squares centred about the centre of the parent of box A. The interacting boxes are those children boxes of the neighbouring boxes of the parent of A which are not neighbouring boxes of A. The multipole expansion associated with box A will be transferred to dual expansions about the centres of the interacting boxes associated with box A.

The full multipole method can now be described as follows:

(a) Starting with those boxes containing only one particle we form the moments from the particle positions and charges. The merging process (point (1) above) is used to obtain moments for boxes at all higher levels of the tree structure. This takes $O(n)$ operations.

(b) The multipole expansions for any box is transferred to dual expansions about the centres of the corresponding interacting boxes (see point(3) above) and the dual expansions are accumulated. For each box this takes $O(1)$ operations, so overall the number of operations is $O(n)$.

(c) Beginning at the coarsest level the dual expansions of parent boxes are distributed (see point(2) above) to their children and the expansions accumulated with any expansion already associated with that child box. When complete, a dual expansion will have been created for each box in the tree, which describes the influence of all particles outside that box's neighbouring boxes. Once again this takes $O(n)$ operations.

(d) Finally the field at each particle position is calculated from the dual expansion of the smallest box which contains the particle plus a direct calculation using the particles in the neighbouring boxes of that smallest box. For each particle this takes $O(1)$ operations, so the number of operations is $O(n)$.

The total amount of work is $O(n)$ where we have assumed that the total number of boxes is $O(n)$.

4. Conclusion

So we have now seen that by a clever combination of analysis and algorithm design that it is possible to reduce the computational complexity of the calculation of the mutual interactions of n particles from $O(n^2)$ operations to the optimal rate of $O(n)$. There are still many questions pertaining to the efficient implementation of these algorithms on vector and parallel machines. There are now standard methods for vectorizing tree codes. The interested reader is referred to the recent articles by Hernquist [11], Barnes [3] and Makino [13] in which three methods for vectorizing tree codes are discussed. The question of efficient parallelization of tree codes is still a very active and fruitful research area.

References

[1] Appel, A. W., "An efficient Program for Many-Body Simulations", *Siam J. Sci. Comput.*, **6**, 85-103, 1985.

[2] Barnes, J., "Encounters of the Disk/Halo galaxies", *The Astrophysical Journal*, **331**, 669, 1988.

[3] Barnes, J., "A Modified Tree Code: Don't Laugh; It Runs", *J. Comput. Phys.*, **87**, 161-170, 1990.

[4] Barnes, J. & Hut, P., "A Heirarchical $O(N \log N)$ Force-Calculation Algorithm", *Nature*, **324**, 446-449, 1986.

[5] Chorin, A. J., "Numerical Study of Slightly Viscous Flows", *J. Fluid Mech.*, **57**, 785-796, 1973.

[6] Dawson, J. M., "Particle Simulations of Plasmas", *Rev. Mod. Phys.*, **55**, 403-447, 1983.

[7] Greengard, L. & Roklin, V., "A Fast Algorithm for Particles Simulations", *J. Comput. Phys.*, **73**, 325-348, 1987.

[8] Greengard, L. & Roklin, V., "The Rapid Evaluation of Potential Fields in Three Dimensions", *Lecture Notes in Mathematics*, **1360**, 121-141, 1987.

[9] Greengard, L., *The Rapid Evaluation of Potential Fields in Particle Systems*, MIT Press, Cambridge, 1988.

[10] Greengard, L., "The Numerical Solution of the N-Body Problem", *Computers in Physics*, **Mar/Apr**, 142-152, 1990.

[11] Hernquist, L., "Vectorization of Tree Traversals", *J. Comput. Phys.*, **87**, 137-147, 1990.

[12] Leonard, A., "Vortex Methods for Flow Simulation", *J. Comput. Phys.*, **37**, 289-

335, 1980.

[13] Makino, J., "Vectorization of a Treecode", *J. Comput. Phys.*, **87**, 148-160, 1990.

[14] Roklin, V., "Rapid Solution of Integral Equations of Classical Potential Theory", *J. Comput. Phys.*, **60**, 187-207, 1985.

[15] Wallace, P. R., *Mathematical Analysis of Physical Problems*, Holt Rinehart & Winston, 1972.