UvA-DARE (Digital Academic Repository)

New Computational Techniques to Simulate Light Scattering from Arbitrary Particles

Hoekstra, A.G.; Sloot, P.M.A.

DOI

[10.1002/ppsc.19940110304](https://doi.org/10.1002/ppsc.19940110304)

Publication date

1994

Published in

Particle & Particle Systems Characterization

[Link to publication](#)

Citation for published version (APA):

Hoekstra, A. G., & Sloot, P. M. A. (1994). New Computational Techniques to Simulate Light Scattering from Arbitrary Particles. *Particle & Particle Systems Characterization*, 11, 189-193. <https://doi.org/10.1002/ppsc.19940110304>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

New computational techniques to simulate Light Scattering from arbitrary particles

Alfons G. Hoekstra and Peter M.A. Sloot

* Faculty of Mathematics and Computer Science, University of Amsterdam,
Kruislaan 403, 1098 SJ Amsterdam, the Netherlands, tel. (+31)20-5257463

Published in

Part. Part. Syst. Charact. Vol 11 (1994), pp. 189-193

Abstract

The Coupled Dipole method, as originally formulated by Purcell and Pennypacker, is a very powerful method to simulate the Elastic Light Scattering from arbitrary particles. This method however has one major drawback: if the size of the particles grows, or if scattering from an ensemble of randomly oriented particles has to be simulated, the computational demands of the Coupled Dipole method soon become too high. In this paper we present two new computational techniques to resolve this problem. First we have implemented the Coupled Dipole method on a Massively Parallel Computer. The parallel efficiency can be very close to one, implying that attained computational speed scales perfectly with the number of processors. Secondly we propose to reduce the computational complexity of the Coupled Dipole method by including ideas from the so-called fast multipole methods (hierarchical algorithms) into the Coupled Dipole method. In this way the calculation time can be decreased with orders of magnitude.

1 Introduction

The possibility to simulate Elastic Light Scattering (ELS) from arbitrary particles is, in many cases, essential in optical particle characterization and sizing applications. This is true if one, for instance, has to verify particle models solely based on ELS information, as is the case for interstellar dust [1]. Another example is offered by our Flowcytometric experiments on human white bloodcells [2], where we have to define an

optimal scattering experiment to detect subtle changes in particle morphology.

The Coupled Dipole (CD) method [3] is a powerful method to calculate Elastic Light Scattering (ELS) from arbitrary particles. However, to be useful in practical applications, a huge computational challenge must be met. Especially if orientational averages have to be calculated, or if the size parameter of the particle grows, the CD method soon needs computing power far beyond the possibilities of desktop computers or even beyond the possibilities of super computers. In this report we present two computational techniques that can make the CD method suited for routine light scattering calculations of arbitrary particles.

The first technique is parallel computing, which is an adaptation at the implementation level. The merits of parallel computing are demonstrated on the basis of CD calculations of systems containing up to 33.000 dipoles. Such large systems require very powerful computers. In our case the CD method is implemented on a 512 node parallel transputer system (a Parsytec GCel-3/512).

Although parallelism is an inevitable concept in modern high performance computing, it will not be powerful enough to meet the computational demands of CD simulations of larger, and/or randomly oriented particles. Therefore we investigate if CD calculations are realistic for even larger systems by adapting the method at the algorithmic level. We demonstrate that ideas from the so-called fast multipole methods (hierarchical algorithms) can decrease the time complexity of the CD method with orders of

magnitude, thus making CD calculations with more than 100,000 dipoles feasible, or making CD calculations with $O(10^4)$ dipoles possible on desktop computers.

We use the CD method to simulate ELS from human white bloodcells. We expect significant biological impact if both computational methods are applied in the simulations. They will enable realistic simulations of ELS from white bloodcells, resulting in optimal definition of scattering experiments, designed to detect subsets of white bloodcells, or small morphological changes of cells, indicating possible pathologies. Furthermore, a rigorous interpretation of scattering experiments of more complex cell samples, such as bone marrow, becomes possible.

2 Computational structure

The CD method divides a particle into N small subvolumes, whose size must be small enough to ensure that it can be viewed as an ideal dipole. Usually one chooses $\lambda/20 < d < \lambda/10$, with d the size of a subvolume, and λ the wavelength of the incident light. From now on we refer to the subvolumes as dipoles. First the electric field on dipole i , $\mathbf{E}(\mathbf{r}_i)$ ($1 \leq i \leq N$), due to the external field $\mathbf{E}^0(\mathbf{r})$ and the field radiated by all other dipoles is calculated. This can be formulated as a matrix equation $\mathbf{A}\mathbf{x} = \mathbf{b}$, with

$$\mathbf{x} = \begin{pmatrix} \mathbf{E}(\mathbf{r}_1) \\ \vdots \\ \mathbf{E}(\mathbf{r}_N) \end{pmatrix}, \quad \mathbf{b} = \begin{pmatrix} \mathbf{E}^0(\mathbf{r}_1) \\ \vdots \\ \mathbf{E}^0(\mathbf{r}_N) \end{pmatrix} \text{ and}$$

$$\mathbf{A} = \frac{\gamma}{4\pi\epsilon_0} \begin{pmatrix} \mathbf{I} & -\mathbf{F}_{12} & \dots & & -\mathbf{F}_{1N} \\ -\mathbf{F}_{21} & \mathbf{I} & & & \\ \vdots & & \ddots & & \vdots \\ & & & \mathbf{I} & -\mathbf{F}_{N-1,N} \\ -\mathbf{F}_{N1} & & & -\mathbf{F}_{N,N-1} & \mathbf{I} \end{pmatrix}.$$

The matrix $\mathbf{A}$ is the $n \times n$ interaction matrix ($n = 3N$), $\mathbf{F}_{ij}$ is a functional describing the field, radiated by dipole j on dipole i , and γ is the polarizability of the dipoles. After solving the matrix equation, the scattered electric field $\mathbf{E}^s$ is calculated by summing the fields, radiated by the dipoles, at the observation point $\mathbf{r}_{obs}$. The dipoles are placed on a cubic grid with grid spacing d . The diameter of the spherical dipoles is equal to the grid spacing d .

Figure 1 gives an estimate of the number of dipoles needed to describe a compact particle, as a function of the size parameter α , with d equal to $\lambda/20$, $\lambda/10$, and $\lambda/5$. Even for modest size parameters the number of dipoles is $O(10^4)$ or larger.

Calculation of the electric field on the dipoles, that is, to solve the system of linear equations $\mathbf{A}\mathbf{x} = \mathbf{b}$ is the computationally most demanding part of the CD

method. Generally speaking linear systems are solved by means of direct or iterative methods [4]. In the past both approaches were applied to solve the coupled dipole equations. For instance, Singham et al. used a direct method (LU factorization) [5], Singham and Bohren described a reformulation of the CD method, which from a numerical point of view is a Jacobi iteration to solve the matrix equation [6], and Draine applied a Conjugate Gradient iteration [7].

Figure 1: Estimation of the number of dipoles needed to model a compact particle.

Direct methods require $O(n^3)$ floating-point operations to find a solution, whereas iterative method require $O(n^2)$ floating-point operations, provided that the number of iterations is much smaller than n . We want to simulate Elastic Light Scattering of particles with $\alpha > 20$, i.e. $N = O(10^5)$. Solely this vast number of dipoles forces us to use iterative methods. Suppose that the implementation can run at a sustained speed 1.0 Gflop/s, and $n = 3.0 \cdot 10^5$. In that case a direct method roughly needs $O(10)$ months to find a solution. An iterative method needs $O(100)$ seconds per iteration. If the number of iterations can be kept small enough, execution times can be acceptable.

The Jacobi iteration is not very well suited for a large number of dipoles; already for a relative small number of dipoles ($N = \sim 500$), the Jacobi iteration becomes non-convergent [8]. A very efficient iterative method is the Conjugate Gradient method [4]. Draine [7] showed that the Conjugate Gradient method is very well suited for solving the coupled dipole equations. The number of iterations needed to find the solution is much smaller than the dimension of the matrix. For instance, for a typical small particle with 2320 dipoles ($n = 6960$) the Conjugate Gradient method only needs 17 iterations to converge. We apply a Conjugate Gradient method, the so-called CGNR method [9], to find the electric field on the dipoles.

3 Parallel Computing

3.1 Parallel calculation of the dipole fields

The dipole fields are calculated in parallel by assigning N/p dipoles to each processor (with p the number of processors of the parallel computer), and each processor calculates the fields on these dipoles, using our parallel implementation of the CGNR method [10]. The CGNR method was implemented on a ring of transputers, with a rowblock decomposition of the matrix. Rowblock decomposition means dividing A in blocks of rows, with every block containing n/p consecutive rows. The CGNR method contains two matrix vector products, three vector updates and three inner products per iteration. Figure 2 schematically shows how these operations are performed in parallel.

$$\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} + [\text{factor}] \times \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \Rightarrow \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$$

2.a: parallel vector update

$$\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \times \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \Rightarrow [1] + [2] + [3] \rightarrow []$$

2.b: the parallel inner product.

$$\begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix} \rightarrow \begin{bmatrix} \\ \\ \end{bmatrix}; \begin{bmatrix} & & \\ & 1 & \\ & 2 & \\ & & 3 \end{bmatrix} \times \begin{bmatrix} \\ \\ \end{bmatrix} \Rightarrow \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}$$

2.c: the parallel matrix vector for a rowblock decomposed matrix.

Figure 2: A schematic drawing of the parallel implementation (here with 3 processors) of the numerical operations. The decomposition of the vector and matrix is symbolized by the dashed lines; a single arrow ($\rightarrow$) means a communication, and the implication mark ($\Rightarrow$) means a (parallel) calculation.

The total computation time is $O(n^2/p)\tau_{calc}$, the communication time is $O(n)\tau_{comm}$. [10]. The parameters τ_{calc} and τ_{comm} are the times to perform one floating-point operation on a processor and to send one byte from a processor to a neighbouring processor. Thus, the efficiency of the parallel CGNR is

$$\varepsilon = \frac{T_1}{pT_p} \approx \left(1 + O\left(\frac{p}{n}\right) \frac{\tau_{comm}}{\tau_{calc}} \right)^{-1}, \quad [1]$$

with T_1 the execution time on 1 processor and T_p the

execution time on p processors. Here we neglected many details of the communications and computations, for this see reference [10]. Still, equation 1 contains the most important conclusion. If n/p is large, the efficiency of the parallel CGNR method can be very close to one. This means that the execution time of the parallel CGNR is almost inversely proportional to the number of processors available in the parallel computer. Performance measurements of the actual implementation support this conclusion [10, 11].

3.2 Parallel calculation of the scattered fields

The scattered electric field is calculated in parallel by calculating the radiated electric fields from the dipoles in parallel, and summing them afterwards. This strategy matches the data decomposition used in the parallel CGNR implementation. After convergence of the CGNR every processor has the electric field on its local dipoles in memory. All processors calculate the scattered fields due to their local dipoles in all observation points (e.g. the scattered field as a function of the scattering angle θ). Finally the results of all processors are accumulated and summed in the root processor, which writes the results to disk for further analysis.

Both the calculation time and the communication time of the parallel calculation of the scattered fields are negligible compared to the calculation - and communication time of the parallel CGNR. Therefore, the efficiency of the parallel CD method will be as good as the efficiency of the parallel CGNR. However, the parallel calculation of the scattered fields also has a very good parallel efficiency on its own right, as a straightforward analysis reveals (data not shown).

3.3 Results

The parallel CD method was implemented on a Parsytec GCel-3/512, a 512 node distributed memory computer, which was recently installed in Amsterdam. The nodes are Inmos T805 transputers. The implementation was carried out in the language C, under Parsytec's parallel programming environment Parix.

Figure 3 show the measured parallel efficiency of the CGNR for some small systems. As the number of dipoles increases the efficiency stays very close to 1 for a larger number of processors. It can be shown that for $N = O(10^4)$ or more, the efficiency is almost 1 for the maximum number of processors ($p = 512$). Most computing time is spent in the CGNR method, therefore these conclusions also hold for the complete parallel implementation of the Coupled Dipole method. This is also supported by actual measurements (data not shown).

Figure 3: The measured parallel efficiency of the CGNR method as a function of the number of processors; the straight line is for 8 dipoles, the short dashed line for 136 dipoles, the long dashed line for 552 dipoles and the point dashed line for 2176 dipoles

Figure 4 shows the result of a CD simulation of scattering by sphere, together with the analytical Mie result. The number of dipoles was 33552, the diameter of the dipoles was $\lambda/10$, resulting in a size parameter $\alpha = 12.6$; the refractive index was 1.05. This calculation required 11 hours on 512 processors. The parallel speedup was estimated to be larger than 500, the computational speed was 250 Mflop/s. As soon as new updates of the C compilers are available, this number is expected to be increased with a factor 2 (due to much better code optimizers).

Figure 4: Coupled Dipole simulation (dots) of scattering by a sphere, and Mie calculations (line). The S_{11} element of the scattering matrix S , as a function of the scattering angle is shown. The size parameter of the sphere is $\alpha = 12.6$, the refractive index is $m = 1.05$. The number of dipoles in the CD simulation was 33552, the size of the dipoles was $\lambda/10$.

The results of the CD simulations are in good agreement with the exact Mie results, except in the backscattering. This is probably due to the relative large

size of the dipoles. The execution times are long, even on massively parallel computers. Note that we simulated a particle in only one orientation. If randomization is required, the execution times are no longer realistic. Furthermore, if the number of dipoles gets even larger, to $O(10^5)$ or $O(10^6)$, the execution time of the CD method for a particle in just one orientation already becomes too high.

4 Hierarchical Many-Body Methods

In its present form the CD method is very promising, but computationally too demanding to calculate ELS from particles with $\alpha > 10$, especially if orientational averages have to be calculated. The execution time of CD simulations, using iterative solvers, scales as N^2 . This is due to the matrix vector products in the CGNR method. From a physical point of view this matrix vector product is a calculation of the electric field on the dipoles, due to radiation from all other dipoles. In this sense the CD method can be viewed as a many-body simulation, which requires to calculate all pairwise interactions between the interacting particles (the dipoles).

Many-body methods possess an algorithmic complexity of $O(N^2/2)$ if all pairwise interactions are calculated (the direct algorithm). For many realistic simulations the number of interacting particles has to be very large. The $O(N^2/2)$ complexity of the direct algorithm is a severe restriction for these large scale many-body simulations. Even on the most powerful (massively parallel) supercomputers the execution times of realistic many-body simulations will soon rise above acceptable (or affordable) values.

The conclusion is that the algorithmic complexity of the direct method must be reduced. Some interaction potentials (e.g. Lennard-Jones) allow the use of cut-off techniques, which can reduce the complexity to $O(N)$. However, for long range interaction potentials, such as the dipolar interaction potential, cut-off techniques cannot be applied. A very important class of "clever" many-body algorithms, which reduce the complexity to $O(N \log N)$ or even to $O(N)$, are the so-called hierarchical tree methods [12, 13]. In these methods the interaction is not calculated for each particle pair directly, but the particles are grouped together in a hierarchical way, and the interaction between single particles and this hierarchy of particle groups is calculated.

Appel [14] introduced the first hierarchical tree method, which relies on using a monopole (center-of-mass) approximation for computing forces over large distances, and on sophisticated data structures to keep track of which particles are sufficiently clustered to make the approximation valid. This method achieves dramatic speedups compared to the direct algorithm, but is less efficient when the distribution of particles is

relatively uniform and the required precision is high. Barnes and Hut applied this method in simulations of interacting galaxies [15]. The next step, which was set by Greengard [12], is the use of multipole expansions to compute interaction potentials or forces. This approach is known as the Fast Multipole Method (FMM), and requires an amount of work proportional to N to evaluate all pairwise interactions to any degree of accuracy. Up till now FMM algorithms are developed for scalar $1/r$ potentials in two and three dimensions [12, 16]. Salmon presents an overview of hierarchical tree methods [13].

We have developed a FMM algorithm for the vector potential of radiating dipoles (in three dimensions). This FMM algorithm replaces the matrix vector products in the iterative solver of the CD simulation. Now the interaction between the dipoles is not calculated for each dipole pair directly, but the dipoles are grouped together in a hierarchical way, and the interaction between single dipoles and this hierarchy of dipolar groups is calculated. In this way the complexity of the complete CD simulation is reduced to $O(N)$. It should be noted that N has to be large to reach a cross over in execution time between the direct algorithm and the FMM algorithm. The FMM algorithm is build along the same lines as Greengard's FMM algorithm for scalar $1/r$ potentials in three dimensions [12]. The algorithm consists of three steps (we ommit the mathematical and algorithmic details here, they will be published elsewhere):

- 1] form multipole expansions for the vector potentials of the hierarchy of dipolar groups (upward pass);
- 2] compute the interactions between all dipoles at the coarsest possible level in the hierarchy; for a given group of dipoles in the hierarchy this is accomplished by including interactions between groups which are well separated from each other, and whose interactions are not accounted for at a higher stage in the hierarchy (downward pass);
- 3] using the resulting vector potential on each dipole, calculate the wanted electric field.

Hierarchical tree methods have proven to be very efficient and accurate, and well suited to be used in realistic many-body simulations. However, efficient implementation on High Performance Computing platforms, specifically massively parallel distributed memory computing systems, is far from obvious. Salmon has successfully implemented the Barnes-Hut method on the Caltech hypercubes [13]. The FMM is implemented on shared memory multicomputers [17], and on the connection machine CM-2 [18]. Furthermore, Leathrum and Board report on FMM implementations on a number of platforms, such as the

Intel Touchstone, transputers, Encore Multimax, and distributed workstations running PVM and Linda [19]. We will implement the CD method, using the FMM algorithm to calculate the dipolar interactions, on the parallel GC-el distributed memory computer.

5 Conclusions

The CD method allows, in principle, simulation of ELS of arbitrary particles. In practice however the calculation times to solve the CD equations soon become unrealistic. We have introduced two computational techniques, parallel computing and hierarchical methods, to meet the computational challenge imposed by the CD method.

The first technique opens the way to perform CD calculations on modern massively parallel computing platforms, but also on networks of workstations. The second technique reduces the complexity of the CD method with orders of magnitude. Now very large CD simulations can be performed on high end systems, and light scattering from smaller arbitrary particles can be simulated on workstations or personal computers.

In the future these advanced innovations will allow routine ELS calculations from arbitrary shaped particles, and serve more detailed optical particle characterizations.

6 References

- [1] J.I. Hage, M. Greenberg, and R.T. Wang, Scattering from arbitrary shaped particles: theory and experiment, *Applied Optics* **30** (1991) 1141-1152.
- [2] P.M.A. Sloot, A.G. Hoekstra, H. van der Liet, and C.G. Figdor, Scattering matrix elements of biological particles measured in a flow through system: theory and practice, *Applied Optics* **28** (1989) 1752-1762.
- [3] E.M. Purcell and C.R. Pennypacker, Scattering and absorption of light by nonspherical dielectric grains, *The Astrophysical Journal* **186** (1973) 705-714.
- [4] G.H. Golub and C.F. van Loan, *Matrix Computations*, (The John Hopkins University Press, Baltimore and London, 1989).
- [5] S.B. Singham and G.C. Salzman, Evaluation of the scattering matrix of an arbitrary particle using the coupled dipole approximation, *J. Chem. Phys.* **84** (1986) 2658-2667.
- [6] S.B. Singham and C.F. Bohren, Light scattering by an arbitrary particle: a physical reformulation of the coupled dipole method, *Optics Letters* **12** (1987) 10-12.
- [7] B.T. Draine, The discrete dipole approximation and its application to interstellar graphite grains, *Astrophys. J.* **333** (1988) 848-872.
- [8] S.B. Singham and C.F. Bohren, Light scattering by an arbitrary particle: the scattering-order formulation of the coupled-dipole method, *J. Opt. Soc. Am. A* **5** (1988) 1867-1872.
- [9] S.F. Ashby, T.A. Manteuffel, and P.E. Saylor, A

taxonomy for conjugate gradient methods, Siam J. Numer. Anal. **27** (1990) 1542-1568.

- [10] A.G. Hoekstra, P.M.A. Sloot, W. Hoffmann, and L.O. Hertzberger Time complexity of a parallel Conjugate Gradient solver for light scattering simulations: theory and SPMD implementation Tech. Rept. CS-92-06, Faculty of Mathematics and Computer Science, University of Amsterdam, 1992.
- [11] A.G. Hoekstra and P.M.A. Sloot, Implementation of a parallel conjugate gradient method for simulation of elastic light scattering, Computers in Physics **6** (1992) 323.
- [12] L.L. Greengard, The Rapid Evaluation of Potential Fields in Particle Systems, (The MIT Press , 1988).
- [13] J.K. Salmon *Parallel Hierarchical N-Body Methods*, Ph.D. dissertation, California Institute of Technology, Pasadena, USA, 1991.
- [14] A.W. Appel, An Efficient Program for many Body Simulations, SIAM J. Sci. Stat. Comput. **6** (1985) 85.
- [15] J. Barnes and P. Hut, A hierarchical $O(N \log N)$ force-calculation Algorithm, Nature **324** (1986) 446.
- [16] K.E. Schmidt and M.A. Lee, Implementing the Fast Multipole Method in Three Dimensions, J. Stat. Phys. **63** (1991) 1223.
- [17] L. Greengard and W.D. Gropp, A Parallel version of the Fast Multipole Method, Computers Math. Applic. **20** (1990) 63.
- [18] F. Zhao and S. Lennart Johnsson, The Parallel Multipole Method on the Connection Machine, SIAM J. Sci. Stat. Comput. **6** (1991) 1420.
- [19] J. Leathrum, James F. and J.A. Board The Parallel Fast Multipole Algorithm in Three Dimensions Tech. Rept. 27706, Duke University, Department of Electrical Engineering, Durham, NC, USA, 1992.