

PROCEEDINGS OF THE 4th INTERNATIONAL CONFERENCE

PHYSICS COMPUTING '92

Editors

Robert A. de Groot and Jaroslav Nadrchal

World Scientific

IMPLEMENTATION OF A PARALLEL CONJUGATE GRADIENT METHOD FOR SIMULATION OF ELASTIC LIGHT SCATTERING

A.G. HOEKSTRA and P.M.A. SLOOT

Department of Computer Systems, Faculty of Mathematics and Computer Science, University of Amsterdam, Kruislaan 409, 1098 SJ Amsterdam, the Netherlands, email alfons@fwi.uva.nl, phone (+31)20-5257543 or (+31)20-5257463

We simulate elastic light scattering with the coupled dipole method. The kernel of this method is a large set of linear equations. The $n \times n$ system matrix is complex, symmetric, full, and diagonally dominant. This application is a typical example of problems arising in computational electromagnetics. The matrix equations are usually solved with (preconditioned) conjugate gradient methods. For realistic problems the size of the matrix is very large ($n \sim 10^4$ to 10^6). In that case sustained calculation speeds in the Gflop/s range are required to keep execution times acceptable. We introduce a methodology to parallelize the conjugate gradient method for this type of problems, with emphasis on coarse grain distributed memory implementations. We present results for an implementation on a transputer network.

1. Introduction

Elastic Light Scattering (ELS) is a powerful particle detection and recognition technique, with important applications e.g. in astrophysics or biophysics. Many exact and approximate theories to calculate ELS from particles are known.¹ Nevertheless, important classes of particles fall outside the range of these theories. The coupled dipole (CD) method, due to Purcell and Pennypacker,² is a method that in principle allows calculation of ELS from any particle.

In the CD method of ELS a particle is divided into N small (\ll the wavelength of the incident light) subvolumes called dipoles. First the electric field on the dipoles, due to the incident field and the fields radiated by the other dipoles are calculated. Subsequently the scattered fields are calculated.²

Calculation of the electric field on the dipoles, the computational most demanding task in the CD method, boils down to solving a system of coupled linear equations

$$Ax = b, \quad (1)$$

with A a $3N \times 3N$ matrix, x a vector containing the unknown dipole fields and b a vector containing the known incident fields.² The matrix A is a large ($N \sim 10^4$ to 10^6), complex, symmetric, full and diagonally dominant matrix. We solve Eq. (1) with a Conjugate Gradient (CG) method, suited for our system matrix. This paper shortly summarizes our implementation of this so-called PCGNR method³ on a transputer network (more details can be found in Hoekstra et al.⁴).

2. Implementation of a Parallel CG Method on a Transputer Network

We parallelize the CG method by domain decomposition of matrix A and mapping this decomposition on an appropriate processor network. A detailed theoretical time complexity analysis revealed that the CG method with a grid decomposition of the matrix, implemented on a cylinder network, or with a row wise decomposition of the matrix, implemented on a ring network, will have comparable execution times for the number of processors and number of dipoles of interest.⁴ Because of the more simple communication structure, we implemented the parallel CG method with a row wise decomposition of the system matrix, on a ring of transputers. Furthermore, the analysis showed that for large values of N the efficiency of this implementation is close to one.

We implemented the parallel CG method on a Meiko Computing Surface with 64 T800 transputers, programmed in Occam2. We measured the execution time of the parallel CG method as a function of the number of processors p ($1 \leq p \leq 64$) and the dimension of the matrix n ($n = 3N$; $60 \leq n \leq 495^\dagger$). The agreement between the theoretical time complexity and the measurements was better than 5%.⁴ For $p = 64$, the efficiency of a relatively small problem ($n=495$) already was 0.85. For a larger problem ($n=3885$) the theoretical efficiency is 0.98. This implies that the implementation exploits the parallelism almost perfectly.

A more important measure is the total execution time of the implementation. One iteration of the CGNR method contains two complex matrix vector products. Assuming a perfect efficiency of one, the execution time t per iteration roughly is

$$t = \frac{16n^2}{p} t_{calc} \quad (2)$$

with t_{calc} the time for one floating point operation. For double precision arithmetic t_{calc} equals $1.57 \mu\text{s}$, as was measured on the T800 transputer. For large matrices t is approximately one hour or more. The number of iterations will be about 100 for particles with very small refractive indices, and (much) larger for more opaque particles. This shows that realistic calculation times for realistic problems can be achieved if more processors are added to the network and if more powerful processors are used. Furthermore, the total execution time can be decreased by suitable preconditioning. We have implemented a parallel polynomial preconditioner.⁴ These adaptations will reduce the calculation time with a certain factor, but the real solution will be a more efficient algorithm. We plan to reduce the $O(n^2)$ matrix vector products in the CG algorithm to $O(n \log(n))$, by using a parallel version of the hierarchical fast multipole algorithm.

References

1. C.F. Bohren and D.R. Huffman, *Absorption and Scattering of Light by Small Particles* (John Wiley and Sons, New York, 1983).
2. E.M. Purcell and C.R. Pennypacker, *The Astrophysical Journal* 186, 705 (1973).
3. S.F. Ashby et al., *Siam J. Numer. Anal.* 27, 1542 (1990).
4. A.G. Hoekstra et al., "Time complexity of a Parallel Conjugate Gradient Solver for Light Scattering Simulations", Technical Report CS-92-06, Faculty of Mathematics and Computer Science, University of Amsterdam, June 1992.

[†] Maximum matrix size to fit in memory of one processor.