PARALLELIZATION OF THE MULTILEVEL FAST MULTIPOLE ALGORITHM (MLFMA) ON HETEROGENEOUS CPU-GPU ARCHITECTURES

Allocation: Illinois/50.0 Knh
PI: Levent Gürel[1]
Co-PI: Prof. Wen-mei Hwu[1]
Collaborators: Weng Cho Chew[1] and Narayan Aluru[1]

[1]University of Illinois at Urbana-Champaign

EXECUTIVE SUMMARY

Our main research goal is to solve extremely large scientific problems involving big-data issues. We have been working on scalable parallel implementations of the multilevel fast multipole algorithm (MLFMA) on massively-parallel supercomputers. We derive our immediate applications from computational electromagnetics, acoustics, optics, elastics, but the methods we develop have applications in a plethora of disciplines, such as quantum mechanics, fluid dynamics, astrophysics, molecular dynamics, structural mechanics, heat equation, potential theory, to name a few. Thanks to the massive parallelization offered by Blue Waters, we have been able to solve larger problems in shorter times. Most recently, we have achieved solutions of electromagnetics scattering problems involving 1.5 billion unknowns. Every phase of such a solution requires the handling of 1,500,000,000 x 1,500,000,000 dense matrices and hence incorporates many aspects of big-data computing. These cutting-edge solutions provide us with new insights into how to improve our solver algorithms and parallelization strategies. We plan to tackle increasingly larger problems and also to benefit from the graphics processing unit resources available on Blue Waters by honing our heterogeneous computing skills.

INTRODUCTION

We have been working on solutions to extremely large problems that are derived mainly from electromagnetics, acoustics, and optics, to name a few. An electromagnetic scattering problem may become very large as either the frequency increases or the target size gets larger. Simulating the scattering of high-frequency radar waves from a large aircraft, for instance, may require the solution of millions or even billions of unknowns. Similarly, medical imaging with microwaves requires solutions of thousands of large problems. For such large-scale simulations, we use our in-house parallel implementation of the multilevel fast multipole algorithm (MLFMA) [1-4].

MLFMA is a reduced-complexity solver, and as such, it enables the solution of unprecedentedly large problems not only in computational electromagnetics but in many other disciplines as well. Parallelization of MLFMA on homogeneous central processing unit (CPU) architectures is not trivial and even more difficult on heterogeneous CPU-graphics processing unit (GPU) architectures [5]. We have been addressing several challenges to obtain a scalable parallel implementation of MLFMA on the CPU nodes of Blue Waters. We plan to overcome many more challenges to achieve a scalable parallel implementation of MLFMA on the heterogeneous CPU-GPU architecture of Blue Waters. This will allow us to address previously unsolvable scientific problems, as an outcome of this exploratory project.

METHODS & RESULTS

Our research revolves around our in-house parallel implementation of the MLFMA to compute solutions to various problems, mostly derived from computational electromagnetics (CEM). Prior to Blue Waters, our computing facilities were limited to 16 computing nodes and up to 256 processes. Using those facilities, the largest problems we were able to solve involved 212 million unknowns, mainly limited by the available memory. We achieved out-of-core solutions by using solid-state drives (SSDs) local to each node in order to overcome the memory limitations. This way we were able to extend our solution capability up to 1.1 billion unknowns, at the expense of runtime.

On Blue Waters, we first had to adapt our code to use more computing nodes and more processes. MLFMA uses an "internode communication map" to determine which node will communicate with whom. These maps had to be updated and generated for the larger number of nodes available. As a result of all of our adaptations, we have so far been able to run simulations on up to 512 nodes and up to 1024 processes.

We have increased the number of unknowns in the problems we can solve, by achieving the solution of an electromagnetic scattering problem with 1.5 billion unknowns. Figure 1 demonstrates the accuracy of the solution, where the target object is a simple but extremely large conducting sphere of 1200 wavelengths in diameter. This canonical object is chosen because of the availability of the semi-analytical Mie-series solutions for spheres so that computational results can be compared and validated. A surface integral equation is used to formulate the physical scattering problem in mathematical terms, and a method-of-moments discretization is used to transform the continuum formulation to a matrix equation, whose solution is well suited to digital computers. The ensuing 1,500,000,000 x 1,500,000,000 dense matrix equation is solved with MLFMA and parallelization on Blue Waters.

We will continue to scale up the problem size that can be efficiently solved by our parallel MLFMA solvers. We are currently investigating the bottlenecks that degrade the efficiency of scaling beyond 512 nodes. Our goal is to scale to at least 4096 nodes and achieve the solution of an electromagnetic scattering problem with 5 billion unknowns.

WHY BLUE WATERS

As the problem size (i.e., the number of unknowns) grows, memory requirement increases to the point where both the total memory needed to solve the problem and also the memory available at a single computing node may become insufficient. Prior to Blue Waters, we had reached this limit and we had been experimenting with out-of-core methods to use the disk storage as additional memory, even though we had to endure a huge penalty in the time required to solve problems. The huge number of computing nodes makes available a much larger total memory, thus enabling the solution of such enormous problems that were **impossible for us to solve before** Blue Waters. Additionally, improving parallelization and using more nodes lead to much shorter solution times.

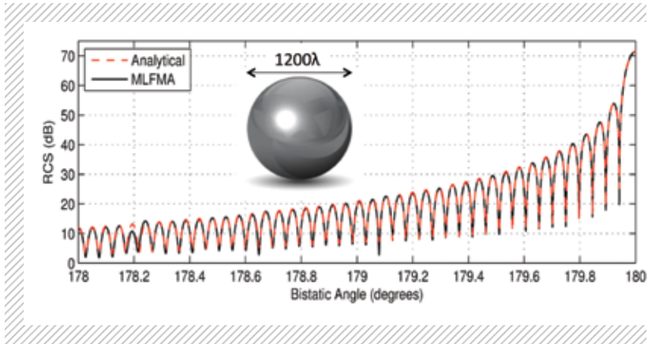

FIGURE 1: Solution of an extremely large electromagnetic scattering problem involving 1.5 billion unknowns. Scattering from a conducting sphere of 1200 wavelengths in diameter requires the solution of a 1,500,000,000 x 1,500,000,000 dense matrix equation. The computed bistatic scattering values are compared to semi-analytical Mie-series solution to demonstrate the high accuracy of the solution.

NEXT GENERATION WORK

We plan to develop scientific solvers for exascale computing. Specifically, we intend to develop parallel implementations of fast-multipole and Fourier-transform software. We consider the software aspect of parallel and heterogeneous computing an essential component for extracting exascale performance from immediately-next-generation supercomputers that will be available within the next 5-10 years. In other words, brute-force use of raw hardware power is neither sufficient nor rational; it will be the intelligent use of software that will unleash the exascale potential of the next-generation supercomputers.

PUBLICATIONS AND DATA SETS

Hidayetoğlu, M., et al., Parallel Solutions of Inverse Multiple Scattering Problems with Born-Type Fast Solvers, *Progress in Electromagnetics Research Symposium (PEIRS2016)*, Shanghai, China, Aug. 2016.

El Hajj, I., KLAP: Kernel Launch Aggregation and Promotion for Optimizing Dynamic Parallelism, *IEEE/ACM International Symposium on Microarchitecture (MICRO'16)*, Taipei, Taiwan, Oct. 2016.