Spatio-temporal Partition of the FMM Interaction Graph

Bo Zhang[a], Jingfang Huang[a], Nikos P. Pitsianis[bc], Xiaobai Sun[b]

[a]Department of Mathematics, University of North Carolina, Chapel Hill, NC 27599, USA
[b]Department of Computer Science, Duke University, Durham, NC 27708, USA
[c]Department of Electrical and Computer Engineering, Aristotle University, Thessaloniki, 54124, Greece

Extended Abstract

The fast multipole method (FMM) has enabled many kinds of important computational simulations to aid scientific discovery or engineering system design, such as for the evaluation of pairwise interactions of N particles in the following simple example,

$$\Phi(x_j) = \sum_{\substack{i=1 \\ i \neq j}}^{N} \frac{e^{-\lambda ||x_j - x_i||}}{||x_j - x_i||} \, q_i, \qquad \lambda > 0, \quad x_j \in R^3, \quad j = 1, \ldots, N, \tag{1}$$

per simulation time step, where $\| \cdot \|$ denotes the Euclidean distance and N is considerably large, from millions to trillions. The number of simulation steps is large also. With the naive and direct evaluation of (1), i.e., an explicit matrix-vector multiplication, the number of arithmetic operations increases quadratically with N. By introducing and utilizing a mathematically compressed approximation structure for the interaction matrix, the FMM calculates the summations in (1) indirectly and approximately in $O(N \log(N))$ arithmetic operations, with respect to any specified accuracy requirement.

A few significant developments took place in the first decade after the initial introduction of the FMM in 1987. First, there were numerous and elaborate attempts and efforts to implement the FMM on parallel computers, by the recognition that further reduction in computation time for such pair-wise interactions would rely heavily on effective utilization of modern computer architectures. Secondly, the FMM became adaptive efficiently and accurately to particle ensembles that are not necessarily uniformly distributed or in a simple cubicle domain. Thirdly, plane-wave expansions were used to make the constant factor associated with the term $N \log(N)$ substantially smaller. However, the results from the early parallelization efforts became very limited or obsolete by the advances in both architectural and algorithmic developments. Recent years see renewed interest and efforts in parallel FMM grow with new and emerging computer architectures and their great promise for faster simulation with larger data sets.

In this paper, we formulate the parallelization of the FMM as a combinatorial problem, namely, the spatio-temporal partition of the interaction graph associated with the FMM on a particular particle ensemble. The FMM with a particular particle ensemble had been associated with a tree structure since the very beginning. In the three dimensional case, the tree root corresponds to the smallest Cartesian box containing all the particles, the tree nodes at level 1 correspond to the sub-boxes partitioned equally along every dimension, the tree at level $k > 1$ correspond to the sub-boxes after the same partition scheme is applied recursively to the sub-boxes k times. The computation process is often overly simplified as consisting of two sweeps along the tree, an upward one for aggregation and a downward one for disaggregation. The parallelization of the FMM was consequently simplified as a partition in the data with certain heuristic criteria, such as the criterion of tree balancing in terms of the particle population. Strictly speaking, the evaluation of (1) by the FMM shall be associated with a graph, with cycles, that has both spatial and temporal attributes. The primary difference lies in the expression of the temporal and spatial interaction relationships in the so-called downward sweep. Specifically, at level k of the tree is a sub-graph $G_k = (V_k, E_k)$, a level-wise interaction sub-graph. The nodes in V_k are the tree nodes at level k. The edges in E_k describe the pairwise interactions among the sub-boxes within a specified neighborhood in a spatially geometric sense. Note also that there is no edge in the associated tree among the tree nodes at the same level. In the FMM, the redundancy in the level-wise interaction is reduced by data sharing and splitting via the sliding-window management, which is a particular form of spatio-temporal arrangement. In the non-adaptive FMM, the level-wise interaction graphs have a regular structure that is common across the levels. In the adaptive FMM, these level-wise sub-graphs do not necessarily have a common or regular structure. This poses a challenge to the partition in data at V_k and in the evaluation of the interaction governed by E_k. Clearly, the challenge can not be answered simply by an equal partition among the nodes in V_k at each level or across all the levels.

With the graph representation of the evaluation of pairwise interactions via the FMM on a specific sample ensemble as described above, we elaborate in the paper on the spatio-temporal partition rules, adaptation rules and metrics with regard to arithmetic complexity, numerical error propagation, computational concurrency degree and memory space limitation. Among other rules, for example, we extract a non-overlapping sub-cliques rule for the level-wise interactions. By this rule, the sub-graph G_k is further partitioned into smaller graphs. The spatially non-overlapping cliques are for concurrent local collections of pairwise interactions, for data sharing, also for reducing accumulated numerical errors below a preset upper bound that is independent of N. In other words, this new rule has a few advantages over the sliding window approach in broad use. We then introduce a spatio-temporal partition approach for parallel adaptive FMM based on the use of spectral graph theory. Experimental results on modern multiprocessor computers, based on the graph-theoretic partition scheme, will be presented.

For extended application of the spatio-temporal graph partition approach, we introduce formally the problem of collaborative carrier network. In a naive way, a carrier from each city travels over all the N cities and makes his all-to-one collection. The all-to-all collections by all carriers can be expressed by the complete graph. The total travel cost by the N carriers is $O(N^2)$. To reduce the cost, the carriers form a hierarchy of regional representatives, and follow certain collection and distribution rules that are similar to that in the FMM, but not necessarily as analytical. The total travel cost by all carriers can then be reduced to $O(N \log(N))$. Following the spatio-temporal partition of the collaboration network, the total travel time for all the all-to-all collections in parallel can be potentially minimized.