Adaptive FMM for fractal sets

key words: *adaptive fast multipole method, fractal set, fractal dimension, linear complexity*

Hadi Pouransari and Eric Darve

In the realm of scientific computing there are variety of situations where calculation of pairwise interaction among N points is of interest. Consider N particles (e.g., N masses) are located at positions $\{x_i\}$ in some metric space $\mathcal{M}$, and the net contribution of these particles at some observation point y is calculated by a sum of the form:

$$f(y) = \sum_{i=1}^{N} K(x_i, y)\sigma_i \tag{1}$$

where K is some $\mathcal{M} \times \mathcal{M} \to \mathbb{R}$ function called kernel, and σ_i is the intensity of the i'th particle.

This well-studied problem has broad application in various fields such as molecular dynamics, fluid dynamics, celestial mechanics, and plasma physics. More broadly, the problem of computing N^2 interactions among N points or N variables appears in the boundary element methods, problems involving radial basis functions or in probability theory to describe dense covariance matrices.

The fast multipole formulation introduced by Greengard and Rokhlin approximates a matrix-vector multiplication of the above form with desired accuracy in $\mathcal{O}(N)$ time. Several works have extended the algorithm by studying different kernels, analyzing the approximation error, introducing parallel implementation techniques, etc. [1]

The adaptive FMM refers to the case where the particle distribution, and the corresponding hierarchical tree, are not uniform. The adaptive FMM and various aspects of its parallel implementation on different machines are an ongoing topic of research. [2]

It is known that the adaptive FMM algorithm maintains the $\mathcal{O}(N)$ complexity irrespective of the point distribution [3]. This requires a modification to the original FMM. We will present a new proof for the linear complexity of the adaptive FMM for any distribution of the points. This also will make it apparent what modifications to the original FMM are required to ensure $\mathcal{O}(N)$ complexity for general particle distributions.

Previous works have limited their analysis to very specific point distributions. The key point essentially is the manner in which points are distributed, in a non-uniform adaptive setting, as N goes to infinity. In the uniform case, the issue of increasing N presents no particular difficulty. We can simply increase the density of points uniformly, and study how accuracy and parameters in the FMM are adjusted as a function of N. However, the non-uniform case is more difficult. One essential point is describing the process of adding points so that $N \to \infty$. The adaptive test cases considered by most previous works fall broadly into the following categories:

1. A small number of subregions are picked (e.g., n spheres) and points are progressively added to each subregion by distributing them with some smooth distribution (e.g., uniform, Gaussian, etc.) inside each region. Then the diameter and distance between regions are varied. [2]

2. Manifolds are considered, that is surfaces or lines. Then points are added on these manifolds again using a randomly uniform distribution.

3. Points are chosen such that they accumulate at some location (e.g., $x_i = 1/i^2$).

Complex non-adaptive cases have also been considered, but in those particular cases N was fixed.

All these cases represent only a small set of possible situations. There are many more ways to create non-uniform distribution of points. We focused on the third case, in which points accumulate.

However we extended this situation to points that essentially accumulate at an infinite number of locations. This naturally leads to fractal sets.

There are several practical problems involving fractal sets. Notoriously models of the universe, and antennas with fractal geometries that take advantage of the space-filling properties of fractal curves. In our numerical benchmarks we have considered the generalized Cantor sets that are constructed based on a recursive definition. The points x_i are generated by going through k iterations of this recursive process. As $k \to \infty$, N goes to infinity in a well-defined manner.

Fractal sets are often characterized in terms of their dimension, for example the fractal dimension, box-counting dimension, or Hausdorff dimension. We studied how parameters in the FMM such as the optimum total number of levels or the maximum number of points per leaf cells can be optimized as a function of the dimension of the set. We considered dimensions ranging continuously from 1 to 3, and exponential dependence of cost on the dimension is presented. Other details of the distribution appear to be less important. Our analysis is based both on mathematical bounds and estimates, as well as numerical benchmarks and investigations.

Theoretical estimates for optimal parameters can be found for uniform distributions, while, for a generic adaptive distribution, not much is known. Most implementations, if not all, manually or heuristically tune parameters to get the optimum values of parameters. In order to analyze arbitrary point distributions, we have characterized and categorized different distributions. We organized all possible fractal point distributions in terms of the fractal dimension of the set. We focused on sets for which the box-counting dimension is defined, which is the case for self-similar fractal sets for example. Note that the box-counting dimension cannot be defined for all sets. The Hausdorff dimension always exists but it cannot be directly related to the FMM (because of the oct-tree decomposition of the FMM) so that the Hausdorff dimension is in general not a good parameter to consider when optimizing FMM parameters. We will discuss these technical points in more details. Specifically in our numerical benchmarks, one of our main examples is a triple tensor product of generalized Cantor sets, which provide all range of box-counting dimensions varying continuously from 1 to 3 (in this case box counting is the same as Hausdorff).

We also present a new strategy to build the adaptive tree. We focused on the criterion used to determine whether a cell needs to be further subdivided or not. The original bisection algorithm uses one threshold value for subdivision, which is the maximum number of particles per leaf node. However, we used two threshold values simultaneously, namely, the maximum level of the tree, and the maximum number of particles per leaf nodes. Essentially, by tuning parameters in the dual threshold method we can transform some expensive operations such as M2P (multipole to particle) and P2L (particle to local) to a cheaper operation M2L (multipole to local). Better performance of the proposed scheme is demonstrated.

The aforementioned particle distributions were studied along with a detailed counting of the number of floating point operations. The calculation begins with some standard cases (e.g., uniform, spiral, etc.), and then extends to general fractal sets.

References

[1] William Fong and Eric Darve, The black-box fast multipole method, *Journal of Computational Physics* (2009) 8712–8725.

[2] Lexing Ying, George Biros, Denis Zorin, A kernel-independent adaptive fast multipole algorithm in two and three dimensions, *Journal of Computational Physics* (2004) 591-625.

[3] K. Nabors, F. T. Korsmeyer, F. T. Leighton, and J White Preconditioned, adaptive, multipole-accelerated iterative methods for three-dimensional first-kind integral equations for potential theory, *SIAM Journal on Scientific Computing*, 15(3), 713-735