

Use of fast multipoles for earthquake modeling

Terry E. Tullis⁽¹⁾, John Salmon⁽²⁾ and Naoyuki Kato^(1,3)

(1) Department of Geological Sciences, Brown University, Providence, RI 02912-1846, USA (e-mail: Terry_Tullis@brown.edu, Naoyuki_Kato@brown.edu, phone: +1-401-863-3829; +1-401-863-1103, fax: +1-401-863-2058). (2) Sandpiper Networks, Inc., 225 W. Hillcrest Dr., Thousand Oaks, CA 91360, USA (email: john@thesalmons.org, phone: +1-805-370-2196). (3) Geological Survey of Japan, 1-1-3 Higashi, Tskuba 305-8567, Japan (email: nkato@gsj.go.jp, phone: +81-298-54-3783, fax: +81-298-54-3697).

Abstract

Numerical earthquake models often use Green's functions for which the compute time increases with N^2 , where N is the number of elements defining the fault(s) used in the model. For realistic models of earthquakes it is necessary to use small elements and thus very large N . Especially for models that require considerable computation in each time step, such as the rate and state friction laws, the limitations of computer time severely restricts the size of the models that can be studied. The Fast Multipole Method involves grouping elements together that are removed from the location of interest and the compute time increases with $N \log N$. Testing of this method on an earthquake problem suggests that it may offer much promise for more realistic modeling of earthquakes.

Introduction

Numerical earthquake models that use the laboratory-based rate and state friction constitutive laws are the most realistic currently in use, based on the features that these models have in common with what is known about earthquake behavior. However, there are still many challenges that prevent these models from reaching their full potential for understanding earthquake mechanics. The difficulties include: 1) the non-linear equations require extensive computation time for each time step, 2) the number of time steps is large because the steps must be small during the dynamic-slip part of the earthquake cycle, but long time intervals must be modeled to simulate several earthquake cycles, 3) we have essentially no knowledge of the appropriate constitutive law for dynamic slip, 4) models that incorporate full dynamics including waves, involve additional complexity and compute-time, but are needed for many reasons, and 5) the number of elements N into which the fault is subdivided must be very large, both to properly represent a continuum and to model microseismicity together with major earthquakes; unfortunately in most computational procedures the compute-time increases as N^2 . Innovative numerical methods and large, fast computers are needed to create models that may overcome all of these difficulties. We are working on adapting Fast Multipole Methods used in other branches of physics to aid in overcoming the last of these difficulties. With such methods the compute time increases with $N \log N$ rather than N^2 , making the use of larger models more feasible.

Background

Previous work on a 3D model of earthquakes at Parkfield California (1,2) have shown behavior that is interesting and useful as we try to understand real Parkfield earthquakes and learn what might be done to predict them. The models use the boundary element method. It employs Green's

functions that give the stress on each element as a function of the slip on every other element. This results in N^2 relations that must be calculated many times in every time step. A fifth-order Runge-Kutta method with adaptive step size (3) is used to extrapolate forward in time. This has the advantage that time step sizes vary over about 10 orders of magnitude, as the model goes through an earthquake cycle. However, the fact that the compute time goes up as N^2 has meant that to date we have been forced to use elements that are much too small to represent a continuum (4), or to model microseismicity. Illustrations of these two points are given in Figures 1 and 2.

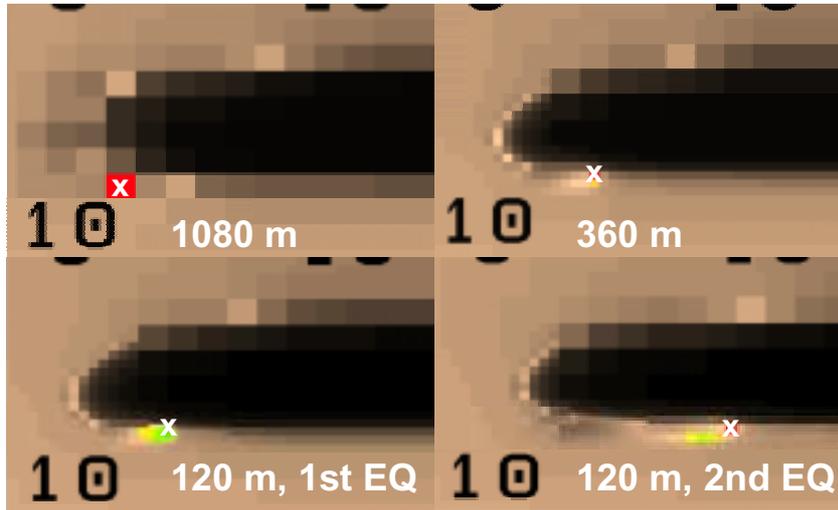


Figure 1. Portions of a Parkfield earthquake model with the smallest elements of different sizes as labeled. The shades of gray relate to the slip velocity in each element. The time steps shown are just prior to the earthquake (that is not modeled correctly and occurs instantaneously). X's show hypocenters. There is a large contrast in velocity between the hypocentral element and those surrounding it, showing that the model is not approximating a continuum. The behavior is modified by the variable element size, as is shown by the fact that for 120 m elements the first two earthquakes nucleate in different places.

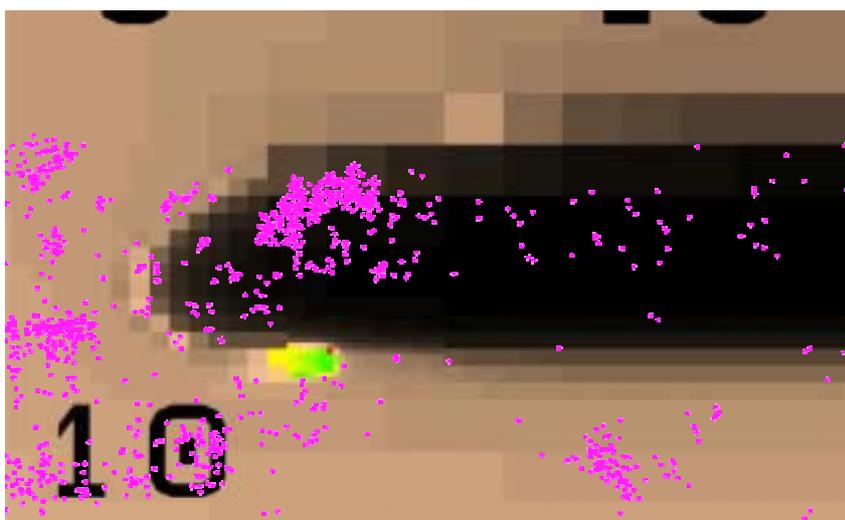


Figure 2. The small dots are the locations of micro-earthquakes at Parkfield, superimposed on the model with the 120 m size for the smallest elements. Even these elements are too large to model this distribution of microseismicity.

Multipole method

General idea

The idea of the multipole method is that elements that are far away from an element of interest can be grouped together and the influence of their slip on the stress at the element of interest can be combined. This reduces the number of Green's functions that need to be evaluated. The idea is a familiar one when we think of gravitational attraction, as is illustrated in Figure 3.

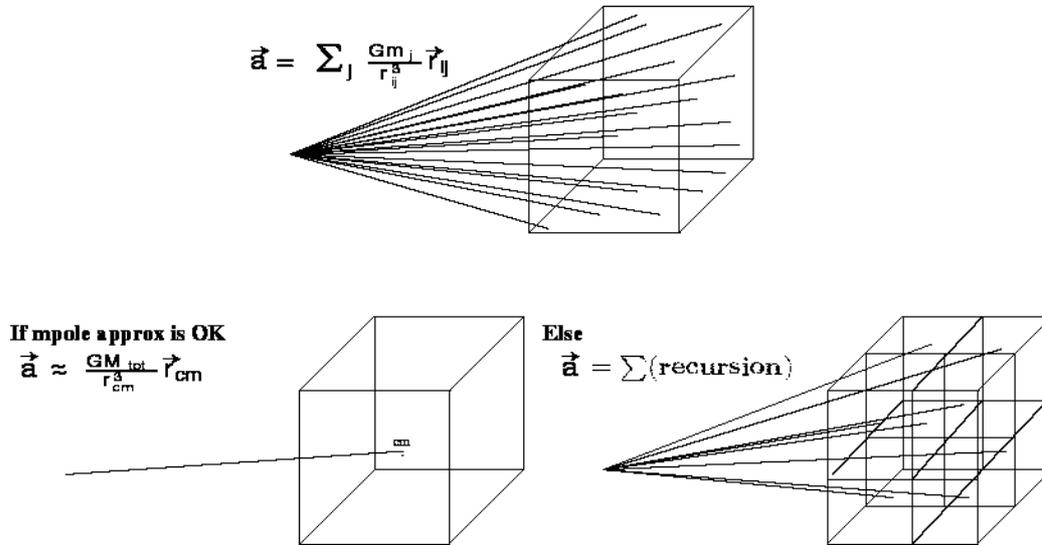


Figure 3. When determining the acceleration due to gravity at a point we may treat distant masses as if they are a single large mass at their center of gravity. If this approximation is inadequate for our uses, a larger number of masses may be used, but still smaller than the actual number.

As the elements get farther and farther away, a coarser grouping of elements may be used. This allows for an hierarchical structure of multipoles of higher order to be used as the distance increases. This is illustrated in figure 4. This can be quite efficient in our problem, since the Green's function we use falls off as $1/r^3$.

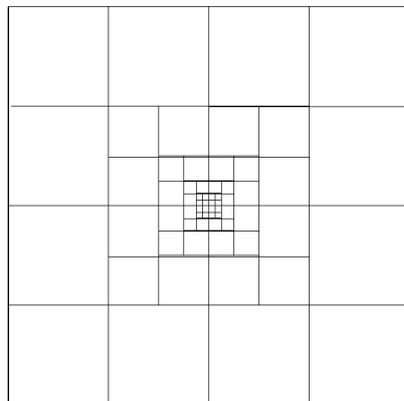


Figure 4. Hierarchical structure of multipoles, centered on the element of interest.

The multiple approach has the advantage over methods that require regular spacing of elements that the elements can be spatially adaptive. Thus, as illustrated in Figure 5, more elements can be assigned in regions where more detail is needed, and they can have appropriate local hierarchies. This can be very advantageous in our earthquake problems as consideration of the distribution of microseismicity illustrated in Figure 2 shows.

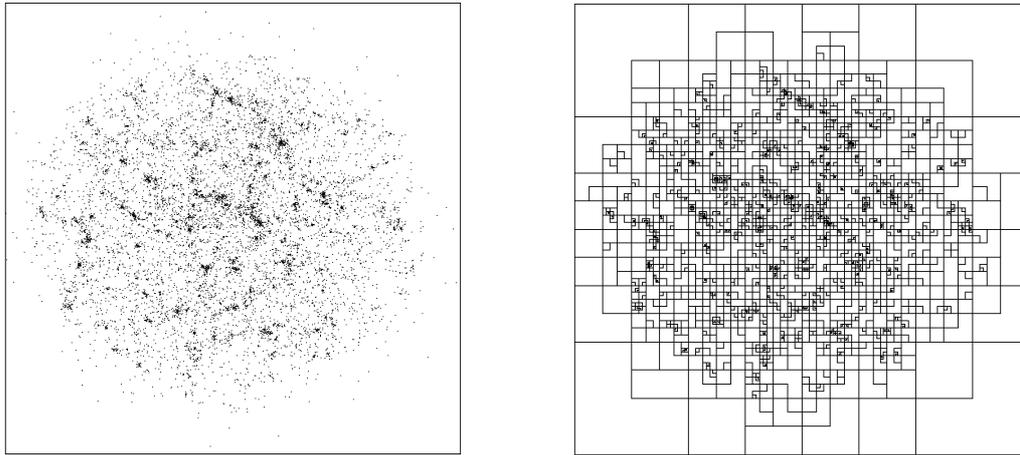


Figure 5. Example of inhomogeneous distribution of bodies (left) and spatially adaptive element distribution to match them (right) from a gravitational problems for which the multiple method has been used.

Ways to make the multipole method faster

There are a number of approaches that can be followed to make the multiple method into a Fast Multipole Method. One is to adjust both the geometry and the order of the multiples used so that they are optimized for the problem being studied. In the approach we use, such optimization is done with the aid of an analytic error bound to make decisions concerning the adequacy of the approximation.

Another approach is to use space filling curves of the type illustrated in Figure 6.

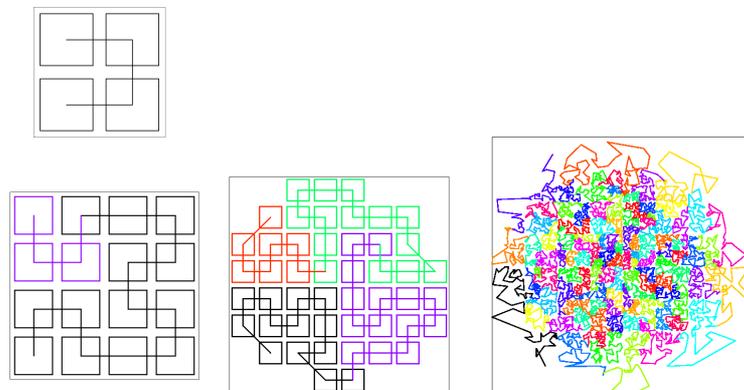


Figure 6. Space filling curves to map between memory and space.

With curves of the type shown in Figure 6 for increasing numbers of elements, both regularly and irregularly spaced, the numbering is done in such a way that nearness in memory implies nearness in space. Thus, an arbitrary region of space covers a small number of “segments” of memory. This allows the computations to be fast because it is possible to map “segments” to: 1) cache lines so there is good cache behavior, 2) to processors so there is a good load balance and minimal communication between processors, and 3) pages so there is a minimum of out-of-core computation. This last point can be important because memory is often more limiting than time.

Multipole Library

In this work we have used a previously created multipole library that can be applied to many types of problems for which the Green’s function approach is appropriate. This library is independent of the physics of the problem. The features and advantages of the library are that it is:

- Portable
- Parallel
- Out-of-core
- Cache friendly
- Space efficient
- Latency tolerant
- Bandwidth frugal
- Written in ANSI C
- Uses POSIX and other standards
- Written to a minimal “subset” of MPI
- Can use existing MPI libraries

Status of applying fast multipole method to our earthquake problem

We have incorporated the fast multipole method into our model for Parkfield earthquakes. We are in the process of testing it to determine how many elements need to be included in the model so that the efficiencies of the method overcome the extra time involved in the overhead of using the method. The result to date show that for our problem, after something in the vicinity of 10,000 elements the multipole method becomes more efficient than the usual N^2 approach. Since millions of elements are needed to satisfy the requirements of models that have sufficiently small elements and model significant sized earthquakes, this crossover number is encouragingly small. We are still examining how well the multipole approach approximates the N^2 results for N small enough that both methods can be used. Although the multipole library is written in conformity with MPI, we have not yet altered our earthquake model code to conform with the MPI standard. Thus, the advantages to be gained by the efficiency of the fast multipole method on parallel processors has not yet been realized, but we see no obstacles to doing this.

Conclusion

The Fast Multipole Method offers much promise for increasing the efficiency of earthquake models that use Green’s functions. We are still in the early stages of testing and using the multipole approach, but so far it appears that the method should allow earthquake models with much larger number of elements to be used than does the standard N^2 approach.

Acknowledgments

The general work on the Fast Multipole Method done by JS was done in collaboration with Michael Warren of Los Alamos National Laboratory.

References

- [1] Stuart, W. D. and Tullis, T. E., 1995, *Fault model for preseismic deformation at Parkfield, California*, J. Geophys. Res., **100**, 24079-24099.
- [2] Tullis, T. E., 1996, *Rock friction and its implications for earthquake prediction examined via models of Parkfield earthquakes*, in Earthquake Prediction: the Scientific Challenge, ed. by Leon Knopoff, Proc. Natl. Acad. Sci. USA, **93**, 3803-3810.
- [3] Press, W. H., Flannery, B. P., Teukolsky, S. A., and Vetterling, W. T., 1986, *Numerical Recipes*, Cambridge Univ. Press, New York.
- [4] Rice, J. R., 1993, *Spatio-temporal complexity of slip on a fault*, J. Geophys. Res., **98**, 9885-9907.