Journal of Computational Physics 328 (2017) 27–45

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

Multiresolution molecular mechanics: Implementation and efficiency

Emre Biyikli, Albert C. To *

Department of Mechanical Engineering and Materials Science, Center for Simulation and Modeling, University of Pittsburgh, Pittsburgh, PA, USA

ARTICLE INFO

Article history:
Received 4 November 2015
Received in revised form 28 September 2016
Accepted 4 October 2016
Available online 7 October 2016

Keywords:
Multiscale
Molecular dynamics
Atomistic/continuum
Implementation
Efficiency
Algorithm

ABSTRACT

Atomistic/continuum coupling methods combine accurate atomistic methods and efficient continuum methods to simulate the behavior of highly ordered crystalline systems. Coupled methods utilize the advantages of both approaches to simulate systems at a lower computational cost, while retaining the accuracy associated with atomistic methods. Many concurrent atomistic/continuum coupling methods have been proposed in the past; however, their true computational efficiency has not been demonstrated. The present work presents an efficient implementation of a concurrent coupling method called the Multiresolution Molecular Mechanics (MMM) for serial, parallel, and adaptive analysis. First, we present the features of the software implemented along with the associated technologies. The scalability of the software implementation is demonstrated, and the competing effects of multiscale modeling and parallelization are discussed. Then, the algorithms contributing to the efficiency of the software are presented. These include algorithms for eliminating latent ghost atoms from calculations and measurement-based dynamic balancing of parallel workload. The efficiency improvements made by these algorithms are demonstrated by benchmark tests. The efficiency of the software is found to be on par with LAMMPS, a state-of-the-art Molecular Dynamics (MD) simulation code, when performing full atomistic simulations. Speed-up of the MMM method is shown to be directly proportional to the reduction of the number of the atoms visited in force computation. Finally, an adaptive MMM analysis on a nanoindentation problem, containing over a million atoms, is performed, yielding an improvement of 6.3–8.5 times in efficiency, over the full atomistic MD method. For the first time, the efficiency of a concurrent atomistic/continuum coupling method is comprehensively investigated and demonstrated.

1. Introduction

A concurrently coupled atomistic/continuum method exploits both highly accurate atomistic method and highly efficient continuum method simultaneously, in order to lower the cost of simulations while retaining high accuracy. Atomistic simulations provide atomistic details of the deformation mechanisms, such as breaking and rearrangement of bonds [1]; however, this method has a high computational expense. On the other hand, a continuum computational framework (such as finite element method (FEM)) is well-established and very efficient, but fails at the atomistic scale, as it completely ignores the

* Corresponding author.
E-mail address: albertto@pitt.edu (A.C. To).

http://dx.doi.org/10.1016/j.jcp.2016.10.010
0021-9991/

inherent discrete structure of the material, leaving the physical underpinnings weak [2]. In addition to the advantages and disadvantages of atomistic and continuum approaches, the deformation and failure of many engineering materials are inherently multiscale, as the observed phenomena occur at many different length and time scales [3]. These macroscopic events are actually the results of microscopic events happening at much smaller length scales. The disparity between length scales is as large as 10^{10}, which is computationally intractable without a multiscale model [4].

In atomistic/continuum multiscale methods, the primary motivation is to systematically reduce the total number of degrees of freedom, while retaining sufficient accuracy. Reduction in the total number of degrees of freedom is best guided by the following insight: only a small portion of atoms actually undergo localized deformation, such as dislocation and stacking fault motion, while the remaining atoms undergo non-localized deformation, such as tension and compression [5]. The regions with localized deformation are physically more important, and hence these localized regions are represented by the more accurate atomistic description, while the much larger remaining portion of the domain with non-localized deformation is handled by an efficient continuum method. For this approach to be effective, the deformation needs to be free from abrupt changes in the continuum region.

To achieve this, the Multiresolution Molecular Mechanics (MMM) method [6] systematically eliminates the degrees of freedom from a full atomistic system by coarse-graining the non-localized regions using finite elements, while retaining an atomistic description for the localized regions. In the MMM method, the energy of a few, carefully selected atoms within the same mesh element are sampled and employed to compute the energy of all ghost atoms (i.e., eliminated degrees of freedom) in the element. The positions of the ghost atoms in each element are interpolated from the positions of the rep atoms (mesh nodes) via finite element shape functions. As a result of the reduction in degrees of freedom, and an efficient energy sampling scheme, calculation of the energy of and the force on the ghost atoms is no longer needed, resulting in higher computational efficiency. The rationale and formulation of the MMM method will be further discussed below. Thus far, the MMM method has been introduced for statics [6] and dynamics [7], followed by convergence and error structure analysis [8], as well as formulation of a unified and consistent framework for general finite element (FE) shape functions [9]. Recently, an adaptivity analysis has also been presented [10]. The accuracy of the MMM method has been well established by our earlier works, but its efficiency has not been investigated beyond the effect of the reduction in degrees of freedom.

Likewise, there have been very few works regarding the implementation and efficiency of other atomistic/continuum multiscale methods. Xiao et al. presented task and data decomposition in atomistic, continuum, and bridging domains of the Bridging Domain (BD) method; however, this demonstration was limited to a one-dimensional wave propagation example [11]. Anciaux et al. presented a parallel implementation of the BD method on a 2-D crack example, which showed load imbalance due to coupling overheads [12]. Fox et al. presented a parallelization of their multiscale method using MPI and demonstrated up to 3.12 times speed-up in simulation time [13]. In addition, there are initiatives to build computational libraries [14] and software infrastructure [15] for multiscale modeling and simulation.

Since the MMM method retains the fundamental nature and functions of the original molecular mechanics method, the details of its implementation will be based on the original method. Molecular mechanics prioritizes proper and efficient sampling of phase space rather than high accuracy, and this can be well achieved by symplectic integrators [16]. Phase space can be explored using minimization algorithms, such as, the conjugate gradient method [17], special relaxation algorithms [18], or using MD [19]. Our focus is on the latter approach, which solves Newton's equations of motion (2nd law) given by the following equation:

$$\begin{aligned} m_i \frac{d\mathbf{v}_i}{dt} &= \sum_j F_2(\mathbf{r}_i, \mathbf{r}_j) + \sum_j \sum_k F_3(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) + \cdots \\ \frac{d\mathbf{r}_i}{dt} &= \mathbf{v}_i \end{aligned} \tag{1}$$

where m is the mass, v is the velocity, r is the position, F_2 is the two-body force function, F_3 is the three-body force function, t is the time, and i, j, k are the atom indices. Usually, the terms beyond the first few are excluded and their effects, along with quantum effects, are represented by the remaining terms [20]. The force term in Equation (1) is equal to the derivative of the potential energy with respect to the position of the atom. The potential energy of an MD system is given by the following equation:

$$E = E_{bonded} + E_{electrostatic} + E_{van\ der\ Waals} \tag{2}$$

where E is potential energy with subscripts denoting the source. In the current work, the first two terms in Equation (2) are neglected, since the last term introduces adequate level of computational complexity to discuss the ideas presented. The last term represents non-bonded interactions, whose effects decay quickly with an increasing distance [21]. Therefore, the interatomic interactions beyond a certain distance (r_c) are truncated, or "cutoff". Due to this reduction, force calculations scale with $O(Nr_c^3\rho)$ instead of N^2, where N is the number of atoms and ρ is the number of atoms per unit volume. In some cases, a correction is used to compensate for the effects of the cutoff region.

An efficient way to numerically integrate the equations of motion is the Velocity-Verlet scheme, which is symplectic, time-reversible, conserves linear and angular momentum, and requires one force evaluation per time step; it is second order accurate (its error is proportional to Δt^2) [16]. Despite the cutoff and efficient integration, the MD method deals with very small length and time scales and hence is highly computationally demanding [20]. The force calculation in MD simulation