Computer Physics Communications 185 (2014) 1857–1864

Contents lists available at ScienceDirect

Computer Physics Communications

journal homepage: www.elsevier.com/locate/cpc

Spatial adaptive sampling in multiscale simulation

Bertrand Rouet-Leduc [a,c], Kipton Barros [a,*], Emmanuel Cieren [a,d], Venmugil Elango [b,h], Christoph Junghans [a], Turab Lookman [a], Jamaludin Mohd-Yusof [b], Robert S. Pavel [b,e], Axel Y. Rivera [b,f], Dominic Roehm [a,g], Allen L. McPherson [b], Timothy C. Germann [a]

[a] Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA
[b] Computer and Computational Sciences Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA
[c] Department of Materials Science and Metallurgy, University of Cambridge, Cambridge CB3 0FS, UK
[d] CEA, DAM, DIF, F-91297 Arpajon, France
[e] Department of Electrical and Computer Engineering, University of Delaware, Newark, DE 19716, USA
[f] School of Computing, University of Utah, Salt Lake City, UT 84112, USA
[g] Institute for Computational Physics, Universität Stuttgart, 70569 Stuttgart, Germany
[h] Department of Computer Science and Engineering, The Ohio State University, Columbus, OH 43210, USA

ARTICLE INFO

Article history:
Received 13 January 2014
Accepted 7 March 2014
Available online 19 March 2014

Keywords:
Multiscale
Adaptive sampling

ABSTRACT

In a common approach to multiscale simulation, an incomplete set of macroscale equations must be supplemented with constitutive data provided by fine-scale simulation. Collecting statistics from these fine-scale simulations is typically the overwhelming computational cost. We reduce this cost by *interpolating* the results of fine-scale simulation over the spatial domain of the macro-solver. Unlike previous adaptive sampling strategies, we do *not* interpolate on the potentially very high dimensional space of inputs to the fine-scale simulation. Our approach is local in space and time, avoids the need for a central database, and is designed to parallelize well on large computer clusters. To demonstrate our method, we simulate one-dimensional elastodynamic shock propagation using the Heterogeneous Multiscale Method (HMM); we find that spatial adaptive sampling requires only $\approx 50 \times N^{0.14}$ fine-scale simulations to reconstruct the stress field at all N grid points. Related multiscale approaches, such as Equation Free methods, may also benefit from spatial adaptive sampling.

1. Introduction

A key challenge in physics and engineering is the development of simplified, coarse-grained descriptions of complex, fine-scale systems. The development of new coarse-graining and multi-scale modeling techniques is a highly active field [1–6]. Coarse-graining approaches relevant to atomistic dynamics include the Mori–Zwanzig [7,8] formalism and its subsequent developments [9,10], "coarse-grained molecular dynamics" to seamlessly refine continuum models down to atomistic detail [11,12], techniques such as "dissipative particle dynamics" [13] and the "multi-scale coarse-grained" method [14,15], among others.

A particularly rich and important topic is the multiscale coupling between these coarse grained models and their fine-scale counterparts. For example, if atomistic detail is desired in only a small spatial region, then domain decomposition may be used to treat the remainder of the system at a coarse-grained, continuum level. Various matching strategies at the atomistic–continuum interface have been employed for dense fluids [16–18] and solids [19–22], with special attention paid to dynamical effects such as phonon propagation and heat transfer [23–28].

Often, however, multiscale coupling is required throughout the *entire* system. Coarse-grained models are necessarily incomplete due to missing information associated with neglected microscopic degrees of freedom. Several multiscale strategies to *augment* the coarse-grained model with data obtained from fine-scale simulations have been proposed [29–33]. These strategies are effective when the problem has a physical "separation of scales", and the statistical information collected at the fine-scale can be meaningfully extended to the coarse-grained level. When this approach is applied to physical simulation, the calculation of fine-scale

* Corresponding author.
E-mail addresses: kbarros@lanl.gov (K. Barros), tcg@lanl.gov (T.C. Germann).

http://dx.doi.org/10.1016/j.cpc.2014.03.011
0010-4655/

statistical quantities typically consumes the vast majority of computational work. By contrast, the integration of the coarse-scale equations themselves is practically free.

In this paper we introduce *spatial adaptive sampling* to reduce the amount of fine-scale simulation required in multiscale modeling. Previously proposed adaptive sampling involved storing the results of all fine-scale simulations in a database; interpolation from this database reduces the need for future fine-scale simulations [34,35]. However, database lookup and interpolation can be expensive, as it occurs within the potentially very high dimensional *input-space* (i.e., the space of inputs to the fine-scale simulations). In our spatial adaptive sampling approach, we instead interpolate in the $d \leq 3$ dimensional spatial domain (i.e., spatial coordinates x, y, and z). All necessary fine-scale simulations are spawned at the current time step, and no database is required. Another advantage of our approach is that we avoid the need to interpolate in the high dimensional input-space. Whereas traditional adaptive interpolation can be likened to building a surrogate model of the fine-scale response at all times [36], our real-space approach specializes to the current time step. The simplicity of our approach yields practical efficiency gains. In particular, interpolation in $d \leq 3$ dimensions is straightforward, and we avoid load balancing difficulties associated with synchronizing the fine-scale response database across large computer clusters [37].

To demonstrate spatial adaptive sampling, we consider a simple multiscale model of elastodynamics [38] based upon the Heterogeneous Multiscale Method (HMM) [32,33]. Here, the coarse-grained (*macroscale*) model is the set of continuum conservation laws for mass, momentum, and energy. These equations are closed by a fine (*micro*) scale, molecular dynamics (MD) model that provides constitutive data in the form of momentum and energy fluxes. We efficiently and accurately resolve shock propagation by devoting the great majority of our microscale computational effort to the immediate vicinity of the propagating shock. In this sense, our approach is analogous to adaptive mesh refinement (AMR, *e.g.*, Refs. [39,40] for conservation laws in two and three dimensions), in which there are more fine-scale model evaluations in regions where the response is more rapidly varying. Unlike AMR, however, our approach can be based upon any discretization scheme for the macro-solver, and directly generalizes to arbitrarily complicated macroscale equations. Our approach also generalizes to other multiscale frameworks, such as Equation Free methods [30].

In developing our spatial adaptive sampling method, we took special care to ensure efficient parallelization across large-scale computer clusters. In particular, the number and the location of fine-scale simulations are completely determined at the beginning of each macroscopic time step, enabling an efficient distribution of the workload. In our simulations of one-dimensional shock propagation with 0.1% target accuracy and $N \gg 1$ grid points, we observe that approximately $4 \times 50 \times N^{0.14}$ fine-scale MD simulations are required at each integration step—a potentially great improvement over naïve HMM, which requires $4 \times N$ fine-scale simulations per time step.

2. Review of the heterogeneous multiscale method

We introduce our spatial adaptive sampling approach in the context of HMM [33], which begins with a microscale dynamical model and then derives a "coarse grained" macroscale representation of this dynamics. We apply MD at the microscale and a system of conservation laws for mass, momentum, and energy at the macroscale [38]. These conservation laws, by themselves, are incomplete. To *close* the macroscale equations, one requires constitutive data from the microscale in the form of fluxes. It is through this closure that the micro and macroscales are coupled. The goal of HMM, then, is to construct a macroscale dynamics on physically relevant length and time scales, while retaining the accuracy of the microscale model.

Real materials are replete with heterogeneous defects including interstitials, vacancies, dislocations, grain boundaries, phase boundaries *etc.* Such defects pose an essential difficulty for HMM, and two approaches are standard [32,33]. In the "type B" approach, one assumes the coarse grained fields are a statistically complete description of the system. In particular, each defect must be faithfully represented macroscopically. This approach has been demonstrated, *e.g.*, for martensitic phase transformation dynamics [38] and combustion front tracking [41]. Alternatively, in the "type A" approach, one uses domain decomposition to simulate isolated defects with full microscopic accuracy. This approach has been applied, *e.g.*, to the propagation of a single crack [28]. In both type A and B approaches, HMM is practical for the study of a small number of isolated defects. A third approach, common in materials science, would posit a phenomenological description of the *coarse grained* defect distribution and dynamics. This third approach is less compatible with the philosophy of HMM because hysteresis and ergodicity breaking makes it difficult to control errors.

In this work, we avoid difficulties associated with material defects and consider a simple demonstration problem: elastodynamic shock propagation in a perfect crystal. The macroscopic conservation laws are closed with momentum and energy fluxes calculated microscopically through an averaged MD stress. HMM makes the key assumption of *separation of scales*. Here, we assume that the microscopic MD simulations are ergodic and that stress is a unique function of the macroscopic inputs, with no history dependence. Under these assumptions, it is valid to extend the microscopically averaged stress to macroscopic scales. The HMM ergodicity assumption would fail for more realistic materials with, *e.g.*, strain hardening due to defects. Such difficulties are absent in our defect-free crystal.

HMM begins with a microscale dynamics, in our case Newton's equations,

$$\dot{\mathbf{x}}_i = \mathbf{v}_i \tag{1}$$

$$m_i \mathbf{v}_i = -\nabla_{\mathbf{x}_i} V, \tag{2}$$

where m_i, $\mathbf{x}_i$, and $\mathbf{v}_i$ are the mass, position, and velocity of particle i. For simplicity we assume that the potential energy is a sum of pairwise interactions,

$$V = \frac{1}{2} \sum_{i \neq j} \phi(\mathbf{x}_{ij}(t)),$$

where $\mathbf{x}_{ij}(t) = \mathbf{x}_i(t) - \mathbf{x}_j(t)$. If we reexpress the mass, momentum, and energy densities as distributions,

$$\rho(\mathbf{x}, t) = \sum_i m_i \delta(\mathbf{x} - \mathbf{x}_i(t)) \tag{3}$$

$$\mathbf{q}(\mathbf{x}, t) = \sum_i m_i \mathbf{v}_i(t) \delta(\mathbf{x} - \mathbf{x}_i(t)) \tag{4}$$

$$e(\mathbf{x}, t) = \sum_i \left[\frac{1}{2} m_i \mathbf{v}_i^2 + \frac{1}{2} \sum_{j \neq i} \phi(\mathbf{x}_{ij}(t)) \right] \delta(\mathbf{x} - \mathbf{x}_i(t)), \tag{5}$$

then Newton's equations become a set of conservation laws [42,38],

$$\partial_t \rho + \nabla_{\mathbf{x}} \cdot \mathbf{q} = 0 \tag{6}$$

$$\partial_t \mathbf{q} + \nabla_{\mathbf{x}} \cdot \tau = 0 \tag{7}$$

$$\partial_t e + \nabla_{\mathbf{x}} \cdot \mathbf{j} = 0. \tag{8}$$

Fluxes are most simply expressed as spatial averages over the domain Ω. The average momentum flux is

$$\tau_\Omega = \frac{1}{|\Omega|} \int_\Omega \tau(\mathbf{x}) \, \mathrm{d}^d \mathbf{x} = -\frac{1}{|\Omega|} \sum_{i,j} c_{ij} \left[\delta_{ij} m_i \mathbf{v}_i \otimes \mathbf{v}_i + \mathbf{f}_{ij} \otimes \mathbf{x}_{ij} \right], \tag{9}$$

and the energy flux is

$$\mathbf{j}_{\Omega} = \frac{1}{2|\Omega|}\sum_{i,j} c_{ij}\left\{\mathbf{v}_i\left[\frac{1}{2}m_i\mathbf{v}_i^2\delta_{ij} + \phi(\mathbf{x}_{ij})\right] - \frac{1}{2}(\mathbf{v}_i+\mathbf{v}_j)\cdot\mathbf{f}_{ij}\mathbf{x}_{ij}\right\}. \tag{10}$$

In the above, $\mathbf{f}_{ij}$ is the force applied on atom i by atom j. The parameter $0 < c_{ij} < 1$ represents the fraction of the line connecting $\mathbf{x}_i$ and $\mathbf{x}_j$ that lies within Ω. For a computational domain with periodic boundary conditions, $c_{ij} = 1$.

We note that the center-of-mass velocity $\bar{\mathbf{v}}$ of the region Ω is not subtracted from the velocities $\mathbf{v}_i$ appearing in the Irving–Kirkwood stress τ. Consequently, τ differs from the virial stress σ (equivalently, the continuum Cauchy stress [43]) by an amount $\tau - \sigma = \rho\bar{\mathbf{v}}\otimes\bar{\mathbf{v}}$.

Eqs. (6)–(10) form the basis of the HMM recipe. The conserved fields ρ, $\mathbf{q}$, and E may be coarse-grained and integrated on macroscopic space and time scales, while the fluxes τ and $\mathbf{J}$ may be estimated as statistical MD averages on microscales. This HMM coupling is valid assuming ergodicity of MD, and assuming that the state (ρ, $\mathbf{q}$, and E) uniquely determines the fluxes (τ and $\mathbf{j}$).

The conservation laws of Eqs. (6)–(8) are expressed in Eulerian coordinates $\mathbf{x}(t) = \mathbf{x}^0 + \mathbf{u}(t)$, where $\mathbf{u}$ is the displacement from the undeformed material position $\mathbf{x}^0$. To model material deformation, it is convenient to work in Lagrangian coordinates $\mathbf{x}^0$, where the equivalent conservation laws are [38]

$$\rho^0\partial_t\mathbf{A} - \nabla_{\mathbf{x}^0}\mathbf{q}^0 = 0 \tag{11}$$

$$\partial_t\mathbf{q}^0 - \nabla_{\mathbf{x}^0}\cdot\tau^0 = 0 \tag{12}$$

$$\partial_t e^0 + \nabla_{\mathbf{x}^0}\cdot\mathbf{j}^0 = 0, \tag{13}$$

where $\mathbf{A} = \nabla_{\mathbf{x}^0}\mathbf{x}$ is the deformation tensor. The mass, momentum, and energy densities are now expressed per unit Lagrangian volume,

$$\rho^0(\mathbf{x}^0) = \rho(\mathbf{x})\det\mathbf{A} \tag{14}$$

$$\mathbf{q}^0(\mathbf{x}^0) = \mathbf{q}(\mathbf{x})\det\mathbf{A} \tag{15}$$

$$e^0(\mathbf{x}^0) = e(\mathbf{x})\det\mathbf{A}. \tag{16}$$

Li and E [38] derived the averaged Lagrangian fluxes as analogues of Eqs. (9) and (10), but lacking kinetic energy terms:

$$\tau^0_{\Omega';\mathrm{L\&E}} = -\frac{1}{|\Omega'|}\sum_{i,j} c_{ij}\mathbf{f}_{ij}\otimes\mathbf{x}^0_{ij} \tag{17}$$

$$\mathbf{j}^0_{\Omega';\mathrm{L\&E}} = -\frac{1}{4|\Omega'|}\sum_{i,j} c_{ij}(\mathbf{v}_i+\mathbf{v}_j)\cdot\mathbf{f}_{ij}\mathbf{x}^0_{ij}, \tag{18}$$

where Ω' denotes a reference volume in Lagrangian coordinates. The flux τ^0 in Eq. (17) is interpreted as the first Piola–Kirchoff stress tensor [38]. These equations are valid when the microscale dynamics is fully Lagrangian [44,45]. However, actual MD simulation is almost always performed in Eulerian coordinates, where the lack of kinetic energy terms in Eqs. (17) and (18) was found to be invalid [43] (contradicting earlier work [46]). In our implementation, we calculate the Lagrangian fluxes by straightforward transformation of the Eulerian ones in the *center-of-mass* frame,

$$\tau^0 = (\mathbf{A}^{-1}\sigma)^T\det\mathbf{A} \tag{19}$$

$$\mathbf{j}^0 = (\mathbf{A}^{-1}\mathbf{j}_{\mathrm{c.o.m.}})^T\det\mathbf{A}. \tag{20}$$

The instantaneous Eulerian fluxes σ and $\mathbf{j}_{\mathrm{c.o.m.}}$ corresponding to an MD state are given by Eqs. (9) and (10) respectively, but with all atom velocities $\mathbf{v}_i$ replaced by $\mathbf{v}_i - \bar{\mathbf{v}}$, where $\bar{\mathbf{v}}$ is the center-of-mass velocity of the region Ω. Proper derivation of these Lagrangian fluxes is subtle, and beyond the scope of this article.

For simplicity, we work in the special case of a one-dimensional deformation. That is, the deformation gradient $\mathbf{A}$ is restricted to a single component A_{11} representing stretch along the x axis. In this case, the Lagrangian equations of motion reduce to [38]

$$\rho^0\partial_t A_{11} - \partial_x q^0_1 = 0 \tag{21}$$

$$\partial_t q^0_1 - \partial_x\sigma_{11} = 0 \tag{22}$$

$$\partial_t e^0 - \partial_x(\sigma_{11}v_1) = 0. \tag{23}$$

Thus, a three-dimensional MD simulation must provide only a single quantity, the (x,x) component of the usual virial stress σ (*i.e.* the continuum Cauchy stress [43].)

In the HMM scheme, we integrate the conserved quantities A_{11}, q^0_1 and e^0 on the macroscale. To complete the dynamical equations, the averaged virial stress σ_{11} must be provided by microscale simulation. At each relevant macroscale point x^0 we perform MD simulation of a perfect reference crystal in a volume deformed by an amount a_{11} along the x component, and with total Lagrangian energy density e^0. Note that the macroscale quantity q^0_1 sets the overall MD center-of-mass velocity, and can be ignored. We evolve the MD simulations until the estimate of virial stress σ_{11} is obtained with sufficient accuracy at each macroscale point x^0. The MD estimated stresses are then used to perform one macroscale integration time step.

3. Numerical integration of conservation laws

The HMM macroscale dynamics is a system of conservation laws, Eqs. (6)–(8) (or in the Lagrangian form, Eqs. (11)–(13)), which may be expressed generically as

$$\partial_t\mathbf{w} + \partial_x\mathbf{f}(\mathbf{w}) + \partial_y\mathbf{g}(\mathbf{w}) + \partial_z\mathbf{h}(\mathbf{w}) = 0, \tag{24}$$

where $\mathbf{w}$ is a collection of conserved fields and $\mathbf{f}$, $\mathbf{g}$ and $\mathbf{h}$ are the corresponding spatial components of flux. Many integration techniques are possible [47]. In one dimension, where $\mathbf{g} = \mathbf{h} = 0$, a particularly simple scheme is due to Nessyahu and Tadmor [48]. This second-order accurate shock-preserving scheme integrates the fields $\mathbf{w}$ on a staggered space–time grid,

$$\mathbf{w}^{n+1/2}_k = \mathbf{w}^n_k - \frac{\Delta t}{2\Delta x}\partial_x\mathbf{f}^n_k \tag{25}$$

$$\mathbf{w}^{n+1}_{k+1/2} = \frac{\mathbf{w}^n_k + \mathbf{w}^n_{k+1}}{2} + \frac{\Delta x}{8}\left(\partial_x\mathbf{w}^n_k - \partial_x\mathbf{w}^n_{k+1}\right) - \frac{\Delta t}{\Delta x}\left(\mathbf{f}^{n+1/2}_{k+1} - \mathbf{f}^{n+1/2}_k\right), \tag{26}$$

where subscripts and superscripts denote space and time indices respectively,

$$\mathbf{w}^n_k = \mathbf{w}(k\Delta x, n\Delta t). \tag{27}$$

The derivatives ∂_x above are defined with slope limiters to suppress unphysical oscillations at a shock. For each conserved field w in the collection $\mathbf{w}$, we use [49,50],

$$\partial_x w^n_k = \mathrm{MM}\left(\frac{w^n_{k+1} - w^n_k}{\Delta x}, \frac{w^n_{k+1} - w^n_{k-1}}{2\Delta x}, \frac{w^n_k - w^n_{k-1}}{\Delta x}\right),$$

and the analogous expression for $\partial_x f^n_k$. The MinMod function is defined as

$$\mathrm{MM}(v_1, v_2, \ldots) = \begin{cases} \min\limits_p\{v_p\} & \text{if } v_p > 0\ \forall p \\ \max\limits_p\{v_p\} & \text{if } v_p < 0\ \forall p \\ 0 & \text{otherwise.} \end{cases}$$

The Nessyahu–Tadmor central, staggered integration scheme is readily generalized to systems of conservation laws in $d > 1$ dimensions [51].

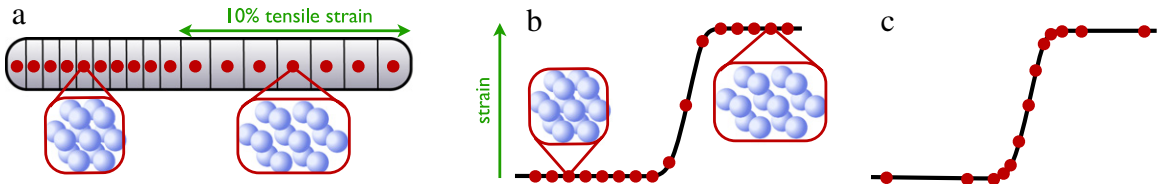

Fig. 1. Schematic illustration of our spatial adaptive sampling scheme applied to the simple problem of elastodynamic shock propagation in one-dimension. (a) Shock propagation dynamics occurs on macroscopic space and time scales. Red dots denote representative *microscopic* regions of the system, each of which is modeled by $\sim 10^3$ atoms of a defect-free fcc Cu crystal. (b) The heterogeneous multiscale method (HMM) integrates conservation laws to determine the dynamical evolution of the macroscopic fields. Constitutive data is provided by the stochastic estimates of *microscopic* molecular dynamics simulations, performed on a regular grid. (c) In spatial adaptive sampling, we dynamically adapt the location of microscopic simulations to increase the accuracy and efficiency of the HMM simulation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4. *Predictive* spatial adaptive sampling for full parallelism

Fig. 1 illustrates our approach in the context of one-dimensional elastodynamics. At a macroscopic level, the system configuration is represented by the A_{11} component of the deformation gradient tensor [Fig. 1(a)], the velocity, and the energy density fields. Macroscopic quantities such as the stress field are represented by microscopic samples of the system. According to HMM, we can use MD simulations at the microscale to estimate the local stress, and to evolve the system at the macroscale [Fig. 1(b)]. These MD simulations are the dominant computational cost of HMM. Our new technique of spatial adaptive sampling greatly improves the efficiency by interpolating the stress field from a small number of MD simulations performed at carefully selected spatial sample points [Fig. 1(c)].

Many types of interpolation are possible. Our experiments with linear interpolation indicate that sharp cusps in the stress tend to destabilize the Nessyahu–Tadmor integration scheme. Instead, we choose Akima splines [52] to smoothly approximate shocks, without introducing unphysical oscillations.

It remains to describe the procedure in which spatial sample points are selected to best reconstruct the stress field. A natural and direct approach is to recursively add sample points until the estimated accuracy of the interpolated stress satisfies an error tolerance. This iterative refinement approach minimizes the number of sample points, but has the disadvantage that the corresponding MD simulations cannot be executed in parallel. This synchronization drawback is especially pronounced for shock dynamics, where the highest resolution is required at a very limited spatial region, the shock front.

We introduce a *predictive* approach to spatial adaptive sampling, in which the number and location of all sample points are selected at the beginning of the macroscale time step. The corresponding MD simulations are then fully independent and can be executed in parallel. To predict important sample points, we use two heuristics: (1) higher resolution is likely to be required in regions where the macroscale fields are rapidly varying, and (2) the important sample points from the previous time step are correlated with the important sample points at the current time step.

Listing 1 presents our predictive sampling algorithm in the context of general conservation law dynamics, equation (24). The conserved fields $\mathbf{w}(\mathbf{x}, t) = \{w_1(\mathbf{x}, t), w_2(\mathbf{x}, t), \ldots\}$ include the deformation gradient, velocity, and energy. To integrate the conservation laws, we require MD-based flux estimates $\mathbf{f}(\mathbf{w}(\mathbf{x}, t)) = \{f_1(\mathbf{x}, t), f_2(\mathbf{x}, t), \ldots\}$. In our one-dimensional shock problem, the fluxes can be reconstructed from the single stress component σ_{11}. We discretize the conserved fields and fluxes on a regular grid, such that $w^n_{i;k}$ denotes the ith conserved field at the kth position and nth macroscopic time step. $f^n_{i;k}$ denotes the analogous flux. We express the sets of field and position indices as $\mathcal{I}$ and $\mathcal{K}$ respectively.

At the beginning of the nth macroscopic time step, the conserved fields $w^n_{i;k}$ are known for all $i \in \mathcal{I}$ and $k \in \mathcal{K}$. Our goal is to accurately and efficiently estimate the fluxes $f^n_{i,k}$ for all $i \in \mathcal{I}$ and $k \in \mathcal{K}$. With spatial adaptive sampling, we interpolate the full flux fields from samples $f^n_{i,k'}$ evaluated at a small set of *sample points* $k' \in \mathcal{S}^n \subset \mathcal{K}$. By construction, $\mathcal{S}^n$ is independent of the field index i.

Listing 1: Algorithm to construct a set of spatial sample points $k \in \mathcal{S}^n$ from which to interpolate the flux fields $\mathbf{f}^n(\mathbf{w}^n)$ (indexed by i) for the current time step n. A sample point k is considered important if any of the conserved fields $\mathbf{w}^n$ are rapidly varying at k (**gradientSampling**), or if k was important for interpolation at the *previous* time step (**holdoutSampling**). The accuracy is controlled by dimensionless threshold parameters $\epsilon_{i;g}$ and $\epsilon_{i;h}$.

```
def gradientSampling(S^n):
  for k ∈ K:
    if Δx|∇w^n_{i;k}/w^0_{i;k}| > ε_{i;g} for any i ∈ I:
      add k to set S^n

def estimatedInterpolationError():
  for k ∈ S^{n-1}, i ∈ I:
      // estimate f̂^{n-1}_{i;k} ≈ f^{n-1}_{i;k} by interpolating from the
      // set {f^{n-1}_{i;j} | j ∈ S^{n-1} and j ≠ k}
      f̂^{n-1}_{i;k} = interpolate(S^{n-1}, k, i)
      // interpolation errors from previous time step
      Δf^{n-1}_{i;k} = f̂^{n-1}_{i;k} − f^{n-1}_{i;k}
  return Δf^{n-1}

def holdoutSampling(Δf^{n-1}, B, S^n):
  if |Δf^{n-1}_{i;k}/f^0_{i;k}| > ε_{i;h} for any i ∈ I, k ∈ B ∩ S^{n-1}:
    add characteristicPoints(B) to set S^n
    // recursively iterate on subdivisions of set B
    for B' ∈ subBlocks(B):
      holdoutSampling(Δf^{n-1}, B', S^n)

def predictiveSampling():
  S^n = emptySet
  // add points k to S^n if w^n_k is rapidly varying
  gradientSampling(S^n)
  if t > 0:
    Δf^{n-1} = estimatedInterpolationError()
    // add points k to S^n if Δf^{n-1}_k is large
    holdoutSampling(Δf^{n-1}, K, S^n)
  return S^n
```

The **predictiveSampling** procedure applies two heuristics to build the sample points $\mathcal{S}^n$. The first heuristic is implemented in the **gradientSampling** procedure, which adds positions k to the set $\mathcal{S}^n$ if $w^n_{i;k}$ is rapidly varying for any field index i. Specifically, we use the condition $\Delta x|\nabla w^n_{i;k}/w^0_{i;k}| > \epsilon_{i;g}$ to test whether the variation of $w^n_{i;k}$ between neighbor grid points, relative to an initial reference value $w^0_{i;k}$, exceeds the dimensionless threshold parameter, $\epsilon_{i;g}$. This criterion helps to identify "interesting points" in the first time step, $n = 0$. However, specification of the thresholds $\epsilon_{i;g}$ supposes prior knowledge of the possible complex functional dependence $\mathbf{f}(\mathbf{w})$. And in any case, what we *actually* seek is a set of sample points $\mathcal{S}^n$ from which we can accurately interpolate $\mathbf{f}(\mathbf{w})$, regard-

less of the variation in **w**. Thus, when $n > 1$ we choose large or infinite $\epsilon_{i;g}$ to de-emphasize the importance of the gradient heuristic.

The **holdoutSampling** procedure implements a second, more robust heuristic. Knowing that the conserved fields **w** change gradually between time steps, we can use the fluxes sampled at $\mathcal{S}^{n-1}$ from the previous time step to predict important sample points $\mathcal{S}^n$ for the current time step. The procedure **estimatedInterpolationError** determines, for each $k \in \mathcal{S}^{n-1}$, how much error $\Delta f_{i;k}^{n-1}$ *would* arise if $f_{i;k}^{n-1}$ were interpolated from the sample points $\mathcal{S}^{n-1}$ excluding k. That is, the error estimate $\Delta f_{i;k}^{n-1}$ measures how important the sample point k was for interpolating the flux fields $f_{i;k}^{n-1}$ at the previous time step. We also use $\Delta f_{i;k}^{n-1}$ as an estimate of the importance of k for the *current* time step. If $|\Delta f_{i;k}^{n-1}/f_{i;k}^{0}|$ exceeds a dimensionless accuracy parameter $\epsilon_{i;h}$, we ensure that the point k is sampled in the next time step by including it in $\mathcal{S}^n$. We go further, and also add some neighbors of k to $\mathcal{S}^n$. The implementation is recursive on spatial blocks $\mathcal{B}$, beginning with the full spatial domain $\mathcal{B} = \mathcal{K}$. Whenever $\mathcal{B}$ contains a point k for which $|\Delta f_{i;k}^{n-1}/f_{i;k}^{0}| > \epsilon_{i;h}$, we add **characteristicPoints** of $\mathcal{B}$ to the set $\mathcal{S}^n$, and also recurse on **subBlocks** of $\mathcal{B}$. In our one-dimensional implementation the blocks $\mathcal{B}$ are line segments, the single characteristic point is the midpoint of the line segment, and sub-blocks of a line segment are its two halves. In this manner, mesh points are recursively added from $\mathcal{K}$ to $\mathcal{S}^n$ in the vicinity of sample points from $\mathcal{S}^{n-1}$ that could not be accurately estimated by interpolation. Our strategy generalizes to higher dimensions in a natural way.

In the next section, we present experiments that demonstrate the efficacy of our **predictiveSampling** algorithm for the one-dimensional shock problem. In particular, we find that spatial adaptive sampling provides high accuracy with a significantly reduced number of sample points compared to standard HMM. Our tests indicate that **holdoutSampling** alone is sufficient to select sample points S^n. Thus, only target accuracies $\epsilon_{i;h}$ are required. However, finite gradient thresholds $\epsilon_{i;g}$ may still be useful to increase efficiency by reducing interpolation requirements.

The algorithm in Listing 1 simplifies somewhat for our one-dimensional shock problem. Specifically, we sample only the single stress component, $f \to \sigma_{11}$, from which the remaining fluxes can be determined. Our algorithm then has only two dimensionless parameters: ϵ_h, the target accuracy in interpolated stress, and ϵ_g, the gradient threshold for the A_{11} component of the deformation gradient tensor.

5. Numerical experiments on shock propagation

We establish the consistency of our HMM approach by comparing to a fully atomistic ("direct-MD") reference simulation. In both cases, we simulate one-dimensional shock propagation in a defect-free fcc crystal of copper using an embedded atom method (EAM) interatomic potential [53]. To generate our initial conditions, we begin with an unstrained system thermalized to the temperature 129 K. This equilibrium configuration has energy density $e = -0.296$ eV/Å^3, fcc lattice constant $a = 3.618$ Å, and zero stress. Next, to generate a shock-wave, we instantaneously apply a 10% tensile strain ($A_{11} = 0.1$) to a subregion of the simulation domain, at fixed kinetic energy. In the strained region, we measure a tensile stress of $\sigma_{11} = 15$ GPa.

The details of our direct-MD reference simulation are as follows. The periodic simulation domain consists of $1440 \times 12 \times 12$ fcc unit cells, yielding an approximate linear system size $L = 0.52\,\mu$m along the x-axis. We integrate for a time $\tau = 19.5$ ps using velocity-Verlet time steps of size $\Delta t = 2.5$ fs. To construct the macroscale stress field, we locally average the virial stress in space and time. We then transform this stress field from Eulerian to Lagrangian coordinates, to enable comparison with the HMM simulation. The stress field at times 0 ps, 6.5 ps, 13 ps and 19.5 ps is shown in Fig. 2.

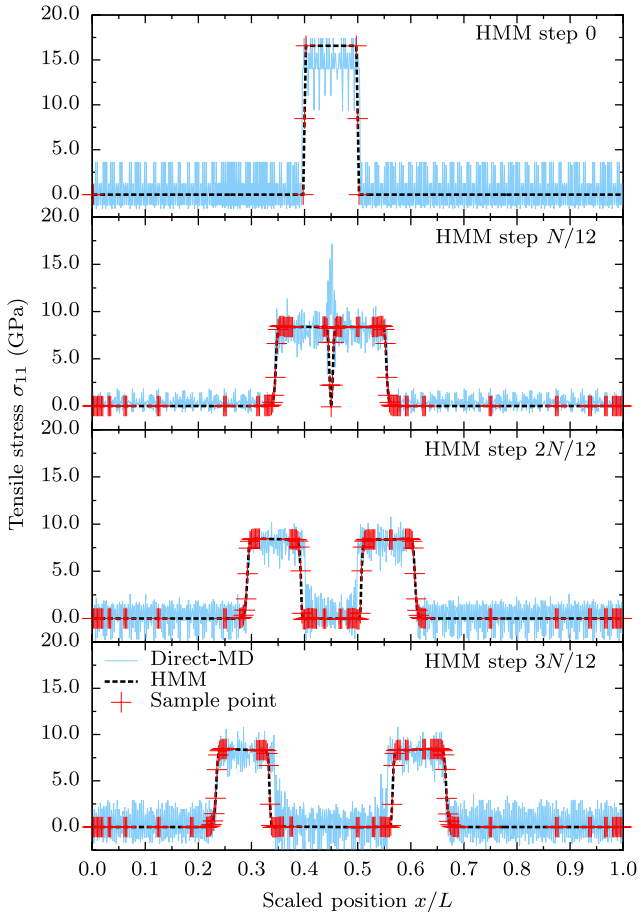

Fig. 2. One-dimensional shock propagation in a defect-free copper crystal. The light blue curve represents a fully atomistic direct-MD simulation of size $L = 0.52\,\mu$m for time 19.5 ps. The black curve represents an equivalent HMM simulation with $N = 400$ grid points using $\epsilon_g = 10^{-2}$ and $\epsilon_h = 10^{-3}$. The red marks are adaptive sample points from which the full stress field is spatially interpolated. The HMM solution is invariant under a rescaling of space and time, and thus also represents simulations on much larger physical scales. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In our corresponding HMM simulations, the macroscale Lagrangian fields (deformation gradient, velocity, and energy) evolve according to the dynamics of Eqs. (11)–(13). These conservation laws, of the general form in Eq. (24), are manifestly scale invariant: if $\mathbf{w}(\mathbf{x}, t)$ is a solution, then so is $\mathbf{w}(\lambda\mathbf{x}, \lambda t)$, for any $\lambda > 0$. Thus, our HMM simulation can be interpreted on any length scale L, including the direct-MD one. We discretize the conserved fields onto a regular lattice of N grid points, and integrate them using the Nessyahu–Tadmor scheme. By increasing N at fixed L, we decrease $\Delta x = L/N$ and improve accuracy. We select integration step size $\Delta t = \Delta x/c$ with $c = 6666$ m/s. At each HMM time step, we use the spatial adaptive sampling algorithm of Listing 1 to selectively spawn fine-scale MD computations, reconstruct the full stress field, and close the HMM macroscale equations. Each fine-scale MD simulation contains $12 \times 12 \times 12$ fcc unit cells, has periodic boundary conditions consistent with the local macroscopic strain, and has energy density fixed to the local macroscopic value. Otherwise, the fine-scale MD details are the same as in the direct-MD simulation. The output of each fine-scale simulation is the estimated virial stress, which we time-average over the second half of our 1.25 ps MD simulation.

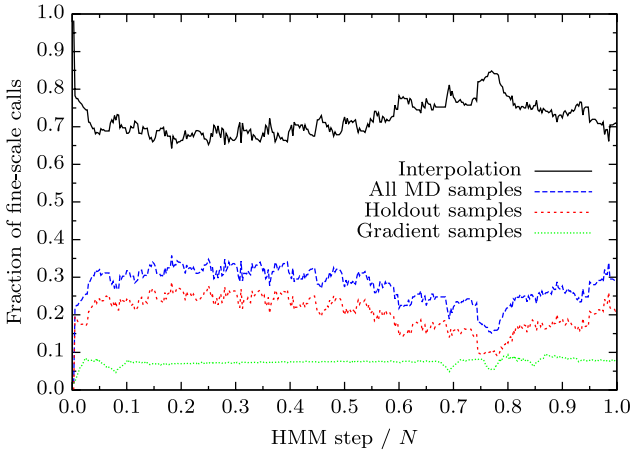

Fig. 3. Reduction in fine-scale simulations enabled by spatial adaptive sampling. For this simulation with $N = 400$ grid points, approximately 70% of the fine-scale MD simulations are replaced by interpolation. That is, we reconstruct the stress field at all 400 grid points using only 0.3×400 fine-scale simulations. Much greater savings is possible with larger N. The majority of sample points are selected by the **holdoutSampling** procedure to achieve target accuracy $\epsilon_h = 10^{-3}$.

Fig. 2 demonstrates the consistency of HMM and direct-MD simulations. Here we use a lattice with $N = 400$ macroscale grid points. For our spatial adaptive sampling scheme, we select gradient threshold $\epsilon_g = 10^{-2}$ and target accuracy $\epsilon_h = 10^{-3}$. To compare to the direct-MD simulation at 19.5 ps, we integrate for $N/4 = 100$ HMM time steps, corresponding to a physical time of $\Delta t N/4 = L/4c$, which is consistent when the HMM simulation is interpreted on the direct-MD length scale, $L = 0.52\ \mu$m.

HMM simulation is significantly accelerated by spatial adaptive sampling, which reduces the need for costly fine-scale MD simulations. Fig. 3 demonstrates that, for our simulation with $N = 400$ grid points, only about 30% of the N grid points require MD simulation. The other 70% are replaced with spatial adaptive sampling based interpolation. Below, we show that larger N leads to more dramatic efficiency gains. In all cases, we integrate for N HMM time steps, corresponding to physical time duration $T = L/c$. The performance of spatial adaptive sampling is roughly independent of integration time. After about $3N/4$ integration steps, the two traveling waves pass through the periodic boundaries and collide with each other. At the time of this collision, the number of required fine-scale simulations actually *decreases*.

To determine the important sample points, our algorithm contains two heuristics: **gradientSampling** and **holdoutSampling**. Fig. 4 demonstrates that the latter is sufficient. Here we use $N = 100$ grid points, fix the **holdoutSampling** target accuracy to $\epsilon_h = 10^{-3}$, and measure how the gradient threshold parameter ϵ_g affects the number of sample points. For large ϵ_g, the number of fine-scale calls spawned by **gradientSampling** drops to zero. We also observe saturation in the average and maximum errors in stress with increasing ϵ_g. These errors are defined as

$$\xi_{\text{avg}} = \frac{1}{LT\sigma_0} \int_0^L \int_0^T \left| \sigma_{11} - \sigma_{11}^{(\text{exact})} \right| \mathrm{d}t\mathrm{d}x \tag{28}$$

$$\xi_{\max} = \frac{1}{\sigma_0} \max_{x,t} \left| \sigma_{11} - \sigma_{11}^{(\text{exact})} \right|, \tag{29}$$

where

$$\sigma_0 = \max_{x,t} |\sigma_{11}^{(\text{exact})}|, \tag{30}$$

and $\sigma_{11}^{(\text{exact})}(\mathbf{x}, t)$ is the result of brute-force HMM simulation *without* the use of adaptive sampling, for which N fine-scale calls are required to construct the full stress field. This observed saturation

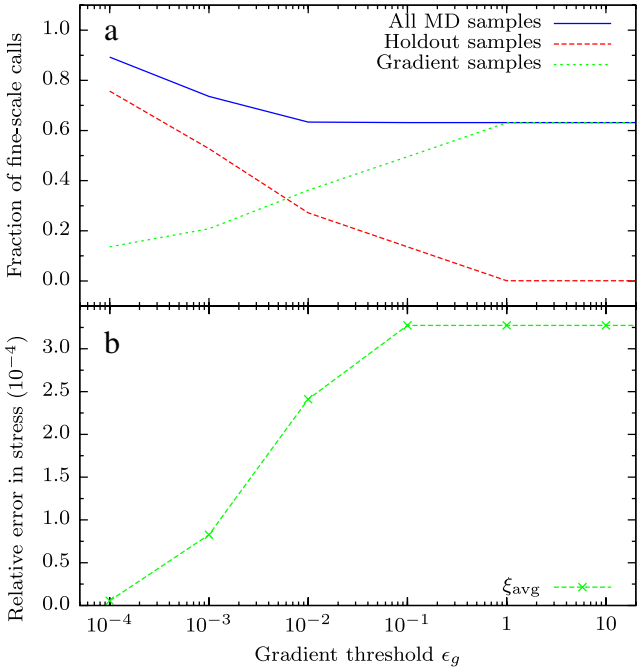

Fig. 4. The gradient threshold ϵ_g does not strongly affect the number of fine-scale sample points or the simulation error. We use $N = 100$ grid points and target accuracy $\epsilon_h = 10^{-3}$. (a) In the limit of large ϵ_g, the number of fine-scale samples is controlled solely by ϵ_h, and (b) the average relative simulation error ξ_{avg} saturates to a value comparable to ϵ_h.

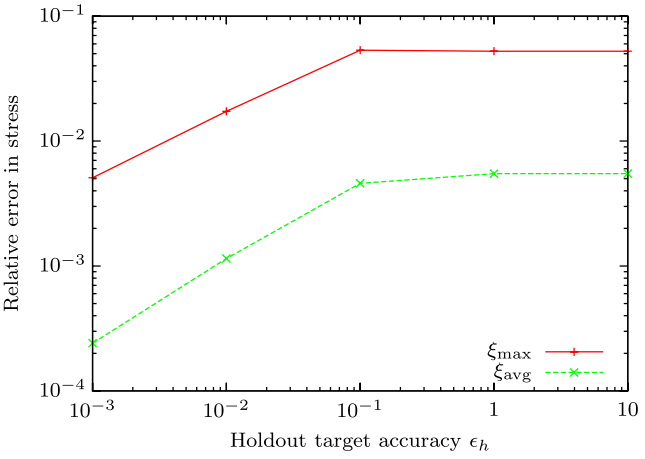

Fig. 5. The relative simulation error, measured with ξ_{avg} and $\xi_{\max}$, is strongly controlled by the target accuracy ϵ_h. We use $N = 100$ and $\epsilon_g = 10^{-2}$. The choice of ϵ_g does not play a strong role (cf. Fig. 4).

in ξ_{avg} confirms that, apart from the first HMM time step, we may take the limit $\epsilon_g \to \infty$ to effectively disable **gradientSampling** with little penalty.

Conversely, Fig. 5 demonstrates that the **holdoutSampling** target accuracy ϵ_h strongly controls the measured simulation errors ξ_{avg} and $\xi_{\max}$. Here we fix $\epsilon_g = 10^{-2}$ and vary ϵ_h, again with $N = 100$ grid points and N HMM time steps. For small ϵ_h, we observe that $\epsilon_h \approx \xi_{\text{avg}}$. The maximum error $\xi_{\max}$ also scales linearly with ϵ_h, but with a larger prefactor.

Fig. 6 illustrates that spatial adaptive sampling becomes extremely effective when the number of grid points N is very large. As before, we use $\epsilon_g = 10^{-2}$ and $\epsilon_h = 10^{-3}$. To extrapolate our measurements to very large N, we also perform HMM simulations with a stochastic linear stress–strain response in place of MD

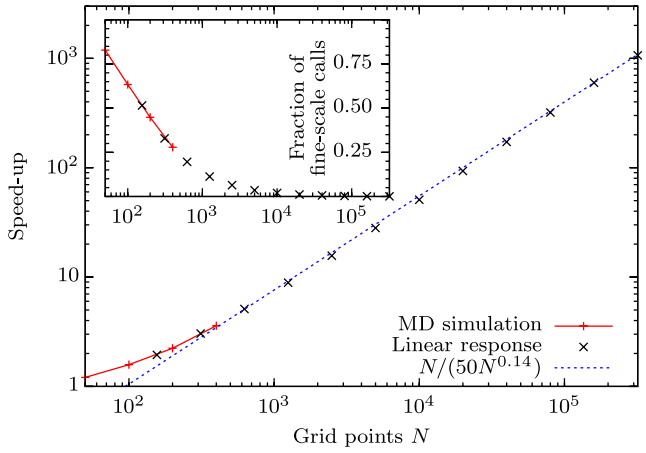

Fig. 6. Speed-up of spatial adaptive sampling compared to brute force HMM, measured by the reduction in fine-scale calls. We use $\epsilon_g = 10^{-2}$ and $\epsilon_h = 10^{-3}$. In our simulations of very large N, we replace fine-scale MD simulation with a stochastic, linear stress–strain response. Tremendous efficiency gains are possible when N grows large, for which the speed-up scales as $N/(50 \times N^{0.14})$.

simulation,

$$\tilde{\sigma}_{11}(A_{11}) = A_{11} \times 165.8\ \text{GPa} + \eta \times 160.0\ \text{kPa}. \quad (31)$$

Here, η represents a Gaussian random variable with zero mean and unit variance. We observe spatial adaptive sampling to perform similarly with both MD and surrogate fine-scale models. In both cases, approximately $50 \times N^{0.14}$ fine-scale sample points are needed to interpolate the stress response at all N grid points, yielding a speed-up factor of $N/(50 \times N^{0.14})$. Because fine-scale simulation is typically the overwhelming computational cost in HMM, spatial adaptive sampling offers tremendous efficiency gains when N is large.

6. Conclusion

We have presented a general technique that substantially accelerates many kinds of multiscale simulations, including HMM and equation-free methods. In these multiscale schemes, the macroscale equations must be completed with constitutive data provided by fine-scale scale simulation. This fine-scale simulation is the overwhelming computational cost. One way to reduce fine-scale simulation is standard adaptive sampling, in which all fine-scale simulations are stored in a database. Future constitutive data can be extracted from this database by interpolating in the high-dimensional space of fine-scale inputs. While attractive in principal, this standard approach to adaptive sampling approach suffers from the difficulty of high dimensional interpolation, and the bottleneck of synchronizing a central database.

In our work, we introduce a *spatial* approach to adaptive sampling. Our method does not use a database. Instead, at each macroscale time step we reconstruct the required constitutive data by interpolating fine-scale simulations on the $d \leq 3$-dimensional simulation domain. We also *predict* the important sample points at the beginning of each macroscale time step, to achieve near perfect parallelism, which we expect to be a major advantage for execution on next-generation supercomputers.

To demonstrate the efficiency of our approach, we used a one-dimensional model of elastodynamic shock propagation based on HMM. For this problem, our spatial adaptive sampling algorithm requires approximately $50 \times N^{0.14}$ sample points to interpolate the stress response at all N grid points, with target accuracy $\epsilon_h = 10^{-3}$. At the highest continuum resolution we tested, $N = 3.2 \times 10^5$, spatial adaptive sampling reduces fine-scale simulation (and thus the total cost of HMM) by three orders of magnitude. The shock problem we selected is particularly suitable to spatial adaptive sampling, but other multiscale simulations should also benefit. Our algorithm requires only a target accuracy parameter ϵ_h, and readily generalizes to more complicated models.

The code we developed for this paper is available from https://github.com/exmatex/CoHMM.

Acknowledgments

This work was supported by the Los Alamos Information Science & Technology Center (IS&T) Co-Design Summer School, the US Department of Energy (DOE), Office of Advanced Scientific Computing Research (ASCR) through the Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx, exmatex.org), and the Center for Nonlinear Studies (CNLS). C.J. acknowledges funding by a Los Alamos National Laboratory Director's Fellowship. Assigned: LA-UR 13-29626. Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the US Department of Energy under contract DE-AC52-06NA25396.

References

[1] D. Givon, R. Kupferman, A. Stuart, Extracting macroscopic dynamics: model problems and algorithms, Nonlinearity 17 (2004) R55–R127.
[2] P. Español, Statistical mechanics of coarse-graining, in: M. Karttunen, A. Lukkarinen, I. Vattulainen (Eds.), Novel Methods in Soft Matter Simulations, in: Lecture Notes in Physics, vol. 640, Springer, Berlin, Heidelberg, 2004, pp. 69–115.
[3] G. Lu, E. Kaxiras, Overview of multiscale simulations of materials, in: M. Rieth, W. Schommers (Eds.), Handbook of Theoretical and Computational Nanotechnology, vol. X, American Scientific Publishers, 2005, pp. 1–33.
[4] S.O. Nielsen, R.E. Bulo, P.B. Moore, B. Ensing, Recent progress in adaptive multiscale molecular dynamics simulations of soft matter, Phys. Chem. Chem. Phys. 12 (2010) 12401–12414.
[5] Y. Chen, J. Zimmerman, A. Krivtsov, D.L. McDowell, Assessment of atomistic coarse-graining methods, Internat. J. Engrg. Sci. 49 (2011) 1337–1349.
[6] S. Fritsch, S. Poblete, C. Junghans, G. Ciccotti, L. Delle Site, K. Kremer, Adaptive resolution molecular dynamics simulation through coupling to an internal particle reservoir, Phys. Rev. Lett. 108 (2012) 170602.
[7] H. Mori, Transport, collective motion, and brownian motion, Progr. Theoret. Phys. 33 (1965) 423–455.
[8] R. Zwanzig, Nonlinear generalized langevin equations, J. Stat. Phys. 9 (1973) 215–220.
[9] C. Hijon, P. Espanol, E. Vanden-Eijnden, R. Delgado-Buscalioni, Mori-zwanzig formalism as a practical computational tool, Faraday Discuss. 144 (2010) 301–322.
[10] X. Li, A coarse-grained molecular dynamics model for crystalline solids, Int. J. Numer. Methods Eng. 83 (2010) 986–997.
[11] R.E. Rudd, J.Q. Broughton, Coarse-grained molecular dynamics and the atomic limit of finite elements, Phys. Rev. B 58 (1998) R5893–R5896.
[12] R.E. Rudd, J.Q. Broughton, Coarse-grained molecular dynamics: Nonlinear finite elements and finite temperature, Phys. Rev. B 72 (2005) 144104.
[13] P. Español, Dissipative particle dynamics, in: S. Yip (Ed.), Handbook of Materials Modeling, Springer, Netherlands, 2005, pp. 2503–2512.
[14] W.G. Noid, J.-W. Chu, G.S. Ayton, V. Krishna, S. Izvekov, G.A. Voth, A. Das, H.C. Andersen, The multiscale coarse-graining method. I. A rigorous bridge between atomistic and coarse-grained models, J. Chem. Phys. 128 (2008) 244114.
[15] W.G. Noid, P. Liu, Y. Wang, J.-W. Chu, G.S. Ayton, S. Izvekov, H.C. Andersen, G.A. Voth, The multiscale coarse-graining method. II. Numerical implementation for coarse-grained molecular models, J. Chem. Phys. 128 (2008) 244115.
[16] S.T. O'Connell, P.A. Thompson, Molecular dynamics–continuum hybrid computations: a tool for studying complex fluid flows, Phys. Rev. E 52 (1995) R5792–R5795.
[17] X.B. Nie, S.Y. Chen, W.N. E, M.O. Robbins, A continuum and molecular dynamics hybrid method for micro- and nano-fluid flow, J. Fluid Mech. 500 (2004) 55–64.
[18] K. Mohamed, A. Mohamad, A review of the development of hybrid atomistic–continuum methods for dense fluids, Microfluid. Nanofluid. 8 (2010) 283–302.
[19] E.B. Tadmor, R. Phillips, M. Ortiz, Mixed atomistic and continuum models of deformation in solids, Langmuir 12 (1996) 4529–4534.
[20] V.B. Shenoy, R. Miller, E.B. Tadmor, R. Phillips, M. Ortiz, Quasicontinuum models of interfacial structure and deformation, Phys. Rev. Lett. 80 (1998) 742–745.

[21] R.E. Miller, E.B. Tadmor, Hybrid continuum mechanics and atomistic methods for simulating materials deformation and failure, MRS Bull. 32 (2007) 920–926.
[22] R.E. Miller, E.B. Tadmor, A unified framework and performance benchmark of fourteen multiscale atomistic/continuum coupling methods, Model. Simul. Mater. Sci. Eng. 17 (2009) 053001.
[23] W. Cai, M. de Koning, V.V. Bulatov, S. Yip, Minimizing boundary reflections in coupled-domain simulations, Phys. Rev. Lett. 85 (2000) 3213–3216.
[24] W. E, Z. Huang, Matching conditions in atomistic–continuum modeling of materials, Phys. Rev. Lett. 87 (2001) 135501.
[25] W. E, Z. Huang, A dynamic atomistic–continuum method for the simulation of crystalline materials, J. Comput. Phys. 182 (2002) 234–261.
[26] G.J. Wagner, W.K. Liu, Coupling of atomistic and continuum simulations using a bridging scale decomposition, J. Comput. Phys. 190 (2003) 249–274.
[27] G. Wagner, R. Jones, J. Templeton, M. Parks, An atomistic-to-continuum coupling method for heat transfer in solids, Comput. Methods Appl. Mech. Eng. 197 (2008) 3351–3365.
[28] X. Li, J.Z. Yang, W. E, A multiscale coupling method for the modeling of dynamics of solids with application to brittle cracks, J. Comput. Phys. 229 (2010) 3970–3987.
[29] C. Theodoropoulos, Y.-H. Qian, I.G. Kevrekidis, Coarse stability and bifurcation analysis using time-steppers: a reaction–diffusion example, Proc. Natl. Acad. Sci. USA 97 (2000) 9840–9843.
[30] I.G. Kevrekidis, G. Samaey, Equation-free multiscale computation: Algorithms and applications, Annu. Rev. Phys. Chem. 60 (2009) 321–344.
[31] A. Brandt, Multiscale scientific computation: review 2001, in: T.J. Barth, T. Chan, R. Haimes (Eds.), Multiscale and Multiresolution Methods, in: Lecture Notes in Computational Science and Engineering, vol. 20, Springer, Berlin, Heidelberg, 2002, pp. 3–95.
[32] W. E, B. Engquist, Z. Huang, Heterogeneous multiscale method: a general methodology for multiscale modeling, Phys. Rev. B 67 (2003) 092101.
[33] W. E, B. Engquist, X. Li, W. Ren, E. Vanden-Eijnden, Heterogeneous multiscale methods: a review, Commun. Comput. Phys. 2 (2007) 367–450.
[34] J. Knap, N.R. Barton, R.D. Hornung, A. Arsenlis, R. Becker, D.R. Jefferson, Adaptive sampling in hierarchical simulation, Int. J. Numer. Methods Eng. 76 (2008) 572–600.
[35] N.R. Barton, J. Knap, A. Arsenlis, R. Becker, R.D. Hornung, D.R. Jefferson, Embedded polycrystal plasticity and adaptive sampling, Int. J. Plasticity 24 (2008) 242–266.
[36] A. Arsenlis, N. Barton, R. Becker, R. Rudd, Generalized in situ adaptive tabulation for constitutive model evaluation in plasticity, Comput. Methods Appl. Mech. Eng. 196 (2006) 1–13.
[37] N.R. Barton, J.V. Bernier, J. Knap, A.J. Sunwoo, E.K. Cerreta, T.J. Turner, A call to arms for task parallelism in multi-scale materials modeling, Int. J. Numer. Methods Eng. 86 (2011) 744–764.
[38] X. Li, W. E, Multiscale modeling of the dynamics of solids at finite temperature, J. Mech. Phys. Solids 53 (2005) 1650–1685.
[39] M. Berger, P. Colella, Local adaptive mesh refinement for shock hydrodynamics, J. Comput. Phys. 82 (1989) 64–84.
[40] J. Bell, M. Berger, J. Saltzman, M. Welcome, Three-dimensional adaptive mesh refinement for hyperbolic conservation laws, SIAM J. Sci. Comput. 15 (1994) 127–138.
[41] W. E, B. Engquist, Y. Sun, The heterogeneous multiscale methods with application to combustion, in: T. Echekki, E. Mastorakos (Eds.), Turbulent Combustion Modeling, in: Fluid Mechanics and Its Applications, vol. 95, Springer, Netherlands, 2011, pp. 439–459.
[42] J.H. Irving, J.G. Kirkwood, The statistical mechanical theory of transport processes IV, J. Chem. Phys. 18 (1950) 817–829.
[43] A.K. Subramaniyan, C. Sun, Continuum interpretation of virial stress in molecular simulations, Int. J. Solids Struct. 45 (2008) 4340–4346.
[44] J. Gao, J.H. Weiner, Excluded-volume effects in rubber elasticity. 1. Virial stress formulation, Macromolecules 20 (1987) 2520–2525.
[45] P.C. Andia, F. Costanzo, G.L. Gray, A classical mechanics approach to the determination of the stress–strain response of particle systems, Model. Simul. Mater. Sci. Eng. 14 (2006) 741.
[46] M. Zhou, A new look at the atomic level virial stress: on continuum-molecular system equivalence, Proc. R. Soc. Lond. Ser. A 459 (2003) 2347–2392.
[47] R. LeVeque, Numerical Methods for Conservation Laws, Birkhäuser, Bosel, 1992.
[48] H. Nessyahu, E. Tadmor, Non-oscillatory central differencing for hyperbolic conservation laws, J. Comput. Phys. 87 (1990) 408–463.
[49] A. Harten, The artificial compression method for computation of shocks and contact discontinuities. I. Single conservation laws, Commun. Pur. Appl. Math 30 (1977) 611–638.
[50] P. Sweby, High resolution schemes using flux limiters for hyperbolic conservation laws, SIAM J. Numer. Anal. 21 (1984) 995–1011.
[51] G. Jiang, E. Tadmor, Nonoscillatory central schemes for multidimensional hyperbolic conservation laws, SIAM J. Sci. Comput. 19 (1998) 1892–1917.
[52] H. Akima, A new method of interpolation and smooth curve fitting based on local procedures, J. ACM 17 (1970) 589–602.
[53] Y. Mishin, M.J. Mehl, D.A. Papaconstantopoulos, A.F. Voter, J.D. Kress, Structural stability and lattice defects in copper: *ab initio*, tight-binding, and embedded-atom calculations, Phys. Rev. B 63 (2001) 224106.