Partial Thermostatting of Coarse-Grained Molecular Dynamics

Simon Gill,[1] Zhidong Jia,[2] Ben Leimkuhler,[3] and Alan Cocks[1]

[1]*Department of Engineering, University of Leicester, University Road, Leicester, UK, LE1 7RH*
[2]*Institute of Computational Mathematics, Chinese Academy of Sciences (CAS), Beijing 100080, CHINA*
[3]*Department of Mathematics, University of Leicester, University Road, Leicester, UK, LE1 7RH*

(Dated: May 26, 2005)

Thermalization issues are studied for a multiscale molecular model based on successive coarse-graining. It is found that high frequency thermal energy can be trapped for long periods in regions with the smallest length scale. As these are precisely the regions where accurate dynamical modelling is generally required, a hybrid strategy is proposed to avoid introduction of artificial thermostatting effects, based on the combination of two novel computational algorithms: Partial Thermostatting Molecular Dynamics (PTMD) is used to thermostat the coarse-grained region while preserving the dynamics of the atomistic region, and Recursive Multiple Thermostats (RMT) provides an effective multiscale dynamic-stochastic heat bath. The combined method is applied to a 1D quasicontinuum model, in which a molecular dynamics simulation is embedded within a dynamic finite element model.

PACS numbers: 02.70.Ns, 02.70.Dh, 45.20.Jj, 65.40.De

The boundary conditions for molecular dynamics (MD) simulations in the condensed phase are a compromise between correct representation of the far field and minimization of the system size due to computational constraints. In recent years, concurrent multiscale methods have been developed for crystalline solids in which the complex response of the far field is represented by a coarse-grained continuum region constructed from finite elements (see [1, 2] for reviews). The requirements of the coarse-grained far field depend on the nature of the simulation, generally either sampling or dynamics. If the purpose of simulation is sampling of equilibrium quantities, then typically only slowly-changing thermodynamical or statistical quantities are of interest and inertial effects are small. Rapid changes occur in truly dynamic situations such as dynamic fracture.

Dynamical simulations are complicated by the reflection of high frequency phonons from the interface between the atomistic and coarse-grained (CG) regions. This leads to energy trapping and localized heating [2]. In most cases, the CG region is only required to provide a (slowly-evolving) statistically accurate (elastostatic) far field representation. Correct transmission of phonons across the interface [1–3] is only necessary if the far boundaries can be seen during the simulation period (e.g. MEMS) or there are two atomistic regions which need to interact dynamically via the CG medium (e.g. two cracks). We assume here that absorption of phonons at the interface is a sufficient requirement.

The ensemble is consequently canonical (constant temperature) as opposed to microcanonical (constant energy). Conventional thermostatting algorithms control the temperature via velocity rescaling. This corrupts the true dynamics of the regions in which it is applied. Liu [4] has recently demonstrated that MD simulations of nanoindentation are very sensitive to global thermostatic control.

This paper aims to address the problem of phonon reflection while preserving the correct dynamics in the atomistic region. We proceed in four stages. First, a finite temperature CG representation of the body is developed and compared with atomistic results. Choice of an appropriate interface between the CG and atomistic regions is then explored. The performance of different local thermostatting algorithms is assessed in the context of correct sampling, and, finally, the more demanding requirements of dynamical simulation are considered.

For the purposes of demonstration, we restrict the analysis to a one-dimensional chain of N atoms. Each atom has mass m and is assigned an index $a = 1, \ldots, N$. Atom a has position q_a and momentum p_a. The atoms interact via the interatomic potential $\phi(r)$ which is a function of the atomic separation r. The dynamics of the above system are described by the following Hamiltonian

$$H(\boldsymbol{Q}, \boldsymbol{P}) = \frac{\boldsymbol{P}^T \boldsymbol{M}^{-1} \boldsymbol{P}}{2} + V(\boldsymbol{Q}), \tag{1}$$

where $\boldsymbol{P} = \{p_a\}$, $\boldsymbol{Q} = \{q_a\}$, $\boldsymbol{M} = m\boldsymbol{I}$, and $\boldsymbol{I}$ is the identity matrix. The potential energy is assumed to have the form:

$$V(\boldsymbol{Q}) = \sum_{a=1}^{N-1} \sum_{b=a+1}^{N} \phi(|q_b - q_a|). \tag{2}$$

The coarse-graining procedure should reduce the number of degrees of freedom of the fully atomistic model defined in (1) while preserving the dynamical and sampling properties of part or all of the system. Following the philosophy of the zero-temperature quasicontinuum method [5], a reduced set of $n(< N)$ atoms are chosen to represent the system. These representative (rep.) atoms form the ends of sub-chains of atoms. The behavior of the slave atoms within each sub-chain (or element) are inferred from the behavior of the rep. atoms at each end.

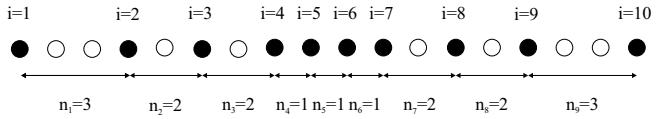

FIG. 1: Coarse-grained geometry of a 1D chain with 10 representative atoms (•) and 8 slave atoms (◦).

The rep. atoms are given the indices $i = 1, \ldots, n$ and each element i is bounded by rep. atoms i and $i+1$. Element i represents n_i atoms, two rep. atoms shared between two elements and $n_i - 1$ slave atoms. An example of the coarse-graining process is illustrated in Figure 1 for N=18 and n=10. This configuration is described by the set $\{3^1, 2^2, 1^3, 2^2, 3^1\}$.

The state of the system is thus characterized by the positions, $\boldsymbol{q}$, and momenta, $\boldsymbol{p}$, of the rep. atoms. To retain correct sampling of quantities in terms of positions and momenta of rep. atoms [6] the partition function must remain constant

$$\begin{aligned} Z &= \frac{1}{h^N}\int_\infty^\infty\int_\infty^\infty \exp\left(-\beta H(\boldsymbol{Q},\boldsymbol{P})\right) d\boldsymbol{Q}d\boldsymbol{P} \qquad (3)\\ &= \frac{1}{h^N}\int_\infty^\infty\int_\infty^\infty \exp\left(-\beta H_{CG}(\boldsymbol{q},\boldsymbol{p})\right) d\boldsymbol{q}d\boldsymbol{p}, \end{aligned}$$

where h is Planck's constant, $\beta = 1/k_BT$, k_B is Boltzmann's constant and T is absolute temperature. The resulting coarse-grained Hamiltonian is

$$H_{CG}(\boldsymbol{q},\boldsymbol{p}) = \frac{\boldsymbol{p}^T\boldsymbol{m}^{-1}\boldsymbol{p}}{2} + V_{CG}(\boldsymbol{q},T), \qquad (4)$$

where the mass matrix $\boldsymbol{m} = m\boldsymbol{I}$ contains only the masses of the rep. atoms (not the slave atoms). As dynamics is only modelled in the atomistic region, the masses of rep. atoms in the CG part are in fact arbitrary. For consistency with other approaches [6], we have utilized a lump mass model (in which the mass of the slave atoms is attributed equally to the element end nodes) for all simulations.

The CG potential $V_{CG}(\boldsymbol{q},T)$ can only be calculated analytically for a harmonic interatomic potential. Assuming only nearest-neighbor interactions with $\phi(r) = \frac{1}{2}k(r-r_0)^2$ gives

$$\begin{aligned} V_{CG}(\boldsymbol{q},T) &= \frac{1}{2}\sum_{i=1}^{n-1}\frac{k}{n_i}(q_{i+1}-q_i-n_ir_0)^2\\ &+ k_BT\sum_{i=1}^{n-1}(n_i-1)\ln\left(\frac{\hbar\omega_i}{k_BT}\right), \qquad (5) \end{aligned}$$

where the frequency $\omega_i = n_i^{\frac{1}{2(n_i-1)}}\sqrt{\frac{k}{m}}$. The first term on the RHS of (5) is the potential energy of the deformed lattice. The second term represents the vibrational energy of the slave atoms and is constant for a harmonic potential. Anharmonic effects such as thermal expansion only arise for nonlinear potentials. We make the approximation that (5) can be extended for a general potential such that the variant contribution is

$$\begin{aligned} V_{CG}(\boldsymbol{q},T) &= \sum_{i=1}^{n-1} n_i\phi(r_i)\\ &+ \frac{1}{2}k_BT\sum_{i=1}^{n-1}(n_i-1)\ln\left(\left.\frac{\partial^2\phi}{\partial r^2}\right|_{r=r_i}\right) \qquad (6) \end{aligned}$$

where $r_i = (q_{i+1}-q_i)/n_i$ is the mean interatomic spacing in element i. This is similar to the result of LeSar et al. [7] for high temperatures. Equation (6) reduces to the atomistic potential for no coarse-graining (all $n_i = 1$) as expected. Note that rep. atoms only truly represent the time-averaged response of an atom (or group of atoms); their oscillation period and amplitude are not those of a normal atom.

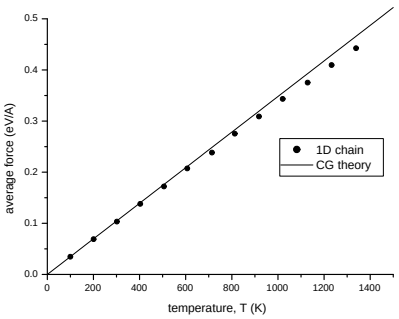

FIG. 2: Average thermal expansion force in a fully atomistic chain with predictions from the coarse-grained potential (6).

For illustration, the Lennard-Jones potential $\phi(r) = 4\epsilon\left(\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^6\right)$ is used. Only nearest-neighbor interactions are considered. In the simulations presented here, we took typical values of ϵ=0.6eV, $\sigma = 2.5\mathring{A}$ and $m = 10^{-25}$kg. However, these results are general and the exact choice of parameters is not critical. The simulations are initialized by assigning rep. atoms momenta from an appropriately weighted Gaussian distribution with zero net momentum. The ends of the chain are fixed so that the average lattice spacing is always the zero Kelvin lattice spacing, $r_0 = \sqrt[6]{2}\sigma$.

For constrained thermal expansion, the average force acting between each atom is predicted to be $\lim_{n_i\to\infty}\left(-\frac{1}{n_i}\left.\frac{\partial V_{CG}}{\partial r}\right|_{r=r_0}\right)$ for a large, single ($n = 2$) element. This is compared with a fully atomistic chain of 160 atoms, $\{1^{160}\}$, in Figure 2. The results are in good agreement for lower temperatures with only a slight deviation at higher temperatures.

For the remainder of this paper, we investigate a partially coarse-grained chain of 1250 atoms of the form {CG9,IF,AR,IF,CG9}. Full atomic resolution is retained in the central atomistic region (AR) which contains

200 atoms, denoted $\{1^{200}\}$. The two outermost regions (CG9) are heavily coarse-grained with 8 slave atoms per element. Each represents 432 atoms and are denoted $\{9^{48}\}$. The two interface (IF) regions each represent 93 atoms and are chosen to increase performance. Without thermostatting it is found that any coarse-graining of the chain drastically slows its convergence to thermal equilibrium, with the effect becoming more pronounced with increasing levels of coarse-graining. This is ameliorated (although not removed) by gradual transition to full coarse-graining via the IF region. An IF region of $\{7^5, 5^5, 3^5, 1^{18}\}$ was found to be appropriate. Overlapping the fully atomistic region into the IF, $\{1^{18}\}$, facilitates thermalization of the AR. The final chain configuration, $\{9^{48}, 7^5, 5^5, 3^5, 1^{18}, 1^{200}, 1^{18}, 3^5, 5^5, 7^5, 9^{48}\}$, contains 362 rep. atoms.

A common approach to temperature control is to apply a thermostat to the whole system. This can be done using the well-known Nosé Hamiltonian [8]

$$\begin{aligned}\mathcal{H}_{Nosé}(\boldsymbol{q},\tilde{\boldsymbol{p}},s,\pi) &= \frac{\tilde{\boldsymbol{p}}^T\boldsymbol{m}^{-1}\tilde{\boldsymbol{p}}}{2s^2}+\frac{\pi^2}{2\mu} \\ &+ V_{CG}(\boldsymbol{q},T)+gk_BT\ln s \end{aligned} \quad (7)$$

where s and π are conjugate position and momentum thermostatting variables, μ is a fictional thermal mass which determines the strength of the thermal coupling to the system, and g is the number of degrees of freedom (including s). The virtual momenta $\tilde{\boldsymbol{p}}$ is related to the actual momenta by $\tilde{\boldsymbol{p}} = s\boldsymbol{p}$. In practice the timescale must be re-parameterized for numerical computation. A symplectic scheme with re-parameterization of time can be realized by the Nosé-Poincaré (NP) Hamiltonian [9] $\mathcal{H}_{NP} = s(\mathcal{H}_{Nosé} - \mathcal{H}_0)$ where $\mathcal{H}_0$ is the initial value of the Nosé Hamiltonian. (The use of symplectic NP-based schemes has stability advantages [9–11] compared to the more common Nosé-Hoover method.)

Thermostatting the entire system is undesirable from a modelling perspective as velocity rescaling can potentially corrupt the dynamics of the AR region. The Partial Thermostatting Molecular Dynamics (PTMD) technique proposed by Jia and Leimkuhler[12] avoids this problem. The set of rep. atoms is divided into two sets, with positions $\boldsymbol{q} = \{\boldsymbol{q_a}, \boldsymbol{q_b}\}$, associated momenta $\boldsymbol{p} = \{\boldsymbol{p_a}, \boldsymbol{p_b}\}$ and masses $\boldsymbol{m} = \{\boldsymbol{m_a}, \boldsymbol{m_b}\}$ where a and b denote thermostatted and unthermostatted degrees of freedom respectively. Applying a Nosé thermostat to set a gives the partial thermostatting Hamiltonian

$$\begin{aligned}\mathcal{H}^{PT}_{Nosé} &= \frac{\tilde{\boldsymbol{p}}_{\boldsymbol{a}}^T\boldsymbol{m}_{\boldsymbol{a}}^{-1}\tilde{\boldsymbol{p}}_{\boldsymbol{a}}}{2s^2}+\frac{\tilde{\boldsymbol{p}}_{\boldsymbol{b}}^T\boldsymbol{m}_{\boldsymbol{b}}^{-1}\tilde{\boldsymbol{p}}_{\boldsymbol{b}}}{2} \\ &+ \frac{\pi^2}{2\mu}+V_{CG}(\boldsymbol{q},T)+gk_BT\ln s \end{aligned} \quad (8)$$

A time transformation is applied to regularize the time variable for the thermostatted (a) variables [12]. The resulting system can be viewed as coupling Newtonian dynamics for the b variables with a Nosé-Poincaré style thermal reservoir. The PTMD model preserves volume and recovers the canonical distribution [12]. This allows the temperature of any part of the chain to be directly regulated, forming an effective heat bath for the remainder of the system. In this paper the IF and CG9 regions are thermostatted with a target temperature of 300K. The temperature in a chain initiated far from thermal equilibrium, with the AR at 600K and the rest (IF and CG9) at 300K, is found to be well-controlled by a single NP thermostat. The NP thermostat provides a reasonable heat bath for the AR and the average temperature in each region approaches the target temperature within a practical time scale.

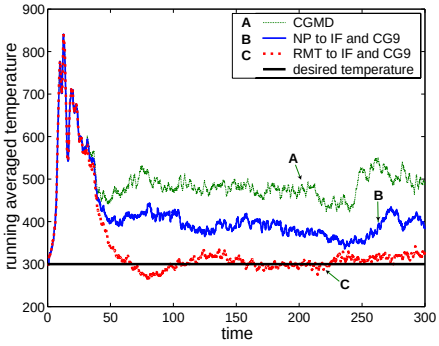

FIG. 3: Running average of temperature in the AR for no thermostatting (CGMD), thermostatting with NP applied to IF and CG9, and thermostatting with RMT applied to IF and CG9. Only RMT absorbs the high frequency phonons at the IF and hence correctly regulate the temperature in the AR.

Thermostatting a more demanding dynamic application is next considered. The chain is initially equilibrated at 300K and then kinetic energy is injected into the AR from time t=0 to t=20 by forcing the central atom in the AR to oscillate with an amplitude of $0.03r_0$ and at 1.2 times the atomic harmonic oscillation frequency. This generates a succession of high frequency travelling waves which rapidly raise the temperature in the AR as shown in Figure 3. The energy packet reaches the IF at t=25 and should have left the AR by t=45. Almost 50% of this energy is trapped in the AR if the simulation is unthermostatted (CGMD). Application of NP to IF and CG9 does not provide a good solution to this problem because a single thermostat cannot respond rapidly to regions operating at different time scales resulting in phonon reflection at the AR/IF interface.

However, any thermostatting technique can be used within the PTMD algorithm. The multiscale thermostatting problem is resolved within a Hamiltonian framework by using the Recursive Multiple Thermostat (RMT) algorithm proposed in [13]. This uses a hierarchy of r thermostatting variables, s_1 to s_r, in which s_1 thermostats s_2 and the original system, and so on in succession. This breaks up the structure of the resulting thermostat and

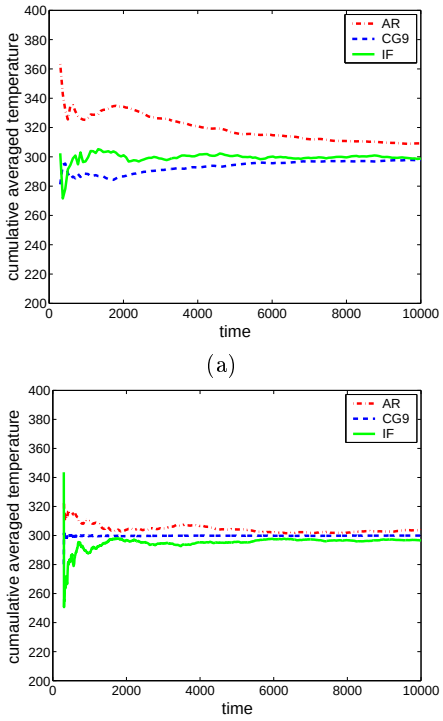

FIG. 4: Cumulative average of temperature in AR, IF and CG9 shows long time response of Fig. 3 for (a) NP applied to IF and CG9 and (b) RMT applied to IF and CG9. RMT provides correct sampling whereas NP does not.

is expected to generate an effective multiscale stochastic heat bath. The following generalized RMT NP Hamiltonian is proposed $\mathcal{H}_{NP}^{RMT} = s_1 \cdots s_r(\mathcal{H}_{Nosé}^{RMT} - \mathcal{H}_0)$ where

$$\begin{aligned}\mathcal{H}_{Nosé}^{RMT} &= \frac{\tilde{\boldsymbol{p}}_{\hat{\boldsymbol{a}}}^T \boldsymbol{m}_{\hat{\boldsymbol{a}}}^{-1} \tilde{\boldsymbol{p}}_{\hat{\boldsymbol{a}}}}{2s_1^2 \cdots s_r^2} + \frac{\tilde{\boldsymbol{p}}_{\bar{\boldsymbol{a}}}^T \boldsymbol{m}_{\bar{\boldsymbol{a}}}^{-1} \tilde{\boldsymbol{p}}_{\bar{\boldsymbol{a}}}}{2s_2^2 \cdots s_r^2} + \frac{\tilde{\boldsymbol{p}}_{\boldsymbol{b}}^T \boldsymbol{m}_{\boldsymbol{b}}^{-1} \tilde{\boldsymbol{p}}_{\boldsymbol{b}}}{2} + \frac{\pi_r^2}{2\mu_r} \\ &+ V_{CG}(\boldsymbol{q}, T) + \sum_{k=1}^{r-1} \frac{\pi_k^2}{2\mu_k s_{k+1}^2 \cdots s_r^2} + \hat{g} k_B T \ln s_1 \\ &+ \sum_{k=2}^{r} \left((g + k - 1) k_B T \ln s_k + \frac{(1 - s_k)^2}{2C_k} \right). \quad (9)\end{aligned}$$

The sets $\hat{a}$ and $\bar{a}$ represent quantities in the IF and CG9 regions respectively, $\hat{g}$ and $\bar{g}$ are the numbers of degrees of freedom in those sets and $g = \hat{g} + \bar{g}$. Importantly, the numerical scheme can be easily and efficiently implemented by Hamiltonian splitting [13]. In our simulations we used r=3, $\mu_1 = m$, $\mu_2 = 1.2m$, $\mu_3 = 1.5m$, $C_2 = 4.1667\beta$ and $C_3 = 0.003333\beta$.

Application of RMT to the IF and CG9 regions for the dynamic problem is shown in Figure 3. It is clearly a great improvement on the NP thermostat. The temperature in the AR is (indirectly) controlled very well by RMT, with practically no phonon reflection. The advantages of RMT are two fold. First, the multiple thermostats can respond to the inherently different frequencies in the different regions. Second, like all Nosé dynamics methods, RMT introduces a feedback control involving average kinetic energy; this control becomes necessarily less sensitive with increasing system size. Modified RMT can respond simultaneously to rapid changes in the small IF and slower changes in the larger CG9 region. The thermostat consequently responds rapidly to the energy entering the IF, absorbing all the high frequency phonons, and effectively thermostatting the low frequency CG9 region. The long time sampling response of the NP and RMT methods for this problem is shown in Fig. 4. RMT equilibrates the chain well with all regions approaching the target temperature within a practical time scale. NP does not achieve this.

In conclusion, PTMD has been used to provide a coarse-grained heat bath for an atomistic simulation. A single Nosé thermostat can adequately thermalize a multiscale chain for slowly changing problems but not dynamic ones. RMT provides a hierarchy of successive thermostatting variables to generate a multiscale stochastic heat bath. Application of RMT to an appropriate coarse-grained region has been shown to effectively thermostat highly dynamic simulations without energy trapping.

Financial support for this work from the EPSRC through the award of grant GR/R24104 is gratefully acknowledged by all the authors. The work of BL was supported by a Leverhulme Foundation Research Fellowship. The first three authors are in alphabetical order.

[1] S. P. Xiao and T. Belytschko, Comput. Methods Appl. Mech. Engrg **193**, 1645 (2004).
[2] W. K. Liu, E. G. Karpov, S. Zhang and H. S. Park, Comput. Methods Appl. Mech. Engrg **193**, 1529 (2004)
[3] J. Q. Broughton, F. F. Abrahams, N. Bernstein and E. Kaxiras, Phys. Rev. B **60**, 2391 (1999); R. E. Rudd and J. Q. Broughton, Phys. Stat. Sol. B **217**, 5893 (2000).
[4] W. K. Liu, E. G. Karpov and H. S. Park, personal communication.
[5] V. B. Shenoy, R. Miller, E. B. Tadmor, R. Phillips and M. Ortiz, Phys. Rev. Lett. **80**, 742 (1998).
[6] S. Curtaloro and G. Ceder, Phys. Rev. Lett. **88** 255504 (2002); L. M. Dupuy, E. B. Tadmor, R. E. Miller and R. Phillips, submitted to Phys. Rev. Lett.
[7] R. LeSar, R. Najafabadi and D. J. Srolovitz, Phys. Rev. Lett. **63**, 624 (1989).
[8] S. Nosé, Molecular Physics **52**, 522 (1984).
[9] S. D. Bond, B. J. Leimkuhler and B. B. Laird, J. Comp. Physics **151** 114 (1999).
[10] E. Hernandez, J. Chem. Phys. **115**, 10282 (2001)
[11] J. B. Sturgeon and B. B. Laird, J. Chem. Phys. **112** , 3474 (2000)
[12] Z. Jia and B. Leimkuhler, J. Multi. Model. Simu., to appear.
[13] B. J. Leimkuhler and C. R. Sweet, SIAM J. of Appl. Dynam. Systems **4**, 187-216 (2005).