Discontinuous molecular dynamics for semiflexible and rigid bodies

Lisandro Hernández de la Peña, Ramses van Zon, and Jeremy Schofield
Chemical Physics Theory Group, Department of Chemistry, University of Toronto, Ontario M5S 3H6, Canada

Sheldon B. Opps
Department of Physics, University of Prince Edward Island, 550 University Avenue, Charlottetown, Prince Edwards Island C1A 4P3, Canada

(Received 20 July 2006; accepted 28 December 2006; published online 20 February 2007)

A general framework for performing event-driven simulations of systems with semiflexible or rigid bodies interacting under impulsive forces is outlined. The method consists of specifying a means of computing the free evolution of constrained motion, evaluating the times at which interactions occur, and determining the consequences of interactions on subsequent motion. Algorithms for computing the times of interaction events and carrying out efficient event-driven simulations are discussed. The semiflexible case and the rigid case differ qualitatively in that the free motion of a rigid body can be computed analytically and need not be integrated numerically. © *2007 American Institute of Physics*. [DOI: 10.1063/1.2434957]

I. INTRODUCTION

There has been an increasing interest over the last decade in performing large-scale simulations of colloidal systems, proteins, micelles, and other biological assemblies. Simulating such systems, and the phenomena that take place in them, typically requires a description of dynamical events that occur over a wide range of time scales. Nearly all simulations of such systems to date are based on following the microscopic time evolution of the system by integration of the classical equations of motion. Usually, due to the complexity of inter- and intramolecular interactions, this integration is carried out in a step-by-step numerical fashion producing a time ordered set of phase-space points (a *trajectory*). This information can then be used to calculate thermodynamic properties, structural functions, or transport coefficients. An alternative approach, which has been employed in many contexts, is to use step potentials to approximate inter- and intramolecular interactions while affording the analytical solution of the dynamics.[1–5] The simplification in the interaction potential can lead to an increase in simulation efficiency since the demanding task of calculating forces is reduced to computing momentum exchanges between bodies at the instant of interaction. This approach is called event-driven or *discontinuous molecular dynamics* (DMD).

In the DMD approach, various components of the system interact via discontinuous interactions, leading to impulsive forces that act at specific moments of time. As a result, the motion of particles is free of inter- and intramolecular forces between impulsive *events* that alter the trajectory of bodies via discontinuous jumps in the momenta of the system at discrete interaction times. To determine the dynamics, the basic interaction rules of how the (linear and angular) momenta of the body are modified by inter- and intramolecular collisions must be specified.

For molecular systems with internal degrees of freedom, it is straightforward to design fully flexible models with discontinuous potentials, but DMD simulations of such systems are often inefficient due to the relatively high frequency of internal motions.[6] This inefficiency is reflected by the fact that most collision events executed in a DMD simulation correspond to intra rather than intermolecular interactions. On the other hand, much of the physics relevant in large-scale simulations is insensitive to details of intramolecular motion at long times. For this reason, methods of incorporating constraints into the dynamics of systems with continuous potentials have been developed that eliminate high frequency internal motion, and thus extend the time scales accessible to simulation. Surprisingly, relatively little work has appeared in the literature on incorporating such constraints into DMD simulations. The goals of this paper are to extend the applicability of DMD methods to include constrained systems and to outline efficient methods that are generally applicable in the simulations of semiflexible and rigid bodies interacting via discontinuous potentials.

In contrast to systems containing only simple spherical particles with single interaction sites,[4–8] the application of DMD methods to constrained systems is complicated by two main challenges. The first challenge is to solve the dynamics of the system so that the position, velocity, or angular velocity of any part of the system can be obtained exactly. In principle, this is possible for a rigid body moving in the absence of forces and torques, even if it does not possess an axis of symmetry which facilitates its motion. However, an explicit solution suitable for numerical implementation seems to be missing in the literature (although partial answers are abundant[9–14]). For this reason, we will present the explicit solution here. Armed with a solution of the dynamics of all bodies in the system, one can calculate the collision times in an efficient manner, and in some instances, analytically.

The second challenge is to determine how the impulsive forces lead to discontinuous jumps in the momenta of the interacting bodies. For complicated rigid or semiflexible

bodies, the rules for computing the momentum jumps are not immediately obvious. It is clear, however, that these jumps in momenta must be consistent with basic conservation laws connected to symmetries of the underlying Lagrangian characterizing the dynamics. Often the basic Lagrangian is invariant to time and space translations, as well as rotations. As a consequence of these symmetries, the rules governing collisions must explicitly obey energy, momentum, and angular momentum constraints. Such conservation laws can be utilized as a guide to derive the proper collision rules.

A first attempt to introduce constraints into an event-driven system was carried out by Ciccotti and Kalibaeva,[15] who studied a system of rigid, diatomic molecules (mimicking liquid nitrogen). Furthermore, nonspherical bodies of a special kind were treated by Donev *et al.*[16,17] by assuming that all rotational motions in between interaction events were that of a spherically symmetric body. More recently, a spherically symmetric hard-sphere model with four tetrahedral short ranged (sticky) interactions (mimicking water) has been studied by De Michele *et al.*[18] with an event-driven molecular dynamics simulation method similar to the most basic scheme presented in this paper. This work primarily focuses on the phase diagram of this "sticky" water model as a prototype of network forming molecular systems. Our purpose, in contrast, is to discuss a general framework that allows one to carry out event-driven DMD simulations in the presence of constraints and, in particular, for general rigid bodies. The methodology is applicable to modeling the correct dynamics of water molecules in aqueous solutions,[19] as well as other many body systems.

The paper is organized as follows: Section II discusses the type of systems considered in this work followed by an analysis of the equations of motions in the presence of constraints in Sec. III. Section IV discusses the calculation and scheduling of collision times. The collision rules are derived in Sec. V. Finally, conclusions are presented in Sec. VI.

II. CONSTRAINED SYSTEMS INTERACTING VIA DISCONTINUOUS POTENTIALS

In this work we focus on molecular systems in which some of the degrees of freedom are constrained under a set of c time-independent holonomic constraints,

$$\sigma_\alpha(\boldsymbol{r}_N) = 0, \tag{1}$$

where the index α runs over all constraints present in the system and $\boldsymbol{r}_N$ is a generalized vector whose components are the set of all Cartesian coordinates of the N total particles in the system. Of particular interest are systems in which the constraints involve clusters (molecules) of particles (atoms or sites). Typically, a constraint equation fixes the distance between sites i and j in each molecule to some value d according to

$$\sigma(r_{ij}) = \tfrac{1}{2}(r_{ij}^2 - d^2) = 0. \tag{2}$$

In general, there are not enough constraints to completely fix the relative positions of all the sites in a molecule; such a molecule is called semiflexible. In later sections, particular attention will be devoted to the special case of rigid body systems, in which *all* intramolecular distances remain fixed.

All interactions in the system are assumed to be expressible in terms of pairwise site-site discontinuous potentials, where a typical interaction between sites i and j can be written in the form

$$\Phi(|\boldsymbol{r}_i - \boldsymbol{r}_j|) = \begin{cases} \Phi_0 & \text{if } |\boldsymbol{r}_i - \boldsymbol{r}_j| \leqslant d_{ij} \\ \Phi_1 & \text{if } |\boldsymbol{r}_i - \boldsymbol{r}_j| > d_{ij}, \end{cases} \tag{3}$$

where d_{ij} is an interaction distance and Φ_0 and Φ_1 are constants. There can be several of these kind of interactions per pair of sites. Although typically the sites i and j are on different molecules, there may be intramolecular site-site interactions for semiflexible molecules. If all interactions are of this form, then the total system evolves freely until the distance between a pair of interaction sites attains the interaction distance.

The simplest example of this kind of system consists of two hard spheres of diameter d located at positions $\boldsymbol{r}_j$ and $\boldsymbol{r}_i$ in a periodic box. If the two spheres are approaching one another, the spheres evolve freely until the time of the system approaches a value t_c, where the distance between the two centers of the hard spheres is equal to the diameter d. At this time, the spheres elastically bounce off one another in a *hard-core collision* in which the momenta of the spheres are altered in a discontinuous fashion. Immediately after time t_c, the spheres once again evolve freely until the next interaction time t_c', as determined by the distance criterion.

The principle of performing event-driven simulations of constrained molecular systems is the same as that of simulating hard-sphere systems. As the system has constant potential energy between interaction events, the motion of the constrained system is free and the trajectory of each molecule is independent of all others. The free propagation of the system between events determines the evolution of the spatial coordinates of the molecules in the system, and the interaction times at which the momenta of the system change discontinuously are determined by identifying the times at which the distance between sites corresponds to a discontinuity in the interaction potential. Even though this interaction time is determined by the free motion of the system, in general, it must be found numerically due to the mathematical complexity of the collision condition. The final ingredient required to perform event-driven simulations of constrained systems consists of specifying the collision rules of how the momenta of the constituents in the molecular system are altered at the interaction times.

In the following sections, we address each of these points individually, paying particular attention to the special case of rigid molecular systems for which a closed-form solution of the free motion is available.

III. EQUATIONS OF MOTION WITH CONSTRAINTS

Typically, the free motion of constrained systems is not exactly solvable and must be treated numerically. An exceptional case is the motion of a rigid body in the absence of forces. Nonetheless, it is possible to develop a systematic scheme for performing DMD simulations based on numerical trajectories of constrained systems. Since the trajectories

of all molecules are independent in between interaction events, it is sufficient to consider the motion of a single molecule. In the following two subsections, the general equations governing the free motion of semiflexible and rigid bodies are presented.

A. Semiflexible dynamics

The equations of motion for a constrained body consisting of sites i with mass m_i and position $\boldsymbol{r}_i$ are[20]

$$m_i \ddot{\boldsymbol{r}}_i = -\lambda_\alpha \frac{\partial \sigma_\alpha}{\partial \boldsymbol{r}_i}, \tag{4}$$

where the parameters λ_α are Lagrange multipliers that enforce the distance constraints σ_α. Here, as well as below, explicit dependence on the time t will not be denoted. For clarity throughout this paper, the Einstein summation convention will be used for sums over repeated *Greek* indices, i.e., $\lambda_\alpha \sigma_\alpha \equiv \Sigma_{\alpha=1}^{c} \lambda_\alpha \sigma_\alpha$, whereas the sum over site indices will be written explicitly.

Note that since intramolecular interactions are treated by discontinuous potentials, they are not evident in Eq. (4), but rather are treated as collision events (see Secs. IV and V). While it is, in principle, possible to introduce smooth intramolecular potentials in the presence of discontinuous intermolecular potentials, such systems will not be considered here. Furthermore, external forces will also not be taken into account in this paper.

These equations of motion must be supplemented by equations for the c Lagrange multipliers λ_α, which are functions of time and follow from the requirement that $\sigma_\alpha = 0$. Although the λ_α are not functions of $\boldsymbol{r}_N$ and $\dot{\boldsymbol{r}}_N$ in a mathematical sense, it will be shown below that once the equations are solved, they can be expressed in terms of $\boldsymbol{r}_N$ and $\dot{\boldsymbol{r}}_N$. Note that the equations of motion show that even in the absence of an interaction potential, the motion of the point masses (atoms) making up a rigid body (molecule) are nontrivial due to the emergence of a *constraint force* $-\lambda_\alpha \partial \sigma_\alpha / \partial \boldsymbol{r}_i$.

In fortuitous cases, the time dependence of the Lagrange multipliers is relatively simple and can be solved for by Taylor expansion of the Lagrange multipliers in time t. To evaluate the time derivatives of the multipliers, one can use time derivatives of the initial constraint conditions, which must vanish to all orders. The result is a hierarchy of equations, which, at order k, is linear in the unknown kth time derivatives $\lambda_\alpha^{(k)}$ but depends on the lower order time derivatives $\lambda_\alpha^{(0)}, \lambda_\alpha^{(1)}, \ldots, \lambda_\alpha^{(k-1)}$. In exceptional circumstances, this hierarchy naturally truncates. For example, for a rigid diatomic molecule with a single bond-length constraint, one finds that the hierarchy truncates at order $k=0$, and the Lagrange multiplier is a constant.[15] However, this is not the typical case.

Alternatively, since the constraints $\sigma_\alpha(\boldsymbol{r}_N) = 0$ are to be satisfied at all times t, and not just at time zero, their time derivatives are zero at all times. From the first time derivative $\dot{\sigma}_\alpha(\boldsymbol{r}_N) = 0$, one sees that the initial velocities $\boldsymbol{v}_i = \dot{\boldsymbol{r}}_i$ must obey

$$\sum_i \boldsymbol{v}_i \cdot \frac{\partial \sigma_\alpha}{\partial \boldsymbol{r}_i} = 0, \tag{5}$$

for each constraint condition α. The Lagrange multipliers can be determined by the condition that the second derivatives of all the constraints vanish, which, using Eq. (4), yields[21]

$$\lambda_\alpha = \mathbf{Z}_{\alpha\beta}^{-1}(\boldsymbol{r}_N) \mathcal{T}_\beta(\boldsymbol{r}_N, \dot{\boldsymbol{r}}_N), \tag{6}$$

where

$$\mathcal{T}_\beta(\boldsymbol{r}_N, \dot{\boldsymbol{r}}_N) = \sum_{i,j} \dot{\boldsymbol{r}}_j \cdot \frac{\partial^2 \sigma_\beta(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_j \partial \boldsymbol{r}_i} \cdot \dot{\boldsymbol{r}}_i, \tag{7}$$

$$\mathbf{Z}_{\alpha\beta}(\boldsymbol{r}_N) = \sum_i \frac{1}{m_i} \frac{\partial \sigma_\alpha(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_i} \cdot \frac{\partial \sigma_\beta(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_i}. \tag{8}$$

As Eq. (6) shows, in general, the Lagrange multipliers are dependent on both the positions $\boldsymbol{r}_N$ and the velocities $\dot{\boldsymbol{r}}_N$ of the particles. To see that this makes the dynamics non-Hamiltonian, the equations of motion can be cast into Hamiltonian-type form using $\boldsymbol{p}_i = m_i \dot{\boldsymbol{r}}_i$, i.e.,

$$\dot{\boldsymbol{r}}_i = \frac{\boldsymbol{p}_i}{m_i}, \quad \dot{\boldsymbol{p}}_i = -\lambda_\alpha \frac{\partial \sigma_\alpha}{\partial \boldsymbol{r}_i}, \tag{9}$$

where it is apparent that the forces in the system depend on the momenta through λ_α in Eq. (6). There exists no Hamiltonian that generates these equations of motion.[22]

Since the underlying dynamics of the system is non-Hamiltonian, the statistical mechanics of the constrained system is potentially more complex. In general, phase-space averages have to be defined with respect to a metric that is invariant to the standard measure of Hamiltonian systems, but $d\boldsymbol{r}_N d\boldsymbol{p}_N$ is not conserved under the dynamics and the standard form of the Liouville equation does not hold.[23,24] In general, there is a phase-space compressibility factor κ associated with the lack of conservation of the measure that is given by the negative of the divergence of the flow in phase space. It may be shown that[24,25]

$$\kappa = -\frac{\partial}{\partial \boldsymbol{r}_i} \dot{\boldsymbol{r}}_i - \frac{\partial}{\partial \boldsymbol{p}_i} \dot{\boldsymbol{p}}_i = \frac{d}{dt} \ln \| \mathbf{Z}(\boldsymbol{r}_N) \|,$$

where $\| \mathbf{Z}(\boldsymbol{r}_N) \|$ is the determinant of the matrix $\mathbf{Z}_{\alpha\beta}(\boldsymbol{r}_N)$ defined in Eq. (8). The compressibility factor is related to the invariant phase-space metric $d\mu = \| \mathbf{Z}(\boldsymbol{r}_N) \| d\boldsymbol{r}_N d\boldsymbol{p}_N$.[25,26] Statistical averages are therefore defined for the non- Hamiltonian system as[27,28]

$$\langle X(\boldsymbol{r}_N, \boldsymbol{p}_N) \rangle = \frac{1}{Q} \int d\boldsymbol{p}_N d\boldsymbol{r}_N \| \mathbf{Z}(\boldsymbol{r}_N) \| X(\boldsymbol{r}_N, \boldsymbol{p}_N) \rho(\boldsymbol{r}_N, \boldsymbol{p}_N) \times \prod_\alpha \delta(\sigma_\alpha(\boldsymbol{r}_N)) \delta(\dot{\sigma}_\alpha(\boldsymbol{r}_N, \boldsymbol{p}_N)), \tag{10}$$

where $\rho(\boldsymbol{r}_N, \boldsymbol{p}_N)$ is the probability density for the unconstrained system and Q is the partition function for the constrained system, given by

$$Q = \int d\boldsymbol{p}_N d\boldsymbol{r}_N \|\mathbf{Z}(\boldsymbol{r}_N)\| \rho(\boldsymbol{r}_N, \boldsymbol{p}_N) \times \prod_\alpha \delta(\sigma_\alpha(\boldsymbol{r}_N)) \delta(\dot{\sigma}_\alpha(\boldsymbol{r}_N, \boldsymbol{p}_N)).$$

Although the invariant metric is nonuniform for many constrained systems, for the special case of entirely rigid systems the $\mathbf{Z}_{\alpha\beta}(\boldsymbol{r}_N)$ matrix is a function only of the point masses and fixed distances. Hence, the term $\|\mathbf{Z}(\boldsymbol{r}_N)\|$ acts as a multiplicative factor, which cancels in the averaging process.

Although the solution of the dynamics of constrained systems via time-independent holonomic constraints is intellectually appealing and useful for developing a formal statistical mechanics for these systems, it is often difficult to analytically solve for the values of the Lagrange multipliers at arbitrary times. One therefore often resorts to numerical solutions of the multipliers in iterative form, using methods such as SHAKE.[29]

Regardless of the method by which the Lagrange multipliers are obtained, in general the free motion of a semiflexible molecule has to be integrated numerically in a step-by-step fashion. Although such an approach is not really consistent with the principles of DMD, it is important to note that the only forces required in the integration of the free motion of a molecule are the constraint forces, which depend only on the coordinates of the molecule itself. Consequently, the biggest bottleneck in standard molecular dynamics, the calculation of inter- and intramolecular forces, is not present here and the numerical computation of the free motion of each molecule can be done independently. This makes the integration of the equations of motion very different and much faster than in standard molecular dynamics simulations. As a result, most of the computational burden will be associated with determining and performing collision events, and the simulation can still rightfully be called event driven.

B. Free motion of rigid bodies

While the above formalism is also applicable to fully constrained, rigid bodies, it is possible to circumvent the constrained Lagrangian formalism by applying other approaches to calculate the evolution of the system analytically in the absence of external forces. The basic simplification in the dynamics of rigid bodies results from the fact that the general motion of a rigid body can be decomposed into a translation of the center of mass of the body plus a rotation about the center of mass. The orientation of the body relative to its center of mass is described by the relation between the so-called body frame, in which a set of axes are held fixed within the body as it moves, and the fixed external laboratory frame. The two frames of reference can be connected by an orthogonal transformation, such that the position of an atom i in a rigid body can be written as

$$\boldsymbol{r}_i = \boldsymbol{R} + \mathbf{A}^\dagger \tilde{\boldsymbol{r}}_i, \tag{11}$$

where $\tilde{\boldsymbol{r}}_i$ is the position of atom i in the body frame (which is independent of time), $\boldsymbol{R}$ is the center of mass, and the matrix $\mathbf{A}^\dagger$ is the orthogonal (rotation) matrix that converts coordinates in the body frame to the laboratory frame. The matrix $\mathbf{A}^\dagger$ is the transpose of $\mathbf{A}$, which is called the attitude matrix and which converts coordinates from the laboratory frame to the body frame. The elements of the rows of the attitude matrix $\mathbf{A}$ are simply the coordinates of the principal axes of the body written in the laboratory frame. Note that Eq. (11) implies that the relative vector $\bar{\boldsymbol{r}}_i$ satisfies

$$\bar{\boldsymbol{r}}_i = \boldsymbol{r}_i - \boldsymbol{R} = \mathbf{A}^\dagger \tilde{\boldsymbol{r}}_i. \tag{12}$$

One sees that in order to determine the location of different parts of the body in the laboratory frame, the attitude matrix $\mathbf{A}$ must be determined. Formally, any rotation matrix $\mathbf{U}$ (such as the attitude matrix) is a so-called special orthogonal matrix that can be specified by a rotation axis $\hat{\boldsymbol{n}} = (n_1, n_2, n_3)$ and an angle θ over which to rotate. Here $\hat{\boldsymbol{n}}$ is a unit vector, so that any nonunit vector $\theta\hat{\boldsymbol{n}}$ can be used to specify a rotation, where its norm is equal to the angle θ and its direction is equal to the axis $\hat{\boldsymbol{n}}$. According to Rodrigues formula, the matrix corresponding to this rotation is[20]

$$\mathbf{U}(\theta\hat{\boldsymbol{n}}) = \mathbf{I} + \mathbf{W}(\hat{\boldsymbol{n}}) \sin\theta + \mathbf{W}(\hat{\boldsymbol{n}})\mathbf{W}(\hat{\boldsymbol{n}})(1 - \cos\theta), \tag{13}$$

where $\mathbf{I}$ is the identity matrix and $\mathbf{W}$ is the skew-symmetric matrix

$$\mathbf{W}(\hat{\boldsymbol{n}}) = \begin{pmatrix} 0 & -n_3 & n_2 \\ n_3 & 0 & -n_1 \\ -n_2 & n_1 & 0 \end{pmatrix}. \tag{14}$$

From these definitions, it can be shown that the attitude matrix $\mathbf{A}$ satisfies the differential equation

$$\dot{\mathbf{A}} = -\mathbf{A}\mathbf{W}(\boldsymbol{\omega}), \tag{15}$$

where $\boldsymbol{\omega}$ is the angular velocity in the laboratory frame.

The equations of motion of the angular velocity are most simply written down by transforming them to the principal axis, or body, frame, i.e., $\tilde{\boldsymbol{\omega}} = \mathbf{A}\boldsymbol{\omega}$. In this frame, the angular velocities obey the well-known Euler equations (here in the absence of forces),

$$\begin{aligned} I_1\dot{\tilde{\omega}}_1 - \tilde{\omega}_2\tilde{\omega}_3(I_2 - I_3) &= 0, \\ I_2\dot{\tilde{\omega}}_2 - \tilde{\omega}_1\tilde{\omega}_3(I_3 - I_1) &= 0, \\ I_3\dot{\tilde{\omega}}_3 - \tilde{\omega}_1\tilde{\omega}_2(I_1 - I_2) &= 0, \end{aligned} \tag{16}$$

where I_1, I_2, and I_3 are the principal components of the moment of inertia tensor. In the principal axis frame, the representation of the components of the angular momentum $\tilde{\boldsymbol{L}} = (\tilde{L}_1, \tilde{L}_2, \tilde{L}_3)$ is particularly simple,

$$\tilde{\boldsymbol{L}} = \tilde{\mathbf{I}}\tilde{\boldsymbol{\omega}} = \begin{pmatrix} I_1\tilde{\omega}_1 \\ I_2\tilde{\omega}_2 \\ I_3\tilde{\omega}_3 \end{pmatrix}. \tag{17}$$

Using some of the properties of rotation matrices, Eq. (15) can be rewritten to relate the time evolution of $\mathbf{A}$ to the angular velocities in the body frame,

$$\dot{\mathbf{A}} = -\mathbf{W}(\tilde{\boldsymbol{\omega}})\mathbf{A}. \tag{18}$$

Thus, once the angular velocity $\tilde{\boldsymbol{\omega}}$ is known, it can be substituted into Eq. (18) to solve for the matrix $\mathbf{A}$. The general solution of Eq. (18) is of the form

$$\mathbf{A} = \mathbf{P}\mathbf{A}(0). \tag{19}$$

Note that the matrices $\mathbf{A}$ and $\mathbf{P}$ depend on t, while $\mathbf{A}(0)$ is the attitude matrix at time zero. Thus, $\mathbf{P}$ is a rotation matrix which "propagates" the orientation $\mathbf{A}(0)$ to the orientation at time t. $\mathbf{P}$ satisfies the same Eq. (18) as $\mathbf{A}$, but with initial condition $\mathbf{P}(0)=\mathbf{I}$. By integrating this equation, one can obtain an expression for $\mathbf{P}$. At first glance, it may seem that $\mathbf{P}$ can only be written as a formal expression containing a time-ordered exponential. However, in the absence of forces, the conservation of angular momentum and energy and the orthogonality of the matrix $\mathbf{P}$ can be used to derive the following explicit expression[30] (implicitly also found in Ref. 14):

$$\mathbf{P} = \mathbf{T}_1\mathbf{T}_2, \tag{20}$$

where $\mathbf{T}_1$ and $\mathbf{T}_2$ are two rotation matrices. The matrix $\mathbf{T}_1$ rotates $\tilde{\boldsymbol{L}}(0)$ to $\tilde{\boldsymbol{L}}$ and can be written as

$$\mathbf{T}_1 = \begin{pmatrix} c_1c_2 - s_1s_2c_3 & c_1s_2 + s_1c_2c_3 & s_1s_3 \\ -s_1c_2 - c_1s_2c_3 & -s_1s_2 + c_1c_2c_3 & c_1s_3 \\ s_2s_3 & -c_2s_3 & c_3 \end{pmatrix}, \tag{21}$$

where

$$s_1 = \frac{\tilde{L}_1}{\tilde{L}_\perp}, \quad c_1 = \frac{\tilde{L}_2}{\tilde{L}_\perp}, \tag{22}$$

$$s_2 = -\frac{\tilde{L}_1(0)}{\tilde{L}_\perp(0)}, \quad c_2 = \frac{\tilde{L}_2(0)}{\tilde{L}_\perp(0)}, \tag{23}$$

$$s_3 = \frac{\tilde{L}_\perp\tilde{L}_3(0) - \tilde{L}_3\tilde{L}_\perp(0)}{L^2}, \quad c_3 = \frac{\tilde{L}_\perp\tilde{L}_\perp(0) + \tilde{L}_3\tilde{L}_3(0)}{L^2}, \tag{24}$$

and $\tilde{L}_\perp = \sqrt{\tilde{L}_1^2 + \tilde{L}_2^2}$ while $L = |\tilde{\boldsymbol{L}}|$.

Using the notation in Eq. (13), the matrix $\mathbf{T}_2$ can be expressed as

$$\mathbf{T}_2 = \mathbf{U}(-\psi L^{-1}\tilde{\boldsymbol{L}}(0)), \tag{25}$$

where the angle ψ is given by

$$\psi = \int_0^t dt'\,\Omega(t'), \tag{26}$$

with

$$\Omega = L\frac{I_1\tilde{\omega}_1^2 + I_2\tilde{\omega}_2^2}{\tilde{L}_\perp^2}. \tag{27}$$

The angle ψ can be interpreted as an angle over which the body rotates. If the body rotates one way, the laboratory frame as seen from the body frame rotates in the opposite way, which explains the minus sign in Eq. (25). For the derivation of Eqs. (20)–(27), we refer to Ref. 30. Similar equations, but in a special reference frame, can be found in Ref. 14.

In the following, the solution of Eqs. (16) and (18) for bodies of differing degrees of symmetry will be analyzed.

1. Spherical rotor

For the case of a spherical rotor in which all three moments of inertia are equal, i.e., $I_1=I_2=I_3$, the form of the Euler equations (16) is particularly simple: $I_1\dot{\tilde{\omega}}_j=0$. Therefore, all components of the angular velocity in the body frame are conserved, as are those of the angular momentum. As a result, $\mathbf{T}_1$ in Eq. (21) is equal to the identity matrix. A second consequence is that Ω in Eq. (27) is constant, so that $\psi=\Omega t$, where Ω may be rewritten, using $I_1=I_2=I_3$, as $\Omega = L/I_1 = |\boldsymbol{\omega}|$. Therefore Eqs. (20) and (25) give

$$\mathbf{P} = \mathbf{U}(-\boldsymbol{\omega}t), \tag{28}$$

corresponding to a rotation by an angle of $-\Omega t$ around the axis $\boldsymbol{\omega}/\Omega$.

2. Symmetric top

For the case of a symmetric top for which $I_1=I_2$, one can solve the Euler equations (16) in terms of simple sines and cosines to obtain

$$\tilde{\boldsymbol{\omega}} = \begin{pmatrix} \tilde{\omega}_1(0)\cos\omega_p t + \tilde{\omega}_2(0)\sin\omega_p t \\ -\tilde{\omega}_1(0)\sin\omega_p t + \tilde{\omega}_2(0)\cos\omega_p t \\ \tilde{\omega}_3(0) \end{pmatrix}, \tag{29}$$

where $\omega_p=(1-(I_3/I_1))\tilde{\omega}_3(0)$ is the precession frequency. Using Eq. (29) and the fact that $\tilde{L}_\perp$ and $\tilde{L}_3$ are conserved in this case, one can easily show that $\mathbf{T}_1$ is given by

$$\mathbf{T}_1 = \mathbf{U}(-\omega_p t\hat{z}), \tag{30}$$

and Ω is given by

$$\Omega = \frac{L[I_1\tilde{\omega}_1^2(0) + I_1\tilde{\omega}_2^2(0)]}{I_1^2\tilde{\omega}_1^2(0) + I_1^2\tilde{\omega}_2^2(0)} = \frac{L}{I_1}. \tag{31}$$

As this Ω is constant, the rotation angle is simply $\psi = (L/I_1)t$, and hence

$$\mathbf{T}_2 = \mathbf{U}\left(-\frac{\tilde{\boldsymbol{L}}(0)t}{I_1}\right). \tag{32}$$

For the symmetric top, the propagation matrix $\mathbf{P}$ is therefore

$$\mathbf{P} = \mathbf{U}(-\omega_p t\hat{z})\mathbf{U}\left(-\frac{\tilde{\boldsymbol{L}}(0)t}{I_1}\right). \tag{33}$$

3. Asymmetric body

If all the principal moments of inertia are distinct, the time dependence of the angular velocity $\tilde{\boldsymbol{\omega}}$ involves Jacobi elliptic functions.[10,30] While this may seem complicated, efficient standard numerical routines exist to evaluate these functions.[31–35] More challenging is the evaluation of the matrix $\mathbf{P}$. While its exact solution has been known for more than 170 years,[9,10] it is formulated—even in more recent texts[11,12]—in terms of undetermined constants and using

complex algebra, which hinders a straightforward implementation of the free motion in a numerical simulation. It is surprisingly difficult to find an explicit formula in the literature for the matrix **P** as a function of the initial conditions, which is the form needed in DMD simulations. For this reason, the explicit general solution for **P**, derived in detail in Ref. 30, will briefly be presented here in terms of general initial conditions.

Following Jacobi,[10] it is useful to adopt the convention that I_2 is the moment of inertia which is intermediate in magnitude (i.e., either $I_1<I_2<I_3$ or $I_1>I_2>I_3$), and one chooses the overall ordering of magnitudes such that

$$I_1 > I_2 > I_3 \quad \text{if } E_R > \frac{L^2}{2I_2},$$
$$I_1 < I_2 < I_3 \quad \text{if } E_R < \frac{L^2}{2I_2}, \tag{34}$$

where E_R is the rotational kinetic energy $E_R=\frac{1}{2}(I_1\tilde{\omega}_1^2+\tilde{\omega}_2^2+I_3\tilde{\omega}_3^2)$ and L is the norm of the angular momentum $L=(I_1^2\tilde{\omega}_1^2+I_2^2\tilde{\omega}_2^2+I_3^2\tilde{\omega}_3^2)^{1/2}$. Without this convention some quantities defined below would be complex valued, which is numerically inconvenient and inefficient. Note that in a simulation, molecules will often be assigned a specific set of physical inertial moments with fixed order, i.e., not depending on the particular values of E_R and L. Nevertheless, a simple way to adopt the convention in Eq. (34) is to introduce internal variables $\tilde{\boldsymbol{\omega}}_{\text{int}}=\mathbf{U}^*\tilde{\boldsymbol{\omega}}$, $\tilde{\mathbf{I}}_{\text{int}}=\mathbf{U}^*\tilde{\mathbf{I}}\mathbf{U}^*$, and $\mathbf{A}_{\text{int}}=\mathbf{U}^*\tilde{\mathbf{A}}\mathbf{U}^*$, which differ when necessary from the physical ones by a rotation given by the rotation matrix

$$\mathbf{U}^* = \begin{pmatrix} 0 & 0 & 1 \\ 0 & -1 & 0 \\ 1 & 0 & 0 \end{pmatrix}. \tag{35}$$

This matrix interchanges the x and z directions and reverses the y direction.

The Euler equations (16) can be solved to obtain[10,11,13]

$$\tilde{\boldsymbol{\omega}} = \begin{pmatrix} \omega_{1m}\,\mathrm{cn}(\omega_p t+\varepsilon|m) \\ \omega_{2m}\,\mathrm{sn}(\omega_p t+\varepsilon|m) \\ \omega_{3m}\,\mathrm{dn}(\omega_p t+\varepsilon|m) \end{pmatrix}, \tag{36}$$

where ε is an integration constant. Here sn, cn, and dn are Jacobi elliptic functions,[31,32,36] while the ω_{im} are the extreme (maximum or minimum) values of the ω_i and are given by

$$\omega_{1m} = \mathrm{sgn}(\tilde{\omega}_1(0))\sqrt{\frac{L^2-2I_3E_R}{I_1(I_1-I_3)}},$$
$$\omega_{2m} = -\,\mathrm{sgn}(\tilde{\omega}_1(0))\sqrt{\frac{L^2-2I_3E_R}{I_2(I_2-I_3)}}, \tag{37}$$
$$\omega_{3m} = \mathrm{sgn}(\tilde{\omega}_3(0))\sqrt{\frac{L^2-2I_1E_R}{I_3(I_3-I_1)}},$$

where $\mathrm{sgn}(x)$ is the sign of x. In Eq. (36), the *precession frequency* ω_p is given by

$$\omega_p = \mathrm{sgn}(I_2-I_3)\mathrm{sgn}(\tilde{\omega}_3(0))\sqrt{\frac{(L^2-2I_1E_R)(I_3-I_2)}{I_1I_2I_3}}. \tag{38}$$

The elliptic functions are periodic functions of their first argument and look very similar to the sine, cosine, and constant functions. They depend on the *elliptic parameter* m (or elliptic modulus $\sqrt{m}$), which determines how closely the elliptic functions resemble their trigonometric counterparts, and which is given by

$$m = \frac{(I_1-I_2)(L^2-2I_3E_R)}{(I_3-I_2)(L^2-2I_1E_R)}. \tag{39}$$

The integration constant ε in Eq. (36) is given by

$$\varepsilon = F(\tilde{\omega}_{20}/\tilde{\omega}_{2m}|m), \tag{40}$$

where F is the incomplete elliptic integral of the first kind[31,32]

$$F(y|m) = \int_0^y \frac{dx}{\sqrt{(1-x^2)(1-mx^2)}}. \tag{41}$$

In fact, $\mathrm{sn}(x|m)$ is simply the inverse of this function. As a result of the ordering convention in Eq. (34), the parameter m in Eq. (39) is guaranteed to be less than 1, which is required in order that $F(y|m)$ in Eq. (41) is not complex valued.

Three more numbers can be derived from the elliptic parameter m which play an important role in determining the properties of elliptic functions. These are the *quarter period* $K=F(1|m)$, the *complementary quarter period* $K'=F(1|1-m)$, and the *nome* $q=\exp(-\pi K'/K)$, which is the parameter in various series expansions.

With the solutions of the Euler equations in hand, we now turn to the solution of Eq. (18) as given by Eqs. (19)–(27). The expression for Ω in Eq. (27) is not a constant in this case but is a periodic function involving elliptic functions. Despite this difficulty, the integral can still be performed using some properties of elliptic functions, with the result[30]

$$\psi = A_1 + A_2 t - \phi. \tag{42}$$

The constants A_1, A_2 and the periodic function ϕ can be expressed using the theta function $H(u|m)$ (Refs. 10, 31, and 32) as

$$\phi = \arg H(\omega_p t+\varepsilon-i\eta|m), \tag{43}$$

$$A_1 = \phi(0) = \arg H(\varepsilon-i\eta|m), \tag{44}$$

$$A_2 = \frac{L}{I_1} + \omega_p \frac{d\log H(i\eta|m)}{d\eta}, \tag{45}$$

where we have used the definition

$$\eta = \mathrm{sgn}(\tilde{\omega}_{30})K' - F\left(\frac{I_3\omega_{3m}}{L}\middle|1-m\right). \tag{46}$$

Equations (43)–(45) involve complex values which are not convenient for numerical evaluation. Using the known

series expansions of the theta function H and its logarithmic derivative[31,32] in terms of the nome q, these equations may be rewritten in a purely real form. In fact, one readily obtains the sine and cosine of ψ, which are all that is needed in Eqs. (13) and (25),

$$\cos\psi = \frac{h_r \cos(A_1 + A_2 t) + h_i \sin(A_1 + A_2 t)}{\sqrt{h_r^2 + h_i^2}},$$
$$\sin\psi = \frac{h_r \sin(A_1 + A_2 t) - h_i \cos(A_1 + A_2 t)}{\sqrt{h_r^2 + h_i^2}}, \tag{47}$$

where h_r and h_i are the real and imaginary parts of the theta function H and are given by

$$h_r = 2q^{1/4}\sum_{n=0}^{\infty}(-1)^n q^{n(n+1)} \cosh\frac{(2n+1)\pi\eta}{2K} \times \sin\frac{(2n+1)\pi(\omega_p t + \varepsilon)}{2K}, \tag{48}$$

$$h_i = -2q^{1/4}\sum_{n=0}^{\infty}(-1)^n q^{n(n+1)} \sinh\frac{(2n+1)\pi\eta}{2K} \times \cos\frac{(2n+1)\pi(\omega_p t + \varepsilon)}{2K}, \tag{49}$$

while the constant A_1 is

$$A_1 = \arctan[h_i(0)/h_r(0)] + n\pi, \tag{50}$$

where $n=0$ if $h_r(0)>0$, $n=1$ if $h_r(0)<0$ and $h_i(0)>0$, and $n=-1$ if $h_r(0)<0$ and $h_i(0)<0$. Finally, the constant A_2 is given by[31,32]

$$A_2 = \frac{L}{I_1} + \frac{\pi\omega_p}{2K}\left[\frac{\xi+1}{\xi-1} - 2\sum_{n=1}^{\infty}\frac{q^{2n}(\xi^n - \xi^{-n})}{1-q^{2n}}\right], \tag{51}$$

where $\xi=\exp(\pi\eta/K)$. The series expansion in q in Eq. (51) converges for $\xi q^2<1$. Because $-K'<\eta<K'$ [cf. Eq. (46)], one has $\xi q^2 \leqslant q<1$, and the series always converges. Since q is typically small, the convergence is rarely very slow [e.g., for convergence up to relative order δ, one needs $O(\log\delta/\log q)$ terms]. Note that since the constants A_1 and A_2 depend only on the initial angular velocities, they only need to be calculated once at the beginning of the motion of a free rigid body. On the other hand, the series expansions in Eqs. (48) and (49), which have to be evaluated any time the positions are desired, have extremely fast convergence due to the $q^{n(n+1)}$ appearing in these expressions [for example, unless $m \gtrsim 0.95$, the convergence of the series up to $O(10^{-15})$ occurs, taking only three terms].

There are efficient routines to calculate the functions cn, sn, dn, and F, see, e.g., Refs. 31–35, and the series in Eqs. (48), (49), and (51) converge; the former two quite rapidly in fact. Therefore, despite an apparent preference in the literature for conventional numerical integration of the equations of motion via many successive small time steps even for torque-free cases, the analytical solution can be used to calculate the same quantities in a computationally more efficient manner requiring only the evaluation of special functions.

IV. INTERACTION EVENTS

As mentioned above, in systems with discontinuous potentials, the particles evolve freely until two sites are an interaction distance apart. In this section the computation for this interaction distance will be shown, and in the next section, the consequences of the interactions will be determined.

The complicated free motion of constrained systems implies that evaluating the distance at arbitrary times between any interaction sites is nontrivial. In contrast to simulations of hard-sphere systems in which it is possible to analytically calculate the time at which a given pair of spheres interact, the interaction times for constrained molecular systems must be computed numerically even if an analytical solution of the free motion is available, as it is for rigid bodies. Nevertheless, a general method of determining the interaction times that does not depend on the details of the dynamics and is equally valid for semiflexible and rigid systems can be constructed.

A. Calculation of interaction event times

The identification of interaction times corresponds to finding the time at which the distance between two sites (on the same or on different molecules) corresponds to a discontinuity in the interaction potential. Assuming that the interactions between sites i and j depend only on the scalar distance r_{ij} between the sites, one can define a *collision indicator function* f_{ij},

$$f_{ij} = \tfrac{1}{2}(r_{ij}^2 - d_{ij}^2), \tag{52}$$

where r_{ij} is the radial distance between sites i and j, and d_{ij} is the distance at which a discontinuity appears in the potential for the given pair of sites. With this definition, an interaction time t_c is a zero of the collision indicator function, so that $f_{ij}(t_c)=0$, and hence the problem of finding the interaction times corresponds to finding the zeros of this typically oscillatory function.

For intermolecular interactions, the search for the earliest interaction or "collision" event time can be facilitated using screening strategies to decide when bodies may overlap.[37,38] Usually, these involve placing the bodies in bounding boxes and using an efficient method to determine when bounding boxes intersect. The simplest way to do this is to place each molecule in the smallest sphere around its center of mass containing all components of the molecule.[15] Note that for semiflexible molecules, this sphere should be large enough to accommodate the largest possible configuration of the molecule. The position of the sphere is determined by the motion of the center of mass, while any change in orientation (and shape in the semiflexible case) of the molecule occurs within the sphere. Collisions between molecules can therefore only occur when their encompassing spheres overlap, and the time at which this occurs can be calculated analytically for any pair of molecules. This time serves as a useful point to begin a more detailed search for

collision events (see below). Similarly, one can also calculate the time at which the spheres no longer overlap, and use these event times to bracket a possible root of the collision indicator function. It is crucial to make the time bracketing as tight as possible in any implementation of DMD with numerical root searches, because the length of the time bracketing interval determines the required number of evaluations of the positions and velocities of the atoms, and therefore, plays a significant role in the efficiency of the overall procedure.

For intermolecular interactions, the intersection of bounding spheres around molecules establishes the largest possible time bracket for which to perform a root search. For intramolecular interactions, such a time bracket is generally not available, but the technique of virtual collisions, as will be explained in the next section, can be used to reduce the root-search interval to a manageable size. To narrow down the root search, this total interval found by bracketing may be subdivided into a set of *grid points* separated by smaller time intervals of length Δt. The simplest, reliable, and reasonably efficient means of detecting a root is thus to perform a *grid search* that looks for changes in the sign of f_{ij} on this set of grid points in time, i.e., one looks at $f_{ij}(t+n\Delta t)$ and $f_{ij}(t+(n+1)\Delta t)$ over a set of integers n. When a time interval in between two grid points is found in which a sign change of f_{ij} occurs, the Newton-Raphson algorithm[33] can be called to numerically determine the root with arbitrary accuracy. Since the Newton-Raphson method requires the calculation of first time derivatives, one must also calculate, for any time t, the derivative $df_{ij}/dt=\boldsymbol{r}_{ij}\cdot\boldsymbol{v}_{ij}$, where the notations $\boldsymbol{r}_{ij}=\boldsymbol{r}_j-\boldsymbol{r}_i$ and $\boldsymbol{v}_{ij}=\boldsymbol{v}_j-\boldsymbol{v}_i$ have been used. Such time derivatives are generally readily evaluated.

Unfortunately, while the Newton-Raphson method is a very efficient algorithm for finding roots, it can be somewhat unstable when one is searching for the roots of an oscillatory function. For translating and rotating molecules, the collision indicator function is often oscillatory due to the periodic or near-periodic motion of the relative orientation of two colliding bodies. It is particularly easy to miss so-called grazing collisions when the grid search interval Δt is too large, in which case the indicator function is positive in two consecutive points of the grid search, yet nonetheless "dips" below zero within the grid interval. It is important that no roots are missed, for a missed root can lead to a different, and possibly, infinite energy. To reduce the frequency of missing grazing collisions to zero, a vanishingly small grid interval Δt would be required. Of course, such a scheme is not practical, and one must balance the likelihood of missing events with practical considerations since several collision indicator functions need to be evaluated at each point of the grid. Clearly, the efficiency of the root-search algorithm significantly depends on the magnitude of the grid interval.

To save computation time, a coarser grid can be utilized if a means of handling grazing collisions is implemented. Since the collision indicator function has a local extremum (maximum or minimum, depending on whether $|\boldsymbol{r}_i-\boldsymbol{r}_j|^2$ is initially smaller or larger than d^2) at some time near the time of a grazing collision, a reasonable strategy to find these kind of collision events is to determine the extremum of the indicator function in cases in which the indicator function f_{ij} itself does not change sign on the interval but its derivative df_{ij}/dt does. Furthermore, since the indicator function at the grid points near a grazing collision is typically small, it is fruitful to search only for extrema when the indicator function at one of the grid points lies below some threshold value. To find the local extrema of the indicator function, any simple routine of locating the extrema of a nonlinear function can be utilized. For example, Brent's minimization method,[33,39] which is based on a parabolic interpolation of the function, is a good choice for sufficiently smooth one-dimensional functions. Once the extremum is found, it is a simple matter to decide whether or not a real collision exists by checking the sign of f_{ij}.

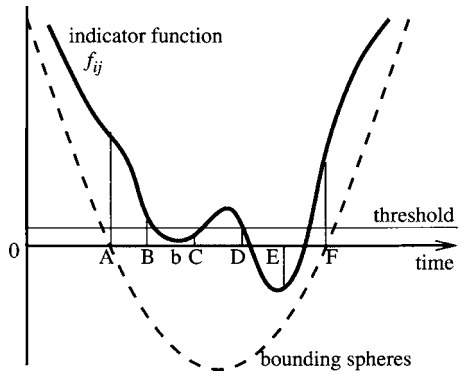

FIG. 1. Example of the root search. The thick line is f_{ij}, the dashed line is a lower bound, and the thin line is the threshold.

We remark that the threshold value in this procedure is found by trial and error; a trial simulation with a very small threshold value is run, which is sure to miss a collision at some point due to a grazing collision. The collision indicator function around the time of this grazing collision is inspected and the threshold value is modified such that this collision will not be missed. New trial simulations are run, and the threshold is adjusted until the frequency of missed grazing collisions is acceptably small.

Once the root has been bracketed (either through a sign change of f_{ij} during the grid search or after searching for an extremum), one can simply use the Newton-Raphson algorithm to find the root to desired accuracy, typically within only a few iterations. The time value returned by the Newton-Raphson routine needs to be in the bracketed interval, and $df_{ij}/dt<0$ if f_{ij} was initially positive and $df_{ij}/dt>0$ if it was initially negative. If those criteria are not satisfied, the Newton-Raphson algorithm has clearly failed and a less efficient but more reliable method is needed to track down the root. For example, the Van Wijngaarden–Dekker–Brent method,[33,39] which combines bisection and quadratic interpolation, is guaranteed to converge if the function is known to have a root in the interval under analysis.

To clarify how this root-search algorithm finds the earliest interaction time for a pair of sites i and j, consider the indicator function f_{ij} shown in Fig. 1. The smallest interaction time t_c, determined by $f_{ij}(t_c)=0$, is found as follows: (1) The intersection of bounding spheres around the molecules containing i and j gives a quadratic lower bound for f_{ij}. The roots of the quadratic are at times A and F, so that the total

search interval is [A, F]. (2) The total interval starting at A is divided into smaller intervals of length Δt. The next grid point in the total interval is at $B=A+\Delta t$, and the successive grid points are at C, D, and E. (3) $f_{ij}(A)$ and $f_{ij}(B)$ are computed. Both are positive and larger than the threshold value, so one concludes that no event takes place between A and B. (4) $f_{ij}(C)$ is computed. Although its value is also positive, it lies below the threshold, and the values of the derivatives at B and C are opposite in sign, so a minimization routine is called that returns the value of the minimum and the time b at which the minimum occurs. In this case, $f_{ij}(b)>0$, so no root exists between B and C. (5) $f_{ij}(D)$ is computed and is found to be positive. $f_{ij}(C)$ is below the threshold value, and df_{ij}/dt at C and D have opposite signs; however, the derivative at C is *positive* so there is no minimum and no root between C and D. (6) $f_{ij}(E)$ is computed and is found to be negative: Therefore t_c must lie between D and E, and is found by the Newton-Raphson method, or, if this fails, by the Van Wijngaarden–Dekker–Brent method.

B. Scheduling events

It was shown above how to determine the time t_{ij} at which two atoms (or sites) i and j interact under the assumption that there is no other earlier interaction. In a DMD simulation, once the possible collision events at times t_{ij} have been computed for all possible collision pairs i and j, the earliest event $t^*_{i^*j^*}=\min_{ij} t_{ij}$ is selected. After the collision event between sites i^* and j^* has been executed (according to the rules that will be derived in the next section), the next earliest collision should be performed. However, because the velocities of atoms of the molecules involved in the collision have changed, the previously computed collision times involving these molecules are no longer valid. The next event in the sequence can be determined and performed only after these collision times have been recomputed.

This process describes the basic strategy of DMD, which, without further improvements, would be needlessly inefficient. For if M is the number of possible collision events, finding the earliest time would require $O(M)$ checks, and $M=O(N^2)$, while the number of invalidated collisions that have to be recomputed after each collision would be $O(N)$. Since the number of collisions in the system per unit of physical time also grows linearly with N, the cost of a simulation for a given physical time would be $O(N^2)$ for the computation of collision times and $O(N^3)$ for finding the first collision event.[40] Fortunately, there are ways to significantly reduce this computational cost.[1,5,41–45] The first technique, also used in molecular dynamics simulations of systems interacting with continuous potentials, reduces the number of possible collision times that have to be computed by employing a *cell division* of the system.[1] Note that while the times of certain interaction events (e.g., involving only the molecule's center of mass) can be expressed in analytical form and thus computed very efficiently, the atom-atom interactions have, in general, an orientational dependence and the possible collision time has to be found by means of a numerical root search as explained in the previous section. As a consequence, the most time consuming task in a DMD simulation is the numerical root search for the collision times. One can, however, minimize the required number of collision time computations by dividing the system into subcells and sorting all molecules into these cells according to the positions of their centers of mass. Each cell has a diameter of at least the largest "interaction diameter" of a molecule (in its largest configuration), as measured from its center of mass. As a result, molecules can only collide if they are in the same cell or in an adjacent cell, so the number of collision events to determine, and to recompute after a collision, is much smaller. In this technique, the sorting of molecules into cells is only done initially, while the cell structure is dynamically updated by introducing a *cell-crossing event* for each molecule that is also stored.[5,42] Since the center of mass of a molecule performs linear motion between collision events, one can express its cell-crossing time analytically and therefore the numerical computation of that time is very fast.

The second technique reduces the cost of finding the earliest event time. It consists of storing possible collision and cell-crossing events in a time-ordered structure called a binary tree. For details we refer to Refs. 5 and 42 (alternative event scheduling algorithms exist,[43,44] but it is not clear which technique is generally the most efficient[45]).

Finally, a third standard technique, which is beneficial if an exact solution for the free motion of a molecule is available, is to update the molecules' positions and velocities only at collisions (and possibly upon their crossing the periodic boundaries), while storing the time of their last collision as a property of the molecule called its local clock.[41] Whenever needed, the positions and velocities at later times can be determined from the solution of force-free and torque-free motion of Sec. III.

The use of cell divisions, a binary event tree to manage the events, and local clocks is a standard practice in DMD simulations and largely improves the simulation's efficiency.[5] To see this, note that one picks the earliest event from the tree in each step of the simulation, which scales as $O(\log N)$ for randomly balanced trees.[5,42] If it is a collision event, it is then performed and subsequently $O(1)$ collisions and cell crossings are recomputed and added to the event tree $[\propto O(\log N)]$. If it is a cell-crossing event, the corresponding molecule is put in its new cell, and new possible collision and crossing events are computed $[O(1)]$ and added to the tree $[O(\log N)]$. Then the program progresses to the next event. Since $O(N)$ real events take place per unit of physical time, one sees that using these techniques, the computational cost per unit of physical time due to the computation of possible collisions and cell-crossing times scales as $O(N)$ instead of $O(N^2)$, while the cost due to the event scheduling is $O(N \log N)$ per unit of physical time instead of $O(N^3)$—a huge reduction.

Contrary to what their scaling may suggest, one often finds that the cost of the computation of collision times greatly dominates the scheduling cost for finite N. This is due to the fact that the computations of many of the collision times require numerical root searches. Thus, to gain further computational improvements, one has to improve upon the efficiency of the numerical search for collision-event times. A nonstandard time-saving technique that we have developed

for this purpose is to use *virtual collision events*. In this case, the grid search (see Sec. IV A) for a possible collision time of atoms i and j is carried out only over a fixed small number of grid points, thus limiting the scope of the root search to a small search interval. If no collisions are detected in this search interval, a virtual collision event is scheduled in the binary event tree, much as if it were a possible future collision at the time of the last grid point that was investigated. If the point at which the grid search is curtailed is rather far in the future, it is likely that this virtual event will not be executed because the atoms i and j probably will have collided with other atoms beforehand. Thus, computational work has been saved by stopping the grid search after a few grid points. Every now and then, however, atoms i and j will not have collided with other atoms at the time at which the grid search was stopped. In this case, the virtual collision event in the tree is executed, which entails continuing the root search from the point at which the search was previously truncated. The continued search again may not find a root in a finite number of grid points, and hence another virtual collision would be scheduled, or it may now find a collision. In either case, the new event is scheduled in the tree. This virtual collision technique avoids the unnecessary computation of a collision time that is so far in the future that it will not be executed in all likelihood anyway, while at the same time ensuring that if, despite the odds, that collision is to happen, it is indeed found and correctly executed. Furthermore, since the virtual collision technique restricts the root-search interval, it solves the problem of finding an initial bracketing for intramolecular interactions. The trade-off of this technique is that the event tree is substantially larger, slowing down the event management. Due to the high cost of numerical root searches, however, the simulations presented in the accompanying paper showed that using virtual collision events yields an increase in efficiency between 25% and 110%, depending mainly on the system size.

V. COLLISION RULES

At each moment of interaction, the impulsive forces lead to discontinuous jumps in the momenta (and angular momenta) of the interacting bodies. The effect of the impulses must be consistent with physical principles, such as the conservation of energy. However, the presence of constraints strongly influences how impulses alter the momenta of the system, since all constraints must be obeyed at all times, including at the moment of interaction. As a result, for semiflexible systems, interactions between a pair of sites on molecules typically lead to discontinuous changes in the momenta of all sites in the interacting molecules.

On the other hand, the effect of the impulsive interactions in fully constrained rigid systems can be analyzed more simply by focusing on the change in the center-of-mass motion and angular momenta rather than the site momenta. In the following subsections, the effect of the impulses at interaction times is analyzed for these two cases.

A. Semiflexible systems

The general collision process in systems with discontinuous potentials can be seen as a limit of the collision process for continuous systems in which the interaction potential becomes infinitely steep. A useful starting point for deriving the collision rules is therefore to consider the effect of a force applied to the overall change in the momentum of any atom k:

$$\boldsymbol{p}_k' \equiv \boldsymbol{p}_k(t_c+\Delta t) = \boldsymbol{p}_k + \int_{t_c-\Delta t}^{t_c+\Delta t} \dot{\boldsymbol{p}}_k(t)dt = \boldsymbol{p}_k + \int_{t_c-\Delta t}^{t_c+\Delta t} \boldsymbol{F}_k(\boldsymbol{r}_N(t))dt, \tag{53}$$

where $\boldsymbol{F}_k$ is the total force acting on atom k and $\boldsymbol{p}_k \equiv \boldsymbol{p}_k(t_c - \Delta t)$. Furthermore, here and below, the pre- and postcollision values of a quantity $\boldsymbol{a}$ are denoted by $\boldsymbol{a}$ and $\boldsymbol{a}'$, respectively.

For discontinuous systems, the forces between sites are impulsive and occur only at an instantaneous collision time t_c. When atoms i and j experience either an inter- or intramolecular collision, the interaction potential Φ depends only on the scalar distance r_{ij} between those atoms, so that the force on an arbitrary atom k is given by (without summation over i and j)

$$-\frac{\partial\Phi(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_k} = -\frac{\partial\Phi(\boldsymbol{r}_N)}{\partial r_{ij}}\frac{\partial r_{ij}}{\partial \boldsymbol{r}_{ij}}\cdot\frac{\partial \boldsymbol{r}_{ij}}{\partial \boldsymbol{r}_k} = -\frac{\partial\Phi(\boldsymbol{r}_N)}{\partial r_{ij}}\frac{1}{r_{ij}}\boldsymbol{r}_{ij}(\delta_{jk}-\delta_{ik}).$$

Note that this is nonzero only for the atoms involved in the collision, as expected. Given that the force is impulsive, it may be written as

$$-\frac{\partial\Phi(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_k} = S\delta(t-t_c)\hat{\boldsymbol{r}}_{ij}(\delta_{jk}-\delta_{ik}), \tag{54}$$

where the scalar S is the magnitude of the impulse (to be determined) on atom a in the collision.

In general, the constraint forces on the right-hand side of Eq. (4) must also have an impulsive component whenever forces are instantaneous in order to maintain the rigid-body constraints at all times. We account for this by writing the Lagrange multipliers as

$$\lambda_\alpha = \nu_\alpha + \mu_\alpha\delta(t-t_c).$$

Because λ_α enters into the equations of motion for all atoms k involved in the constraint σ_α, there is an effect of this impulsive constraint force on all those atoms. Thus, one can write for the force on an atom k when atoms i and j collide:

$$\boldsymbol{F}_k(\boldsymbol{r}_N) = -\nu_\alpha\frac{\partial\sigma_\alpha(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_k} + \delta(t-t_c)\times\left[S\hat{\boldsymbol{r}}_{ij}(\delta_{jk}-\delta_{ik}) - \mu_\alpha\frac{\partial\sigma_\alpha(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_k}\right]. \tag{55}$$

Substituting Eq. (55) into Eq. (53), one finds that the term proportional to ν_α vanishes in the limit that the time interval

Δt approaches zero, so that the postcollision momenta $\boldsymbol{p}_k'$ are related to the precollision momenta $\boldsymbol{p}_k$ by

$$\boldsymbol{p}_k' = \boldsymbol{p}_k - \mu_\alpha \frac{\partial \sigma_\alpha(\boldsymbol{r}_N)}{\partial \boldsymbol{r}_k} + S\hat{\boldsymbol{r}}_{ij}(\delta_{jk} - \delta_{ik}). \tag{56}$$

Note that at the instant of collision $t=t_c$, the positions of all atoms $\boldsymbol{r}_N$ remain the same (only their momenta change) so that there is no ambiguity in the right-hand side of Eq. (56) as to whether to take the $\boldsymbol{r}_N$ before or after the collision. It is straightforward to show that due to the symmetry of the interaction potential, the total linear momentum and angular momentum of the system are conserved by the collision rule [Eq. (56)] for arbitrary values of the unknown scalar functions S and μ_α. In addition to these constants of motion, the collision rule must also conserve the total energy and preserve the constraint conditions, $\sigma_\alpha(\boldsymbol{r}_N)=0$ and $\dot{\sigma}_\alpha(\boldsymbol{r}_N)=0$, before and after the collision. The first constraint condition is trivially satisfied at the collision time, since the positions are not altered at the moment of contact. The second constraint condition allows the scalars μ_α to be related to the value of S using Eq. (5) before and after the collision, since we must have

$$\dot{\sigma}_\alpha = \sum_k \frac{\boldsymbol{p}_k}{m_k} \cdot \frac{\partial \sigma_\alpha}{\partial \boldsymbol{r}_k} = 0,$$
$$\dot{\sigma}_\alpha' = \sum_k \frac{\boldsymbol{p}_k'}{m_k} \cdot \frac{\partial \sigma_\alpha}{\partial \boldsymbol{r}_k} = 0. \tag{57}$$

Inserting Eq. (56) into Eq. (57), one gets

$$0 = \sum_k \frac{1}{m_k}\left(S\hat{\boldsymbol{r}}_{ij}(\delta_{jk} - \delta_{ik}) - \mu_\beta \frac{\partial \sigma_\beta}{\partial \boldsymbol{r}_k}\right) \cdot \frac{\partial \sigma_\alpha}{\partial \boldsymbol{r}_k}. \tag{58}$$

Solving this linear equation for μ_α gives

$$\mu_\alpha = \mathbf{Z}^{-1}_{\alpha\beta}\mathcal{F}_\beta,$$
$$\mathcal{F}_\beta = S\hat{\boldsymbol{r}}_{ij} \cdot \left(\frac{1}{m_j}\frac{\partial \sigma_\beta}{\partial \boldsymbol{r}_j} - \frac{1}{m_i}\frac{\partial \sigma_\beta}{\partial \boldsymbol{r}_i}\right), \tag{59}$$

where the $\mathbf{Z}_{\alpha\beta}$ matrix was defined in Eq. (8). Note that for intermolecular interactions, atoms i and j are on different bodies, so that a given constraint σ_β involves either one or the other atom (or neither), and at least one of the two terms on the right-hand side of Eq. (59) is zero. For intramolecular interactions, both terms can be present. Equation (56) can now be written as

$$\boldsymbol{p}_k' = \boldsymbol{p}_k + S\Delta \boldsymbol{p}_k,$$
$$\Delta \boldsymbol{p}_k = \hat{\boldsymbol{r}}_{ij}(\delta_{jk} - \delta_{ik}) - \mu_\alpha^* \frac{\partial \sigma_\alpha}{\partial \boldsymbol{r}_k}, \tag{60}$$

where $\mu_\alpha^* = \mu_\alpha / S$ is a function of the phase-space coordinate as determined by Eq. (59) and is independent of S.

Finally, the scalar S can be determined by employing energy conservation,

$$\sum_k \frac{\boldsymbol{r}_k' \cdot \boldsymbol{p}_k'}{2m_k} + \Delta\Phi = \sum_k \frac{\boldsymbol{p}_k \cdot \boldsymbol{p}_k}{2m_k}, \tag{61}$$

where $\Delta\Phi = \Phi' - \Phi$ denotes the discontinuous change in the potential energy at the collision time. Inserting the expression in Eq. (60) into Eq. (61) and using Eq. (5), one gets a quadratic equation for the scalar S,

$$aS^2 + bS + \Delta\Phi = 0,$$
$$a = \sum_k \frac{\Delta \boldsymbol{p}_k \cdot \Delta \boldsymbol{p}_k}{2m_k},$$
$$b = \sum_k \frac{\boldsymbol{p}_k \cdot \Delta \boldsymbol{p}_k}{m_k} = \hat{\boldsymbol{r}}_{ij} \cdot \boldsymbol{v}_{ij}. \tag{62}$$

For finite values of $\Delta\Phi$, the value of S is therefore

$$S = \frac{-b \pm \sqrt{b^2 - 4a\Delta\Phi}}{2a}, \tag{63}$$

where the physical solution corresponds to the positive (negative) root if $b>0$ ($b<0$), provided $b^2 > 4a\Delta\Phi$. If this latter condition is not met, there is not enough kinetic energy to overcome the discontinuous barrier, and the system experiences a hard-core scattering, with $\Delta\Phi=0$, so that Eq. (63) gives $S=-b/a$. Once the value of S has been computed, the discrete changes in momenta or velocities are easily computed using Eq. (60).

B. Rigid-body approach

The solution method outlined above can be applied to semiflexible as well as rigid molecular systems, but is not very suitable for rigid, continuous bodies composed of an infinite number of point particles. For perfectly rigid bodies, a more convenient approach is therefore to analyze the effect of impulsive collisions on the center of mass and angular coordinates of the system, which are the minimum number of degrees of freedom required to specify the dynamics of rigid systems. The momentum of the center of mass $\boldsymbol{P}_a$ and the angular momentum $\boldsymbol{L}_a$ of rigid molecule a are affected by the impulsive collision via

$$\boldsymbol{P}_a' = \boldsymbol{P}_a + \Delta \boldsymbol{P}_a,$$
$$\boldsymbol{L}_a' = \boldsymbol{R}_a \times \boldsymbol{P}_a' + \mathbf{I}_a \boldsymbol{\omega}_a' = \boldsymbol{L}_a + \boldsymbol{R}_a \times \Delta \boldsymbol{P}_a + \mathbf{I}_a \Delta \boldsymbol{\omega}_a$$
$$= \boldsymbol{L}_a + \Delta \boldsymbol{L}_a, \tag{64}$$

where $\mathbf{I}_a$ and $\boldsymbol{\omega}_a$ are the moment of inertia tensor and the angular velocity of body a in the laboratory frame, respectively. Note that they are related to their respective quantities in the principal axis frame (body frame) via the matrix $\mathbf{A}_a$ (now associated with the body a):

$$\mathbf{I}_a = \mathbf{A}_a \tilde{\mathbf{I}}_a \mathbf{A}_a^\dagger,$$
$$\boldsymbol{\omega}_a = \mathbf{A}_a \tilde{\boldsymbol{\omega}}_a. \tag{65}$$

Note that intramolecular interactions cannot have any effect on fully rigid bodies, hence, the momentum and angu-

lar momentum of a body a can only change due to an intermolecular interaction with another body b. To derive specific forms for the impulsive changes $\Delta \boldsymbol{P}_a$, $\Delta \boldsymbol{\omega}_a$, $\Delta \boldsymbol{P}_b$, and $\Delta \boldsymbol{\omega}_b$, one may calculate the impulsive force and torque acting on the center of mass and angular momentum, leading to $\Delta \boldsymbol{P}_a = -\Delta \boldsymbol{P}_b = -S\hat{\boldsymbol{r}}$ and $\Delta \boldsymbol{L}_a = -\Delta \boldsymbol{L}_b = \boldsymbol{r}_a \times \Delta \boldsymbol{P}_a$, or

$$\Delta \boldsymbol{\omega}_a = -S\mathbf{I}_a^{-1}((\boldsymbol{r}_a - \boldsymbol{R}) \times \hat{\boldsymbol{r}}), \tag{66}$$

where $\boldsymbol{r}_a$ and $\boldsymbol{r}_b$ are the points at which the forces are applied on bodies a and b, respectively, while $\hat{\boldsymbol{r}} = (\boldsymbol{r}_b - \boldsymbol{r}_a)/r_{ab}$ with $r_{ab} = |\boldsymbol{r}_b - \boldsymbol{r}_a|$.

To determine the scalar S, one again uses the conservation of total energy ($E' = E$) to obtain a quadratic equation for S of the form of Eq. (63), with

$$a = \frac{1}{2M_a} + \frac{1}{2M_b} + \frac{\Delta E_\omega^a + \Delta E_\omega^b}{2},$$

$$b = \boldsymbol{v} \cdot \hat{\boldsymbol{r}},$$

where $\boldsymbol{v}$ is the relative velocity of the interaction sites, i.e.,

$$\boldsymbol{v} = \boldsymbol{V}_b - \boldsymbol{V}_a + \boldsymbol{\omega}_b \times (\boldsymbol{r}_b - \boldsymbol{R}_b) - \boldsymbol{\omega}_a \times (\boldsymbol{r}_a - \boldsymbol{R}_a),$$

and

$$\Delta E_\omega^a = \boldsymbol{n}_a \cdot \mathbf{I}_a^{-1} \boldsymbol{n}_a,$$

$$\Delta E_\omega^b = \boldsymbol{n}_b \cdot \mathbf{I}_b^{-1} \boldsymbol{n}_b,$$

with

$$\boldsymbol{n}_a = (\boldsymbol{r}_a - \boldsymbol{R}_a) \times \hat{\boldsymbol{r}},$$

$$\boldsymbol{n}_b = (\boldsymbol{r}_b - \boldsymbol{R}_b) \times \hat{\boldsymbol{r}}.$$

VI. CONCLUSIONS

In this paper we have shown how to carry out discontinuous molecular dynamics simulations for semiflexible and rigid molecules with arbitrary discontinuous intermolecular and intramolecular interaction potentials in the absence of external potentials. For semiflexible bodies, the dynamics and collision rules have been derived from the principles of constrained mechanics. The implementation of an efficient DMD method for semiflexible systems is hindered by the fact that in almost all cases the equations of motion must be propagated numerically in an event searching algorithm so that the constraints are enforced at all times. Nonetheless, such a scheme can be realized using the SHAKE (Ref. 29) or RATTLE (Ref. 46) algorithms in combination with the root-searching methods outlined here.

The dynamics of a system of completely rigid molecules interacting through discontinuous potentials is more straightforward. For such a system, the Euler equations for rigid-body dynamics can be used to calculate the free evolution of a general rigid object. This analytical solution enables the design of efficient numerical algorithms for the search for collision events. In addition, the collision rules for calculating the discontinuous changes in the components of the center of mass velocity and angular momenta have been obtained for arbitrary bodies with single-site interactions based on conservation principles.

From an operational standpoint, the difference between the method of DMD and molecular dynamics using continuous potentials in rigid systems lies in the fact that the DMD approach does not require the calculation of forces and sequential updating of phase-space coordinates at discrete (and short) time intervals since the response of the system to an impulse can be computed analytically. Instead, the computational effort focuses on finding the precise time at which such impulses exert their influence. The basic building block outlined here for the numerical computation of collision times is a grid search, for which the positions of colliding atoms on a given pair of molecules need to be computed at equally spaced points in time. As outlined in Sec. IV, this can be done efficiently starting with a completely explicit analytical form of the motion of a torque-free rigid body. An efficient implementation of the DMD technique to find the collision-event times should make use of (a) a large grid step combined with a threshold scenario to catch pathological cases, (b) sophisticated but standard techniques such as binary event trees, cell divisions, and local clocks, and (c) a new technique of finding collision times numerically that involves truncating the grid search and scheduling virtual collision events.

On a fundamental level, it is natural to wonder whether the "stepped" form of a discontinuous potential could possibly model any realistic interaction. Such concerns are essentially academic, since it is always possible to approximate a given interaction potential with as many (small) steps as one would like in order to approximate a given potential to any desired level of accuracy.[6] Of course, the drawback to mimicking a smooth potential with a discontinuous one with many steps is that the number of collision events that occur in the system per unit time scales with the number of steps in the potential. Hence, one would expect that the efficiency of the simulation scales roughly inversely with the number of steps in the interaction potential. Nonetheless, the issue is a practical one: How small can the number of steps in the interaction potential be such that one still gets a good description of the physics under investigation? In the accompanying paper,[47] we will see for benzene and methane that it takes surprisingly few steps (e.g., a hard core plus a square-well interaction) to get results which are very close to those of continuous molecular dynamics. Additionally, we will compare the efficiency of such simulations to simulations based on standard molecular dynamics methods.

ACKNOWLEDGMENT

The authors would like to acknowledge support by grants from the Natural Sciences and Engineering Research Council of Canada (NSERC).

[1] B. J. Alder and T. E. Wainwright, J. Chem. Phys. **31**, 459 (1959).
[2] B. J. Alder and T. E. Wainwright, J. Chem. Phys. **33**, 1439 (1960).
[3] M. P. Allen and D. J. Tildesley, *Computer Simulation of Liquids* (Clarendon, New York, 1987).
[4] M. P. Allen, D. Frenkel, and J. Talbot, Comput. Phys. Rep. **9**, 301 (1989).
[5] D. C. Rapaport, *The Art of Molecular Dynamics Simulation* (Cambridge

University Press, Cambridge, 2004).
[6] G. A. Chapela, H. T. Davis, and L. E. Sciven, Chem. Phys. **129**, 201 (1989).
[7] G. A. Chapela, S. E. Martínez-Casas, and J. Alejandre, Mol. Phys. **53**, 139 (1984).
[8] R. van Zon and J. Schofield, Phys. Rev. E **65**, 011107 (2002).
[9] A. S. Rueb, Ph.D. thesis, Utrecht, Netherlands, 1834.
[10] C. G. J. Jacobi, Crelle J. Reine Angew. Math. **39**, 293 (1849).
[11] E. T. Whittaker, *A Treatise on the Analytical Dynamics of Particles and Rigid Bodies*, 4th ed. (Cambridge University Press, Cambridge, 1937).
[12] L. D. Landau and E. M. Lifshitz, *Mechanics*, 3rd ed. (Pergamon, New York, 1976).
[13] J. E. Marsden and T. S. Ratiu, *Introduction to Mechanics and Symmetry: A Basic Exposition of Classical Mechanical Systems*, 2nd ed. (Springer, New York, 2002).
[14] Y. Masutani, T. Iwatsu, and F. Miyazaki, *Proceedings of the IEEE International Conference on Robotics and Automation* (IEEE Computer Society Press, San Diego, 1994), Vol. 2, pp. 1066–1072.
[15] G. Ciccotti and G. Kalibaeva, J. Stat. Phys. **115**, 701 (2004).
[16] A. Donev, S. Torquato, and F. H. Stillinger, J. Comput. Phys. **202**, 737 (2005).
[17] A. Donev, S. Torquato, and F. H. Stillinger, J. Comput. Phys. **202**, 765 (2005).
[18] C. De Michele, S. Gabrielli, P. Tartaglia, and F. Sciortino, J. Phys. Chem. B **110**, 8064 (2006).
[19] L. Hernández de la Peña, R. van Zon, J. Schofield, and S. B. Opps (unpublished).
[20] H. Goldstein, *Classical Mechanics* (Addison-Wesley, Reading, MA, 1980).
[21] D. Frenkel and B. Smit, *Understanding Molecular Dynamics*, 2nd ed. (Academic, San Diego, CA, 2004).
[22] To be more precise, no Hamiltonian exists with $\boldsymbol{r}_i$ and $\boldsymbol{p}_i$ as conjugate variables. Otherwise, $\dot{\boldsymbol{r}}_i$ should be equal to $\partial\mathcal{H}/\partial\boldsymbol{p}_i$ and $\dot{\boldsymbol{p}}_i$ should be equal to $-\partial\mathcal{H}/\partial\boldsymbol{r}_i$, whence $\partial\dot{\boldsymbol{r}}_i/\partial\boldsymbol{r}_i=-\partial\dot{\boldsymbol{p}}_i/\partial\boldsymbol{p}_i$, but Eq. (9) violates this relation.
[23] J. D. Ramshaw, Phys. Lett. A **116**, 110 (1986).
[24] M. E. Tuckerman, C. J. Mundy, and G. J. Martyna, Europhys. Lett. **45**, 149 (1999).
[25] M. E. Tuckerman, Y. Liu, G. Ciccotti, and G. J. Martyna, J. Chem. Phys. **115**, 1678 (2001).
[26] S. Melchionna, Phys. Rev. E **61**, 6165 (2000).
[27] M. Fixman, Proc. Natl. Acad. Sci. U.S.A. **71**, 3050 (1974).
[28] E. A. Carter, G. Ciccotti, J. T. Hynes, and R. Kapral, Chem. Phys. Lett. **156**, 472 (1989).
[29] J.-P. Ryckaert, G. Ciccotti, and H. J. C. Berendsen, J. Comput. Phys. **23**, 327 (1977).
[30] R. van Zon and J. Schofield, J. Comput. Phys. (in press).
[31] E. T. Whittaker and G. N. Watson, *A Course of Modern Analysis*, 4th ed. (Cambridge University Press, Cambridge, 1927).
[32] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables* (Dover, New York, 1965).
[33] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in Fortran, The Art of Scientific Computing*, 2nd ed. (Cambridge University Press, Cambridge, 1992).
[34] S. L. Moshier, *Methods and Programs for Mathematical Functions* (Harwood, New York, 1989).
[35] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, M. Booth, and F. Rossi, *GNU Scientific Library Reference Manual*, revised 2nd ed. (Network Theory Ltd., Bristol, 2005).
[36] K. Knopp, *Theory of Functions* (Dover, Mineola, NY, 1947), Pt.II.
[37] D. Baraff, Ph.D. thesis, Cornell University, 1992.
[38] D. Baraff, Comput. Graphics **23**, 223 (1989).
[39] R. P. Brent, *Algorithms for Minimization Without Derivatives* (Prentice-Hall, Englewood Cliffs, NJ, 1973).
[40] This assumes no optimization whatsoever. It is easy to reduce the cost of finding the first possible collision times to $O(N^2)$ per unit physical time, however, by storing only the first collision for every given particle (Ref. 3).
[41] J. J. Erpenbeck and W. W. Wood, in *Statistical Mechanics*, edited by B. J. Berne (Plenum, New York, 1977), Pt. B.
[42] D. C. Rapaport, J. Comput. Phys. **34**, 184 (1980).
[43] B. D. Lubachevsky, J. Comput. Phys. **94**, 255 (1991).
[44] M. Marin, D. Risso, and P. Cordero, J. Comput. Phys. **109**, 306 (1993).
[45] M. Marin and P. Cordero, Comput. Phys. Commun. **92**, 214 (1995).
[46] H. C. Andersen, J. Comput. Phys. **52**, 24 (1983).
[47] L. Hernández de la Peña, R. van Zon, J. Schofield, and S. B. Opps, J. Chem. Phys. **126**, 074106 (2007), following paper.