A semi-analytical approach to molecular dynamics

Dominik L. Michels [a,b,*], Mathieu Desbrun [c,d]

[a] Computer Science Department, Stanford University, 353 Serra Mall, MC 9515, Stanford, CA 94305, USA
[b] Max Planck Institute for Informatics, Campus E1 4, 66123 Saarbrücken, Germany
[c] Department of Computing and Mathematical Sciences, California Institute of Technology, 1200 E. California Blvd., MC 305-16, Pasadena, CA 91125, USA
[d] Centre de Recherche INRIA Sophia-Antipolis Méditerranée, 2004 Route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex, France

A R T I C L E   I N F O

A B S T R A C T

Despite numerous computational advances over the last few decades, molecular dynamics still favors explicit (and thus easily-parallelizable) time integrators for large scale numerical simulation. As a consequence, computational efficiency in solving its typically stiff oscillatory equations of motion is hampered by stringent stability requirements on the time step size. In this paper, we present a semi-analytical integration scheme that offers a total speedup of a factor 30 compared to the Verlet method on typical MD simulation by allowing over three orders of magnitude larger step sizes. By efficiently approximating the exact integration of the strong (harmonic) forces of covalent bonds through matrix functions, far improved stability with respect to time step size is achieved without sacrificing the explicit, symplectic, time-reversible, or fine-grained parallelizable nature of the integration scheme. We demonstrate the efficiency and scalability of our integrator on simulations ranging from DNA strand unbinding and protein folding to nanotube resonators.

1. Introduction

With an ever-increasing need to understand complex behavior at the molecular level comes a high demand for computational methods that can simulate the macroscopic properties of systems from models describing the geometry and the interactions of their molecules. Molecular Dynamics (MD), in particular, has been proven helpful in nanomaterials and bio-engineering, as it allows to understand observed phenomena and predict observations that would be difficult or costly to make experimentally. The drastic improvements in computational power witnessed in recent years have allowed to investigate the structure, dynamics, and even thermodynamics of increasingly complex biological molecules. Yet, only modest progress has been made in extending the size of the time step used in numerical schemes: since explicit time integrators are favored for their ease of parallelization, the time step size remains restricted by the highest frequency components of the intrinsic dynamics of the nuclei—typically arising from stiff bonding between atoms.

This paper tackles this long-standing issue, and proposes the use of a semi-analytical integrator derived from [16,11] to reliably produce three orders of magnitude larger time steps in MD than regular integrators [60], for a resulting 30-fold speedup on average. We demonstrate the efficiency and scalability of our explicit, structure-preserving, and easily

---

* Corresponding author at: Computer Science Department, Stanford University, 353 Serra Mall, MC 9515, Stanford, CA 94305, USA.
E-mail addresses: michels@cs.stanford.edu (D.L. Michels), mathieu@caltech.edu (M. Desbrun).

parallelizable approach on various typical MD simulation runs, varying from DNA unfolding and protein folding to nanotube resonators.

1.1. Related work

Classical versus Quantum Molecular Dynamics. While the motion of atoms and molecules can be obtained in principle by solving the time-dependent Schrödinger equation simultaneously for both electrons and nuclei, such a quantum mechanics approach remains too computationally prohibitive in practice to investigate large molecular systems (see [29] for recent advances). Instead, classical Molecular Dynamics uses Newtonian mechanics: it treats the nuclei (which are much heavier than electrons) as point particles in a force field that accounts for both their mutual interaction as well as the electronic interactions. This force field derives from a potential energy function that is formulated either from expectation values of the quantum system (see Section 2), or using empirical laws.

Numerical integration schemes. Given the typically large number of molecules involved in MD simulations, a fully analytical solution of the resulting Newtonian mechanical system is out of reach. Consequently, numerical methods that evaluate the position of each nucleus at (fixed or adaptive) time intervals are used to find computational approximations to the solutions, given proper initial conditions [28]. Established molecular dynamics simulators (e.g., *LAMMPS* [36]) often make use of the "velocity Verlet" integration scheme [60]—a simple, explicit integrator that can easily be parallelized and whose symplectic nature provides numerical benefits including exact momenta preservation and excellent energy behavior. However, solving for this initial value MD problem is particularly challenging due to the strong covalent bonds between nuclei, requiring painfully small time steps to properly capture the dynamics without generating instability. Turning these covalent bonds into rigid constraints (using RATTLE [4], SHAKE [51], or through internal variables [58] for instance) alleviates the most stringent time step restrictions, but at the cost of having to solve non-linear systems and a significant decrease in parallelizability. Computational efficiency has further increased over the past few years, either through algorithmic improvements (e.g., by computing the more distant interactions less often), or by leveraging specialized hardware for parallel computing (GPU computing). However, the ability to achieve longer time steps efficiently remains a major computational challenge.

1.2. Contributions

In this paper, we present a practical approach to computational molecular dynamics. We propose an efficient (explicit, second-order, and linear-time), structure-preserving (symplectic), and semi-analytical (exponential, or Gautschi-type[1]) integration scheme that allows the use of significantly larger time steps than usual methods through its closed-form treatment of the strongest bonding forces deriving from harmonic potentials.

The key ingredient of our approach, i.e., the use of an exponential integrator combined with a Krylov-based evaluation of matrix functions, has been proposed in the context of quantum molecular dynamics [29]; but hasn't gained acceptance in classical MD so far. Yet, we will demonstrate that its use is particularly appropriate for Hamiltonian systems involving a potential energy with a strong harmonic part, such as the harmonic potentials of bounded atoms in combination with regular non-bonded potentials in MD, or in coarse-graining methods [34]. Moreover, our algorithm scales linearly with the number of atoms and exhibits excellent long-term preservation of momenta and energy, which are hallmarks of symplectic and time-reversible integrators. Finally, our scheme is easily parallelizable as it involves sparse matrix and vector operations, rendering it particularly attractive as a basis for large-scale MD computations.

2. Background

For completeness, we begin our exposition by reviewing the forces involved in molecular dynamics, then discussing the resulting stiffness of the differential equations that we wish to numerically simulate.

2.1. Ehrenfest potential

To bypass the difficulty of solving the time-dependent Schrödinger equation simultaneously for both electrons and nuclei, Ehrenfest was the first to describe the evolution in time of a molecular system using Newtonian (classical) dynamics where nuclei are simply submitted to a force field deriving from a single, effective potential, given as the mean expectation value

---

[1] The history of so-called *exponential integrators* [30] goes back to the late 1950s, when Joseph Hersch pointed out that traditional numerical integrators do not compute the correct solution even if the differential equation is analytically solvable. He then proposed an exact integration scheme for linear ordinary differential equations with constant coefficients, see [24]. Three years later, Walter Gautschi proposed the integration of the non-linear part of the variation of parameters formula with trigonometric polynomials, see [16]. Almost 20 years later, this approach was combined with the trapezoidal rule [11] and a significant improvement was achieved in the late 1990s by Bosco García-Archilla, who introduced *filter functions* for the non-linear part arguments in this context [15]. Today, there exists a variety of exponential integrators, like multistep integrators of this type (cf. [9]), Rosenbrock methods (cf. [50]), and Runge–Kutta-based methods of simple (cf. [37]) and higher orders (cf. [14]) aiming at finding an efficient approximation of the solutions of stiff differential equations.
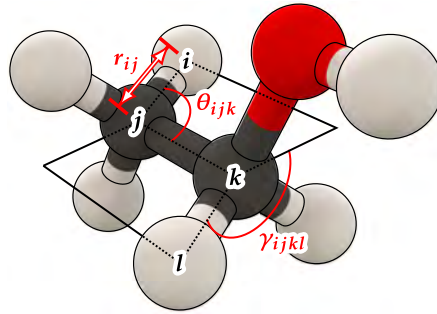
Fig. 1. Example scenario of distance r_{ij}, angle θ_{ijk}, and torsional angle γ_{ijkl} for $(i, j, k, l) \in T$ in the case of a so-called *trans* configuration illustrated with an ethanol molecule.

of the electronic Hamiltonian. (The reader can refer to Appendix A for a recap of the various approximations needed to derive this simplified model.) This Ehrenfest potential V is of the form

$$V(\boldsymbol{R}_1(t), \dots, \boldsymbol{R}_N(t)) := \sum_{i=1}^{N} \sum_{j=1}^{N} u_{ij}(r_{ij})$$

with distances $r_{ij} := \|\boldsymbol{R}_j - \boldsymbol{R}_i\|$ and pairwise potentials u_{ij} between nucleus i and nucleus j for $i, j = 1, \dots, N$, see [49]. As it will become useful in the formal expressions of the various parts of this potential, we denote the set of bonded nuclei $\boldsymbol{B} \subset \{1, \dots, N\}^2$ through pairs of nucleus indices, the set of bond angles $\boldsymbol{A} \subset \{1, \dots, N\}^3$ through triplets of nucleus indices, and the set of torsional angles $\boldsymbol{T} \subset \{1, \dots, N\}^4$ through its four associated nucleus indices. Finally, we denote by

$$\theta_{ijk} := \sphericalangle \left(\boldsymbol{R}_i - \boldsymbol{R}_j, \boldsymbol{R}_k - \boldsymbol{R}_j \right)$$

the angle for $(i, j, k) \in \boldsymbol{A}$ and by

$$\gamma_{ijkl} := \pi \pm \arccos \left(\frac{\langle \boldsymbol{r}_{ij} \wedge \boldsymbol{r}_{jk} | \boldsymbol{r}_{jk} \wedge \boldsymbol{r}_{kl} \rangle}{\|\boldsymbol{r}_{ij} \wedge \boldsymbol{r}_{jk}\| \, \|\boldsymbol{r}_{jk} \wedge \boldsymbol{r}_{kl}\|} \right)$$

the torsional angle[2] for $(i, j, k, l) \in \boldsymbol{T}$ (see Fig. 1), where $\langle \cdot | \cdot \rangle$ denotes the Euclidean inner product.

2.2. Bonded and non-bonded potentials

The potential V is typically written as the sum of four distinct potentials $V = V_{\boldsymbol{B}} + V_{\boldsymbol{A}} + V_{\boldsymbol{T}} + V_{LJC}$, where the covalent bond potential $V_{\boldsymbol{B}}$ is given by a simple harmonic potential

$$V_{\boldsymbol{B}} := \sum_{(i,j) \in \boldsymbol{B}} \frac{k_{\boldsymbol{B}_{ij}}}{2} (r_{ij} - \bar{r}_{ij})^2, \tag{1}$$

the angle potential $V_{\boldsymbol{A}}$ is expressed as

$$V_{\boldsymbol{A}} := \sum_{(i,j,k) \in \boldsymbol{A}} \frac{k_{\boldsymbol{A}_{ijk}}}{2} (\theta_{ijk} - \bar{\theta}_{ijk})^2,$$

the torsional potential $V_{\boldsymbol{T}}$ is written as

$$V_{\boldsymbol{T}} := \sum_{(i,j,k,l) \in \boldsymbol{T}} \sum_{s=1}^{3} k_{\boldsymbol{T}_{ijkl,s}} \left(\cos \left(s \gamma_{ijkl} - \delta_s \right) + 1 \right),$$

and finally, the so-called *Lennard-Jones and Coulomb* potential is expressed as

---

[2] We make use of the equivalent so-called *scalar product definition*

$$\gamma_{ijkl} := \pi \pm \arccos \langle \boldsymbol{r}_{ij} - \langle \boldsymbol{r}_{ij} | \boldsymbol{r}_{kj} / r_{kj} \rangle \boldsymbol{r}_{kj} | \boldsymbol{r}_{lk} - \langle \boldsymbol{r}_{lk} | \boldsymbol{r}_{kj} / r_{kj} \rangle \boldsymbol{r}_{kj} \rangle,$$

since it leads to more efficient computations [19]. Note that in both cases, the sign is determined by $\operatorname{sgn}(\det(\boldsymbol{r}_{ij}, \boldsymbol{r}_{jk}, \boldsymbol{r}_{kl})) = \operatorname{sgn} \langle \boldsymbol{r}_{ij} | \boldsymbol{r}_{jk} \wedge \boldsymbol{r}_{kl} \rangle$.

$$V_{LJC} := \sum_{i=1}^{N} \sum_{j=1}^{i-1} 4\epsilon \left(\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^{6} \right) + \frac{e^2}{4\pi\epsilon_0} \frac{Z_i Z_j}{r_{ij}}$$

and accounts for non-bonded Van der Waals interactions, see [23,38,39]. These potentials are parameterized by the harmonic bending constants $k_{B_{ij}}$ for $(i,j) \in B$, the initial distances $\overline{r}_{ij}$ between i and j for $(i,j) \in B$, the angular constants $k_{A_{ijk}}$ for $(i,j,k) \in A$, the initial angles $\overline{\theta}_{ijk}$ between $(i,j,k) \in A$, the torsional constants $k_{T_{ijkl,s}}$ and phase shifts δ_s for $(i,j,k,l) \in T$ and $s \in \{1,2,3\}$, the well-depth ϵ and finite distance values σ_{ij} for all $(i,j) \in \{1,\ldots,N\}^2$ of the Lennard-Jones potential. Note that in order to account for the anharmonicity of the bond for large length deviation, the covalent bond potential can be expressed as a Morse potential instead, as the steep "walls" of a quadratic potential can prevent bond dissociation. However, a compromise between accuracy and computational speed often dictates the use of mostly quadratic (and sometimes cubic and quartic) terms in practice, with no discernably advert effects on large molecules.

2.3. Stiffness of molecular dynamics

The typical coefficients $k_{B_{ij}}$ of covalent bond potentials are of the order of $1 \text{ MJ mol}^{-1}\text{nm}^{-2}$, whereas the coefficients $k_{A_{ijk}}, k_{T_{ijkl,s}}$, and ϵ associated to angle, torsional, and Lennard-Jones potentials are around $1 \text{ kJ mol}^{-1}\text{nm}^{-2}$. Bonded potentials thus differ by approximately three orders of magnitude from the remaining potentials, and render molecular dynamics *stiff* [10]. Stiffness spells trouble for numerical simulation as it typically requires unduly small time step sizes, particularly for explicit integrators: the presence of significant fast motions in MD can limit the time step to below one femtosecond. While implicit integration can partially remedy this restriction, it can also introduce artificial viscosity to the system, damping high frequencies and violating energy conservation. Even symplecticity does not help with stiff problems: as the authors of [22] point out, "for differential equations where the high frequencies of the oscillations remain nearly constant along the solution, [...] symmetry of the methods is still essential, but symplecticity plays no role in the analysis and in the numerical experiments [...]." Constraint algorithms have also been proposed for molecular dynamics simulations to directly satisfy the stiff bond-length constraints, thus removing the strongest forces in the system [25,51,58]. However, these methods are typically difficult to parallelize and involve tradeoffs between speed, stability, and range of applicability. Recent efforts towards efficient and parallelizable integrators (e.g., using constant Jacobian approximations [13]) are still limited to time steps of the order of one femtosecond for typical MD simulations.

3. Semi-analytical time integration

In order to handle the stiff nature of the differential equations at play in an efficient and structure-preserving manner, we base our approach to molecular dynamics on a time-reversible, symplectic, and semi-analytical integration scheme. Our method consists in reexpressing the equation of motion for MD by singling out distinctively stiff forces that are linear in nuclei displacements. We then provide a closed-form solution of the resulting differential equation using trigonometric matrix functions. We finally evaluate this analytical solution over a large time step through Krylov subspace reduction for computational efficiency, without affecting the numerical behavior expected from a symplectic and time-reversible integrator. This semi-analytical treatment of time integration of MD allows time steps that are over three orders of magnitude larger than the conventional velocity Verlet scheme for the same level of accuracy, resulting in a net speedup of a factor 30 on average. The overall numerical treatment has a computational complexity that is linear in the size of the molecular structures.

3.1. Reexpressing the equation of motion

Molecular dynamics is often written as a second order ordinary differential equation derived from the Hamiltonian form (see Eq. (A.7)) through:

$$M\ddot{R}(t) = -\nabla_R \left(V_B(R(t)) + V_A(R(t)) + V_T(R(t)) + V_{LJC}(R(t)) \right)$$

where R stores the positions of all nuclei. We can thus rewrite the equation of motion as

$$M\ddot{R}(t) + KR(t) + G(R(t)) = 0$$

where $KR(t)$ corresponds to the bonding forces $\nabla_R V_B(R(t))$ in the previous equation, and $G(R(t)) := \nabla_R(V_A(R(t)) + V_T(R(t)) + V_{LJC}(R(t)))$ regroups all the weaker forces deriving from, respectively, the angle, torsional, and Lennard-Jones and Coulomb potentials. We now transform the term $KR(t)$ using a co-rotational formulation originally introduced in [6, 61] for finite element simulations, where geometrically non-linear deformations are decomposed into rigid rotations and the remaining, small contributions due to pure strain: from a rotation field (expressed through a sparse matrix Q) easily determined using basic geometric operations [47], we turn $KR(t)$ into $QK(Q^{-1}R(t) - R_0)$ using the rest positions R_0. We now have isolated the (instantaneous) forces that are linear in $R(t)$, and we can treat the (non-linear, but small) extra term $-QKR_0$ with the other weak terms. Note that Q is orthogonal and K is symmetric as it is a Hessian. Since the mass matrix M is invertible, we reexpress the equation of motion as

$$\ddot{\boldsymbol{R}}(t) + \boldsymbol{A}\boldsymbol{R}(t) + \boldsymbol{M}^{-1}\tilde{\boldsymbol{G}}(\boldsymbol{R}(t)) = \boldsymbol{0}$$

with $\tilde{\boldsymbol{G}}(\boldsymbol{R}(t)) := \boldsymbol{G}(\boldsymbol{R}(t)) - \boldsymbol{Q}\,\boldsymbol{K}\boldsymbol{R}_0$ and where now $\boldsymbol{A} := \boldsymbol{M}^{-1}\boldsymbol{Q}\,\boldsymbol{K}\,\boldsymbol{Q}^{\mathsf{T}}$ may no longer be symmetric for nuclei with different atomic masses. In order to maintain the symmetry of the linear term, we use instead the equivalent equation

$$\ddot{\boldsymbol{\chi}}(t) + \boldsymbol{\Omega}^2\boldsymbol{\chi}(t) + \boldsymbol{\Lambda}(\boldsymbol{\chi}(t)) = \boldsymbol{0} \tag{2}$$

using the modified coordinates $\boldsymbol{\chi}(t) := \sqrt{\boldsymbol{M}}\boldsymbol{R}(t)$, $\boldsymbol{\Omega}^2 := \sqrt{\boldsymbol{M}}\boldsymbol{A}\sqrt{\boldsymbol{M}}^{-1}$, and $\boldsymbol{\Lambda}(\boldsymbol{\chi}(t)) := \sqrt{\boldsymbol{M}}^{-1}\tilde{\boldsymbol{G}}(\sqrt{\boldsymbol{M}}^{-1}\boldsymbol{\chi}(t))$. The matrix $\boldsymbol{\Omega}^2$ is symmetric and positive definite by definition, and thus diagonalizable, which implies that its square root $\boldsymbol{\Omega}$ is well defined.

3.2. Formal solution

We now seek to find the analytical solution of Eq. (2) for a given initial condition $\boldsymbol{\chi}(t_0) = \boldsymbol{\chi}_0$. We begin by considering the equivalent first-order Hamiltonian system

$$\begin{bmatrix} \dot{\boldsymbol{\chi}}(t) \\ \dot{\boldsymbol{v}}(t) \end{bmatrix} = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{1} \\ -\boldsymbol{\Omega}^2 & \boldsymbol{0} \end{bmatrix}\begin{bmatrix} \boldsymbol{\chi}(t) \\ \boldsymbol{v}(t) \end{bmatrix} + \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{\Lambda}(\boldsymbol{\chi}(t)) \end{bmatrix}$$

with $\boldsymbol{v}(t) := \dot{\boldsymbol{\chi}}(t)$ and identity matrix $\boldsymbol{1}$ of dimension $3N$. For conciseness, we rewrite this system of equations as

$$\dot{\boldsymbol{X}}(t) = \mathfrak{A}\boldsymbol{X}(t) + \boldsymbol{\Gamma}(\boldsymbol{X}(t)) \quad \text{where} \quad \boldsymbol{X}(t) := \begin{bmatrix} \boldsymbol{\chi}(t) \\ \boldsymbol{v}(t) \end{bmatrix}, \ \mathfrak{A} = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{1} \\ -\boldsymbol{\Omega}^2 & \boldsymbol{0} \end{bmatrix}, \ \boldsymbol{\Gamma}(\boldsymbol{X}(t)) = \begin{bmatrix} \boldsymbol{0} \\ \boldsymbol{\Lambda}(\boldsymbol{\chi}(t)) \end{bmatrix}. \tag{3}$$

We can now apply Lagrange's variation of parameters approach[3] by considering the general solution of the corresponding homogeneous equation $\dot{\boldsymbol{X}}(t) = \mathfrak{A}\boldsymbol{X}(t)$, i.e.:

$$\boldsymbol{X}(t) = \exp(t\mathfrak{A})\,\boldsymbol{\eta}(t). \tag{4}$$

Since the time derivative of this general solution form is $\dot{\boldsymbol{X}}(t) = \mathfrak{A}\,\boldsymbol{X}(t) + \exp(t\mathfrak{A})\,\dot{\boldsymbol{\eta}}(t)$, we find using Eq. (3) that the function $\boldsymbol{\eta}(t)$ must satisfy:

$$\dot{\boldsymbol{\eta}}(t) = \exp(-t\mathfrak{A})\,\boldsymbol{\Gamma}(\boldsymbol{X}(t)). \tag{5}$$

Consequently, we deduce the closed-form solution of Eq. (3) with initial condition $\boldsymbol{X}(t_0)$ to be:

$$\boldsymbol{X}(t) = \exp(t\mathfrak{A}) \int_{t_0}^{t} \exp(-\tau\mathfrak{A})\,\boldsymbol{\Gamma}(\boldsymbol{X}(\tau))\,\mathrm{d}\tau + \exp\left(-(t - t_0)\mathfrak{A}\right)\boldsymbol{X}(t_0)$$

or, equivalently:

$$\boldsymbol{X}(t + \Delta t) = \exp(\Delta t\mathfrak{A})\boldsymbol{X}(t) + \int_{t}^{t+\Delta t} \exp\left((t + \Delta t - \tau)\mathfrak{A}\right)\boldsymbol{\Gamma}(\boldsymbol{X}(\tau))\,\mathrm{d}\tau. \tag{6}$$

The original first order system Eq. (3) thus admits the following solution:

$$\begin{bmatrix} \boldsymbol{\chi}(t) \\ \boldsymbol{v}(t) \end{bmatrix} = \int_{t_0}^{t} \begin{bmatrix} \boldsymbol{\Omega}^{-1}\sin\left((t - \tau)\,\boldsymbol{\Omega}\right) \\ \cos\left((t - \tau)\,\boldsymbol{\Omega}\right) \end{bmatrix}\boldsymbol{\Lambda}(\boldsymbol{\chi}(\tau))\,\mathrm{d}\tau + \begin{bmatrix} \cos\left((t - t_0)\,\boldsymbol{\Omega}\right) & \boldsymbol{\Omega}^{-1}\sin\left((t - t_0)\,\boldsymbol{\Omega}\right) \\ -\boldsymbol{\Omega}\sin\left((t - t_0)\,\boldsymbol{\Omega}\right) & \cos\left((t - t_0)\,\boldsymbol{\Omega}\right) \end{bmatrix}\begin{bmatrix} \boldsymbol{\chi}(t) \\ \boldsymbol{v}(t) \end{bmatrix} \tag{7}$$

where we used the identity

$$\exp(t\mathfrak{A}) = \exp\left(t\begin{bmatrix} \boldsymbol{0} & \boldsymbol{1} \\ -\boldsymbol{\Omega}^2 & \boldsymbol{0} \end{bmatrix}\right) = \begin{bmatrix} \cos(t\boldsymbol{\Omega}) & \boldsymbol{\Omega}^{-1}\sin(t\boldsymbol{\Omega}) \\ -\boldsymbol{\Omega}\sin(t\boldsymbol{\Omega}) & \cos(t\boldsymbol{\Omega}) \end{bmatrix}.$$

---

[3] The variation of parameters approach was introduced by Joseph-Louis Lagrange to solve ordinary differential equations: one employs the general formula of the solution of the associated homogeneous system, and replace the arbitrary constant with an initially unknown function that has to be determined such that the inhomogeneous equation is satisfied [35].

3.3. Discrete formulation

From the analytical expression of the solution given in Eq. (7), we follow [21] to discretize it into a one-step integrator in $\boldsymbol{\chi}$ and $\boldsymbol{v}$ through:

$$\boldsymbol{\chi}_{n+1} = \cos\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\chi}_n + \boldsymbol{\Omega}^{-1} \sin\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{v}_n + \frac{1}{2} \Delta t^2 \psi\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\Lambda}\left(\phi\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\chi}_n\right), \tag{8}$$

$$\boldsymbol{v}_{n+1} = -\boldsymbol{\Omega} \sin\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\chi}_n + \cos\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{v}_n + \frac{1}{2} \Delta t \left(\psi_0\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\Lambda}\left(\phi\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\chi}_n\right) + \psi_1\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\Lambda}\left(\phi\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\chi}_{n+1}\right)\right) \tag{9}$$

where $\boldsymbol{\chi}_n := \boldsymbol{\chi}(t_0 + n\Delta t)$, $\boldsymbol{v}_n := \boldsymbol{v}(t_0 + n\Delta t)$, and the so-called filter functions $\phi, \psi, \psi_0, \psi_1 : \mathbb{C}^{3N \times 3N} \to \mathbb{C}^{3N \times 3N}$ are even functions that depend smoothly on the square of their arguments (we discuss their exact expressions in Section 3.4). If $\phi(0) = \psi(0) = \psi_0(0) = \psi_1(0) = 1$, the resulting integrator is consistent, and of order two, see [21]. Moreover, the integrator is time-reversible for the choice $\psi(\cdot) = \text{sinc}(\cdot)\psi_1(\cdot)$ (where the cardinal sine matrix function sinc is defined through $\text{sinc}(\cdot) := (\cdot)^{-1}\sin(\cdot)$ for any symmetric and positive definite matrix argument) and $\psi_0(\cdot) = \cos(\cdot)\psi_1(\cdot)$, as easily verified by interchanging $n \leftrightarrow (n+1)$ and $\Delta t \leftrightarrow (-\Delta t)$.

If a two-step method is preferred, one can readily derive it by incorporating the term $\boldsymbol{\Omega}^{-1}\sin\left(\Delta t \boldsymbol{\Omega}\right)\boldsymbol{v}_{n+1}$ of Eq. (8) into Eq. (9) after multiplication by $\boldsymbol{\Omega}^{-1}\sin\left(\Delta t \boldsymbol{\Omega}\right)$, leading to:

$$\boldsymbol{\chi}_{n+1} = 2\cos\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\chi}_n - \boldsymbol{\chi}_{n-1} + \Delta t^2 \psi\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\Lambda}\left(\phi\left(\Delta t \boldsymbol{\Omega}\right) \boldsymbol{\chi}_n\right). \tag{10}$$

The integration scheme in Eq. (10) is semi-analytic as it combines the analytical closed-form solution of the linear part with a numerical integration of the non-linear part. Because the stiff homogeneous part is treated exactly, we will show in Section 4 that one can handle step sizes over three orders of magnitude larger than commonly used integration schemes without running the risk of becoming unstable or losing accuracy.

3.4. Filter functions, symplecticity, and long-term behavior

The choice of filter functions ψ and ϕ in Eq. (10) can greatly impact the long-term properties of our integration method. In effect, our integrator combines the analytic solution of the stiff linear part with a ψ-filtered version of the numerically computed solution of the non-linear part. To allow the error function to smoothly decrease as the step size becomes smaller, the nonlinear argument is also premultiplied by an additional filter function ϕ.

The one-step formulation given by Eq. (8) and Eq. (9) defines a discrete flow

$$\boldsymbol{\Phi}_{\Delta t} : (\boldsymbol{\chi}_n, \boldsymbol{v}_n)^\mathsf{T} \mapsto (\boldsymbol{\chi}_{n+1}, \boldsymbol{v}_{n+1})^\mathsf{T}$$

that is a symplectic transformation if

$$\boldsymbol{\Phi}'_{\Delta t}{}^\mathsf{T} \boldsymbol{J} \boldsymbol{\Phi}'_{\Delta t} = \boldsymbol{J} \tag{11}$$

where $\boldsymbol{\Phi}'_{\Delta t}$ represents the Jacobian of $\boldsymbol{\Phi}_{\Delta t}$ (see Appendix B, Eq. (B.1)) and $\boldsymbol{J}$ an anti-diagonal matrix (see Appendix A, Eq. (A.7)). We already saw that the discrete flow is time reversible if and only if $\psi(\cdot) = \text{sinc}(\cdot)\psi_1(\cdot)$ and $\psi_0(\cdot) = \cos(\cdot)\psi_1(\cdot)$. In this case, direct calculation shows that Eq. (11) holds if and only if $\psi(\cdot) = \text{sinc}(\cdot)\phi(\cdot)$. Therefore the simplest symplectic choice for the filter functions would be the pair $(\psi, \phi) = (\text{sinc}, \mathbf{1})$. In order to improve numerical stability in stiff problems, it is strongly recommended [15] to filter/mollify the position argument of the non-linear term $\boldsymbol{\Lambda}$. Hence we use the symplectic pair $(\psi, \phi) = (\text{sinc}^2, \text{sinc})$: this choice of filters guarantees a smooth evolution of the error function.[4] The resulting integrator in Eq. (10) is thus time-reversible and symplectic, and is able to accurately and efficiently integrate stiff molecular systems without artificial damping. Additionally, linear and angular momenta are exactly preserved, while energy does not drift even over long runs. Note finally that one can trivially handle dissipative systems as well: a damping force of the form $\boldsymbol{D}\dot{\boldsymbol{\chi}}(t)$ can be incorporated if need be by simply adding it, as an external force, to the non-linear forces $\boldsymbol{\Lambda}$ with a ϕ-pre-filtered velocity computed via $(\boldsymbol{\chi}_{n+1} - \boldsymbol{\chi}_n)/\Delta t$.

3.5. Efficient evaluation of matrix functions

The price to pay for a semi-analytical treatment of stiff linear forces is the need to evaluate functions $\mathsf{f}(\boldsymbol{\Omega}^2)$ of the matrix $\boldsymbol{\Omega}^2$ that occur in Eq. (10), such as $\cos(\Delta t \boldsymbol{\Omega})$, $\sin(\Delta t \boldsymbol{\Omega})$, $\psi(\Delta t \boldsymbol{\Omega})$, or $\phi(\Delta t \boldsymbol{\Omega})$. This is a non-trivial computational task, that we address using a Krylov-based reduction [27,52,54] for efficiency.[5]

---

[4] In contrast, without filtering the position argument of the non-linearity, the error oscillates with an amplitude that grows with the step size.

[5] An excellent survey of suitable methods for the particular purpose of matrix exponential computation was published by [45]. Twenty-five years later, the same authors presented a revised version of their survey, cf. [46]. An enormous improvement has been achieved through the development of Krylov-type methods.

The key idea is to prevent the direct computation of the matrix functions, and instead, calculate the action of the matrix $f(\boldsymbol{\Omega}^2)$ on a vector $\boldsymbol{x} \in \mathbb{C}^{3N}$ using a Krylov subspace projection method as described in [59]. A simple approach to approximate $f(\boldsymbol{\Omega}^2)\boldsymbol{x}$ would be the use of the truncated Taylor series of f of degree $m-1$ that is given by $\sum_{k=0}^{m-1} a_k \boldsymbol{\Omega}^{2k}\boldsymbol{x}$. However, the Taylor coefficients a_k are not optimal due to truncation, and a more accurate evaluation is desirable.

Krylov subspace and Arnoldi iteration. One can, instead, find the optimal $\mathcal{L}^2$ approximation to $f(\boldsymbol{\Omega}^2)\boldsymbol{x}$ within the m-dimensional *Krylov subspace* $\mathcal{K}_m$ (with $m \ll n$) generated by $\boldsymbol{\Omega}^2$ and $\boldsymbol{x}$:

$$\mathcal{K}_m(\boldsymbol{\Omega}^2, \boldsymbol{x}) := \operatorname{span}\left(\boldsymbol{x}, \boldsymbol{\Omega}^2\boldsymbol{x}, \ldots, \boldsymbol{\Omega}^{2(m-1)}\boldsymbol{x}\right)$$

as introduced by Aleksey N. Krylov in 1931. Note that, as exploited in the power method, the term $\boldsymbol{\Omega}^{2(m-1)}\boldsymbol{x}$ approximates the dominant eigenvector, and thus, the Krylov subspace is an approximation of the subspace formed by the eigenvectors corresponding to the m largest eigenvalues [44]. An orthonormal basis $\boldsymbol{X}_m$ for the Krylov subspace $\mathcal{K}_m$ is efficiently found via *Arnoldi iteration* [5], which uses a stabilized Gram–Schmidt process [53] to compute $\boldsymbol{X}_m := (\boldsymbol{x}_0, \ldots, \boldsymbol{x}_{m-1}) \in \mathbb{C}^{3N \times m}$ as well as $\boldsymbol{H}_m := \boldsymbol{X}_m^* \boldsymbol{\Omega}^2 \boldsymbol{X}_m \in \mathbb{C}^{m \times m}$. The matrix $\boldsymbol{H}_m$ is in upper Hessenberg form, and describes the projection of $\boldsymbol{\Omega}^2$ onto $\mathcal{K}_m(\boldsymbol{\Omega}^2, \boldsymbol{x})$ relative to the basis $\boldsymbol{X}_m$. Moreover, this Hessenberg matrix satisfies $g(\boldsymbol{\Omega}^2)\boldsymbol{x} = \boldsymbol{X}_m^* g(\boldsymbol{H}_m)\boldsymbol{e_1}$ where $\boldsymbol{e_1} = (1, 0 \ldots, 0)^t$ for any polynomial g of order less than m (see [52], Lemma 3.1). One thus obtain the approximation:

$$f(\boldsymbol{\Omega}^2)\boldsymbol{x} \approx \boldsymbol{X}_m^* f(\boldsymbol{H}_m)\boldsymbol{e_1}$$

where the initial problem of the evaluation of a function with $3N \times 3N$ range has been reduced to the evaluation of a function with $m \times m$ range, where $m \ll 3N$ is the length of the Krylov basis.

Final evaluation. Finally, the remaining function of a low-dimensional matrix can be efficiently evaluated by $f(\boldsymbol{H}_m) = \mathcal{T}^\mathsf{T} f(\mathcal{D})\mathcal{T}$ via an unitary eigendecomposition of $\boldsymbol{H}_m = \mathcal{T}^\mathsf{T} \mathcal{D}\mathcal{T}$. Since $\mathcal{D}$ is diagonal, the evaluation of $f(\mathcal{D})$ can be done trivially, per diagonal element, with only m evaluations.

Automatic determination of Krylov dimension. To determine a conservative value of m such as the Krylov-based evaluations of matrix functions remain below a given error threshold, we follow [27]: the absolute error

$$\epsilon_m := \left\| f(\boldsymbol{\Omega}^2)\boldsymbol{x} - \boldsymbol{X}_m^* f(\boldsymbol{H}_m)\boldsymbol{e_1} \right\|$$

can be bounded from above by $\epsilon_m \leq 10 e^{-m^2/(5\rho)}$ if $\sqrt{4\rho} \leq m \leq 2\rho$ and by $\epsilon_m \leq 10/\rho\, e^{-\rho}\left(\frac{e\rho}{m}\right)^m$ if $m \geq 2\rho$ for $f(\cdot) = \exp(\cdot)$ if the spectrum of the symmetric positive-definite system matrix $\boldsymbol{\Omega}^2$ is a subset of $[0, 4\rho]$. (Note that we approximate ρ with the use of a linear time Gershgorin circle approach.[6]) Since other functions f can trivially be derived from matrix exponentials and trivial algebraic operations, it allows us to determine the basis length m for a given maximum desirable error.

3.6. Long-range potential evaluations

A final issue to address is the efficient evaluation of the long-range Lennard-Jones and Coulomb potentials for non-bonded interactions. While a naive evaluation of the forces on all pairs of nuclei is of quadratic complexity, we employ the *Fast Multipole Method* (FMM, originally introduced in [18]) to reduce potential evaluation to a linear number of operations. Legendre polynomials are used in the spherical series expansion of the potential, see Appendix C for a brief overview of the FMM approach in our context. Two types of expansions are distinguished: the outer expansion, expressing the potential of a charged particle located outside a sphere of a fixed radius due to all charged particles inside it, and the so-called inner expansion, describing the potential due to all charged particles outside the sphere, see Fig. 2. The approximations of the potentials are used in conjunction with a hierarchical space partitioning technique—in our case, an octree. The outer expansions are computed for all particles at the leaf level of the octree and are then accumulated in a bottom-up pass by applying concatenation, translation, and rotation operations. In a subsequent top-down pass over each level, the inner expansions are computed from the outer expansions. We thus obtain a linear-time evaluation of all pairwise long-range interactions.

---

[6] For a matrix $\boldsymbol{A} := \boldsymbol{\Omega}^2 \in \mathbb{C}^{3N \times 3N}$, the Gershgorin circles $\boldsymbol{G}_1, \ldots, \boldsymbol{G}_{3N}$ named after Semyon A. Gershgorin are defined by $\boldsymbol{G}_i := \boldsymbol{S}(a_{ii}, \sum_{j=1..3N} |a_{ij}|)$ for $i \in \{1, \ldots, 3N\}$. We denote by $\boldsymbol{S}(x, r)$ the circle of center $x \in \mathbb{C} \cong \mathbb{R}^2$ with radius $r \in \mathbb{R}$. According to Gershgorin's circle theorem, the spectrum of $\boldsymbol{A}$ is a subset of the union of all Gershgorin circles $\boldsymbol{G}^\circ = \bigcup_{i=1}^{3N} \boldsymbol{G}_i$. Furthermore, the number of Gershgorin circles included in each connected component of $\boldsymbol{G}^\circ$ are equal to the numbers of eigenvalues included in the same components, see [17]. This is used to easily determine approximations for the extremal eigenvalues of the matrix $\boldsymbol{A} = \boldsymbol{\Omega}^2$. Because $\boldsymbol{\Omega}^2$ is symmetric, the spectrum of $\boldsymbol{A}^\mathsf{T}$ is identical and a second estimation can be realized using columnwise evaluated Gershgorin circles. Since $\boldsymbol{A}$ is sparse, the whole process has linear complexity.
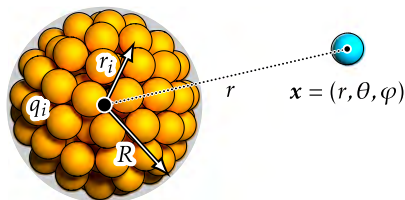
Fig. 2. Illustration of the Fast Multipole Method (FMM): instead of computing all pairwise interactions, the potential outside a sphere of radius R due to all particles inside the sphere is expressed by its outer multipole expansion.
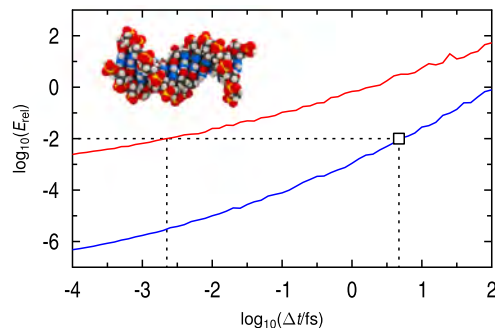

Fig. 3. Unfolding DNA experiment: relative error E_{rel} as a function of the step size Δt in log–log scale for our semi-analytical algorithm (blue curve) and the velocity Verlet-based integration scheme (red curve), both using a FMM-based evaluation of non-bonded potentials. The step size $\Delta t = 4.7$ fs leading to an error $E_{rel} = 1.0\%$ using our semi-analytical algorithm is labeled. Velocity Verlet method requires a step size three orders of magnitude smaller to achieve the same relative error. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.7. Discussion

We note here that the integration scheme we described is valid for a broad range of applications: any simulation involving very strong harmonic potentials can benefit from our semi-analytical approach, including coarse graining methods where pseudo-atoms are used to approximate entire groups of atoms and pseudo-forces are typically chosen to be harmonic [34]. For a serial implementation of our algorithm, the most time consuming part is the multiple evaluations of matrix functions. Since these evaluations only require matrix and vector operations during the execution of the polynomial Arnoldi iteration, an efficient low grain parallelism can be achieved trivially by parallelizing these elementary linear algebra operations. Therefore, our approach combines the two desirable properties of being easy to parallelize and being stable for large time steps. Moreover, our integrator inherits the velocity Verlet algorithm's simplicity with which it can be combined with various more advanced MD techniques, such as multiple time steps or thermostats.

4. Simulation experiments

In this section, we demonstrate the scalability, stability, and long-term behavior of our approach to molecular dynamics for a series of simulation runs. Force field parameters were chosen from the *Assisted Model Building with Energy Refinement* (AMBER) data set [3]. Simulation of fullerenes and nanotubes, this time using *Brenner* bond order potentials, are also presented to exemplify the versatility of our method. When comparing our approach with existing integration methods in molecular dynamics, we use a fully serial execution (i.e., no parallelization) of both of our semi-analytical integrator and of the classical velocity Verlet [60] (also known as Newmark-β integrator for $\beta = 0$ and $\gamma = 1/2$; its simplicity and frequent use by practitioners makes it an ideal reference implementation), as parallelization could render execution times less obvious to analyze. All our timings were measured on a 2.7 GHz Intel® Core™ i7 computer with 16 GB DDR3 SDRAM.

4.1. DNA unfolding: large time steps

As a first evaluation of our approach, we use a simulation of the unfolding of DNA, using two complementary strands with a nucleic acid sequence forming the genetic palindrome CGCGAATTCGCG.[7] The DNA helix is unfolded by breaking the hydrogen bonds between the strands (as is typically happening in, e.g., protein biosynthesis via the helicase enzyme, see [2]). Each CGCGAATTCGCG strand contains 387 nuclei ($N = 774$) as shown in Fig. 4. We simulate the whole unfolding

---

[7] DNA nucleotides pair in the following way: Adenine (A) with Thymine (T) and vice versa; Cytosine (C) with Guanine (G) and vice versa. A nucleotide sequence is denoted a *genetic palindrome* if and only if it is equal to its pair complement [2].
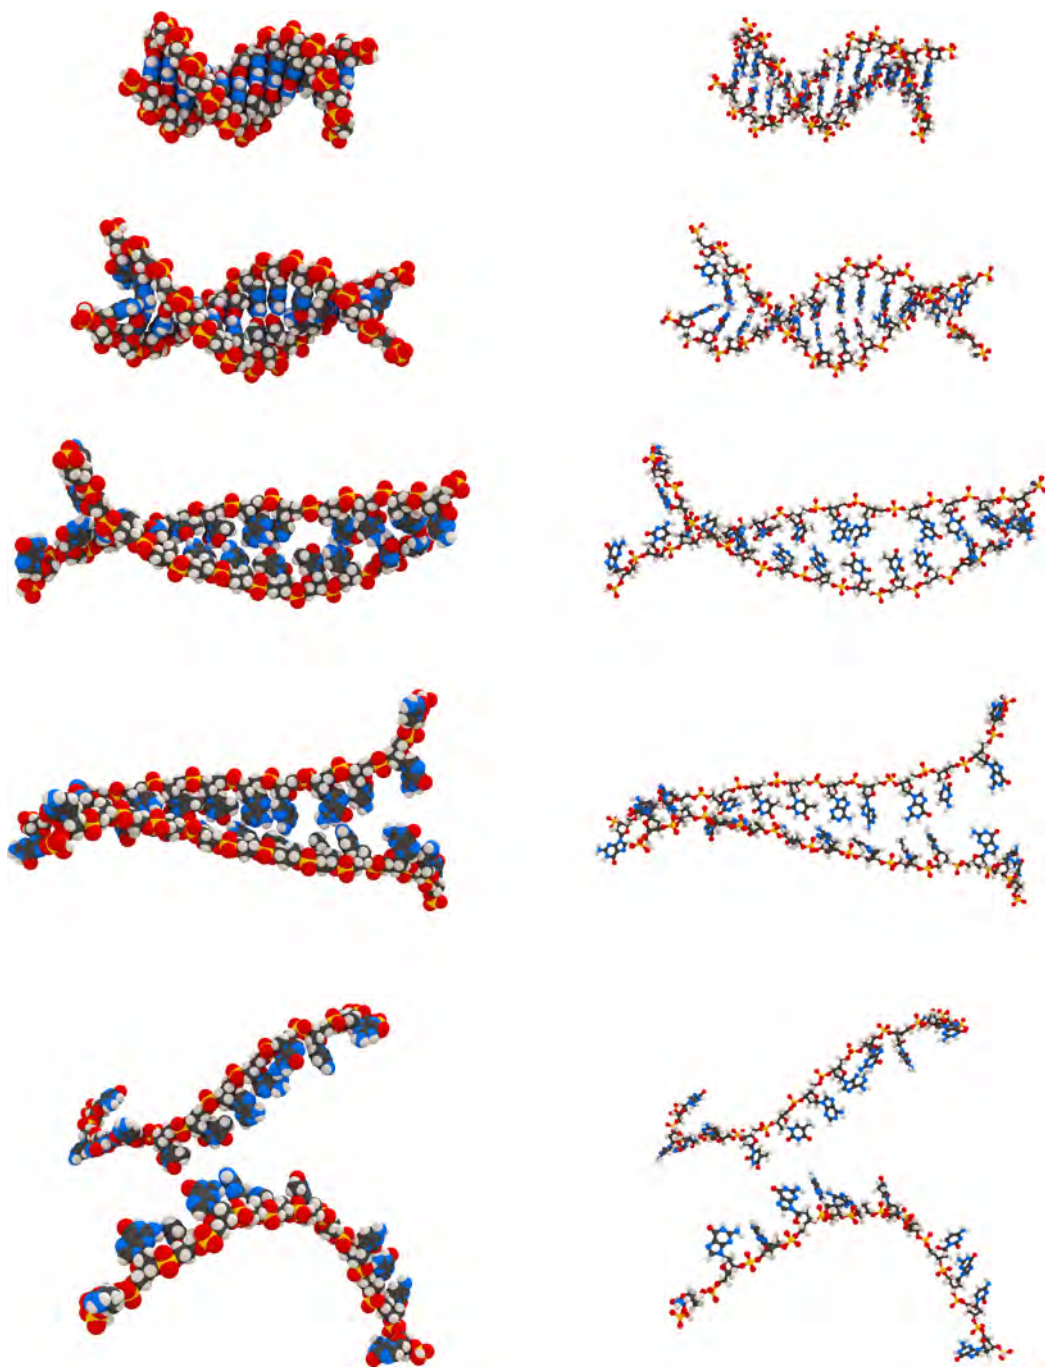
Fig. 4. Two unfolding CGCGAATTCGCG strands simulated with our approach, visualized using a space-filling (left) and a ball-and-stick model (right).

process, which lasts around 0.5 ps. Initial velocities v_0 are set according to the Maxwell–Boltzmann distribution $d(v_0) = 4\pi (m/(2\pi kT))^{1.5} v_0^2 \exp(-m v_0^2/(2kT))$, where m is the nucleus mass, k is Boltzmann's constant, and $T = 296.15$ K is the thermodynamic temperature [42,43].

We begin by comparing the accuracy of the semi-analytical scheme versus velocity Verlet, both using the same FMM-based evaluation of non-bonded potentials. We measure the relative error E_{rel} in position space after simulating a time duration of 0.5 ps, i.e., the $\mathcal{L}^2$ error in position of the nuclei are measured compared to a reference velocity Verlet simulation with step size $\Delta t_{\mathrm{ref}} := 10^{-8}$ fs. Fig. 3 illustrates the dependence of E_{rel} on the step size Δt for both the semi-analytical algorithm and the velocity Verlet integrator. Our numerical tests systematically show asymptotic quadratic accuracy (cf. Fig. 12). The error is smoothly decreasing with decreasing step size for both schemes as expected from second-
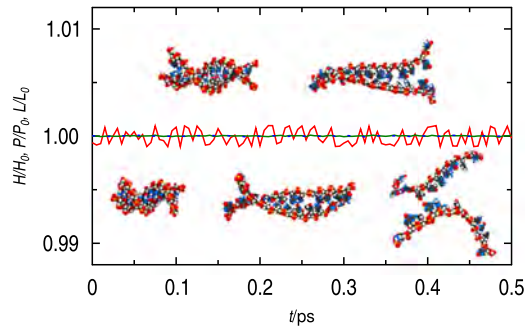
Fig. 5. Illustration of the relative change of the Hamiltonian (red curve), linear momentum (green curve), and angular momentum (blue dashed curve) over the course of the DNA unfolding example for our semi-analytical integration scheme with step size $\Delta t = 4.7$ fs. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)
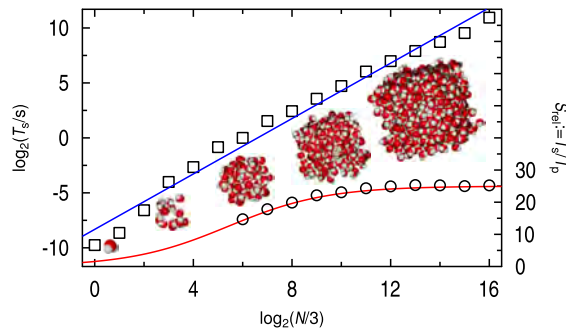

Fig. 6. Average run-time T_s (plotted with squares; their regression line is shown in blue) per iteration for the fully-sequential simulation of interacting water molecules over 0.5 ps as a function of the number of H_2O molecules (shown in log–log scale), using a step size of $\Delta t = 5.0$ fs. If a direct, non-optimized CPU parallelization using pthreads is used, the speedup T_s/T_p between the serial execution time T_s and parallel execution time T_p (plotted in circles; their sigmoidal fit is shown in red) levels off around 24. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

order accurate schemes. However, one can notice that our approach achieves a significantly lower error for a given time step size. In particular, if a maximum relative error of $E_{rel} = 1.0\%$ is acceptable, one can use a step size of $\Delta t = 4.7$ fs with our semi-analytical approach (the size of the Krylov basis is automatically found to be $m = 36$ according to the estimation of ϵ_m for this time step size); to reach the same maximum relative error, the velocity Verlet integrator requires a three orders of magnitude smaller step size of 3.9 as. However, the runtime per iteration of our semi-analytic integrator is 40 times longer than for the velocity Verlet integrator due to the need for estimating matrix functions. Therefore, our scheme achieves a typical speed-up of 30 for a given accuracy, taking around 5 minutes using the semi-analytical method compared to 2.5 hours using the velocity Verlet algorithm.

Finally, we confirm in Fig. 5, that the linear and angular momenta are numerically preserved over the course of the unfolding, and the Hamiltonian does not drift over long runtimes ($> 10^5$ time steps), remaining within 0.2% of its original value. This result, representative of all our other tests, gives further proof that the accuracy threshold used in our Krylov subspace reduction does *not* affect the numerical properties that are expected from a time-reversible and symplectic integrator. The approximation errors introduced by our fast evaluation of matrix functions are, indeed, of the same order as the tolerance thresholds used in non-linear solvers called by implicit (symplectic and/or time-reversible) geometric integrators. The numerical properties of the basic integrator are thus preserved as we will further demonstrate in our subsequent tests.

4.2. Water simulation: scalability, parallelization, and long-term behavior

Our semi-analytical integration scheme (Eq. (10)) is supposed to have a time complexity linear in the number of nuclei at play due to our use of a Krylov subspace projection for our matrix function evaluations and of the Fast Multipole Method for the efficient evaluation of non-bonded potentials. To support this claim, we simulate an increasing number of interacting H_2O molecules—from $2^0 = 1$ to $2^{16} = 65,536$ H_2O molecules (i.e., $N = 3$ to $N = 196,608$ nuclei). Again, initial velocities are set up according to the Maxwell–Boltzmann distribution, with the temperature set to $T = 296.15$ K. For each power-of-two number of molecules, we run a simulation over 100 time steps with a step size of $\Delta t = 5.0$ fs. The runtimes exhibit a clear linear trend as illustrated in Fig. 6, confirming that our combination of a semi-analytical integration and a FMM-based non-bonded potential evaluation scales linearly as expected. The same figure shows that a typical speedup behavior is witnessed when a POSIX thread based CPU parallelization is used instead of a fully sequential implementation.
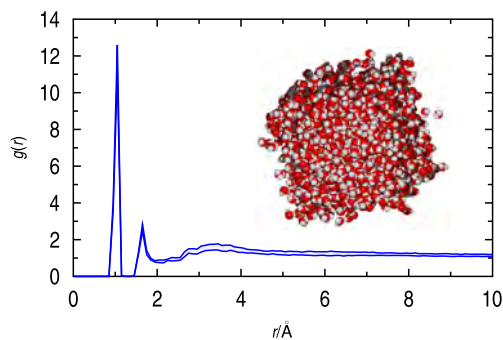
Fig. 7. Range of the radial distribution function (RDF) for a 2 ns long-term simulation of water with 2500 molecules. The two curves delineate the region within which the RDF remains throughout the simulation. The resulting shape of the function is in agreement with known statistics of water molecules.
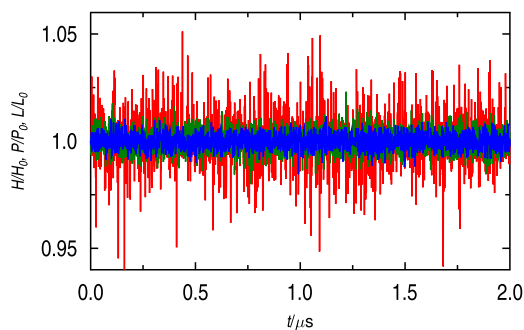

Fig. 8. Plot of the relative changes of the Hamiltonian (red curve), linear momentum (green curve), and angular momentum (blue dashed curve) over a long-term water simulation of 2μ s using our semi-analytical integration scheme with step size $\Delta t = 30$ fs (i.e., 66 million steps). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Additionally, we computed the radial distribution function (also called pair correlation function) $g(r)$ during simulation of $2,500$ H$_2$O molecules over 2 ns using time steps of 260 fs to test the validity of our results. This classical measure of molecule distributions from statistical mechanics describes the histogram of the distances between the two kind of atoms (hydrogen and oxygen). The fact that $g(r)$ is in agreement with theoretical results (see, e.g., [41]) and is numerically preserved over > 7000 steps (see Fig. 7) offers yet another indicator of the stability of our integration method. We also measured the impact of the time step size – and thus, the runtime – on the relative error (compared, as earlier, to a reference implementation with three orders of magnitude smaller time steps) after 2 ns of simulation. Fig. 9 shows a total speedup factor of 30 compared to velocity Verlet for an equivalent relative error, and confirms a graceful degradation of the error for large time steps.

Finally, the excellent long-term properties of the semi-analytical integration scheme are demonstrated by visualizing the Hamiltonian and both linear and angular momenta of the simulation of 128 H$_2$O molecules over a time interval of 2μ s using time steps of 30 fs. Even after over 66 million steps of time integration, no drift is present, and the energy oscillations are limited to below 5% of the original energy. Here again, the approximation of our matrix functions does not hamper the numerical properties of our integrator. (See Fig. 8.)

4.3. Protein folding: long-term behavior

To further underline the practicability of our approach, we ran a folding simulation of the *tryptophan cage* (Trp-Cage) protein sequence NLYIQWLKDGGPSSGRPPPS with $N = 301$ atoms. Starting from a linear chain, this protein should fold into its native conformation within 4 to 10 ns at a temperature of $350K$ [48]. Our simulation for 10 ns of this folding process, carried out in about 90 minutes, is shown in Fig. 10. Accurate preservation of the Hamiltonian as well as linear and angular momenta are demonstrated in Fig. 11, offering more evidence that our fast evaluation of matrix functions does not interfere with the symplectic and time-reversible nature of our basic integration method. Finally, we also compared time step size and runtime vs. positional errors (based, once again, on a reference implementation with three orders of magnitude smaller time steps) in Fig. 9, exhibiting the same graceful degradation as a function of the time step size.
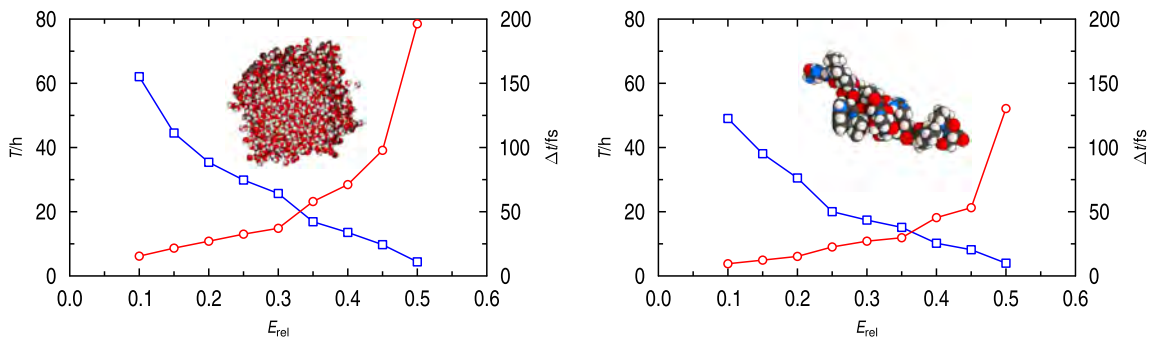
Fig. 9. Plot of the time step size (red curve with circles) and the total runtime (blue curve with squares) as a function of the $\mathcal{L}^2$ relative error after a 2 ns water simulation (left) and a 10 ns protein folding simulation (right). All computations were carried out using a single thread.
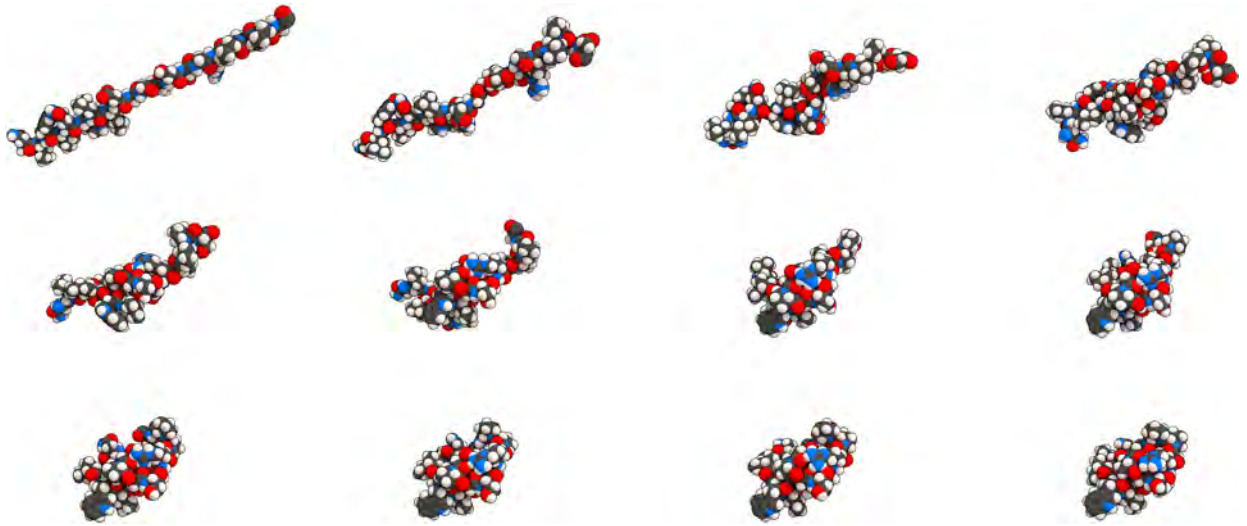

Fig. 10. Folding of the NLYIQWLKDGGPSSGRPPPS sequence simulated using our approach over 10 ns, visualized using a space-filling model. Using naive parallelization, we achieve a speed-up of about a factor of 30, so that the protein folding can be carried out in about 90 minutes, for a $\mathcal{L}^2$ position error of 10%.
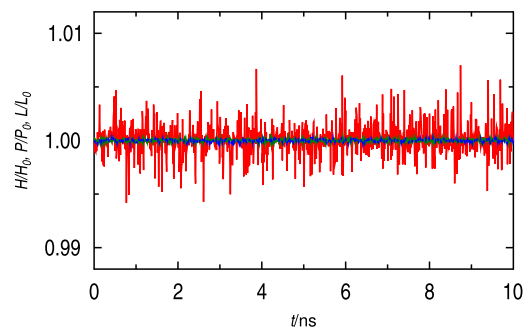

Fig. 11. Plot of the relative changes of the Hamiltonian (red curve), linear momentum (green curve), and angular momentum (blue dashed curve) over a protein folding simulation of 10 ns with our semi-analytical integration scheme. The relative $\mathcal{L}^2$ error in position is about 10% whereas the Hamiltonian oscillates within 0.6% from its original value. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

4.4. Fullerenes and nanotubes: bond order potentials

Semi-analytical integration can efficiently be performed for more general potentials as well. We tested our integrator on the simulation of nanotubes and fullerenes, using *Brenner* bond order potentials.[8] In contrast to classical molecular

---

[8] Bond order potentials are based on Linus C. Pauling's bond order concept, which defines the bond order as the half difference of binding and anti-binding valence electrons in the molecular orbital [1,57].
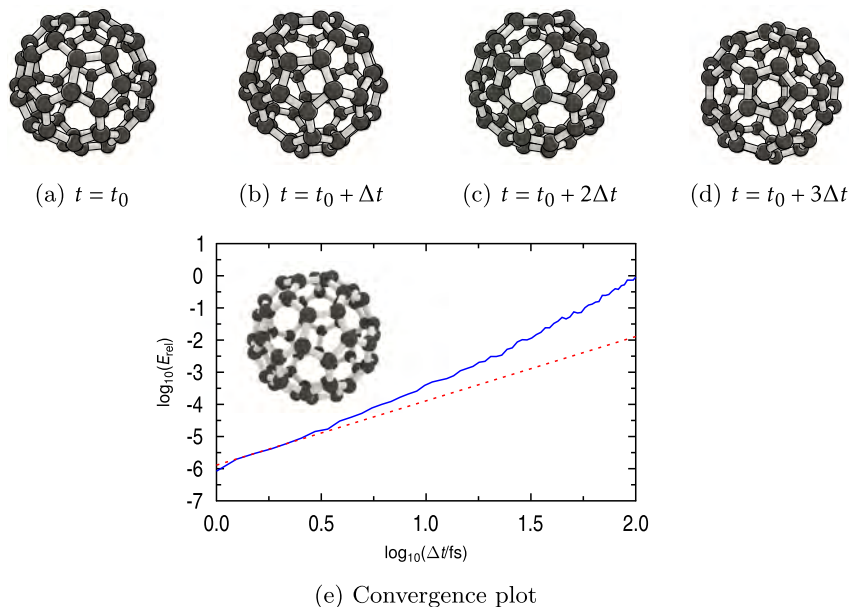
(a) $t = t_0$ (b) $t = t_0 + \Delta t$ (c) $t = t_0 + 2\Delta t$ (d) $t = t_0 + 3\Delta t$

(e) Convergence plot

Fig. 12. Simulation of a rotating buckminsterfullerene C_{60} [33] (named after the famous architect R. Buckminster Fuller) for $\Delta t = 5.0$ fs using our semi-analytical algorithm. The log–log error plot shows a quadratic rate of convergence (red dotted line is of slope 2) of the integrator with decreasing time step size (compared to a reference solution), and an error increasing very slowly as the time step grows.

dynamics force fields, this type of potential can describe several bonding states, and is an example of empirical many-body potential that models intramolecular chemical bonding in several hydrocarbon molecules as well as in graphite and diamond lattices [7]. It is of the form

$$V_{Br} := \sum_{i=1}^{N} \sum_{j=1}^{i-1} f_{ij}\left(r_{ij}\right) \left(U_R\left(r_{ij}\right) - B_{ij} U_A\left(r_{ij}\right)\right)$$

with a repulsive part $U_R(r_{ij})$ and an attractive part $U_A(r_{ij})$. The short range cut-off function $f_{ij}(r_{ij})$ limits the interaction of the nuclei to their nearest neighbors [8].

A simple simulation example of a rotating buckminsterfullerene C_{60} is illustrated in Fig. 12. Similar to the DNA unfolding test, over three orders of magnitude larger step sizes can be used with our integration scheme compared to velocity Verlet for a total speedup factor of 30. We also computed the $\mathcal{L}^2$ error in position of the carbon atoms compared to a reference velocity Verlet simulation with step size $\Delta t_{\text{ref}} := 10^{-8}$ fs, and observed second-order accuracy for small time steps, with an error growing slowly for larger time steps.

We also tested the accuracy of our semi-analytical method in the high frequency range by simulating the atomic-resolution nanomechanical mass sensor introduced in [31]: the idea is to use a nanomechanical resonator as a high precision mass sensor, since the resonant frequency shifts if a particle adsorbs to the resonator. The authors of [31] reach a sensitivity in the range of one Au (gold) atom using a nanotube resonator without destructive ionization effects typically found in traditional mass spectrometers. The relation indicating how a change of mass Δm influences the resonant frequency ν_{res} is highly dependent on the geometry: a position dependent responsivity function can be used to describe the ratio of $\Delta \nu_{\text{res}}$ and Δm. Assuming that the adsorbed mass is distributed along the resonator, this relation can be expressed by $\Delta \nu_{\text{res}} = -\nu_0/(2m_0)\Delta m$, in which ν_0 denotes the initial resonant frequency of the resonator with initial mass m_0.

We set up a nanotube resonator as illustrated in Fig. 13 composed of $N = 651$ C (carbon) atoms and measure the frequency shift $\Delta \nu_{\text{res}}$ depending on the change of mass Δm. As illustrated in Fig. 14 the linear decrease of $\Delta \nu_{\text{res}}$ is successfully reproduced with a mass sensitivity of only one atomic mass unit. For the whole measurement a step size $\Delta t = 20.0$ as is used, which is comparatively large for this accuracy. In contrast, such a high fidelity simulation cannot be carried out efficiently using velocity Verlet-based integration since 10^4 times smaller step sizes are needed in order to capture the linear trend of Fig. 14.

5. Conclusions

We have introduced an exponential integrator for molecular dynamics. Our integration scheme is explicit, which makes it easy to implement either serially or with fine-grained parallelism. Far field approximations are used to render the integrator linear in complexity with the number of nuclei. We also presented ample numerical evidence of second-order accuracy in
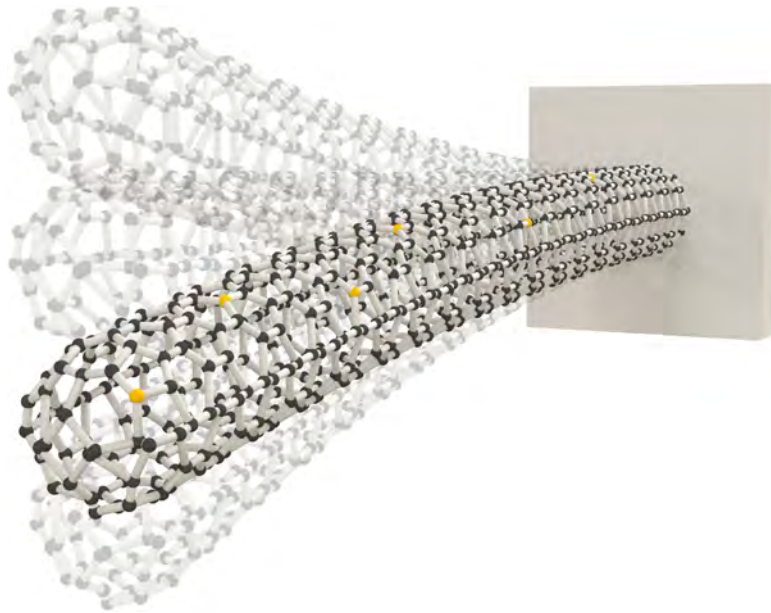
Fig. 13. Simulation of the oscillating nanotube resonator composed of $N = 651$ carbon atoms. A few carbon atoms are replaced by gold (Au) atoms to illustrate the adsorption of gold atoms in carbon nanostructures.
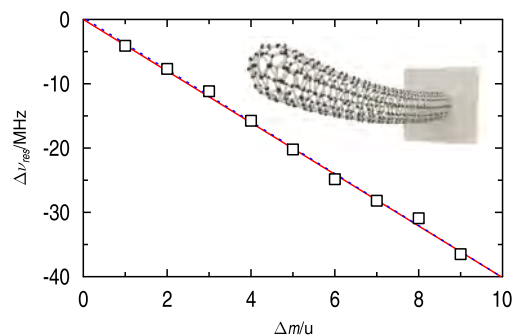

Fig. 14. Illustration of the frequency shift $\Delta \nu_{res}$ depending on the change of mass Δm in a range varying from 1 u to 9 u. The linear reference given by $\Delta \nu_{res}(\Delta m) = -\nu_0/(2m_0)\Delta m$ with $\nu_0 = 318.72$ MHz is shown (red line), as well as the regression line (blue dashed line).

time, as well as an excellent long-term behavior in terms of momenta and energy preservation inherited from the symplectic and time-reversible nature of the integrator. Most distinguishing from previous integrators with similar properties is its ability in extending the length of the time step by over three orders of magnitude compared to the classical velocity Verlet, a property that is obtained by treating the strongest (covalent bond) harmonic forces in a semi-analytical manner to bypass the usual time step size restrictions of explicit integrators. As a result, our approach offers a speedup of at least a factor 30 compared to classical MD algorithms on all the examples we tested. We believe that the efficiency of its treatment of linear forces makes this alternative approach interesting and very valuable for practical simulation of large scale molecular systems and other dynamical systems.

Limitations. Our approach will only offer a significant speedup in systems with strong harmonic potentials. However, this case is particularly common: not only molecular systems are typically exhibiting much stronger covalent bonds than all other forces, but even applications such as coarse graining make heavy use of strong harmonic potentials between pseudo-nuclei.

Future work. Our work can be extended in various ways. First and foremost, a fine-grained parallel implementation on GPU is likely to bring much improved performance, since carefully hand-coded routines can dramatically enhance speedup of explicit integrators (see, e.g., [26]). Second, our algorithm can easily be improved as well. For instance, the use of multiple time steps in the Verlet-1/r-RESPA/Impulse MTS method can be incorporated as well to better deal with the non-linear part of the force field—although the savings will be unlikely to affect the efficiency of our approach in a significant way since we already addressed the stiffest part. Similarly, Langevin dynamics [32] could be used to replace solvent forces by

homogenized and stochastic forces while preserving symplecticity [56]. Finally, we also plan to test out an approach to obtain coarse-grained parallelism using *rational* Krylov subspaces instead: the associated rational Arnoldi method, in which a linear system is solved independently for each pole of the polynomial denominator [20], could offer a coarse-grained way to parallelize the matrix function computations involved in the integrator.

Acknowledgements

Appendix A. Ehrenfest's molecular dynamics

In quantum mechanics, a system S of N nuclei (i.e., protons and neutrons) and K electrons is defined by its wave function

$$\Psi := \Psi_S(\boldsymbol{R}, \boldsymbol{r}, t)$$

which is the solution of the *Schrödinger equation*.[9] In our notation, $\boldsymbol{R} := (\boldsymbol{R}_1, \ldots, \boldsymbol{R}_N)^\mathsf{T}$ denotes the set of positions $\boldsymbol{R}_1, \ldots, \boldsymbol{R}_N \in \mathbb{R}^3$ of the nuclei and $\boldsymbol{r} := (\boldsymbol{r}_1, \ldots, \boldsymbol{r}_K)^\mathsf{T}$ the set of positions $\boldsymbol{r}_1, \ldots, \boldsymbol{r}_K \in \mathbb{R}^3$ of the electrons. All these positions are dependent on time (denoted by t).

Heisenberg's uncertainty principle[10] states that one cannot determine the exact state of S. Hence, the wave function does not provide a classical description of the system and its observables, but instead serves as a means to predict probabilities of observing specific measurements[11]: the probability that S is at time t in the volume element $\prod_{i=1}^{N+K} \mathrm{d}V_i$ of the space of all possible configurations centered at $(\boldsymbol{R}, \boldsymbol{r})$ is given by

$$\langle \Psi(\boldsymbol{R}, \boldsymbol{r}, t) | \Psi(\boldsymbol{R}, \boldsymbol{r}, t) \rangle \prod_{i=1}^{N+K} \mathrm{d}V_i.$$

Analogous to the Hamiltonian in classical mechanics, the Hamilton operator $\mathcal{H}$ corresponds to the total energy of the system, and its spectrum is the set of possible measurements of the total energy. The time-independent operator $\mathcal{H} := \mathcal{H}(\boldsymbol{R}, \boldsymbol{r})$ in the case of negligible relativistic effects [19] is given by

$$\mathcal{H} := -\frac{\hbar^2}{2m_e} \sum_{k=1}^{K} \Delta_{\boldsymbol{r}_k} - \frac{\hbar^2}{2} \sum_{k=1}^{N} \frac{1}{M_k} \Delta_{\boldsymbol{R}_k} - \frac{e^2}{4\pi\epsilon_0} \sum_{k=1}^{K} \sum_{j=1}^{N} Z_j \left\| \boldsymbol{r}_k - \boldsymbol{R}_j \right\|^{-1}$$

$$+ \frac{e^2}{4\pi\epsilon_0} \sum_{k<j}^{K} \left\| \boldsymbol{r}_k - \boldsymbol{r}_j \right\|^{-1} + \frac{e^2}{4\pi\epsilon_0} \sum_{k<j}^{N} Z_k Z_j \left\| \boldsymbol{R}_k - \boldsymbol{R}_j \right\|^{-1}$$

in which the first two terms correspond to the operators T_e and T_n representing the kinetic energy of the electrons and of the nuclei respectively, while the last three terms are the Coulomb potential operators V_{en}, V_{ee}, and V_{nn}, representing the electrostatic potential energies due to interactions between, respectively, electrons and nuclei, electrons only, and nuclei only. Z_k and M_k denotes the atomic number of the k-th nucleus and its mass, while e denotes the elementary electric charge, m_e the electron mass, ϵ_0 the vacuum permittivity, and $\hbar$ the reduced Planck constant.

The Schrödinger equation can now be written out as

$$i\hbar \, \partial_t \Psi(\boldsymbol{R}, \boldsymbol{r}, t) = \mathcal{H} \Psi(\boldsymbol{R}, \boldsymbol{r}, t) \tag{A.1}$$

where we use $\Delta_{\boldsymbol{r}_k} \Psi(\boldsymbol{R}, \boldsymbol{r}, t) := \Delta_{\boldsymbol{y}} \Psi(\boldsymbol{R}, \boldsymbol{r}_1, \ldots, \boldsymbol{r}_{k-1}, \boldsymbol{y}, \boldsymbol{r}_{k+1}, \ldots, \boldsymbol{r}_K, t)|_{\boldsymbol{r}_k}$ and similarly for $\Delta_{\boldsymbol{R}_j} \Psi$. Moreover, since the Hamiltonian is not an explicit function of time, the Schrödinger equation is separable into a product of spatial and temporal parts. Indeed, the wavefunction takes generally the "product" form $\Psi(\boldsymbol{R}, \boldsymbol{r}, t) = \psi(\boldsymbol{R}, \boldsymbol{r}) \cdot f(t)$, so Eq. (A.1) can then be reexpressed as:

---

[9] This equation describes the evolution in time of a non-relativistic quantum systems. It was formulated by Erwin R. Schrödinger in 1926 as a wave equation and successfully used in its first application to explain the spectrum of the hydrogen atom.

[10] This uncertainty principle discovered by Werner K. Heisenberg in 1927 formulates a fundamental limit to the precision with which certain pairs of physical properties (such as position/momentum) can be known simultaneously.

[11] This view is present in the *Copenhagen interpretation* of quantum mechanics formulated by Niels H. Bohr and Werner K. Heisenberg during a meeting in Copenhagen in 1927. It is based on Max Born's probability interpretation of the wave function.

$$i\hbar\, f(t)^{-1}\mathrm{d}_t f(t) = \psi(\boldsymbol{R}, \boldsymbol{r})^{-1}\mathcal{H}\psi(\boldsymbol{R}, \boldsymbol{r}) \tag{A.2}$$

for $\psi(\boldsymbol{R}, \boldsymbol{r}) \cdot f(t) \neq 0$. Since the left-hand (resp., right-hand) side of Eq. (A.2) is only dependent on time (resp., on the positions), we can split the equation in two:

$$i\hbar\, f(t)^{-1}\mathrm{d}_t f(t) = E \tag{A.3}$$

and

$$\psi(\boldsymbol{R}, \boldsymbol{r})^{-1}\mathcal{H}\psi(\boldsymbol{R}, \boldsymbol{r}) = E \tag{A.4}$$

where a constant E is introduced [40]. Eq. (A.3) has closed-form solutions of the form $f(t) = c\exp(-iEt/\hbar)$ with integration constant c scaling the time evolution of the wave equation. Eq. (A.3) can be formulated as an eigenvalue problem $\mathcal{H}\psi(\boldsymbol{R}, \boldsymbol{r}) = E\psi(\boldsymbol{R}, \boldsymbol{r})$ of the operator $\mathcal{H}$, and is known as the *stationary Schrödinger equation*. Hence, to every (eigen) energy value E_k correspond a spatial function ψ_k and a temporal function f_k. Consequently, the solution of Eq. (A.1) is given by a superposition of eigenfunctions ψ_k and corresponding f_k:

$$\Psi(\boldsymbol{R}, \boldsymbol{r}, t) = \sum_k c_k \exp\left(-\frac{i}{\hbar}E_k t\right)\psi_k(\boldsymbol{R}, \boldsymbol{r})$$

with weights $c_k = \int \langle \psi_k(\boldsymbol{R}, \boldsymbol{r})|\Psi(\boldsymbol{R}, \boldsymbol{r}, 0)\rangle\, \mathrm{d}\boldsymbol{R}\, \mathrm{d}\boldsymbol{r}$.

Now we apply the so-called *single configuration ansatz* as presented in [19]. For that we define the decomposition $\mathcal{H} = \mathcal{H}_e + \mathcal{H}_K$ with $\mathcal{H}_e := T_e + V_{ee} + V_{eK} + V_{KK} = T_e + V_e$ and $V_e := V_{ee} + V_{eK} + V_{KK}$. It leads to the approximation

$$\Psi(\boldsymbol{R}, \boldsymbol{r}, t) \approx \tilde{\Psi}(\boldsymbol{R}, \boldsymbol{r}, t) := \chi(\boldsymbol{R}, t)\phi(\boldsymbol{r}, t)\exp\left(\frac{i}{\hbar}\int_{t_0}^{t}\tilde{E}_e(\tau)\,\mathrm{d}\tau\right) \tag{A.5}$$

with nucleus wave function χ, electron wave function ϕ, and phase vector $\tilde{E}_e(t) = \int \langle \chi(\boldsymbol{R}, t)\phi(\boldsymbol{r}, t)|\mathcal{H}_e\phi(\boldsymbol{r}, t)\chi(\boldsymbol{R}, t)\rangle\, \mathrm{d}\boldsymbol{R}\, \mathrm{d}\boldsymbol{r}$. Note that χ and ϕ are normalized, i.e. $\int \langle \chi(\boldsymbol{R}, t)|\chi(\boldsymbol{R}, t)\rangle\, \mathrm{d}\boldsymbol{R} = \int \langle \phi(\boldsymbol{r}, t)|\phi(\boldsymbol{r}, t)\rangle\, \mathrm{d}\boldsymbol{r} = 1$. We substitute $\tilde{\Psi}(\boldsymbol{R}, \boldsymbol{r}, t) \approx \Psi(\boldsymbol{R}, \boldsymbol{r}, t)$ from Eq. (A.5) in Eq. (A.1), multiply with $\phi^*(\boldsymbol{r}, t)$ and $\chi^*(\boldsymbol{R}, t)$, and integrate over $\boldsymbol{R}$ and $\boldsymbol{r}$. Due to energy conservation, i.e.

$$\mathrm{d}_t \int \left\langle \tilde{\Psi}|\mathcal{H}\tilde{\Psi}\right\rangle \mathrm{d}\boldsymbol{R}\, \mathrm{d}\boldsymbol{r} = 0,$$

we end up with the following system of coupled equations[12]:

$$i\hbar\,\partial_t \phi(\boldsymbol{r}, t) + \sum_k \frac{\hbar^2}{2m_e}\Delta_{\boldsymbol{r}_k}\phi(\boldsymbol{r}, t) = \int \langle \chi(\boldsymbol{R}, t)|V_e(\boldsymbol{R}, \boldsymbol{r})\chi(\boldsymbol{R}, t)\rangle\, \mathrm{d}\boldsymbol{R}\,\phi(\boldsymbol{r}, t),$$

$$i\hbar\,\partial_t \chi(\boldsymbol{R}, t) + \sum_k \frac{\hbar^2}{2M_k}\Delta_{\boldsymbol{R}_k}\chi(\boldsymbol{R}, t) = \int \langle \phi(\boldsymbol{R}, t)|\mathcal{H}_e(\boldsymbol{R}, \boldsymbol{r})\phi(\boldsymbol{r}, t)\rangle\, \mathrm{d}\boldsymbol{r}\,\chi(\boldsymbol{R}, t).$$

We approximate the nucleus wave function using classical particles:

$$\chi(\boldsymbol{R}, t) = A(\boldsymbol{R}, t)\exp\left(\frac{i}{\hbar}S(\boldsymbol{R}, t)\right)$$

with amplitude $A(\boldsymbol{R}, t) \in \mathbb{R}_{>0}$ and phase factor $S(\boldsymbol{R}, t) \in \mathbb{R}$. Separating the real and the imaginary parts leads to

$$\partial_t S(\boldsymbol{R}, t) + \sum_{k=1}^{N}\frac{1}{2M_k}\left(\nabla_{\boldsymbol{R}_k}S(\boldsymbol{R}, t)\right)^2 + \int \langle \phi(\boldsymbol{r}, t)|\mathcal{H}_e\phi(\boldsymbol{r}, t)\rangle\, \mathrm{d}\boldsymbol{r} = \hbar^2\sum_{k=1}^{N}\frac{1}{2M_k}\frac{\Delta_{\boldsymbol{R}_k}A(\boldsymbol{R}, t)}{A(\boldsymbol{R}, t)}$$

and

$$\partial_t A(\boldsymbol{R}, t) + \sum_{k=1}^{N}\frac{1}{M_k}\left(\nabla_{\boldsymbol{R}_k}A(\boldsymbol{R}, t)\right)\left(\nabla_{\boldsymbol{R}_k}S(\boldsymbol{R}, t)\right) + \sum_{k=1}^{N}\frac{1}{2M_k}A(\boldsymbol{R}, t)\left(\Delta_{\boldsymbol{R}_k}S(\boldsymbol{R}, t)\right) = 0.$$

Since $\hbar$ is close to zero, the upper equation can be rewritten as

$$\partial_t S(\boldsymbol{R}, t) + \sum_{k=1}^{N}\frac{1}{2M_k}\left(\nabla_{\boldsymbol{R}_k}S(\boldsymbol{R}, t)\right)^2 + \int \left\langle \tilde{\phi}(\boldsymbol{r}, t)|\mathcal{H}_e\tilde{\phi}(\boldsymbol{r}, t)\right\rangle \mathrm{d}\boldsymbol{r} = 0$$

---

[12] This equations are also called TDSCF system, since Paul A. Dirac's time-dependent self-consistent field approach (TDSCF) is based upon it, see [12].

for a function $\tilde{\phi}(\mathbf{r}, t) \approx \phi(\mathbf{r}, t)$. With the shorthand notation $\nabla_{\mathbf{R}} = (\nabla_{\mathbf{R}_1}, \ldots, \nabla_{\mathbf{R}_N})^{\mathsf{T}}$, this last equation has a Hamilton–Jacobi form

$$\partial_t S(\mathbf{R}, t) + H(\mathbf{R}, \nabla_{\mathbf{R}} S(\mathbf{R}, t)) = 0$$

in generalized coordinates $\mathbf{R}$ for a Hamiltonian $H(\mathbf{R}, \mathbf{P}) = T(\mathbf{P}) + V(\mathbf{R})$, kinetic energy $T(\mathbf{P}) = \sum_{k=1}^{N} \|\mathbf{P}_k\|^2 / (2M_k)$, and conjugated momenta $\mathbf{P}(t) := \nabla_{\mathbf{R}} S(\mathbf{R}(t), t)$. Using the classical Hamilton–Jacobi formalism, we get the equations of motion

$$\mathrm{d}_t \mathbf{P}_k = -\nabla_{\mathbf{R}_k} \int \left\langle \tilde{\phi}(\mathbf{r}, t) | \mathcal{H}_e \tilde{\phi}(\mathbf{r}, t) \right\rangle \mathrm{d}\mathbf{r} =: -\nabla_{\mathbf{R}_k} V_e^E(\mathbf{R}(t)),$$

respectively

$$M_k \ddot{\mathbf{R}}_k(t) = -\nabla_{\mathbf{R}_k} V(\mathbf{R}(t)), \tag{A.6}$$

for $k = 1, \ldots, N$, or in abbreviated first-order notation

$$\dot{\mathbf{U}} = \mathbf{J}^{-1} \nabla_{\mathbf{U}} H(\mathbf{U}), \text{ with } \mathbf{U}(t) := \begin{bmatrix} \mathbf{P}(t) \\ \mathbf{R}(t) \end{bmatrix} \text{ and } \mathbf{J} := \begin{bmatrix} \mathbf{0} & \mathbf{1} \\ -\mathbf{1} & \mathbf{0} \end{bmatrix}. \tag{A.7}$$

Eq. (A.7) describes a classical Hamiltonian system acting in a phase space $\mathcal{U} \subset \mathbb{R}^{6N}$. Hence the nuclei are moving according to the equations from classical mechanics in an effective so-called *Ehrenfest*[13] potential V given by the electrons resulting from averaging over the degrees of freedom of the electrons weighted with $\mathcal{H}_e$.

Appendix B. Hamiltonian systems and symplecticity

Symplecticity is a key issue [55] for the integration of Hamiltonian systems like Eq. (A.7), we thus briefly illustrate this concept in this appendix.

The *flow* of Eq. (A.7) is defined by the mapping $\boldsymbol{\varphi}_t : \mathcal{U} \to \mathbb{R}^{6N}$, which maps $\mathbf{U}_0$ to $\boldsymbol{\varphi}_t(\mathbf{U}_0) := \mathbf{U}(t)$ if and only if $\mathbf{U}_0 = \mathbf{U}(t_0)$. We will point the relevance of symplecticity out according to [22]: a differentiable mapping $\mathbf{g} : \mathcal{U} \to \mathbb{R}^{6N}$ acting on an open subset $\mathcal{U} \subset \mathbb{R}^{6N}$ is called *symplectic*, if and only if

$$\omega\left(\mathbf{g}'(\mathbf{U})\boldsymbol{\xi}, \mathbf{g}'(\mathbf{U})\boldsymbol{\mu}\right) = \omega(\boldsymbol{\xi}, \boldsymbol{\mu}) \tag{B.1}$$

holds for all $\boldsymbol{\xi}, \boldsymbol{\mu} \in \mathbb{R}^{6N}$. In Eq. (B.1), $\mathbf{g}'$ denotes the Jacobian of $\mathbf{g}$, whereas the mapping $\omega : \mathbb{R}^{6N} \to \mathbb{R}^{6N}$ maps the pair $(\boldsymbol{\xi}, \boldsymbol{\mu})$ to $\omega(\boldsymbol{\xi}, \boldsymbol{\mu}) := \boldsymbol{\xi}^{\mathsf{T}} \mathbf{J} \boldsymbol{\mu}$.[14] One can think about a symplectic mapping as one which preserves the structure of the phase space, in the sense that it preserves the surface areas resulting from the projections onto planes of momenta and positions. In order to make this statement clear, consider a two-dimensional submanifold M of a set $\mathcal{U} \subset \mathbb{R}^2$ given by the image $M = \boldsymbol{\zeta}(C)$ of a compact set $C \subset \mathbb{R}^2$ under a continuous differentiable function $\boldsymbol{\zeta} : \mathbb{R}^2 \to \mathbb{R}^2$. One can regard M as the limit of the unions of the parallelograms, which are spanned by the vectors $\partial_p \boldsymbol{\zeta}(p, q) \, \mathrm{d}p$ and $\partial_q \boldsymbol{\zeta}(p, q) \, \mathrm{d}q$ with areas $\omega\left(\partial_p \boldsymbol{\zeta}(p, q) \, \mathrm{d}p, \partial_q \boldsymbol{\zeta}(p, q) \, \mathrm{d}q\right)$. Integration over all parallelograms leads to the area

$$\Omega(M) = \int_C \omega\left(\partial_p \boldsymbol{\zeta}(p, q), \partial_q \boldsymbol{\zeta}(p, q)\right) \mathrm{d}p \, \mathrm{d}q \tag{B.2}$$

of the whole surface of M. After applying a mapping $\mathbf{g} : \mathcal{U} \to \mathbb{R}^2$ on M, the surface area reads

$$\Omega(\mathbf{g}(M)) = \int_C \omega\left(\partial_p(\mathbf{g} \circ \boldsymbol{\zeta})(p, q), \partial_q(\mathbf{g} \circ \boldsymbol{\zeta})(p, q)\right) \mathrm{d}p \, \mathrm{d}q. \tag{B.3}$$

With $(\mathbf{g} \circ \boldsymbol{\zeta})'(p, q) = (\mathbf{g}' \circ \boldsymbol{\zeta})(p, q) \, \boldsymbol{\zeta}'(p, q)$, it can be easily proven that the integrands in Eq. (B.2) and Eq. (B.3) are equal and hence $\Omega(\mathbf{g}(M)) = \Omega(M)$ holds iff $\mathbf{g}$ is symplectic.

Furthermore, it can also be proven that for each fixed $t \in \mathbb{R}_{\geq 0}$, the flow $\boldsymbol{\varphi}_t$ of a Hamiltonian system is a symplectic transformation.[15] Conversely, it can be shown that a system is *locally Hamiltonian* if its flow is symplectic.[16] Moreover,

---

[13] There exists a more general theorem named after Paul Ehrenfest, which renders the connection between classical and quantum mechanics. It states that in special cases, classical equations of motions can be used to compute average values in quantum-mechanical systems.

[14] To motivate the definition of ω, consider the case of a single particle moving in the one-dimensional space $\mathbb{R}$. Let $\boldsymbol{\xi}, \boldsymbol{\mu} \in \mathcal{U}$ be two-dimensional vectors of the phase space $\mathcal{U} \subset \mathbb{R}^2$. The area of the parallelogram spanned by $\boldsymbol{\xi}$ and $\boldsymbol{\mu}$ is given by $\det(\boldsymbol{\xi}, \boldsymbol{\mu}) = \boldsymbol{\xi}^{\mathsf{T}} \mathbf{J} \boldsymbol{\mu} = \omega(\boldsymbol{\xi}, \boldsymbol{\mu})$.

[15] This statement was proved by J. Henri Poincaré in the 19th century, if the Hamiltonian is twice continuously differentiable. It follows for $t = 0$ directly and for all $t \in \mathbb{R}_{>0}$ by calculation with $\mathrm{d}_t \omega\left(\partial_{\mathbf{U}_0} \boldsymbol{\varphi}_t(\mathbf{U}_0), \partial_{\mathbf{U}_0} \boldsymbol{\varphi}_t(\mathbf{U}_0)\right) = 0$ using the fact, that $\partial_{\mathbf{U}_0} \boldsymbol{\varphi}_t(\mathbf{U}_0)$ is a solution of the variational equation $\mathbf{J} \dot{\mathbf{V}} = \nabla^2 H(\boldsymbol{\varphi}_t(\mathbf{U}_0)) \mathbf{V}$.

[16] A system $\mathbf{U} = \mathbf{F}(\mathbf{U})$ is called *locally Hamiltonian*, iff for every $\mathbf{U}_0 \in \mathcal{U}$ there exists a neighborhood with $\mathbf{F}(\mathbf{U}) = \mathbf{J}^{-1} \nabla H(\mathbf{U})$ for some analytical function H.

the set of all Hamiltonian systems is closed under transformations of coordinates with symplectic functions, and every function which maps a Hamiltonian system to a Hamiltonian system is symplectic.[17] Therefore, the notions of *canonical* transformations (i.e. transformations described by Hamiltonian systems) and symplectic transformations are equivalent.

Therefore it is evident that integration methods for Hamiltonian systems, which are structure preserving in the sense of symplecticity, are naturally of special interest. By definition, an integration method is symplectic if and only if the discrete flow $\Phi_{\Delta t} : \boldsymbol{U}_0 = \boldsymbol{U}(t_0) \mapsto \boldsymbol{U}_1 \approx \boldsymbol{U}(\Delta t)$ with fixed step size Δt is symplectic. Such methods accurately preserve linear and angular momentum, and energy does not drift, but oscillates instead around the initial energy.[18]

Appendix C. Fast Multipole Method

For completeness, we illustrate in this appendix the Fast Multipole Method (FFM) using the notation from Fig. 2 for the Coulomb potentials $\Phi^i(\boldsymbol{x})$ resulting from charges $Q_i := Z_i e$ in terms of their distances r^i to the origin. These potentials can be written in polynomial terms

$$\Phi^i(\boldsymbol{x}) = \sum_{k=0}^{\infty} \hat{r}_i^k r_i^{-k-1} \mathsf{P}_k \left(\cos \left(\sphericalangle \left(\boldsymbol{x}_i, \boldsymbol{x}_i^0 \right) \right) \right)$$

after the transformation to spherical coordinates $\boldsymbol{x}_i \mapsto (r_i, \theta_i, \varphi_i)$, $\boldsymbol{x}_i^0 \mapsto (\hat{r}_i, \hat{\theta}_i, \hat{\varphi}_i)$ and the use of Legendre polynomials $\mathsf{P}_0, \mathsf{P}_1, \ldots$ defined by

$$\mathsf{P}_k(\psi) = \frac{1}{2^k k!} \mathsf{d}_{\psi}^k ((\psi^2 - 1)^k).$$

For a configuration of N charges of strengths $q_1, q_2, \ldots, q_N$, the potential at a given point $\boldsymbol{x} = (r, \theta, \varphi)$ outside is then given by

$$\Phi(\boldsymbol{x}) = \sum_{k=0}^{\infty} \sum_{l=-k}^{k} r^{-k-1} \mathsf{M}_k^l \mathsf{Y}_k^l(\theta, \varphi)$$

in which

$$\mathsf{M}_k^l = \sum_{j=1}^{N} q_j \hat{r}_j^k \mathsf{Y}_k^{-l}(\hat{\theta}_j, \hat{\varphi}_j)$$

denotes the expansion moment accounting for the charge distribution and

$$\mathsf{Y}_k^l(\theta, \varphi) = \sqrt{\frac{(k - |l|)!}{(k + |l|)!}} \, \mathsf{P}_k^{|l|} \cos(\theta) \exp(il\varphi)$$

the spherical harmonic polynomials of degree k. For the evaluation process of the potentials, an octree is built over the set of particles. The multipole expansion is evaluated for all particles at a leaf node. These potentials are then propagated throughout the octree in a two-phase process: a bottom-up step, in which the multipole expansions of all cells on a single level are accumulated and propagated upwards, and a top-down step, in which the influence of all particles outside a cell at a certain level onto the particles inside this cell are computed.

References

[1] G.C. Abell, Empirical chemical pseudopotential theory of molecular and metallic bonding, Phys. Rev. B 31 (1985) 6184–6196.
[2] B. Alberts, A. Johnson, J. Lewis, M. Raff, K. Roberts, P. Walter, Molecular Biology of the Cell: Reference Edition, vol. 1, Garland Science, 2009.
[3] *AMBER* molecular mechanical force fields, http://ambermd.org/, May 2014.
[4] H.C. Andersen, Rattle: a velocity version of the Shake algorithm for molecular dynamics calculations, J. Comput. Phys. 52 (1983) 24–34.
[5] W.E. Arnoldi, The principle of minimized iterations in the solution of the matrix eigenvalue problem, Q. Appl. Math. 9 (1951) 17–29.
[6] T. Belytschko, L.W. Glaum, Applications of higher order corotational stretch theories to nonlinear finite element analysis, Comput. Struct. 10 (1–2) (1979) 175–182.
[7] D.W. Brenner, Empirical potential for hydrocarbons for use in simulating the chemical vapor deposition of diamond films, Phys. Rev. B 42 (1990) 9458–9471.
[8] D.W. Brenner, O.A. Shenderova, J.A. Harrison, S.J. Stuart, B. Ni, S.B. Sinnott, A second-generation reactive empirical bond order (rebo) potential energy expression for hydrocarbons, J. Phys. Condens. Matter 14 (4) (2002) 783–802.
[9] J. Certaine, The solution of ordinary differential equations with large time constants, in: A. Ralston, H. Wilf (Eds.), Mathematical Methods for Digital Computers, Wiley, 1960, pp. 128–132.

---

[17] This holds for continuous differentiable functions and is known as *Theorem X* proven by Carl G. Jacobi in the 19th century.
[18] Of course, the intensity of this oscillations depends on the step size Δt used for the integration.

[10] C.F. Curtiss, J.O. Hirschfelder, Integration of Stiff equations, Proc. Natl. Acad. Sci. USA 38 (3) (1952) 235–243.
[11] P. Deuflhard, A study of extrapolation methods based on multistep schemes without parasitic solutions, J. Appl. Math. Phys. 30 (1979) 177–189.
[12] P.A. Dirac, Note on exchange phenomena in the Thomas atom, Math. Proc. Camb. Philos. Soc. 26 (1930) 376–385.
[13] P. Eastman, V.S. Pande, Constant constraint matrix approximation: a robust, parallelizable constraint method for molecular simulations, J. Chem. Theory Comput. 6 (2) (2010) 434–437.
[14] A. Friedli, Verallgemeinerte Runge–Kutta Verfahren zur Lösung steifer Differentialgleichungen, Lect. Notes Math. 631 (1978) 214–219.
[15] B. García-Archilla, J.M. Sanz-Serna, R.D. Skeel, Long-time-step methods for oscillatory differential equations, SIAM J. Sci. Comput. 20 (1999) 930–963.
[16] W. Gautschi, Numerical integration of ordinary differential equations based on trigonometric polynomials, Numer. Math. 3 (1961) 381–397.
[17] S.A. Gershgorin, Über die Abgrenzung der Eigenwerte einer Matrix, Bull. Acad. Sci. URSS, Cl. Sci. Math. Nat. 6 (1931) 749–754.
[18] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, J. Comput. Phys. 73 (2) (1987) 325–348.
[19] M. Griebel, S. Knapek, G. Zumbusch, Numerical Simulation in Molecular Dynamics: Numerics, Algorithms, Parallelization, Applications, Springer, 2007.
[20] S. Güttel, Rational Krylov approximation of matrix functions: numerical methods and optimal pole selection, GAMM-Mitt. 36 (1) (2013) 8–31.
[21] E. Hairer, C. Lubich, Long-time energy conservation of numerical methods for oscillatory differential equations, SIAM J. Numer. Anal. 38 (2) (2000) 414–441.
[22] E. Hairer, C. Lubich, G. Wanner, Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations, Springer Series in Computational Mathematics, vol. 31, Springer, 2006.
[23] G.G. Hall, J.E. Lennard-Jones, The molecular orbital theory of chemical valency. III. Properties of molecular orbitals, Proc. R. Soc. Lond. Ser. A, Math. Phys. Sci. 202 (1069) (1950) 155–165.
[24] J. Hersch, Contribution à la méthode des équations aux différences, Z. Angew. Math. Phys. 9 (1958) 129–180.
[25] B. Hess, H. Bekker, H.J.C. Berendsen, J.G.E.M. Fraaije, Lincs: a linear constraint solver for molecular simulations, J. Comput. Chem. 18 (12) (1997) 1463–1472.
[26] B. Hess, C. Kutzner, D. van der Spoel, E. Lindahl, Gromacs 4: algorithms for highly efficient, load-balanced, and scalable molecular simulation, J. Chem. Theory Comput. 4 (3) (2008) 435–447.
[27] M. Hochbruck, C. Lubich, On Krylov subspace approximations to the matrix exponential operator, SIAM J. Numer. Anal. 34 (5) (1997) 1911–1925.
[28] M. Hochbruck, C. Lubich, A bunch of time integrators for quantum/classical molecular dynamics, in: P. Deuflhard, J. Hermans, B. Leimkuhler, A. Mark, S. Reich, R.D. Skeel (Eds.), Algorithms for Macromolecular Modelling, Springer, 1999, pp. 421–432.
[29] M. Hochbruck, C. Lubich, Exponential integrators for quantum-classical molecular dynamics, BIT Numer. Math. 39 (4) (1999) 620–645.
[30] M. Hochbruck, C. Lubich, H. Selfhofer, Exponential integrators for large systems of differential equations, SIAM J. Sci. Comput. 19 (5) (1998) 1552–1574.
[31] K. Jensen, K. Kim, A. Zettl, An atomic-resolution nanomechanical mass sensor, Nat. Nanotechnol. 3 (9) (2008) 533–537.
[32] P.E. Kloeden, E. Platen, Numerical Solution of Stochastic Differential Equations, Stochastic Modelling and Applied Probability, vol. 23, Springer, 1993.
[33] H.W. Kroto, J.R. Heath, S.C. O'Brien, R.F. Curl, R.E. Smalley, C_{60}: buckminsterfullerene, Nature 318 (6042) (1985) 162–163.
[34] Y. Kulkarni, Coarse-graining of atomistic description at finite temperature, PhD thesis, California Institute of Technology, 2006.
[35] J.-L. Lagrange, Supplément au Mémoire sur la théorie générale de la variation des constantes arbitraires, dans tous les problèmes de la mécanique, Mém. Cl. Sci., Math. Phys. (1809).
[36] *LAMMPS* molecular dynamics simulator, http://lammps.sandia.gov/, May 2014.
[37] J.D. Lawson, Generalized Runge–Kutta processes for stable systems with large Lipschitz constants, SIAM J. Numer. Anal. 4 (3) (1967) 372–380.
[38] J.E. Lennard-Jones, On the determination of molecular fields. I. From the variation of the viscosity of a gas with temperature, Proc. R. Soc. Lond. Ser. A, Math. Phys. Sci. 106 (738) (1924) 441–462.
[39] J.E. Lennard-Jones, On the determination of molecular fields. II. From the equation of state of a gas, Proc. R. Soc. Lond. Ser. A, Math. Phys. Sci. 106 (738) (1924) 463–477.
[40] D. Marx, J. Hutter, Ab Initio Molecular Dynamics: Basic Theory and Advanced Methods, Cambridge University Press, 2009.
[41] E. Matteoli, G.A. Mansoori, A simple expression for radial distribution functions of pure fluids and mixtures, J. Chem. Phys. 103 (11) (1995) 4672–4677.
[42] J.C. Maxwell, Illustrations of the dynamical theory of gases, Philos. Mag. 19 (1860) 19–32.
[43] J.C. Maxwell, Illustrations of the dynamical theory of gases, Philos. Mag. 20 (1860) 21–37.
[44] R. Mises, H. Pollaczek-Geiringer, Praktische Verfahren der Gleichungsauflösung, Z. Angew. Math. Mech. 9 (1929) 152–164.
[45] C. Moler, C.V. Loan, Nineteen dubious ways to compute the exponential of a matrix, SIAM Rev. 20 (4) (1978) 801–836.
[46] C. Moler, C.V. Loan, Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later, SIAM Rev. 45 (1) (2003) 3–49.
[47] M. Müller, J. Dorsey, L. McMillan, R. Jagnow, B. Cutler, Stable real-time deformations, in: Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation, SCA '02, ACM, New York, NY, USA, 2002, pp. 49–54.
[48] J.W. Neidigh, R.M. Fesinmeyer, N.H. Andersen, Designing a 20-residue protein, Nat. Struct. Mol. Biol. 9 (6) (2002) 425–430.
[49] D.C. Rapaport, The Art of Molecular Dynamics Simulation, Cambridge University Press, 1996.
[50] H.H. Rosenbrock, Some general implicit processes for the numerical solution of differential equations, Comput. J. 5 (4) (1963) 329–330.
[51] J.-P. Ryckaert, G. Ciccotti, H.J.C. Berendsen, Numerical integration of the Cartesian equations of motion of a system with constraints: molecular dynamics of n-alkanes, J. Comput. Phys. 23 (3) (1977) 327–341.
[52] Y. Saad, Analysis of some Krylov subspace approximations to the matrix exponential operator, SIAM J. Numer. Anal. 29 (1992) 209–228.
[53] Y. Saad, Iterative Methods for Sparse Linear Systems, 2nd edition, SIAM, 2003.
[54] R.B. Sidje, EXPOKIT: a software package for computing matrix exponentials, ACM Trans. Math. Softw. 24 (1) (1998) 130–156.
[55] R. Skeel, Integration schemes for molecular dynamics and related applications, in: M. Ainsworth, J. Levesley, M. Marletta (Eds.), The Graduate Student's Guide to Numerical Analysis, in: Springer Series in Computational Mathematics, vol. 26, Springer, Berlin, Heidelberg, 1999, pp. 119–176.
[56] M. Tao, H. Owhadi, J.E. Marsden, Nonintrusive and structure preserving multiscale integration of stiff odes, sdes, and Hamiltonian systems with hidden slow dynamics via flow averaging, Multiscale Model. Simul. 8 (4) (2010) 1269–1324.
[57] J. Tersoff, New empirical approach for the structure and energy of covalent systems, Phys. Rev. B 37 (1988) 6991–7000.
[58] N. Vaidehi, A. Jain, W.A. Goddard, Constant temperature constrained molecular dynamics: the Newton–Euler inverse mass operator method, J. Phys. Chem. 100 (25) (1996) 10508–10517.
[59] H.A. Van der Vorst, An iterative solution method for solving $f(A)x = b$, using Krylov subspace information obtained for the symmetric positive definite matrix A, J. Comput. Appl. Math. 18 (2) (1987) 249–263.
[60] L. Verlet, Computer experiments on classical fluids. I. Thermodynamical properties of Lennard–Jones molecules, Phys. Rev. 159 (1967) 98–103.
[61] G.A. Wempner, Finite elements, finite rotations and small strains of flexible shells, Int. J. Solids Struct. 5 (2) (1969) 117–153.