

RESEARCH ARTICLE | AUGUST 28 2024

GPU-accelerated on-the-fly nonadiabatic semiclassical dynamics

Special Collection: [Algorithms and Software for Open Quantum System Dynamics](#)

Christopher A. Myers ; Ken Miyazaki ; Thomas Trepl ; Christine M. Isborn  ; Nandini Ananth 



J. Chem. Phys. 161, 084114 (2024)

<https://doi.org/10.1063/5.0223628>



Articles You May Be Interested In

Comparing (stochastic-selection) *ab initio* multiple spawning with trajectory surface hopping for the photodynamics of cyclopropanone, fulvene, and dithiane

J. Chem. Phys. (March 2021)

Benchmarking various nonadiabatic semiclassical mapping dynamics methods with tensor-train thermo-field dynamics

J. Chem. Phys. (July 2024)

Photoinduced quantum dynamics of *ortho*- and *para*-fulvene: Hindered photoisomerization due to mode selective fast radiationless decay via a conical intersection


J. Chem. Phys. (March 2009)

Challenge us.

What are your needs for periodic signal detection?



[Find out more](#)



GPU-accelerated on-the-fly nonadiabatic semiclassical dynamics

Cite as: J. Chem. Phys. 161, 084114 (2024); doi: 10.1063/5.0223628

Submitted: 16 June 2024 • Accepted: 11 August 2024 •

Published Online: 28 August 2024



View Online



Export Citation



CrossMark

Christopher A. Myers,¹  Ken Miyazaki,²  Thomas Treppl,³  Christine M. Isborn,^{1,a)} 
and Nandini Ananth^{4,b)} 

AFFILIATIONS

¹Department of Chemistry and Biochemistry, University of California Merced, Merced, California 95343, USA

²Department of Chemistry, Northwestern University, Evanston, Illinois 60208, USA

³Theoretical Physics IV, University of Bayreuth, 95440 Bayreuth, Germany

⁴Department of Chemistry and Chemical Biology, Baker Laboratory, Cornell University, Ithaca, New York 14853, USA

Note: This paper is part of the JCP Special Topic on Algorithms and Software for Open Quantum System Dynamics.

^{a)}Author to whom correspondence should be addressed: cisborn@ucmerced.edu

^{b)}E-mail: ananth@cornell.edu

ABSTRACT

GPU-accelerated on-the-fly nonadiabatic dynamics is enabled by interfacing the linearized semiclassical dynamics approach with the TeraChem electronic structure program. We describe the computational workflow of the “PySCES” code interface, a Python code for semiclassical dynamics with on-the-fly electronic structure, including parallelization over multiple GPU nodes. We showcase the abilities of this code and present timings for two benchmark systems: fulvene solvated in acetonitrile and a charge transfer system in which a photoexcited zincphthalocyanine donor transfers charge to a fullerene acceptor through multiple electronic states on an ultrafast timescale. Our implementation paves the way for an efficient semiclassical approach to model the nonadiabatic excited state dynamics of complex molecules, materials, and condensed phase systems.

Published under an exclusive license by AIP Publishing. <https://doi.org/10.1063/5.0223628>

I. INTRODUCTION

Nonadiabatic molecular dynamics (NAMD), where the motion of nuclei and electrons cannot be uncoupled and the dynamics occur beyond the Born–Oppenheimer regime, occur in many key chemical processes, such as photo-initiated molecular dissociation, photocatalytic reactions that use electronic excited states to drive chemical reactions, and photoinduced charge separation dynamics in photovoltaics.^{1–5} However, such dynamics are challenging to model with a fully quantum treatment of nuclei and electrons, especially for complex and condensed phase systems.

Widely considered the gold standard for NAMD simulations, the multiconfigurational time-dependent Hartree (MCTDH) approach is highly accurate but only computationally feasible for systems with a handful of nuclear degrees of freedom and generally requires pre-computed potential energy surfaces. Despite the recent development of multi-layer (ML) MCTDH methods⁶ that make possible applications of MCTDH to more complex systems, at present, on-the-fly dynamics for complex molecular systems with

many degrees of freedom.⁷⁹ Ab initio multiple spawning (AIMS) is an approximate implementation of a formally exact spawning Gaussian wavepacket approach that can account for quantum effects. However, this method is expensive, requiring some remedies to the computational cost.^{3,7} Of the approximate methods available, the fewest switches surface hopping (FSSH) and Ehrenfest dynamics are, arguably, the most popular. Both employ classical nuclear trajectories and have been used in *ab initio* on-the-fly simulations. Unfortunately, neither approach accounts for nuclear quantum effects, and both suffer from strong coherence and slow decoherence, necessitating additional approximations.^{8–11}

Nonadiabatic semiclassical implementations for on-the-fly simulations are in their infancy, with the *ad hoc* symmetrical quasiclassical (SQC) method showing some promise.^{12,13} The linearized semiclassical (LSC) approach^{14,15} is a rigorous alternative to the SQC approximation that two of us recently demonstrated can be used for on-the-fly simulations of excited state dynamics.¹⁶ LSC offers two key advantages: (i) it directly approximates the quantum correlation function, allowing us to compute experimental observables;

and (ii) the implementation involves evolving independent classical trajectories from an initial quantum Wigner distribution and does not involve any parameters unlike FSSH (decoherence corrections) or AIMS (spawning thresholds). LSC is highly parallelizable with a computational expense that scales almost classically with system dimensions, making it a very promising candidate for on-the-fly simulations of excited state dynamics in condensed phase environments.

Beyond the method chosen for the dynamics, the main determination of computational cost for NAMD is the electronic structure calculations of the ground and excited state gradients as well as the nonadiabatic coupling (NAC) vectors. For many NAMD simulations of bond-breaking reactions, a multi-reference method such as state-averaged complete active space self-consistent field (SA-CASSCF) or CAS with configuration interaction (CAS-CI) is employed.^{17–19} In other nonadiabatic processes, more affordable single reference excited state methods such as linear response time-dependent density functional theory (TDDFT) are used.^{20–22} TDDFT, especially when employing a long-range corrected hybrid functional, can often provide accurate vertical excitation energies near the ground state minimum,^{23–26} but due to its single reference nature, it will generally poorly describe bond breaking, biradical intermediates, and crossings between ground and excited state potential energy surfaces.²⁷ TDDFT has recently been shown to be accurate in modeling conical intersections between excited states,²⁸ and variants of excited state single reference methods, e.g., with references with different spins, particle numbers, or excitation levels, have shown promise in capturing non-dynamic correlation^{29–33} but have not seen widespread use for NAMD simulations.

The TeraChem software package is built to take advantage of GPU parallelization^{34,35} and offers significant speed ups in quantum electronic structure calculations compared to CPU-based packages. A variety of excited state electronic structure methods have been implemented within TeraChem, including SA-CASSCF,³⁶ full-CI,³⁷ CASCI,³⁸ and TDDFT.^{39,40} These GPU-accelerated implementations have been used with both AIMS and surface hopping NAMD simulations.^{19,41–44} The recent extension of TeraChem to a protocol buffer client–server model has allowed for a straightforward interface between TeraChem’s electronic structure capabilities and any external code for on-the-fly dynamics.⁴⁵

Motivated by the success of the recent on-the-fly nonadiabatic LSC simulations, here we introduce the PySCES code (Python code for Semiclassical Dynamics with on-the-fly Electronic Structure) that interfaces the LSC method with the TeraChem electronic structure code. We explain the interface between the LSC Python code and TeraChem, describing how NAMD can be run in parallel over multiple GPU nodes and how we account for the phase difference in wave function that leads to sign swapping of the NAC vector. We compare, for a single NAMD trajectory, the excited state occupations for ethylene using the GAMESS electronic structure code⁸⁰ and TeraChem, then showcase the performance of the code for two larger test systems: solvated fulvene and the charge transfer in a zinc phthalocyanine-C₆₀ donor–acceptor complex.

II. METHODS

Here, we provide a brief overview of the LSC method for non-adiabatic dynamics, introduce the relevant equations of motion,

and describe the distributions from which initial conditions are sampled.

A. Linearized semiclassical approximation to real-time correlation functions

We start with a general real-time quantum correlation function between two operators, \hat{A} and \hat{B} ,

$$C_{AB}(t) = \text{Tr} \left[\hat{A} e^{i\hat{H}t/\hbar} \hat{B} e^{-i\hat{H}t/\hbar} \right], \quad (1)$$

where \hat{H} is the Hamiltonian. The path integral formulation recasts a real-time quantum propagator into a sum over all possible paths in phase space, rendering Eq. (1) a double integral over forward and backward paths. The LSC approximation is obtained by linearizing the difference in action between the forward and backward paths to obtain^{15,46}

$$C_{AB}^{\text{LSC}}(t) = \frac{1}{(2\pi\hbar)^N} \iint d\mathcal{X}_0 d\mathcal{P}_0 A_W(\mathcal{X}_0, \mathcal{P}_0) B_W(\mathcal{X}_t, \mathcal{P}_t), \quad (2)$$

where A_W and B_W are the phase space functions resulting from the Wigner transforms⁴⁷ of \hat{A} and \hat{B} , respectively. In Eq. (2), all the system degrees of freedom, electronic and nuclear, are represented by their corresponding phase space variables, $(\mathcal{X}, \mathcal{P})$. The LSC approximation to the quantum correlation function is obtained by sampling trajectory initial conditions from the function A_W at time $t = 0$, propagating trajectories under the classical analog Hamiltonian for a time t , and using the time-evolved phase space variables to evaluate function B_W .

We note that the LSC dynamics is entirely classical, and trajectories are independent of each other. Therefore, the parallelization of LSC dynamics is simple and intuitive, in contrast to more elaborate quantum dynamics methods such as MCTDH⁴⁸ and AIMS⁴⁹ and less approximate semiclassical methods with non-negligible interactions between trajectories.⁵⁰ However, the lack of trajectory interference in LSC dynamics means that the method fails to describe nuclear quantum coherence effects and deep tunneling, although coherences between electronic states are captured owing to the nature of the electronic mapping variables and the associated nonadiabatic equations of motion discussed in Sec. II B.^{51–53} Given that the present work is aimed at enabling simulations of high-dimensional systems where nuclear quantum coherences are not expected to play a significant role, we predict that the limitations of LSC dynamics will not significantly affect the resulting accuracy.

B. Mapping approach and nonadiabatic equations of motion

To express the quantum Hamiltonian in phase space and subsequently derive the associated equations of motion, we adopt the so-called mapping approach. Here, we limit our discussion to the Meyer Miller Stock Thoss (MMST) mapping,^{54–56} although several mapping protocols have been introduced recently that aim to better preserve the electronic phase space.^{57,58} For a comparative study of different types of mapping approaches in the context of on-the-fly LSC dynamics, we refer readers to an earlier study.¹⁶

In the MMST mapping, the occupation of electronic states is expressed using the ladder operators of a quantum harmonic oscillator,

$$\begin{aligned} |n\rangle &\mapsto \hat{a}_n^\dagger \\ |n\rangle\langle m| &\mapsto \hat{a}_n^\dagger \hat{a}_m, \end{aligned} \quad (3)$$

where $\hat{a}_n = (\hat{x}_n + i\hat{p}_n)/\sqrt{2}$ and $\hbar = 1$. The general Hamiltonian for F states is then

$$\hat{H} = \frac{1}{2} \mathbf{P}^\dagger \boldsymbol{\mu}^{-1} \mathbf{P} + \sum_{n,m}^F \varepsilon_{nm}(\mathbf{R}) |n\rangle\langle m|, \quad (4)$$

which can thus be reformulated in terms of nuclear and electronic phase space operators,

$$H = \frac{1}{2} \mathbf{P}^\dagger \boldsymbol{\mu}^{-1} \mathbf{P} + \sum_{n,m}^F \frac{1}{2} (x_n x_m + p_n p_m - \delta_{nm}) \varepsilon_{nm}(\mathbf{R}), \quad (5)$$

where (\mathbf{R}, \mathbf{P}) and $\boldsymbol{\mu}$ are the nuclear phase space operators and the associated mass vector, respectively, and $\varepsilon_{nm}(\mathbf{R})$ is the element of the diabatic potential energy matrix. In the adiabatic representation more suitable for interfacing with *ab initio* electronic structure programs, Eq. (5) becomes^{54,59}

$$\begin{aligned} H &= \frac{1}{2} \mathbf{P}_{\text{kin}}^\dagger \boldsymbol{\mu}^{-1} \mathbf{P}_{\text{kin}} + \sum_n^F \frac{1}{2} (x_n^2 + p_n^2 - 1) E_n(\mathbf{R}) \\ &= \frac{1}{2} \mathbf{P}_{\text{kin}}^\dagger \boldsymbol{\mu}^{-1} \mathbf{P}_{\text{kin}} + \frac{1}{F} \sum_n^F E_n(\mathbf{R}) \\ &\quad + \frac{1}{4F} \sum_{n,m}^F (x_n^2 + p_n^2 - x_m^2 - p_m^2) (E_n(\mathbf{R}) - E_m(\mathbf{R})), \end{aligned} \quad (6)$$

where $E_n(\mathbf{R})$ is the energy of the n th adiabatic state, and \mathbf{P}_{kin} is the kinematic momentum,⁵⁹

$$\mathbf{P}_{\text{kin}} = \mathbf{P} + \sum_{n,m}^F x_n p_m \mathbf{d}_{nm}(\mathbf{R}), \quad (7)$$

and $\mathbf{d}_{nm}(\mathbf{R}) = \langle n | \frac{\partial}{\partial \mathbf{R}} | m \rangle$ is the NAC vector between adiabatic states $|n(\mathbf{R})\rangle$ and $|m(\mathbf{R})\rangle$. The second equality in Eq. (6) is the result of an exact transformation called symmetrization,^{12,59} which has been found to yield more accurate classical analog dynamics.

The equations of motion for the dynamical variables of the system are now derived from the classical analog of the Hamiltonian in Eq. (6). The complete forms of the equations of motion are presented in the Appendix. The force terms necessary to integrate the equations of motion require input from electronic structure $[E_n(\mathbf{R}), \mathbf{g}_n = \frac{\partial E_n(\mathbf{R})}{\partial \mathbf{R}}$, and $\mathbf{d}_{nm}(\mathbf{R})]$, and these are obtained from on-the-fly *ab initio* calculations. It bears reiterating that these equations of motion effectively describe nonadiabatic dynamics in multiple coupled electronic states, despite the fact that time evolution is entirely classical in both the electronic and nuclear degrees of freedom. The ability to describe electronic coherences and the efficiency resulting from using classical dynamics render LSC a compelling option for large-scale *ab initio* NAMD simulations.

C. Sampling of initial conditions

The initial conditions for an LSC simulation span $2(F + N)$ variables, two phase space variables for every F electronic states and N nuclear modes. Those variables are sampled from A_W , where $\hat{A} = e^{-\beta \hat{H}} / Z$ is the thermal density operator with $\beta = 1/k_B T$ and $Z = \text{Tr} [e^{-\beta \hat{H}}]$. We assume that initially the nuclear and electronic variables can be treated independently, which allows us to sample the nuclear variables from $[\hat{\rho}_n]_W(\mathbf{R}, \mathbf{P})$ and the electronic ones from $[\hat{\rho}_e]_W(\mathbf{x}, \mathbf{p})$, where $\hat{\rho}_n$ and $\hat{\rho}_e$ are the nuclear and electronic density operators, respectively (i.e., $\text{Tr}_n[\hat{\rho}_n] = \text{Tr}_e[\hat{\rho}_e] = 1$), and the notation $[\dots]_W$ indicates a Wigner transform. Making a harmonic approximation for the N nuclear vibrational normal modes yields an analytic expression for the nuclear Wigner distribution,

$$[\hat{\rho}_n]_W = \prod_i^N \frac{1}{2\pi} \exp \left[-\tanh \left(\frac{\beta \omega_i}{2} \right) \left(\frac{P_i^2}{\mu_i \omega_i} + \mu_i \omega_i R_i^2 \right) \right], \quad (8)$$

where ω_i and μ_i are the frequency and the reduced mass of the i th vibrational mode, respectively. For the purpose of excited state dynamics simulations following an impulsive photo-excitation and within the Franck–Condon approximation, nuclear modes are assumed to be in thermal equilibrium on the electronic ground state for which $\{\omega_i\}$ and $\{\mu_i\}$ are calculated.

Similarly, the initial electronic state population is assumed to be 1 in the initial photo-excited state i and zero for all other states,

$$|\alpha\rangle\langle\alpha| = \delta_{i\alpha}, \quad (9)$$

where $\delta_{i\alpha}$ is a Kronecker delta. The sampling functions for each state thus form the system of equations

$$\begin{aligned} [|\alpha\rangle\langle\alpha|]_W &= 2^{F+1} \left(x_\alpha^2 + p_\alpha^2 - \frac{1}{2} \right) \exp \left[-\sum_k^F (x_k^2 + p_k^2) \right] \\ &= \delta_{i\alpha}. \end{aligned} \quad (10)$$

Solving Eq. (10), we obtain equations describing circles in phase space with radii, r_i ,

$$\begin{aligned} x_i^2 + p_i^2 &= r_i^2, \\ x_\alpha^2 + p_\alpha^2 &= \frac{1}{2} \quad \text{for } \alpha \neq i, \end{aligned} \quad (11)$$

where the values of the radii vary with F . For some F values, there exist two possible roots for the value of r_i^2 from Eq. (10). Following previous work, we chose the larger one of the two for the numerical stability of subsequent trajectory propagation.⁶⁰

III. COMPUTATIONAL WORKFLOW AND IMPLEMENTATION

The computational expense of the LSC nonadiabatic dynamics stems from the computation of the electronic structure variables, $E_n(\mathbf{R})$, \mathbf{g}_n , and $\mathbf{d}_{nm}(\mathbf{R})$, for all electronic states under consideration at a given nuclear geometry, \mathbf{R} . Here, we describe how the workload of the TeraChem electronic structure calculations is distributed among the available GPU compute nodes used in a simulation and how they communicate with our LSC code to propagate nuclear motion. The PySCES code is available open-source on GitHub,

along with appropriate documentation and input files to test the code.⁶¹

Figure 1 depicts the workflow of running an LSC simulation with the PySCES code, starting with an optimized geometry. The user provides PySCES with the location of a TeraChem normal mode analysis of the optimized structure, for which initial nuclear and electronic conditions are assigned from the Wigner distributions in Eqs. (8) and (11), respectively. Starting from this new randomized geometry and for all subsequent geometries throughout the simulation, both the gradients \mathbf{g}_i of each adiabatic energy surface and the NAC vectors \mathbf{d}_{ij} between two surfaces need to be computed. The signs of the NACs are then checked and, if necessary, corrected for random sign-flips between time steps that occur due to a change in phase of the wave function (see Sec. III A below). We then obtain updated nuclear and electronic variables

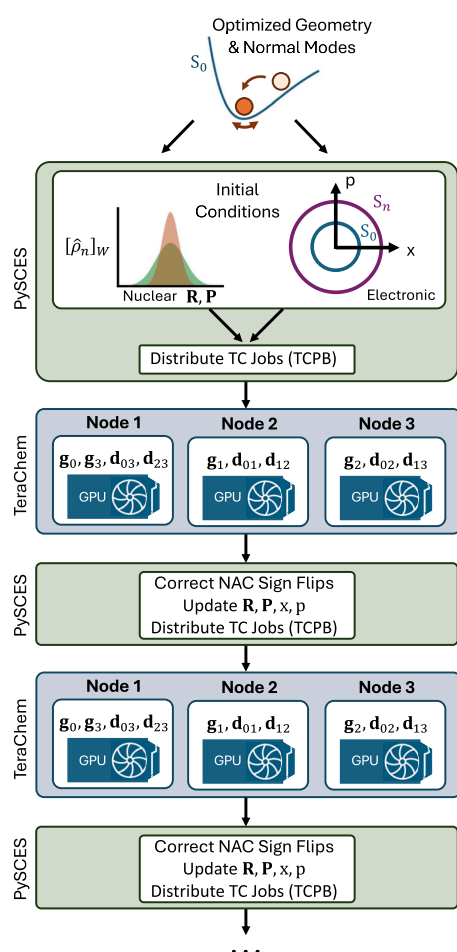


FIG. 1. Workflow of the PySCES code, an LSC–TeraChem interface. Initial nuclear positions and momenta of the normal modes are sampled from a Wigner distribution, and electronic mapping variables are initialized to populate the desired state as discussed in the text. The TeraChem gradient and NAC jobs are then distributed across GPUs. The NAC vectors are assessed, with a sign flip assigned as necessary, and then the nuclear and electronic positions and momenta are updated for the next step of the trajectory.

(position and momenta) via a Runge–Kutta fourth-order integrator as implemented in SciPy.^{62,63} The updated nuclear geometries are then used once again as inputs for electronic structure calculations, and the process continues. All relevant information, including the nuclear and electronic variables, gradients, NAC vectors, timings of the individual electronic structure calculations, and checkpoint files needed to restart the trajectory, is logged throughout the simulation as well.

Since the number of NAC vectors grows quadratically with the number of electronic states used in the simulation, the amount of electronic structure information needed for a single time step of a large molecular system can, very quickly, become prohibitively expensive to compute unless the work is parallelized across multiple computational nodes. Because each gradient and NAC computation can be performed independently from one another, we utilized the TeraChem Protocol Buffer (TCPB) library⁴⁵ to efficiently distribute the workload across all available GPUs within a computing network. As opposed to running TeraChem with multiple input files, TeraChem is instead executed in a server mode that receives and transmits “job” information as requested by PySCES via the TCPB library. A request to perform a particular job (gradient and NAC) is sent to the server along with the atomic coordinates, the basis set, the electronic structure method, and any setting otherwise used in a traditional TeraChem input file. Once all of the required gradients and NAC vectors are computed (energies E_i are also computed with each job), their results are summarized and used to update the molecular geometry.

One of the benefits of using TCPBs is that the electronic structure jobs can be efficiently parallelized in environments that have a limited or varying number of GPUs per compute node. For parallel jobs, multiple TeraChem servers can be run on a single or separate compute node, and each server can utilize any number of GPUs available on its respective node. PySCES will then distribute all computations to each of the servers on the same network. As the performance of TeraChem varies based on the hardware environment, such as the GPU model or the number of GPUs per compute node, it is left to the user to allocate TeraChem servers themselves and simply provide the IP addresses and port numbers to PySCES via the input file. PySCES will then distribute the gradient and NAC jobs in a round-robin manner across all specified servers, as depicted in each of the labeled “Node” blocks in Fig. 1.

Each job keeps track of the previously computed coefficients, both for the ground and excited states, and uses these files as inputs for the next job to speed up subsequent calculations. Completed jobs transmit minimal information from TeraChem through the network (only the energies, gradients, and NACs), and the remaining electronic structure properties of each job are copied from their local output files at every time step. Furthermore, because each computation is run as a separate job in TeraChem, GPUs only need to be initialized once with the server, a task that would otherwise introduce several seconds of inactive computation time at the start of every job.

For certain systems or geometries, electronic structure calculations of the ground or excited state can sometimes fail to converge, and then the NAMD simulation will need to be restarted. Additional TeraChem options, e.g., modified convergence or integral thresholds, can help limit or prevent these failures. We have implemented the ability to specify additional TeraChem runtime options for spe-

cific gradient/NAC jobs or for a specified number of initial time steps at the start of a simulation. These options can also request the calculation of additional electronic structure properties, such as ground and excited state point charges, throughout a simulation. We note that computing dynamical observables in the LSC method requires an ensemble of classical trajectories that can be evolved as independent simulations in parallel. The current implementation of PySCES does not parallelize this step.

A. Sign of the nonadiabatic coupling vector

We evaluate the NAC vectors using the states' adiabatic wave function representation as computed by TeraChem. Adiabatic wave functions are, however, defined only up to a phase factor, which can lead to phase inconsistencies between different time steps.⁶⁴ This phase inconsistency can also lead to a change in the sign of the calculated NACs, significantly affecting the resulting dynamics. In practice, TeraChem always uses real-valued state functions. Thus, during the propagation, the state functions (and the sign of the NAC vectors) can either keep their phase or change by a factor of -1 .

In order to preserve NAC sign consistency, we use two-level decision criteria during run time to estimate whether such a sign swap occurred for the NAC vectors between the current and previous time steps. The first level is to check if the NAC itself flipped the sign unexpectedly, while the second level is to check if the transition dipole moment (TDM), computed with the same adiabatic states, also flips the sign. We are able to combine both criteria by leveraging the property that if one of the adiabatic states flips a sign, then both the NAC and the TDM will contain the same sign flip.

Figure 2(a) illustrates the first level of the procedure in one dimension. In this scheme, we use the NAC from the two previous time steps to linearly extrapolate an estimation for the next time step. We then check whether the direction of the NAC from TeraChem matches our expectations by calculating the dot product of the TeraChem computed NAC vector with our extrapolated one. A positive dot product indicates that a sign-flip has not occurred between the current and previous time steps, and we continue the dynamics using the NAC from TeraChem. If instead, the direction of the NAC does not match our expectations (with a negative dot product), then we perform the second-level check by using the same extrapolation scheme to determine whether the TDM for the two corresponding states also flipped sign. If the transition dipole does not show a sign swap, we accept the sign swap of the computed NAC. However, if both the TDM and the NAC have negative dot products, then we correct the TeraChem computed NAC by a factor of -1 and continue the dynamics with the correction. We note that the extrapolated NAC is only used as a check for sign flips and is never used in the dynamics itself.

Figure 2(b) shows the first component of the NAC vector for photo-excited ethylene with and without the extrapolation correction scheme for a single (2, 2) SA-CASSCF trajectory with identical initial conditions. Without the correction, the NAC clearly undergoes discontinuous jumps due to the sign swaps, whereas the NAC with the correction smoothly evolves along the trajectory. Figures 2(c) and 2(d) show the corresponding S_2 and S_1 occupations along the trajectory, with substantial deviations observed after 10 fs.




FIG. 2. (a) For a single trajectory, we show the linear extrapolation scheme implemented in PySCES to determine the sign of the NAC vector components, d_{ij} . (b) The value of the first component of the NAC vector between S_0 and S_1 for ethylene shown without correction (as computed) and with correction. (c) Corresponding populations of the S_2 and (d) S_1 states with and without the sign swap for the NAC vector.

IV. RESULTS AND DISCUSSION

We next show results for LSC trajectories generated using PySCES with the TeraChem TCPB interface. First, we validate our implementation through comparison to the GAMESS (version 2023 R1) results using identical initial conditions for a single ethylene trajectory. Next, we report timings for larger systems.

All TeraChem jobs in this paper were performed with version 1.9-2023.09-dev and ran as a docker image using Singularity 3.8.3. Computation nodes utilized a pair of NVIDIA A100 GPUs with 40 GB of VRAM in conjunction with dual 28 core Intel Xeon Gold 6330 CPUs. All TeraChem servers were initiated with both A100 GPUs on a single node.

A. Comparison of ethylene dynamics with GAMESS and TeraChem

To verify our implementation within PySCES with the TeraChem interface, we compared the dynamics of ethylene using (2, 2) SA-CASSCF with both GAMESS and TeraChem drivers. An

initial nuclear geometry was generated from a Wigner distribution over normal mode coordinates and frequencies determined from a ground state CAM-B3LYP/6-31G(d) computation. Electronic phase space variables were sampled from the circles in phase space defined in Eq. (11).⁶⁰ The S_0 , S_1 , and S_2 states were used in this assignment with an initial photoexcitation to S_1 , i.e., $F = 3$, and the radii, $r_{S_1} = 1.559$, $r_{S_2} = r_{S_0} = 0.707$. Both trajectories with GAMESS and TeraChem used the exact same initial conditions for each trajectory.

The resulting state energies and state populations from a single trajectory are shown in Fig. 3. The dynamics are in excellent agreement, both in energy and population, until 30 fs, when small numerical differences in the electronic structure output lead to visible differences in their trajectories. We verified that these differences remain even when employing full double precision and an integral threshold of 10^{-20} in TeraChem. Our testing revealed that, at least for ethylene, small differences in nuclear position can lead to large differences in NAC vectors. Take, for example, the avoided crossing that occurs around 46 fs when the S_1 and S_2 energies approach each other in Fig. 3. At this simulation frame, the minimum RMSD⁶⁵ between the geometry produced by the GAMESS and TeraChem driven trajectories is less than 0.03 Å, whereas the magnitude of their NAC vectors differed by over 30%, thus sending the TeraChem and GAMESS simulations on slightly different trajectories that eventually diverge from each other. However, we expect that, when averaged over multiple trajectories with varying initial conditions, the average rate of population change resulting from simulations utilizing both drivers should be very similar.

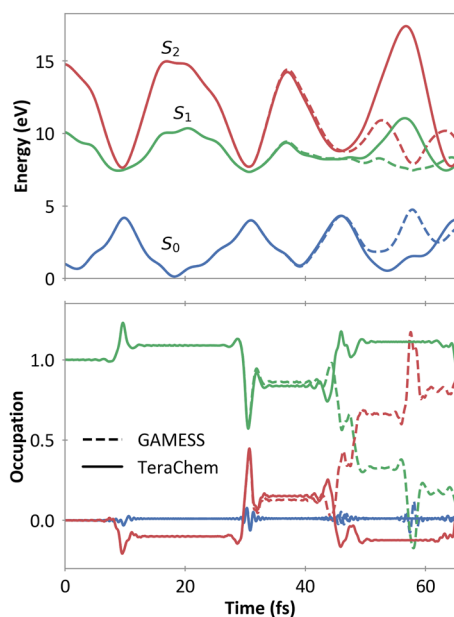


FIG. 3. Energy and occupations of the S_0 , S_1 , and S_2 states of ethylene for a single LSC trajectory with identical initial conditions generated by PySCES. The deviation between GAMESS and TeraChem due to a buildup of small numerical differences becomes apparent after 35 fs.

B. Computational timings

We next showcase the performance of PySCES for two larger systems using the TeraChem TCPB library. First, we compare timings for 1 and 2 nodes for solvated fulvene as we increase the size of the explicit solvent environment and demonstrate the scaling of our NAMD implementation as a function of system size. Next, we model the charge transfer in an excited zinc phthalocyanine- C_{60} donor-acceptor complex and provide insights into the computational effort needed to run NAMD LSC simulations for systems consisting of over 100 heavy atoms.

1. Fulvene in acetonitrile

The fulvene molecule, which is a well-studied benchmark system within the nonadiabatic dynamics community, possesses a sloped S_0/S_1 conical intersection along the methylene torsion coordinate that facilitates population transfer between the S_0 and S_1 surfaces after photo exciting to S_1 .^{45,66,67} In most studies of the conical intersections and nonadiabatic dynamics of fulvene, a multi-configurational approach with at least the π system is employed, i.e., a CASSCF(6, 6) wave function description. However, spin-flip TDDFT has also been employed to characterize the conical intersection,⁶⁸ as well as TDDFT within the Tamm-Dancoff approximation (TDA) to explore the region of potential around a conical intersection of fulvene as it rearranges to benzene.⁶⁹ Here, we do not perform a detailed study of the nonadiabatic dynamics of fulvene, and so we are not particularly concerned with the choice of electronic structure method and employ TD-CAM-B3LYP/6-31G(d) within the TDA to model the S_0 , S_1 , and S_2 states of fulvene in acetonitrile within the LSC approach.

The character of a conical intersection, as well as the nonadiabatic dynamics of chromophores through such intersections, will be affected by a solvent environment. However, such effects are rarely considered within NAMD simulations due to the computational expense of including explicit solvent molecules in on-the-fly dynamics. With the increased computational efficiency of TeraChem for large-scale electronic structure calculations, solvated chromophore NAMD, although still computationally formidable, becomes feasible within PySCES. We next show the performance of PySCES employing TeraChem as we increase the amount of explicit solvent included in the electronic structure calculations.

Individual fulvene and acetonitrile molecules were geometry optimized in isolation with TeraChem using CAM-B3LYP/6-31G(d). The PACKMOL⁷⁰ program was then used to generate a 20 \AA^3 solvated box of 200 acetonitrile molecules with the fulvene at the center, from which the closest 20 solvent molecules to the fulvene solute were selected to generate configurations for timings. Solvent molecules were then sorted by their distance from the fulvene center to generate configurations with fewer solvent molecules. A final ground state geometry optimization and normal-mode analysis were performed on each of the fulvene-solvent clusters, for which the geometries and normal mode properties of each cluster were used to generate initial conditions for the NAMD simulations according to Eq. (8). The smallest system, consisting of a single acetonitrile solvent molecule and fulvene, as well as the largest 20 solvent molecule-fulvene subsystem, are depicted as insets in Fig. 4(a).

In total, 15 chromophore-solvent subsystems were constructed and optimized, and the LSC NAMD was simulated for at least 1 fs

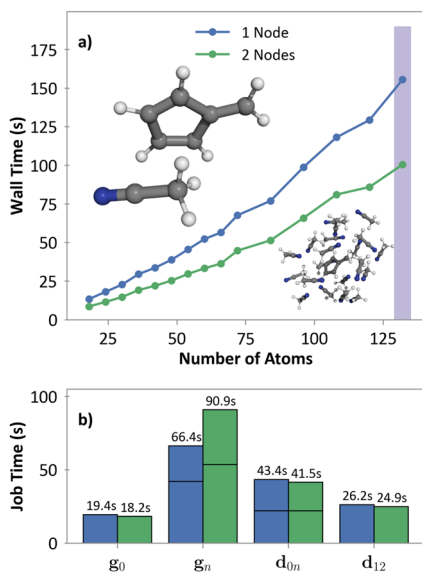


FIG. 4. (a) Average wall time for a single dynamics step for solvated fulvene using S_0 , S_1 , and S_2 states computed with CAM-B3LYP/6-31G(d). The increasing number of atoms is due to the increase in the number of acetonitrile solvent molecules included in the calculation. (b) Timings for the fully solvated system (136 atoms) broken down into the different gradient g_n and nonadiabatic coupling d_{mn} calculations over one (blue) and two (green) nodes. The total job time for a particular type of calculation is shown atop each set of bars.

with a 1 a.u. time step. CAM-B3LYP/6-31G(d) was also used for all simulations, and S_1 and S_2 excited states were computed with TDDFT within the TDA. The last 40 time steps of all simulations were used in our benchmarks, averaging the computational time needed to compute individual gradients and NAC vectors between the three states as well as the total wall time for a NAMD time step. Each time step within a simulation uses the ground state molecular orbitals and excited state guess coefficients of the previous geometry as the initial guess for the next TeraChem calculation. To ensure that our statistics are not influenced by time steps that do not have these initial guesses available, we eliminated the first few sets of timings of each simulation from our averaging. We compared the performance of using TCPBs to distribute work over multiple computers on the same network, running all 15 simulations on both a single and dual node setup.

The resulting wall clock timings for one and two compute nodes are shown in Fig. 4(a), where we show that as we increase the number of atoms in the solvent region, the computational time increases linearly, at least up to system sizes of 136 atoms in total. Figure 4(b) further decomposes the average time spent by TeraChem to compute each of the gradient and NAC vectors for the largest 20 solvent molecule-fulvene system, with the wall timings highlighted in purple in Fig. 4(a). When two nodes are used, all gradient and NAC “job times” do not sum to the “wall time” as indicated by the green line in Fig. 4(a); the wall time measures the real-world time as seen by the user while all of the TeraChem jobs are simultaneously running, whereas the job time is the cumulative time measured by the individual nodes themselves.

We find that for the job timings for solvated fulvene, shown in Fig. 4(b), excited state gradients take the longest time to compute, with gradients for the S_1 state taking the longest as subsequent calculations of higher lying states can use previous excited state vectors to reduce job time. As discussed in Sec. III, each gradient and NAC vector are computed as separate jobs within TeraChem, with the ground state orbital and excited state coefficients of the previous job used in subsequent calculations and saved to the server’s local storage. Therefore, each TeraChem job that is performed after the S_1 gradient job uses these guess orbitals and coefficients. The NAC jobs are not much more expensive than higher energy (S_2 or greater) excited state gradient jobs. In Fig. 4(b), we find that NAC vectors involving two excited states take longer time to compute than those involving the ground state, a trend also observed for the charge transfer system discussed later.

Figure 4(a) also shows that using two compute nodes can decrease the total wall time for a single simulation time step, although not perfectly with double the performance. To avoid having to save the ground and excited state coefficient files to a network shared storage system for every time step and every TeraChem excited state job, each server maintains its own local copy of the coefficient files. This strategy further ensures that future users of PySCES, who may wish to distribute their TeraChem jobs over dissimilar hardware, are not faced with the possibility that the (binary) coefficient files are incompatible between the different computing environments. As a result, each TeraChem server generates its own ground and excited state coefficient files.

For the dual node set up, the S_1 gradient (computed on node 2) will not have a ground state orbital file since it is the first job that is computed on this node. As evident in Fig. 4(b), the average S_1 gradient job time using two nodes is indeed larger than that of the equivalent S_1 gradient job using only one compute node. Similarly, the S_2 gradient (computed on node 1) does not have the excited state coefficient file since the previous gradient job on the same node was only for the S_0 state, also leading to a larger job time in Fig. 4(b) compared to the single node setup. Nevertheless, a performance gain of 153% was observed when using two compute nodes over just one. We expect that the relative performance gains using two nodes should get closer to the 200% limit when more than two excited states are involved, such as we describe below for the zinc-phthalocyanine and fullerene charge transfer systems.

2. Zinc-phthalocyanine and fullerene charge transfer

Another system that we applied the PySCES LSC-TeraChem interface to is the charge transfer at the interface of zinc phthalocyanine (ZnPc) and fullerene (C_{60}), a system of potential interest in organic photovoltaics.^{71–73} The ultrafast charge transfer dynamics of this system have previously been simulated employing TDDFT with the Ehrenfest and FSSH approaches.^{74,75} For the face-on configuration, Liu *et al.* found that charge transfer from a higher energy locally excited state on the ZnPc to lower energy charge transfer states with the electron transferred to the C_{60} acceptor is thermodynamically favorable, with exciton charge transfer completed within 100 fs.⁷⁵

For the face-on configuration of ZnPc- C_{60} , we optimized the geometry with CAM-B3LYP/6-31G and performed a normal mode analysis on the full 117 atom system. TDDFT calculations were performed with the same level of theory within the TDA. Our excited

state characterization showed the transition from the ground state to the S_4 state was the brightest and localized on ZnPc, in agreement with the analysis of Liu *et al.* We began the LSC NAMD trajectory with full occupation of the S_4 state at the optimized geometry, initiating the dynamics with only the momenta sampled from a Wigner distribution and including NACs between excited states S_0 through S_4 . In order to demonstrate the ability of LSC to simulate charge transfer between the two species and to measure the computational cost of performing the necessary electronic structure calculations, removing the Wigner sampling over nuclear positions allowed us to achieve these goals with fewer time steps. Initialization of electronic variables was performed to occupy the bright S_4 state only, and the trajectory was then propagated with a 1 a.u. step size.

For our single LSC trajectory depicted in Fig. 5, we found the charge transfer to occur from the S_4 state to the S_1 state within 25 fs. Figure 5(a) shows the excited state and ground state density difference plots at different times during the trajectory, where we weight the density difference according to the occupation of the excited

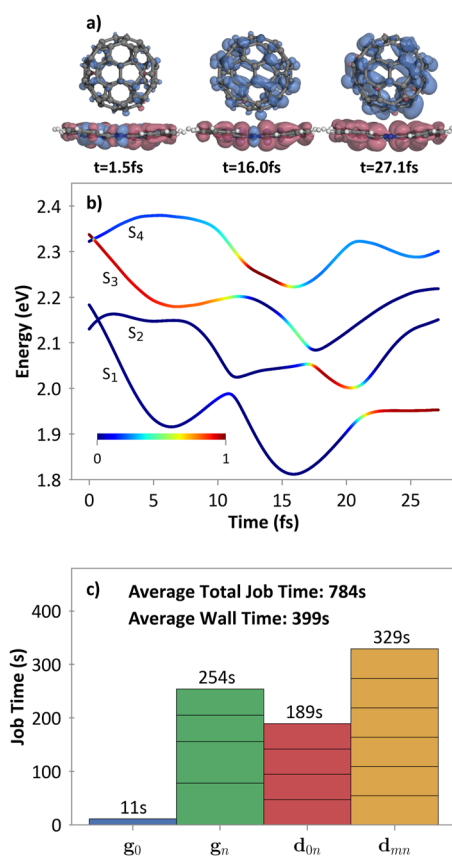


FIG. 5. (a) ZnPc-C₆₀ excited state and ground state density differences at three different times in the LSC NAMD trajectory generated with PySCES. C₆₀ accepts the charge after a local, bright photoexcitation of the zinc phthalocyanine to the higher lying S_4 state. (b) TD-CAM-B3LYP/6-31G excitation energies computed during the LSC trajectory with the corresponding occupation shown in the color bar. (c) Timings shown for electronic structure gradient \mathbf{g}_n and nonadiabatic coupling \mathbf{d}_{mn} calculations performed with CAM-B3LYP/6-31G averaged over all steps of LSC dynamics that included five electronic states: S_0 through S_4 .

state, as shown in Fig. 5(b). For example, at $t = 1.5$ fs, where a crossing between the S_4 and S_3 surfaces occurred, the population was predominantly on the S_3 surface, so the density difference shown is primarily between S_3 and S_0 . Similarly, at the end of the trajectory, the population is mostly in the S_1 state, and the density difference shown is mostly that for S_1 and S_0 . At $t = 16$ fs shortly after the S_4 and S_3 avoided surface crossing, this state has mixed locally excited and charge transfer character, and the occupation is shared between S_4 and S_3 . Figure 5(b) shows the TDDFT energies of the four lowest lying excited states, colored according to the occupation during the trajectory. Surface crossings and avoided surface crossings occur throughout the trajectory, with state mixings often leading to population transfer between states, notably at $t \approx 14$ fs, $t \approx 17$ fs, and $t \approx 21$ fs. Note that these LSC NAMD dynamics are for a single trajectory, and an ensemble of LSC NAMD trajectories would be required to compute the LSC correlation function.

For the ZnPc-C₆₀ donor-acceptor system, modeled with 661 basis functions and four excited states, we show the job time for electronic structure contributions averaged across the trajectory in Fig. 5(c). The calculation of the excited state gradient and the NAC vectors between excited states takes up the majority of the computation time. Because of the number of states used throughout the simulation, about 66% of the computational job time was spent computing the 10 NAC vectors. We note that for this trajectory, we see little participation from the S_0 state in the dynamics, suggesting that we may be able to eliminate this state from future simulations of the dynamics, thereby reducing the number of NAC vectors to be computed. The total job time for a single LSC step averaged across the trajectory was 784 s, but parallelization over two GPU nodes brought the average wall time per step to 399 s, a performance gain of 196%. Similarly, with the fulvene system in Fig. 4(b), which was also parallelized over two compute nodes, the first two excited state gradients, \mathbf{g}_1 and \mathbf{g}_2 , take the longest to compute as they are the first and only jobs that solve for the TDDFT excited state coefficients. However, because there are many more jobs after the S_1 and S_2 gradient calculations that use the excited state coefficients than in the fulvene system, the computational performance is nearly double that of a single node. We conclude with a reminder that the simulations reported here demonstrate the abilities of the PySCES code and document timings for a single trajectory; however, in an actual investigation of excited state dynamics, an ensemble of independent trajectories would be evolved to compute the corresponding LSC correlation function.

V. CONCLUSIONS

We introduce PySCES, an LSC-TeraChem Python interface code to enable GPU-accelerated NAMD simulations using classical dynamics trajectories within an extended phase space of nuclear and electronic variables. Our interface makes use of the TeraChem protocol buffer client-server model and can be parallelized over multiple GPU nodes, opening the door for semiclassical NAMD simulations of excited state dynamics in the condensed phase.

We detail our computational workflow and implementation, including a linear extrapolation scheme for determining the sign of the NAC vector. Computational timings are presented for two NAMD benchmark systems. The first system is fulvene solvated in acetonitrile, where we show a linear scaling in compute time as

we increase the amount of solvent included in the electronic structure calculations, and we show the performance of the code over one and two GPU nodes. The second benchmark system is a zinc-phthalocyanine donor face-on with a fullerene acceptor. We show the state occupations during the course of a single LSC NAMD trajectory, which shows population transfer from the locally excited bright state to lower energy dark charge transfer states over 25 fs. The corresponding timings show that the computation of the NAC vectors between excited states takes up a large portion of the total computation time.

Moving forward, we will continue to benchmark LSC NAMD against molecular systems studied using other on-the-fly nonadiabatic methods such as AIMS and surface hopping, with an emphasis on accuracy and numerical efficiency. We also plan to explore new and more efficient integrators and approaches to compute the NAC,⁷⁶ including its replacement with time-derivative coupling followed by its approximation by wavefunction interpolation, a commonly employed strategy in AIMS and surface hopping programs.^{77,78} Implementing these capabilities could further increase the efficiency of LSC NAMD simulations for complex excited state molecular phenomena in the condensed phase on an ultrafast timescale.

ACKNOWLEDGMENTS

C.I. and C.M. acknowledge the support from a MURI award through the Air Force Office of Scientific Research, Grant No. FA9550-22-1-0317. T.T. acknowledges the support from the Elite Study Program “Biological Physics” of the Elite Network of Bavaria. K.M. acknowledges the support from the Mark A. Ratner Postdoctoral Fellowship and the Northwestern University International Institute for Nanotechnology (IIN). N.A. acknowledges the support from NSF Grant No. CHE-2404809.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Christopher A. Myers: Conceptualization (equal); Data curation (equal); Formal analysis (equal); Software (equal); Writing – original draft (equal); Writing – review & editing (equal). **Ken Miyazaki:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Software (equal); Writing – original draft (equal); Writing – review & editing (equal). **Thomas Trepl:** Conceptualization (equal); Data curation (equal); Formal analysis (equal); Funding acquisition (supporting); Software (equal); Writing – original draft (equal). **Christine M. Isborn:** Conceptualization (equal); Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Writing – original draft (equal); Writing – review & editing (equal). **Nandini Ananth:** Conceptualization (equal); Funding acquisition (equal); Project administration (equal); Resources (equal); Supervision (equal); Writing – original draft (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request, along with files and software available on the GitHub repository for the PySCES code, <http://doi.org/10.5281/zenodo.11985696>.

APPENDIX: NONADIABATIC EQUATIONS OF MOTION

The equations of motion for the nuclear phase space variables, as well as the electronic mapping variables, are obtained from the classical analog mapping Hamiltonian in Eq. (6),

$$\dot{x}_n = \frac{p_n}{F} \sum_m^F (E_n(\mathbf{R}) - E_m(\mathbf{R})) + \sum_m^F x_m \mathbf{d}_{mn}(\mathbf{R}) \cdot \frac{\mathbf{P}_{\text{kin}}}{\boldsymbol{\mu}}, \quad (\text{A1})$$

$$\dot{p}_n = -\frac{x_n}{F} \sum_m^F (E_n(\mathbf{R}) - E_m(\mathbf{R})) + \sum_m^F p_m \mathbf{d}_{nm}(\mathbf{R}) \cdot \frac{\mathbf{P}_{\text{kin}}}{\boldsymbol{\mu}}, \quad (\text{A2})$$

$$\dot{\mathbf{R}} = \frac{\mathbf{P}_{\text{kin}}}{\boldsymbol{\mu}}, \quad (\text{A3})$$

$$\begin{aligned} \dot{\mathbf{P}}_{\text{kin}} &= \sum_{n,m}^F (\dot{x}_n p_m + x_n \dot{p}_m) \mathbf{d}_{nm}(\mathbf{R}) \\ &\quad - \frac{1}{4F} \sum_{n,m}^F (x_n^2 + p_n^2 - x_m^2 - p_m^2) (\mathbf{g}_n(\mathbf{R}) - \mathbf{g}_m(\mathbf{R})) - \frac{1}{F} \sum_n^F \mathbf{g}_n(\mathbf{R}), \end{aligned} \quad (\text{A4})$$

where $\mathbf{g}_n(\mathbf{R}) = \frac{\partial E_n(\mathbf{R})}{\partial \mathbf{R}}$ is the gradient of the adiabatic state $|n\rangle$. In deriving Eq. (15), the derivative of the nonadiabatic coupling vector $\mathbf{d}_{nm}(\mathbf{R})$ was tactically removed by utilizing the chain rule.³⁹ By recognizing $\Delta \mathbf{P} = \sum_{n,m}^F x_n p_m \mathbf{d}_{nm}(\mathbf{R}) \equiv f(\mathbf{x}, \mathbf{p}, \mathbf{R})$ and using V_{eff} to represent the electronic part of our Hamiltonian, we obtain

$$\begin{aligned} \dot{\mathbf{P}}_{\text{kin}} &= \dot{\mathbf{P}} + \frac{\partial \Delta \mathbf{P}}{\partial t} \\ &= -\frac{\mathbf{P}_{\text{kin}}}{\boldsymbol{\mu}} \cdot \frac{\partial f}{\partial \mathbf{R}} - \frac{\partial V_{\text{eff}}}{\partial \mathbf{R}} + \dot{\mathbf{x}} \cdot \frac{\partial f}{\partial \mathbf{x}} + \dot{\mathbf{p}} \cdot \frac{\partial f}{\partial \mathbf{p}} + \dot{\mathbf{R}} \frac{\partial f}{\partial \mathbf{R}} \\ &= -\frac{\partial V_{\text{eff}}}{\partial \mathbf{R}} + \sum_{n,m}^F (\dot{x}_n p_m + x_n \dot{p}_m) \mathbf{d}_{nm}(\mathbf{R}), \end{aligned} \quad (\text{A5})$$

which is in agreement with Eq. (15).

REFERENCES

- J. C. Tully, “Perspective: Nonadiabatic dynamics theory,” *J. Chem. Phys.* **137**, 22A301 (2012).
- E. Tapavicza, G. D. Bellchambers, J. C. Vincent, and F. Furche, “*Ab initio* non-adiabatic molecular dynamics,” *Phys. Chem. Chem. Phys.* **15**, 18336–18348 (2013).
- B. F. E. Curchod, W. J. Glover, and T. J. Martínez, “SSAIMS—Stochastic-Selection *ab initio* multiple spawning for efficient nonadiabatic molecular dynamics,” *J. Phys. Chem. A* **124**, 6133–6143 (2020).
- F. Agostini and B. F. E. Curchod, “Different flavors of nonadiabatic molecular dynamics,” *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **9**, e1417 (2019).

- ⁵T. R. Nelson, A. J. White, J. A. Bjorggaard, A. E. Sifain, Y. Zhang, B. Neugebauer, S. Fernandez-Alberti, D. Mozyrsky, A. E. Roitberg, and S. Tretiak, "Non-adiabatic excited-state molecular dynamics: Theory and applications for modeling photophysics in extended molecular materials," *Chem. Rev.* **120**, 2215–2287 (2020).
- ⁶H. R. Larsson, "A tensor network view of multilayer multiconfiguration time-dependent Hartree methods," *Mol. Phys.* **122**, e2306881 (2024).
- ⁷B. F. E. Curchod and T. J. Martínez, "Ab initio nonadiabatic quantum molecular dynamics," *Chem. Rev.* **118**, 3305–3336 (2018).
- ⁸M. D. Hack and D. G. Truhlar, "A natural decay of mixing algorithm for non-born–oppenheimer trajectories," *J. Chem. Phys.* **114**, 9305 (2001).
- ⁹M. J. Bedard-Hearn, R. E. Larsen, and B. J. Schwartz, "Mean-field dynamics with stochastic decoherence (MF-SD): A new algorithm for nonadiabatic mixed quantum/classical molecular-dynamics simulations with nuclear-induced decoherence," *J. Chem. Phys.* **123**, 234106 (2005).
- ¹⁰A. Jain, E. Alguire, and J. E. Subotnik, "An efficient, augmented surface hopping algorithm that includes decoherence for use in large-scale simulations," *J. Chem. Theory Comput.* **12**, 5256–5268 (2016).
- ¹¹R. Tempelaar and D. R. Reichman, "Generalization of fewest-switches surface hopping for coherences," *J. Chem. Phys.* **148**, 102309 (2018).
- ¹²S. J. Cotton and W. H. Miller, "Symmetrical windowing for quantum states in quasi-classical trajectory simulations: Application to electronically non-adiabatic processes," *J. Chem. Phys.* **139**, 234112 (2013).
- ¹³J. J. Talbot, M. Head-Gordon, and S. J. Cotton, "The symmetric quasi-classical model using on-the-fly time-dependent density functional theory within the Tamm–Dancoff approximation," *Mol. Phys.* **121**, e2153761 (2023).
- ¹⁴W. H. Miller, "The semiclassical initial value Representation: A potentially practical way for adding quantum effects to classical molecular dynamics simulations," *J. Phys. Chem. A* **105**, 2942–2955 (2001).
- ¹⁵Q. Shi and E. Geva, "A relationship between semiclassical and centroid correlation functions," *J. Chem. Phys.* **118**, 8173 (2003).
- ¹⁶K. Miyazaki and N. Ananth, "Nonadiabatic simulations of photoisomerization and dissociation in ethylene using *ab initio* classical trajectories," *J. Chem. Phys.* **159**, 124110 (2023).
- ¹⁷P. Slavíček and T. J. Martínez, "*Ab initio* floating occupation molecular orbital-complete active space configuration interaction: An efficient approximation to CASSCF," *J. Chem. Phys.* **132**, 234102 (2010).
- ¹⁸J. Olsen, "The CASSCF method: A perspective and commentary," *Int. J. Quantum Chem.* **111**, 3267–3272 (2011).
- ¹⁹D. Hollas, L. Šišťák, E. G. Hohenstein, T. J. Martínez, and P. Slavíček, "Nonadiabatic *ab initio* molecular dynamics with the floating occupation molecular orbital-complete active space configuration interaction method," *J. Chem. Theory Comput.* **14**, 339–350 (2018).
- ²⁰B. F. E. Curchod, U. Rothlisberger, and I. Tavernelli, "Trajectory-based nonadiabatic dynamics with time-dependent density functional theory," *ChemPhysChem* **14**, 1314–1340 (2013).
- ²¹L. Yue, Y. Liu, and C. Zhu, "Performance of TDDFT with and without spin-flip in trajectory surface hopping dynamics: *cis*–*trans* azobenzene photoisomerization," *Phys. Chem. Chem. Phys.* **20**, 24123–24139 (2018).
- ²²L. D. M. Peters, J. Kussmann, and C. Ochsenfeld, "Combining graphics processing units, simplified time-dependent density functional theory, and finite-difference couplings to accelerate nonadiabatic molecular dynamics," *J. Phys. Chem. Lett.* **11**, 3955–3961 (2020).
- ²³A. D. Laurent and D. Jacquemin, "TD-DFT benchmarks: A review," *Int. J. Quantum Chem.* **113**, 2019–2039 (2013).
- ²⁴F. Santoro and D. Jacquemin, "Going beyond the vertical approximation with time-dependent density functional theory," *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **6**, 460–486 (2016).
- ²⁵S. Kümmel, "Charge-transfer excitations: A challenge for time-dependent density functional theory that has been met," *Adv. Energy Mater.* **7**, 1700440 (2017).
- ²⁶J. Liang, X. Feng, D. Hait, and M. Head-Gordon, "Revisiting the performance of time-dependent density functional theory for electronic excitations: Assessment of 43 popular and recently developed functionals from rungs one to four," *J. Chem. Theory Comput.* **18**, 3460–3473 (2022).
- ²⁷J. Q. B. G. Levine, C. Ko, and T. J. Martínez, "Conical intersections and double excitations in time-dependent density functional theory," *Mol. Phys.* **104**, 1039–1051 (2006).
- ²⁸J. T. Taylor, D. J. Tozer, and B. F. E. Curchod, "On the description of conical intersections between excited electronic states with LR-TDDFT and ADC(2)," *J. Chem. Phys.* **159**, 214115 (2023).
- ²⁹N. Minezawa and M. S. Gordon, "Optimizing conical intersections by spin-flip density functional theory: Application to ethylene," *J. Phys. Chem. A* **113**, 12749–12753 (2009).
- ³⁰D. Casanova, "Avoided crossings, conical intersections, and low-lying excited states with a single reference method: The restricted active space spin-flip configuration interaction approach," *J. Chem. Phys.* **137**, 084105 (2012).
- ³¹S. Gozem, A. I. Krylov, and M. Olivucci, "Conical intersection and potential energy surface features of a model retinal chromophore: Comparison of EOM-CC and multireference methods," *J. Chem. Theory Comput.* **9**, 284–292 (2013).
- ³²Y. Yang, L. Shen, D. Zhang, and W. Yang, "Conical intersections from particle–particle random phase and Tamm–Dancoff approximations," *J. Phys. Chem. Lett.* **7**, 2407–2411 (2016).
- ³³C. Bannwarth, J. K. Yu, E. G. Hohenstein, and T. J. Martínez, "Hole–hole Tamm–Dancoff-approximated density functional theory: A highly efficient electronic structure method incorporating dynamic and static correlation," *J. Chem. Phys.* **153**, 024110 (2020).
- ³⁴I. S. Ufimtsev and T. J. Martínez, "Quantum chemistry on graphical processing units. 3. Analytical energy gradients, geometry optimization, and first principles molecular dynamics," *J. Chem. Theory Comput.* **5**, 2619–2628 (2009).
- ³⁵S. Seritan, C. Bannwarth, B. S. Fales, E. G. Hohenstein, C. M. Isborn, S. I. L. Kokkila-Schumacher, X. Li, F. Liu, N. Luehr, J. W. Snyder, Jr. *et al.*, "TeraChem: A graphical processing unit-accelerated electronic structure package for large-scale *ab initio* molecular dynamics," *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **11**, e1494 (2021).
- ³⁶J. W. Snyder, Jr., E. G. Hohenstein, N. Luehr, and T. J. Martínez, "An atomic orbital-based formulation of analytical gradients and nonadiabatic coupling vector elements for the state-averaged complete active space self-consistent field method on graphical processing units," *J. Chem. Phys.* **143**, 154107 (2015).
- ³⁷B. S. Fales and B. G. Levine, "Nanoscale multireference quantum chemistry: Full configuration interaction on graphical processing units," *J. Chem. Theory Comput.* **11**, 4708–4716 (2015).
- ³⁸E. G. Hohenstein, M. E. F. Bouduban, C. Song, N. Luehr, I. S. Ufimtsev, and T. J. Martínez, "Analytic first derivatives of floating occupation molecular orbital-complete active space configuration interaction on graphical processing units," *J. Chem. Phys.* **143**, 014111 (2015).
- ³⁹C. M. Isborn, N. Luehr, I. S. Ufimtsev, and T. J. Martínez, "Excited-state electronic structure with configuration interaction singles and Tamm–Dancoff time-dependent density functional theory on graphical processing units," *J. Chem. Theory Comput.* **7**, 1814–1823 (2011).
- ⁴⁰C. M. Isborn, A. W. Götz, M. A. Clark, R. C. Walker, and T. J. Martínez, "Electronic absorption spectra from MM and *ab initio* QM/MM molecular dynamics: Environmental effects on the absorption spectrum of photoactive yellow protein," *J. Chem. Theory Comput.* **8**, 5092–5106 (2012).
- ⁴¹J. W. J. Snyder, B. F. E. Curchod, and T. J. Martínez, "GPU-accelerated state-averaged complete active space self-consistent field interfaced with *ab initio* multiple spawning unravels the photodynamics of provitamin D3," *J. Phys. Chem. Lett.* **7**, 2444–2449 (2016).
- ⁴²A. Sisto, C. Stross, M. W. van der Kamp, M. O'Connor, S. McIntosh-Smith, G. T. Johnson, E. G. Hohenstein, F. R. Manby, D. R. Glowacki, and T. J. Martínez, "Atomistic non-adiabatic dynamics of the LH2 complex with a GPU-accelerated *ab initio* exciton model," *Phys. Chem. Chem. Phys.* **19**, 14924–14936 (2017).
- ⁴³L. M. Ibele, Y. Lassmann, T. J. Martínez, and B. F. E. Curchod, "Comparing (stochastic-selection) *ab initio* multiple spawning with trajectory surface hopping for the photodynamics of cyclopropanone, fulvene, and dithiane," *J. Chem. Phys.* **154**, 104110 (2021).
- ⁴⁴B. Rana, E. G. Hohenstein, and T. J. Martínez, "Simulating the excited-state dynamics of polaritons with *ab initio* multiple spawning," *J. Phys. Chem. A* **128**, 139–151 (2024).

- ⁴⁵V. W. D. Cruzeiro, Y. Wang, E. Pieri, E. G. Hohenstein, and T. J. Martínez, “TeraChem protocol buffers (TCPB): Accelerating QM and QM/MM simulations with a client–server model,” *J. Chem. Phys.* **158**, 044801 (2023).
- ⁴⁶H. Wang, X. Sun, and W. H. Miller, “Semiclassical approximations for the calculation of thermal rate constants for chemical reactions in complex molecular systems,” *J. Chem. Phys.* **108**, 9726–9736 (1998).
- ⁴⁷E. Wigner, “On the quantum correction for thermodynamic equilibrium,” *Phys. Rev.* **40**, 749 (1932).
- ⁴⁸H.-D. Meyer, U. Manthe, and L. S. Cederbaum, “The multi-configurational time-dependent Hartree approach,” *Chem. Phys. Lett.* **165**, 73–78 (1990).
- ⁴⁹M. Ben-Nun and T. J. Martínez, “*Ab initio* quantum molecular dynamics,” in *Advances in Chemical Physics* (John Wiley & Sons, Ltd., 2002), pp. 439–512.
- ⁵⁰S. Malpathak, M. S. Church, and N. Ananth, “A semiclassical framework for mixed quantum classical dynamics,” *J. Phys. Chem. A* **126**, 6359–6375 (2022).
- ⁵¹N. Ananth, C. Venkataraman, and W. H. Miller, “Semiclassical description of electronically nonadiabatic dynamics via the initial value representation,” *J. Chem. Phys.* **127**, 84114 (2007).
- ⁵²M. S. Church, T. J. Hele, G. S. Ezra, and N. Ananth, “Nonadiabatic semiclassical dynamics in the mixed quantum-classical initial value representation,” *J. Chem. Phys.* **148**, 102326 (2018).
- ⁵³S. Malpathak and N. Ananth, “A linearized semiclassical dynamics study of the multiquantum vibrational relaxation of NO scattering from a Au(111) surface,” *J. Phys. Chem. Lett.* **15**, 794–801 (2024).
- ⁵⁴H. D. Meyer and W. H. Miller, “A classical analog for electronic degrees of freedom in nonadiabatic collision processes,” *J. Chem. Phys.* **70**, 3214 (1979).
- ⁵⁵G. Stock and M. Thoss, “Semiclassical description of nonadiabatic quantum dynamics,” *Phys. Rev. Lett.* **78**, 578 (1997).
- ⁵⁶M. Thoss and G. Stock, “Mapping approach to the semiclassical description of nonadiabatic quantum dynamics,” *Phys. Rev. A* **59**, 64 (1999).
- ⁵⁷J. Liu, “A unified theoretical framework for mapping models for the multi-state Hamiltonian,” *J. Chem. Phys.* **145**, 204105 (2016).
- ⁵⁸J. E. Runeson and J. O. Richardson, “Spin-mapping approach for nonadiabatic molecular dynamics,” *J. Chem. Phys.* **151**, 044119 (2019).
- ⁵⁹S. J. Cotton, R. Liang, and W. H. Miller, “On the adiabatic representation of Meyer–Miller electronic–nuclear dynamics,” *J. Chem. Phys.* **147**, 064112 (2017).
- ⁶⁰J. R. Duke and N. Ananth, “Simulating excited state dynamics in systems with multiple avoided crossings using mapping variable ring polymer molecular dynamics,” *J. Phys. Chem. Lett.* **6**, 4219–4223 (2015).
- ⁶¹C. Myers, K. Miyazaki, T. Trepl, N. Ananth, and C. Isborn (2024), “PySCES – Python code for linearized semi-classical dynamics with on-the-fly electronic structure,” Zenodo. <https://doi.org/10.5281/zenodo.11985696>
- ⁶²L. F. Shampine, “Some practical Runge-Kutta formulas,” *Math. Comput.* **46**, 135–150 (1986).
- ⁶³J. Dormand and P. Prince, “A family of embedded Runge-Kutta formulae,” *J. Comput. Appl. Math.* **6**, 19–26 (1980).
- ⁶⁴A. V. Akimov, “A simple phase correction makes a big difference in nonadiabatic molecular dynamics,” *J. Phys. Chem. Lett.* **9**, 6096–6102 (2018).
- ⁶⁵D. L. Theobald, “Rapid calculation of RMSDs using a quaternion-based characteristic polynomial,” *Acta Crystallogr., Sect. A: Found. Crystallogr.* **61**, 478–480 (2005).
- ⁶⁶D. Mendive-Tapia, B. Lasorne, G. A. Worth, M. J. Bearpark, and M. A. Robb, “Controlling the mechanism of fulvene S_1/S_0 decay: Switching off the stepwise population transfer,” *Phys. Chem. Chem. Phys.* **12**, 15725–15733 (2010).
- ⁶⁷L. M. Ibele and B. F. E. Curchod, “A molecular perspective on tully models for nonadiabatic dynamics,” *Phys. Chem. Chem. Phys.* **22**, 15183–15196 (2020).
- ⁶⁸M. Winslow, W. B. Cross, and D. Robinson, “Comparison of spin-flip TDDFT-based conical intersection approaches with XMS-CASPT2,” *J. Chem. Theory Comput.* **16**, 3253–3263 (2020).
- ⁶⁹B. K. Carpenter, G. B. Ellison, M. R. Nimlos, and A. M. Scheer, “A conical intersection influences the ground state rearrangement of fulvene to benzene,” *J. Phys. Chem. A* **126**, 1429–1447 (2022).
- ⁷⁰L. Martínez, R. Andrade, E. G. Birgin, and J. M. Martínez, “PACKMOL: A package for building initial configurations for molecular dynamics simulations,” *J. Comput. Chem.* **30**, 2157–2164 (2009).
- ⁷¹O. Esenturk, J. S. Melinger, P. A. Lane, and E. J. Heilweil, “Relative photon-to-carrier efficiencies of alternating nanolayers of zinc phthalocyanine and C60 films assessed by time-resolved terahertz spectroscopy,” *J. Phys. Chem. C* **113**, 18842–18850 (2009).
- ⁷²W. Zeng, K. S. Yong, Z. M. Kam, F. Zhu, and Y. Li, “Effect of blend layer morphology on performance of ZnPc: C60-based photovoltaic cells,” *Appl. Phys. Lett.* **97**, 133304 (2010).
- ⁷³B. P. Rand, D. Cheyng, K. Vasseur, N. C. Giebink, S. Mothy, Y. Yi, V. Coropceanu, D. Beljonne, J. Cornil, J.-L. Brédas, and J. Genoe, “The impact of molecular orientation on the photovoltaic properties of a phthalocyanine/fullerene heterojunction,” *Adv. Funct. Mater.* **22**, 2987–2995 (2012).
- ⁷⁴E. J. G. Santos and W. L. Wang, “Ultrafast charge-transfer in organic photovoltaic interfaces: Geometrical and functionalization effects,” *Nanoscale* **8**, 15902–15910 (2016).
- ⁷⁵X.-Y. Liu, Z.-W. Li, W.-H. Fang, and G. Cui, “Nonadiabatic exciton and charge separation dynamics at interfaces of zinc phthalocyanine and fullerene: Orientation does matter,” *J. Phys. Chem. A* **124**, 7388–7398 (2020).
- ⁷⁶Y. Shu, L. Zhang, D. Wu, X. Chen, S. Sun, and D. G. Truhlar, “New gradient correction scheme for electronically nonadiabatic dynamics involving multiple spin states,” *J. Chem. Theory Comput.* **19**, 2419–2429 (2023).
- ⁷⁷G. A. Meek and B. G. Levine, “Evaluation of the time-derivative coupling for accurate electronic state transition probabilities from numerical simulations,” *J. Phys. Chem. Lett.* **5**, 2351–2356 (2014).
- ⁷⁸S. Mai, D. Avagliano, M. Heindl, P. Marquetand, M. F. S. J. Menger, M. Oppel, F. Plasser, S. Polonius, M. Ruckebauer, Y. Shu, D. G. Truhlar, L. Zhang, P. Zobel, and L. Gonzalez, “SHARC3.0: Surface hopping including arbitrary couplings—Program package for non-adiabatic dynamics,” <https://sharc-md.org/> (2023).
- ⁷⁹H. Wang, “Multilayer multiconfiguration time-dependent Hartree theory,” *J. Phys. Chem. A* **119**(29), 7951–7965 (2015).
- ⁸⁰G. M. J. Barca, C. Bertoni, L. Carrington, D. Datta, N. De Silva, J. E. Deustua, D. G. Fedorov, J. R. Gour, A. O. Gunina, E. Guidez, T. Harville, S. Irle, J. Ivanic, K. Kowalski, S. S. Leang, H. Li, W. Li, J. J. Lutz, I. Magoulas, J. Mato, V. Mironov, H. Nakata, B. Q. Pham, P. Piecuch, D. Poole, S. R. Pruitt, A. P. Rendell, L. B. Roskop, K. Ruedenberg, T. Sattasathuchana, M. W. Schmidt, J. Shen, L. Slipchenko, M. Sosonkina, V. Sundriyal, A. Tiwari, V. Galvez, L. Jorge, B. Westheimer, M. Wloch, P. Xu, F. Zahariev, and M. S. Gordon, “Recent developments in the general atomic and molecular electronic structure system,” *J. Chem. Phys.* **152**(15), 154102 (2020).