A hybrid stochastic hierarchy equations of motion approach to treat the low temperature dynamics of non-Markovian open quantum systems

A hybrid stochastic hierarchy equations of motion approach to treat the low temperature dynamics of non-Markovian open quantum systems

Jeremy M. Moix and Jianshu Cao

Citation: J. Chem. Phys. **139**, 134106 (2013); doi: 10.1063/1.4822043
View online: http://dx.doi.org/10.1063/1.4822043
View Table of Contents: http://jcp.aip.org/resource/1/JCPSA6/v139/i13
Published by the AIP Publishing LLC.

A hybrid stochastic hierarchy equations of motion approach to treat the low temperature dynamics of non-Markovian open quantum systems

Jeremy M. Moix and Jianshu Cao[a)]
Department of Chemistry, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA

(Received 2 August 2013; accepted 10 September 2013; published online 2 October 2013)

The hierarchical equations of motion technique has found widespread success as a tool to generate the numerically exact dynamics of non-Markovian open quantum systems. However, its application to low temperature environments remains a serious challenge due to the need for a deep hierarchy that arises from the Matsubara expansion of the bath correlation function. Here we present a hybrid stochastic hierarchical equation of motion (sHEOM) approach that alleviates this bottleneck and leads to a numerical cost that is nearly independent of temperature. Additionally, the sHEOM method generally converges with fewer hierarchy tiers allowing for the treatment of larger systems. Benchmark calculations are presented on the dynamics of two level systems at both high and low temperatures to demonstrate the efficacy of the approach. Then the hybrid method is used to generate the exact dynamics of systems that are nearly impossible to treat by the standard hierarchy. First, exact energy transfer rates are calculated across a broad range of temperatures revealing the deviations from the Förster rates. This is followed by computations of the entanglement dynamics in a system of two qubits at low temperature spanning the weak to strong system-bath coupling regimes.
 [http://dx.doi.org/10.1063/1.4822043]

I. INTRODUCTION

Despite a long history and intense interest, there are still relatively few methods capable of generating the numerically exact dynamics of open quantum systems across a wide range of system parameters. Among the most successful numerical methods are those derived from the path integral formalism, such as the hierarchical equations of motion (HEOM),[1,2] the quasi-adiabatic path integral (QUAPI) approach,[3,4] or direct integration of the path integral expression through Monte Carlo methods.[5,6] The widespread success of the path integral formalism in open quantum systems lies in the ability to analytically integrate out the (Gaussian) environmental degrees of freedom. This procedure leads to the Feynman-Vernon influence functional which accounts for all of the effects of the bath on the system.[7,8] The price to be paid for this simplification, however, is that the influence functional is nonlocal in time, with a correlation time that depends on the intrinsic relaxation time of the bath, the temperature, and the system-bath coupling strength. The key insight in developing the QUAPI formalism was that, in many cases, the bath correlation function, and hence the influence functional itself, decays relatively quickly so that the propagation may be obtained through a small number of deterministic tensor multiplications. Alternatively the HEOM replaces the influence functional by a set of auxiliary density matrices that account for the non-Markovian effects of the bath, provided that the spectral density may be represented in Drude-Lorentz form. Essentially, both the HEOM and QUAPI represent expansions of the influence functional in terms of the memory time of the environment, and as a result they often become prohibitively expensive for strong system-bath interactions, highly non-Markovian environments, or low temperatures.

The main result of this work is a numerical algorithm that significantly extends the parameter regimes accessible to the HEOM, although the general approach outlined here is equally applicable to QUAPI. As is well known, the quantum bath correlation function that appears in the influence functional is a complex function with real and imaginary parts. The real part is responsible for the fluctuations from the environment with a magnitude that increases with temperature, whereas the imaginary part is temperature independent and accounts for the dissipative effects of the environment. All of the methods mentioned above treat these two components on the same footing, but this restriction is not necessary or indeed desirable. Here we show that it can be advantageous to treat the real and imaginary parts of the correlation function by different methods. The central idea is to perform a stochastic unraveling of the real part of the bath correlation function appearing in the influence functional. The remaining imaginary term will then be treated by the HEOM.

The stochastic unraveling of the influence functional has recently formed the basis of a new approach for computing the exact dynamics of open quantum systems.[9–12] The unraveling is achieved through a series of Hubbard-Stratonovich transformations that ultimately result in two coupled linear stochastic differential equations describing the respective forward- and backward-time propagators. As a Monte Carlo, wavefunction-based approach, this scheme appears quite promising for computing the exact dynamics of large

a)jianshu@mit.edu

systems. Performing a Hubbard-Stratonovich transformation only on the real part of the bath correlation function in the influence functional generates an auxiliary field that is purely real. If one stops at this point and ignores the remaining imaginary part of the bath correlation function, then this procedure leads to an evolution equation identical to that of the Haken-Strobl model, except that the Gaussian noise is colored instead of white. In fact, numerical simulations in this case are no more difficult than the Haken-Strobl model itself, and can be readily applied to very large systems.[13,14] However, due to the neglect of the dissipative aspects of the environment, detailed balance is recovered only at infinite temperature.

Correctly treating the imaginary part of the correlation function appearing in the influence functional by a stochastic unraveling is more complicated. It requires a complex auxiliary field that introduces a corresponding complex noise term into the stochastic evolution equations.[15] Consequently the norm of the individual realizations of the wavefunction is not conserved during the propagation. Of course the ensemble average is perfectly normalized, but this norm loss at the realization level greatly degrades the convergence properties of the Monte Carlo procedure. In principle, these difficulties may be circumvented through the use of a Girsanov transformation and other numerical schemes, but the resulting stochastic Schrödinger equations are generally nonlinear.[9,10]

Here, we present an approach that essentially combines the strengths of the stochastic and deterministic methods, while mitigating their respective difficulties. The central idea is to perform a stochastic unraveling only of the real part of the bath correlation function appearing in the influence functional. The remaining imaginary term can then be treated by QUAPI or the HEOM. In fact, similar ideas have been previously proposed, although generally coupled with further approximations. For example, Stockburger and Mak have applied this approach in the context of QUAPI, although the numerical implementation was restricted to an environment that was nearly Markovian.[16] In subsequent work, they treated the real part of the influence functional exactly through a Hubbard-Stratonovich transformation, as is also used here, but a Markovian approximation was made for the dissipative term.[17] Although approximate, this approach leads to a single, closed equation of motion for the reduced density matrix that resembles Kubo's stochastic Liouville equation. Additionally, Tanimura outlined such an approach within the HEOM formalism in Ref. 1 which was referred to as the Fokker-Planck equation with Langevin forces, but the proposed algorithm and preliminary numerical results were only valid in the high temperature limit. A similar result was also derived from a purely stochastic perspective in Ref. 18. While the formal results in Refs. 1 and 18 are exact in principle, the simple description of the real and imaginary parts of the correlation function by stochastic and deterministic schemes, respectively, leads to an unstable numerical algorithm which has limited their practical application. Through the introduction of a reference temperature, here we propose an improved decomposition scheme that leads to a substantial improvement over the previous methods. It is shown that a suitable choice of the reference temperature can provide an optimal balance between the stochastic sampling and deterministic evolution. We make no approximations and treat the non-Markovian characteristics of the bath exactly, demonstrating that the hybrid stochastic hierarchy approach is widely applicable across a broad range of system parameters including, in particular, the low temperature, strong coupling regime.

As with the standard HEOM, the spectral density of the bath is restricted to be of Drude-Lorentz form. However, this restriction offers many advantages. In this case, the imaginary part of the bath correlation function consists of a single, temperature-independent exponential term. Thus, the only convergence parameter with respect to the depth of hierarchy is determined by the reorganization energy and the time scale of the bath—the infinite summation over the Matsubara frequencies is performed exactly through a Monte Carlo sampling of the auxiliary stochastic field. The real power of this approach over the standard HEOM is that it is valid for arbitrary temperatures with a numerical cost that is nearly independent of temperature. In addition, because the imaginary part of the bath correlation function is generally of smaller magnitude than the real part and also decays more rapidly, the stochastic hierarchy can typically be truncated at a much lower tier than with the standard approach. Thus, in principle, one can treat larger systems with the hybrid approach. The only drawback is that a stochastic average over the independent HEOM evolutions is required. However, our calculations so far have indicated that the stochastic average generally converges rapidly, and of course, Monte Carlo algorithms can be trivially parallelized.

In Sec. II, the hybrid stochastic hierarchy equations of motion (sHEOM) are formalized. Following this, benchmark calculations of two-level systems are presented in Sec. III for which standard hierarchy results can be obtained. Then energy transfer rates in a model donor-acceptor system are computed across a broad range of temperatures. Finally, the entanglement dynamics in a system of two qubits are presented at near zero temperature spanning the weak to strong system-bath coupling regimes.

II. FORMALISM

A. Preliminaries

For clarity we consider only one-dimensional systems with continuous degrees of freedom. For the case of discrete systems, the path integral formalism is most readily developed either through the use of Grassmann variables[2] or in the mapping representation.[19] However, our final working expressions are valid for either continuous or discrete systems and numerical results will be presented for the latter. We consider generic system-bath Hamiltonians of the form, $\hat{H} = \hat{H}_0 + \hat{H}_{sb}$, where

$$\hat{H}_{sb} = \frac{1}{2}\sum_j \left[\hat{p}_j^2 + \omega_j^2 \left(\hat{x}_j^2 - \frac{c_j}{\omega_j^2}\hat{V}(\hat{Q}) \right)^2 \right]. \quad (1)$$

The specific form of the system Hamiltonian, $\hat{H}_0$, is irrelevant at this point and need not be specified. The thermal bath is

composed of independent harmonic oscillators characterized by their respective mass-weighted coordinates, $\hat{x}_j$, and momenta, $\hat{p}_j$, with frequencies ω_j and coupling constants c_j to the system degrees of freedom through the interaction potential $\hat{V}(\hat{Q})$, where $\hat{Q}$ denotes a system coordinate.

Assuming that the bath remains in thermal equilibrium and that the initial state of the composite system factorizes such that $\rho(0) = \rho_s(0)\rho_b(0)$, then the environmental degrees of freedom may be integrated out analytically.[7,8] As a result the path integral expression for the forward-backward propagation of the reduced density matrix is succinctly given by

$$U(x_f, y_f, t; x_0, y_0, 0) = \int_{x(0)=x_0}^{x(t)=x_f} \mathcal{D}[x] \int_{y(0)=y_0}^{y(t)=y_f} \mathcal{D}[y] \exp\left(\frac{i}{\hbar}(S_0[x] - S_0[y])\right) \times F[x, y], \tag{2}$$

where $S_0[x]$ denotes the classical action function associated with the bare system Hamiltonian, H_0, computed along the path starting at x_0 at $t = 0$ and ending at x_f at time t. All of the effects of the environment on the system dynamics are accounted for by the Feynman-Vernon influence functional,[7] which, for future convenience, is decomposed as $F[x, y] = F_r[x, y]F_i[x, y]F_b[x, y]$ where the respective terms are associated with the real part of the bath correlation function ($F_r[x, y]$), the imaginary part ($F_i[x, y]$), and the static bath renormalization ($F_b[x, y]$). Each term is given explicitly as

$$F_r[x, y] = \exp\left(-\frac{1}{\hbar}\int_0^t dt' \int_0^{t'} dt'' V^\times(t') K_r(t'-t'') V^\times(t'')\right),$$

$$F_i[x, y] = \exp\left(-\frac{i}{\hbar}\int_0^t dt' \int_0^{t'} dt'' V^\times(t') K_i(t'-t'') V^\circ(t'')\right),$$

$$F_b[x, y] = \exp\left(-\frac{i}{\hbar}\int_0^t dt' \, \lambda V^\times(t') V^\circ(t')\right), \tag{3}$$

where the notation $V^\times(t) = V(x(t)) - V(y(t))$ and $V^\circ(t) = V(x(t)) + V(y(t))$ has been introduced, along with the decomposition of the bath correlation function into its real, K_r, and imaginary, K_i, parts:

$$K(t) = K_r(t) + i K_i(t) \tag{4}$$

$$= \int_0^\infty \frac{d\omega}{\pi} J(\omega) \left[\coth\left(\frac{\hbar\beta\omega}{2}\right)\cos(\omega t) - i \sin(\omega t)\right]. \tag{5}$$

The spectral density function,

$$J(\omega) = \frac{\pi}{2}\sum_{j=1} \frac{c_j^2}{\omega_j}\delta(\omega - \omega_j), \tag{6}$$

contains all of the relevant features in terms of the system dynamics.

The HEOM requires that the bath is described by the Drude-Lorentz spectral density,

$$J_D(\omega) = 2\lambda\omega_c \frac{\omega}{\omega^2 + \omega_c^2}, \tag{7}$$

where ω_c is the cutoff frequency of the bath and the reorganization energy λ is defined as

$$\lambda = \frac{1}{2}\sum_{j=1} \frac{c_j^2}{\omega_j^2} = \frac{1}{\pi}\int_0^\infty d\omega \, \frac{J(\omega)}{\omega}, \tag{8}$$

where the second equality is obtained in the continuum limit. For the Drude spectrum, the corresponding bath correlation function may be computed analytically and the time dependence is of multi-exponential form,

$$K_r(t) = \lambda\omega_c \cot(\hbar\beta\omega_c/2) e^{-\omega_c t} + \frac{4\lambda\omega_c}{\hbar\beta}\sum_{j=1}^\infty \frac{\nu_j}{\nu_j^2 - \omega_c^2} e^{-\nu_j t},$$
$$K_i(t) = -\lambda\omega_c e^{-\omega_c t}, \tag{9}$$

where the Matsubara frequencies are defined as $\nu_j = \frac{2\pi j}{\hbar\beta}$.

With these preliminaries, the formulation of the sHEOM is rather straightforward. In Sec. II B, a Hubbard-Stratonovich transformation is first performed on $F_r[x, y]$, the influence functional containing the real part of the bath correlation function.[12,17,20] Then in Sec. II C, the remaining imaginary part is developed in a hierarchical expansion. Ultimately the expressions for the time evolution of the reduced density matrix are formally identical to those of the standard HEOM, except that here the Hamiltonian is stochastic.[1,18]

B. Partial stochastic unraveling

The Hubbard-Stratonovich transformation is a Gaussian integral identity. Physically, it allows for the decomposition of an interacting system into separate, non-interacting systems but subject to the influence of a common auxiliary field.[21] We have recently employed such an approach in the case of the equilibrium reduced density matrix which led to a highly efficient numerical algorithm.[12,22,23] For real time dynamics, this transformation was applied to the entire influence functional in Ref. 9 leading to a set of stochastic Schrödinger equations. Here we apply the transformation only to the real part of the correlation function in the influence functional as in Refs. 1 and 17, which introduces the corresponding real, auxiliary field, $\xi(t)$. As such, the term $F_r[x, y]$ is equivalently represented as

$$\begin{aligned} F_r[x, y] &= \int \mathcal{D}[\xi] \, N_\xi \exp\left(-\frac{1}{2\hbar}\int_0^t dt' \int_0^t dt'' \, \xi(t') \times K_r^{-1}(t'-t'')\xi(t'') + \frac{i}{\hbar}\int_0^t dt' \xi(t') V^\times(t')\right) \\ &= \int \mathcal{D}[\xi] \, P[\xi] \exp\left(\frac{i}{\hbar}\int_0^t dt' \xi(t') V^\times(t')\right) \\ &= \left\langle \exp\left(\frac{i}{\hbar}\int_0^t dt' \xi(t') V^\times(t')\right)\right\rangle_\xi, \end{aligned} \tag{10}$$

where N_ξ denotes the normalization of the Gaussian functional integral. The Hubbard-Stratonovich transformation is valid provided that $K_r(t)$ is positive semi-definite, which is ensured in this case since it is defined by a covariance function of the bath. Performing the Gaussian integration over $\xi(t)$ clearly leads back to the original expression for $F_r[x, y]$ in

Eq. (3). The normalization of the functional integral implies that the Gaussian term in Eq. (10) can serve as a true probability measure, $P[\xi]$, which leads to the subsequent equalities in Eq. (10). These features allow for the interpretation of $\xi(t)$ as a colored noise driving the dynamics with one- and two-time correlation functions that obey the relations

$$\begin{aligned} \langle \xi(t) \rangle &= 0, \\ \langle \xi(t)\xi(t') \rangle &= \hbar K_r(t-t'). \end{aligned} \tag{11}$$

Since the unraveled influence functional in Eq. (10) and the bath renormalization in Eq. (3) are both local in time, they can be written as additional action terms such that the partially unraveled propagator is given by

$$\begin{aligned} &U(x_f, y_f, t; x_0, y_0, 0) \\ &= \left\langle \int_{x(0)=x_0}^{x(t)=x_f} \mathcal{D}[x] \int_{y(0)=y_0}^{y(t)=y_f} \mathcal{D}[y] \exp\left(\frac{i}{\hbar}(S_0[x] - S_0[y])\right) \right. \\ &\quad \left. \times \exp\left(\frac{i}{\hbar}(S_1[x] - S_1[y])\right) F_i[x, y] \right\rangle_\xi , \end{aligned} \tag{12}$$

where S_0 is the classical action function associated with the bare system Hamiltonian H_0. The noise and potential renormalization are included in the additional action term,

$$S_1[x] = \int_0^t dt'\, \xi(t') V(x(t')) - \lambda V(x(t'))^2 . \tag{13}$$

Alternatively, it is also apparent that a modified action $\tilde{S}[x] = S_0[x] + S_1[x]$ can be defined that is generated from the corresponding time-dependent (stochastic) Hamiltonian,

$$\tilde{H}(\xi; t) = \hat{H}_0 - \xi(t)\hat{V}(\hat{x}) + \lambda \hat{V}(\hat{x})^2 . \tag{14}$$

This interpretation will be used in the ensuing developments.

In a similar manner one may perform an additional Hubbard-Stratonovich transformation on the remaining portion of the influence functional, $F_i[x, y]$, that ultimately leads to the stochastic Schrödinger equations of Ref. 9. However, in this case the auxiliary stochastic fields are required to be complex which leads to serious numerical difficulties due to the associated loss of norm of the individual wavefunction realizations.[15] Therefore, we stop here while the auxiliary field is well behaved and develop $F_i[x, y]$ in a hierarchical expansion following Ref. 1. It should be mentioned, however, that the QUAPI formalism could be equally applied to the imaginary part of the bath correlation as proposed in Ref. 16.

C. Hierarchical expansion

Specifying to the Drude-Lorentz spectral density, then the time dependence of the bath correlation function is of a purely exponential form (see Eq. (9)). For such correlation functions, the hierarchical equations of motion have proved to be a highly efficient numerical approach to compute the exact dynamics of open quantum systems at moderately high temperature. The standard approach, however, becomes extremely costly at low temperature. The source of this difficulty is that the temperature dependence of the real part of the bath correlation is expressed as an infinite summation over Matsubara terms. At low temperatures, many terms of the series must be retained which gives rise to a very deep hierarchy of equations before truncation is acceptable, although some progress has been made recently by utilizing a Padé decomposition of the bath correlation function.[24,35] The stochastic approach developed in Sec. II B circumvents this difficulty entirely by replacing the temperature-dependent terms in the influence functional by a noise sampling procedure that is independent of the number of Matsubara frequencies. The remaining imaginary part of the correlation function appearing in $F_i[x, y]$ consists of only a single exponential term and is independent of temperature.

At this point, one may follow the standard derivations of the hierarchy expansion, leading to nearly identical results for the equations of motion of the auxiliary density matrices.[1,2] For conciseness, this procedure will not be reproduced here and we only present the final result. For a given realization of the stochastic noise, $\xi(t)$, the hierarchy equations of motion are given by

$$\begin{aligned} \frac{\partial}{\partial t}\hat{\rho}_n(\xi; t) = &-\left(\frac{i}{\hbar}\tilde{H}(\xi; t)^\times + n\omega_c\right)\hat{\rho}_n(\xi; t) \\ &-\frac{i}{\hbar}\hat{V}^\times\hat{\rho}_{n+1}(\xi; t) - n\frac{\lambda\omega_c}{\hbar}\hat{V}^\circ\hat{\rho}_{n-1}(\xi; t), \end{aligned} \tag{15}$$

where the hyper-operator notations $\hat{V}^\times\hat{O} = \hat{V}\hat{O} - \hat{O}\hat{V}$ and $\hat{V}^\circ\hat{O} = \hat{V}\hat{O} + \hat{O}\hat{V}$ denote commutators and anti-commutators, respectively. The stochastic Hamiltonian is given in Eq. (14), and the associated noise statistics are given in Eq. (11). The numerically exact reduced density matrix is obtained after preforming the stochastic average, $\rho(t) = \langle \rho_0(\xi; t) \rangle_\xi$. The generalization to systems containing multiple independent baths is identical to those in Refs. 1 and 2 except that the stochastic Hamiltonians must also include multiple independent noise terms. Equation (15) has been previously derived by Tanimura in Ref. 1 (cf. Eq. (6.35) therein), and referred to as the Fokker-Planck equation with Langevin forces. Similarly it was also derived from a purely stochastic approach in Ref. 18 and used to study the low temperature dynamics of the spin-boson model. The closed stochastic equation for the density matrix obtained by Stockburger and Mak is recovered when a Markovian approximation is made to the imaginary part of the influence functional, in which case the hierarchy truncates at the lowest level.[17]

D. Reference temperature

Unfortunately there is a formidable numerical instability associated with the hierarchy as given in Eq. (15) that has limited its practical application. If one considers the noiseless case with $\xi(t) = 0$, then the resulting deterministic hierarchy equations of motion describe a purely dissipative quantum bath. There is no physical limit in which such a case occurs since thermal fluctuations are always present, even at zero temperature. As a result the reduced density matrix obtained from the noiseless version of Eq. (15) is not necessarily positive semi-definite, although it is both Hermitian and norm-preserving. This divergence is shown in the inset of

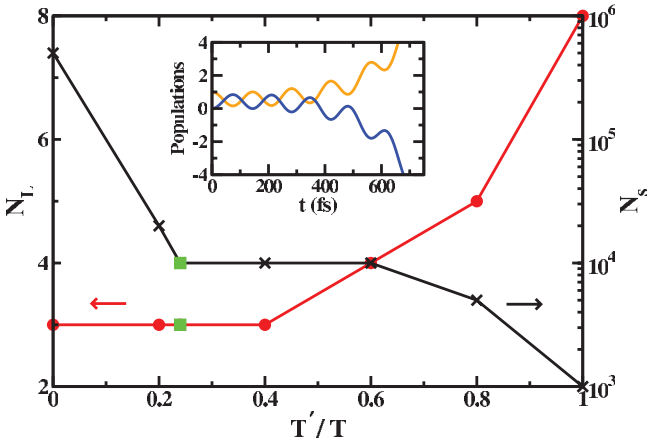

FIG. 1. The convergence properties of the sHEOM as a function of the reference temperature T' in an unbiased two level system coupled to two independent, identical baths at $T = 300$ K with electronic coupling, reorganization energy, and cutoff frequency of $J = 100$ cm^{-1}, $\lambda = 2J$, and $\omega_c = J$, respectively. The axis labels N_L and N_s refer to the number of hierarchy levels and Monte Carlo samples required for convergence, respectively. The green square symbols denote the choice $\beta' = 2/\omega_c$ used here. The inset displays the loss of positivity in the noiseless dynamics of Eq. (15) ($T'/T = 0$), which causes the sharp increase in the required number of Monte Carlo samples.

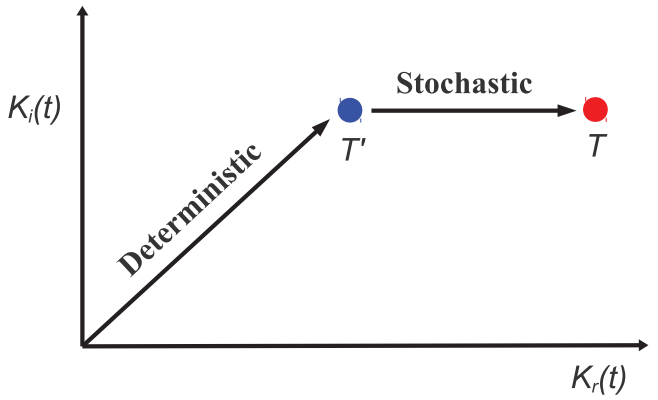

FIG. 2. In the sHEOM scheme, the propagation to a fictitious temperature T' is accomplished deterministically by the standard hierarchy, while the remaining portion of the real part of the bath correlation function is described by the noise. The standard hierarchy describes the entire bath correlation function deterministically and the original proposal in Eq. (15) describes the entire real part stochastically.

Fig. 1. Upon including the stochastic terms, positivity is recovered on average, but the numerical convergence of the Monte Carlo calculations is greatly retarded, particularly for long-time simulations. Here we demonstrate that this problem can be largely mitigated by including a classical portion of the real part of the influence functional in the hierarchy.[17] Then the reduced density matrix is described by the standard high-temperature hierarchy equations,

$$\frac{\partial}{\partial t}\hat{\rho}_n(\xi;t) = -\left(\frac{i}{\hbar}\tilde{H}(\xi;t)^{\times} + n\omega_c\right)\hat{\rho}_n(\xi;t) - \frac{i}{\hbar}\hat{V}^{\times}\hat{\rho}_{n+1}(\xi;t) - \frac{n\lambda}{\hbar}\left(\omega_c\hat{V}^{\circ} + i\frac{2}{\beta'}\hat{V}^{\times}\right)\hat{\rho}_{n-1}(\xi;t), \tag{16}$$

and the weakened noise obeys the modified autocorrelation,

$$\begin{aligned}\langle\xi(t)\rangle &= 0,\\ \langle\xi(t)\xi(t')\rangle &= \hbar[K_r(t-t') - \gamma(t-t')],\end{aligned} \tag{17}$$

where the classical friction function is defined by

$$\gamma(t) = \frac{2}{\pi\hbar\beta'}\int_0^{\infty} d\omega\,\frac{J(\omega)}{\omega}\cos(\omega t) = \frac{2\lambda}{\hbar\beta'}e^{-\omega_c t}\,. \tag{18}$$

In this approach, the temperature of the classical bath, $\beta' = 1/(k_b T')$, is a free parameter that can take on any value in the physically allowable range of $0 \le T' \le T$. Outside of this range the noise becomes complex, which is highly undesirable from a numerical standpoint. The general scheme is depicted in Fig. 2 illustrating that one may tune the reference temperature, T', in the sHEOM method to combine the strengths of the deterministic and stochastic approaches in order to achieve optimal numerical convergence. Setting $T' = 0$ recovers the original Eq. (15) proposed in Refs. 1 and 18 where the entire real part of the bath correlation function is described by the noise. In the opposite limit, when $T' = T$ then only the quantum corrections to the correlation function are treated stochastically. At sufficiently high temperature where $K_r(t) = \gamma(t)$, this case reduces to the standard deterministic HEOM. Equations (16) and (17) constitute the central formal developments of this work.

The impact of the choice of reference temperature on the convergence properties of the sHEOM is shown in more detail in Fig. 1 for an unbiased two level system coupled to two independent, identical baths at 300 K. For $T'/T = 1$, the temperature dependent term in the deterministic hierarchy (last term of Eq. (16)) accounts for almost all of $K_r(t)$ so that convergence requires many levels of the hierarchy, but only few samples of the noise. As the reference temperature is lowered, the temperature-dependent deterministic component of the hierarchy carries less weight as more of the bath correlation function is described by the stochastic sampling. The numerical convergence steadily requires more Monte Carlo samples, but fewer hierarchy tiers. This behavior continues until the temperature dependent, deterministic term becomes negligible compared with the dissipative term such that the required number hierarchy tiers is constant for $T'/T \lesssim 0.4$. However, near $T' = 0$ (Eq. (15)) the noiseless hierarchy becomes unstable as seen in the inset of Fig. 1 and the required number of Monte Carlo samples increases rapidly. We have found that the choice $\beta' = 2/\omega_c$ (provided, of course, that $T' \le T$), which equates the prefactors in the last term of Eq. (16), generally preserves the positivity of the reduced density matrix and allows for quite stable simulations over long timescales, while also keeping the depth of the hierarchy to a minimum. For example, this choice allows for converged results with nearly two orders of magnitude fewer Monte Carlo samples than in Eq. (15) without increasing the required number of hierarchy levels. The quantitative behavior seen in Fig. 1 is somewhat system dependent, but we have found a similar qualitative trend for all cases studied so far. We will use the value $\beta' = 2/\omega_c$ in all of the computations presented below.

The numerical integration of the sHEOM requires only a minor modification to the standard numerical approach to the HEOM. First a realization of the stochastic noise trajec-

tory is generated that is consistent with the autocorrelation function given in Eq. (17). The numerical generation of Gaussian noise with an arbitrary correlation function is discussed in the Appendix. Following this, the stochastic Hamiltonian (Eq. (14)) is formed and the integration of the hierarchy equations proceeds as usual. This procedure is then repeated until the dynamics of the reduced density matrix is converged to acceptable accuracy. Typically this is reached within 10^4–10^6 realizations of the noise.

III. NUMERICAL RESULTS

A. Two level systems

In order to demonstrate the efficacy of the sHEOM approach, we first present benchmark studies of two level systems for which results from the standard hierarchy can be independently computed.[25] We consider the population dynamics in the biased two level system studied in Ref. 26, where each of the sites is coupled to an independent, identical bath. The system Hamiltonian is given by $\hat{H}_0 = \Delta\hat{\sigma}_z + J\hat{\sigma}_x$, where $\Delta = 50$ cm^{-1} and $J = \lambda = 100$ cm^{-1}, with the cutoff frequency $\omega_c = 53$ cm^{-1}. Fig. 3(a) displays the standard hierarchy dynamics at 300 K where the high temperature approximation is completely valid and no Matsubara terms are required. The sHEOM results in Fig. 3(b) are very well behaved, and already at $L = 0$ most of the short time dynamics are accurately captured. At this level, the dynamics are similar to the Haken-Strobl model except that the noise is colored and constructed so as to correctly describe the real part of the bath correlation function. Since the dissipative effects of the bath are completely neglected at $L = 0$, the sites are equally populated at equilibrium. Nevertheless, this approach has recently been used to study diffusion processes in large unbiased systems containing hundreds of sites,[13,14] and is valid as long as the temperature is larger than the bandwidth of the system. In Fig. 3, it is readily seen that the $L = 2$ results in the stochastic approach are almost completely converged while

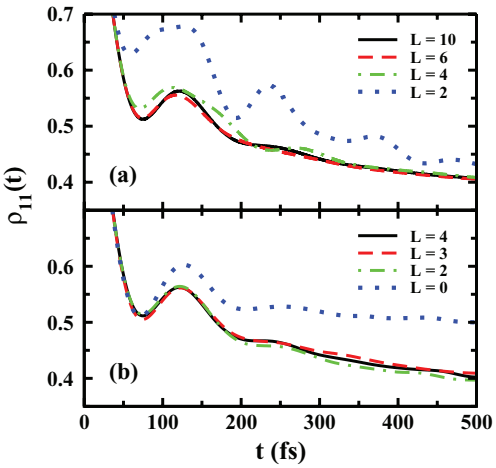

FIG. 3. The convergence of the standard (a) and stochastic (b) HEOM with respect to the number of tiers, L, at a temperature of $T = 300$ K. The bias, electronic coupling, and system-bath coupling are 100 cm^{-1} while the cutoff frequency is 53 cm^{-1}. The sHEOM results are computed with $N = 10^4$ trajectories at each tier.

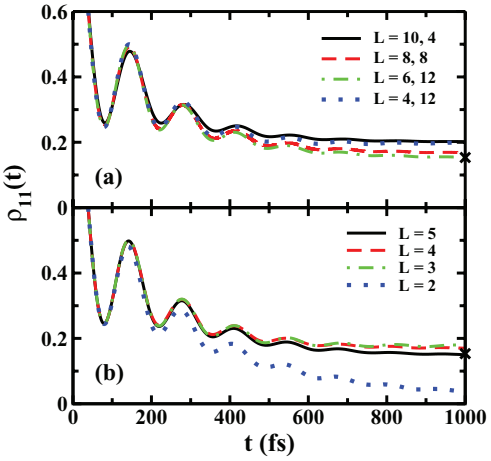

FIG. 4. The convergence of the standard (a) and stochastic (b) HEOM at a temperature of $T = 10$ K with respect to the number of tiers and Matsubara terms (in (a)). The remaining parameters are the same as in Fig. 3. The stochastic calculations required 10^5 Monte Carlo samples for convergence. The cross-symbols on the ordinate denote the exact equilibrium values.

the corresponding standard HEOM results display highly erratic behavior. At least $L = 6$ is required for the standard hierarchy to obtain results of comparable accuracy to the sHEOM approach at $L = 3$. Indeed, small differences are seen in the standard hierarchy results up to $L = 10$. It should be noted that even for $L = 10$, the standard HEOM is more efficient than the stochastic approach at any level, but this advantage is quickly lost as the system size increases or the temperature is lowered. Additionally, one could set the temperature of the classical bath to the physical temperature ($T' = T$ in Eq. (16)) in which case the standard hierarchy is recovered since the low temperature corrections are nearly negligible for these parameters.

A more interesting benchmark case is displayed in Fig. 4 for the same system as in Fig. 3 except with the temperature lowered from 300 K to 10 K. Fully converged results for the standard HEOM with respect to the number of Matsubara terms were possible for 4 and 6 hierarchy tiers, but memory requirements limited convergence at level $L = 8$ to 8 Matsubara terms, and at $L = 10$ to only 4 Matsubara terms. As can be seen by comparing the long-time behavior of the dynamics to the exact equilibrium value obtained from imaginary time path integral calculations,[12] most of these results are not fully converged. The closest result to the exact limit for the standard hierarchy is $L = 6$ with 12 Matsubara terms, although without being able to converge the $L = 8$ calculations it is impossible to know *a priori* if this is an acceptable result. In contrast, the sHEOM approach is completely free of such demands. Fully converged results are readily obtained as shown in Fig. 4(b). In addition as seen previously in Fig. 3, the sHEOM calculations are nearly converged at a lower tier than in the standard approach. In this case, the computational cost of the standard hierarchy with Matsubara terms at $L = 6$ becomes comparable to that of the hybrid approach. In the former case, one is required to integrate a single deterministic system of equations containing $\sim 10^6$ density matrices, while in the latter one must compute 10^5 Monte Carlo samples, but the system of equations contains only ~ 20 matrices.

This scaling becomes even more favorable as the system size increases.

B. Energy transfer rates

As seen above, the sHEOM approach is capable of generating the numerically exact dynamics at both high and low temperatures, as well as the correct equilibrium limit. Now we turn to simulations that are difficult, if not impossible, to carry out with the standard hierarchy, but become straightforward when using the hybrid approach. We first consider energy transfer in the two level donor-acceptor system considered in Ref. 26. The bias, electronic coupling, and cutoff frequency are $\Delta = 50\ \text{cm}^{-1}$, $J = 20\ \text{cm}^{-1}$, and $\omega_c = 53\ \text{cm}^{-1}$, respectively. For this relatively weak value of the electronic coupling, the Förster rates provide a reasonable approximation to the exact rates (see Fig. 5). In Ref. 26 the temperature was fixed at 300 K and the energy transfer rate from the higher lying electronic state to the lower state was computed as a function of the reorganization energy. However, the reorganization energy is generally not an experimentally tunable parameter. Here, we fix the reorganization at the maximum energy transfer rate observed in Ref. 26 ($\lambda = 20\ \text{cm}^{-1}$) and scan the temperature from 1 to 1000 K. In every case, the results were converged with $L = 2$ hierarchy tiers and 10^4 Monte Carlo samples. The rates, k, are computed by fitting the long-time population dynamics to a kinetic model for the energy transfer dynamics between the two sites.[27] In this model, the population of the initially excited state is given by

$$P_1(t) = \frac{\chi_{\text{eq}} + e^{-(1+\chi_{\text{eq}})kt}}{1 + \chi_{\text{eq}}}, \tag{19}$$

where $\chi_{\text{eq}} = P_1^{\text{eq}}/P_2^{\text{eq}}$ is the ratio of the equilibrium populations computed from imaginary time path integral calculations.[12] In the results presented in Ref. 26, the accuracy of the Förster rates improved for very large or very small values of the reorganization energy. Here, the Förster rates are seen to systematically overestimate the exact results for all temperatures. This error can be largely corrected by including the fourth order correction in the electronic coupling to the Förster rates, as shown in previous studies of diffusion-limited electron transfer.[28,29]

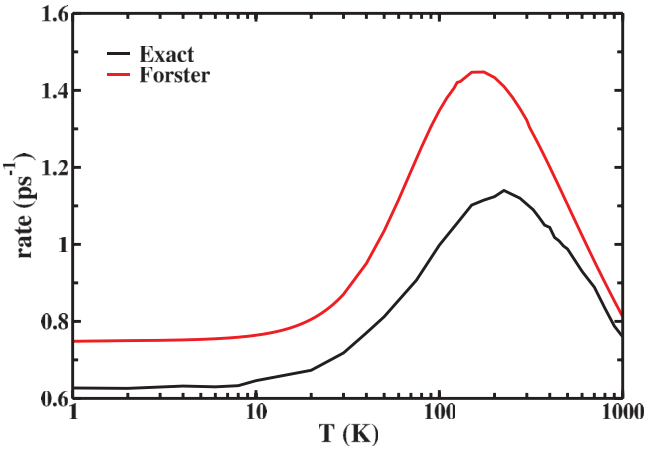

FIG. 5. The energy transfer rates computed from the exact sHEOM simulations (black) and the standard Förster rates (red). The bias is $\Delta = 50\ \text{cm}^{-1}$ and the electronic coupling, $J = 20\ \text{cm}^{-1}$, with the cutoff frequency $\omega_c = 53\ \text{cm}^{-1}$ and reorganization energy $\lambda = 20\ \text{cm}^{-1}$. The exact results are computed with $L = 2$ levels of the hierarchy and 10^4 Monte Carlo samples.

C. Entanglement dynamics

Finally, we present results on the entanglement dynamics of two qubits at near zero temperature across a wide range of system-bath coupling strengths. The system Hamiltonian is given by $\hat{H}_0 = \omega_0(\hat{\sigma}_1^z + \hat{\sigma}_2^z) + J\hat{\sigma}_1^x\hat{\sigma}_2^x$ and the two qubits are coupled to identical, independent baths through their respective $\hat{\sigma}^x$ operators. The system frequency and bath cutoff are $\omega_0 = 1.5J$ and $\omega_c = 3J$, respectively, in units where $J = 1$ sets the energy scale. This model was considered in Ref. 30, and, although not shown, we have reproduced the high temperature dynamics shown therein. The initial state is chosen to be a completely entangled state, $\rho(0) = I + \hat{\sigma}_1^x\hat{\sigma}_2^x + \hat{\sigma}_1^y\hat{\sigma}_2^y - \hat{\sigma}_1^z\hat{\sigma}_2^z$, where I denotes the identity matrix. Fig. 6 displays results for the concurrence with the low temperature, $\beta J = 50$ converged with 10^6 Monte Carlo samples. At very weak coupling, $\lambda = 0.01J$ the results of the secular Redfield equation are in excellent agreement with the exact numerical results. However, the equilibrium state generated by the Redfield dynamics is always that given by the Boltzmann distribution computed with respect to $\hat{H}_0$, regardless of the system-bath coupling strength. Additionally, in the scaled units of time, λt, the secular Redfield dynamics are independent of λ. Nevertheless, for weak system-bath coupling, $\lambda = 0.1J$, the Redfield dynamics still provide an accurate approximation to the exact results. The entanglement displays an initial death followed by a subsequent reappearance and slow equilibration.[31] In contrast to the Redfield results, the exact dynamics correctly demonstrate that the environment-induced decoherence steadily destroys the entanglement between the qubits as the

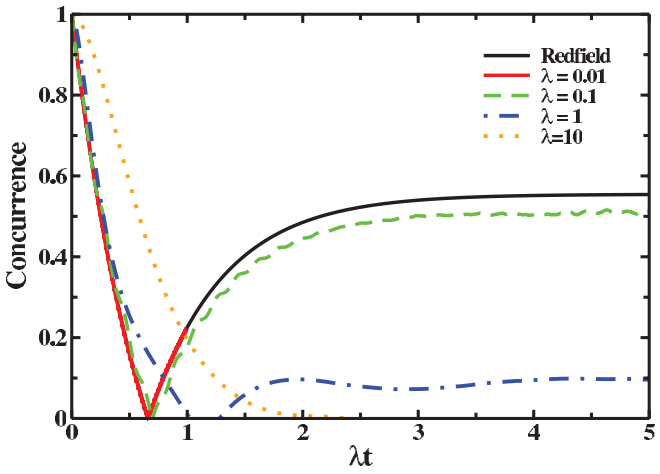

FIG. 6. The exact concurrence between two qubits compared with the results from the secular Redfield equation. The solid red, dashed green, dot-dashed blue, and dotted orange lines display the results for reorganization energies of $\lambda/J = 0.01$, 0.1, 1, and 10, respectively. The exact $\lambda = 0.01J$ results are indistinguishable from the Redfield results on this scale and for clarity are shown only until $\lambda t = 1.10^6$ Monte Carlo samples were used in each of the stochastic hierarchy calculations. In the scaled units of time, λt, the Redfield dynamics are independent of λ.

system-bath coupling strength increases. In fact, for $\lambda = 10J$ there is no reappearance of the concurrence after the initial death. In this case the system-bath coupling is sufficiently strong so as to destroy all of the equilibrium entanglement.

IV. CONCLUSIONS

In summary, a hybrid stochastic hierarchy equation of motion approach has been proposed that substantially extends the parameter regimes accessible to the hierarchy formalism. The method combines many of the strengths of the standard hierarchy with those of recently proposed stochastic methods.[1,9] By performing a Hubbard-Stratonovich transformation only on the real part of the bath correlation function, the problematic temperature dependent terms in the hierarchy are exchanged for a Monte Carlo average over real noise trajectories. This extension eliminates the hierarchy tiers that are required to account for the low temperature corrections in the standard approach. Consequently the numerical cost of the sHEOM simulations is nearly independent of the temperature. Additionally, treating the imaginary part of the bath correlation function by the hierarchy cures many of the normalization difficulties associated with the purely stochastic methods.[9] The introduction of the reference temperature in Eq. (16) allows one to achieve an optimal balance between the stochastic and deterministic components of the evolution that substantially improves the convergence properties of the algorithm. Numerical results were presented across a broad range of parameters, including both the high and low temperature limits, as well as the strong to weak system-bath coupling regimes. The computations of the energy transfer rates and the low temperature entanglement dynamics in two qubits presented in Sec. III are difficult, if not impossible, to obtain with most other methods. However, these calculations are very straightforward with the sHEOM approach presented here. If the noiseless hierarchy is well-behaved then the Monte Carlo procedure generally converges rapidly, but even in cases where positivity is not ensured convergence is generally achieved within 10^6 samples. Additionally, the use of the hybrid approach allows for a much lower truncation of the hierarchy tiers which, in principle, allows one to simulate larger systems than is possible with the standard hierarchy. Although not discussed here, we have easily performed calculations on the light harvesting system LH2 across a broad range of temperatures.

The Hubbard-Stratonovich transformation of the real part of the bath correlation function appearing in the influence functional is valid for arbitrary spectral densities. It is only the hierarchy treatment of the remaining imaginary part that demands a Drude-Lorentz form. In principle, it should be possible to employ other spectral densities such as that of the underdamped Brownian oscillator[32] or by decomposing more complicated spectral densities into a sum of Drude-Lorentz terms.[33] However, the hierarchy of equations will be more complicated. In these cases, it may be more advantageous to adapt the procedure outlined above to QUAPI,[16] which suffers many of the same problems as the standard HEOM in the low temperature regime. The extension of the hybrid approach to the QUAPI formalism would allow for numerically exact simulations of the non-Markovian dynamics of open quantum systems for arbitrary spectral densities across a broad range of the parameter space.

ACKNOWLEDGMENTS

This work was supported by the National Science Foundation (NSF) (Grant No. CHE-1112825) and Defense Advanced Research Projects Agency (DARPA) (Grant No. N99001-10-1-4063). J. Moix has been supported by the Center for Excitonics, an Energy Frontier Research Center funded by the U.S. Department of Energy, Office of Science, Office of Basic Energy Sciences under Award No. DE-SC0001088. We thank Jian Ma and Arend Dijkstra for useful discussions, as well as the secular Redfield code used to generate the results in Fig. 6.

APPENDIX: NOISE GENERATION

Apart from the integration of the HEOM, which is achieved through standard Runge-Kutta methods, the only additional computational requirement of the sHEOM approach is the generation of the stochastic process, $\xi(t)$. There are numerous methods to generate Gaussian colored noise. In this appendix, two approaches are briefly described that we have found to be particularly useful. The most straightforward approach originates from signal processing techniques and is based on filtering white noise through an appropriate kernel:

$$\xi(t) = \int_0^t dt' k(t-t')\zeta(t') \,. \tag{A1}$$

The noise, $\zeta(t)$, is a standard Wiener process with zero mean, $\langle \zeta(t) \rangle = 0$ and autocorrelation, $\langle \zeta(t)\zeta(t') \rangle = \delta(t-t')$. The filtering kernel, $k(t)$, is determined from the factorization of the autocorrelation function:

$$K_r(t-t') = \int dt'' k(t-t'')k(t''-t') \,. \tag{A2}$$

In practice, $k(t)$ is most efficiently constructed through a Cholesky decomposition of the discretized kernel matrix $\bar{K}_{ij} = K_r(t_i - t_j)$ with $\bar{k}_{ij}$ defined accordingly, such that $\bar{K} = \bar{k}^T \bar{k}$. A discretized sample of the desired noise sequence is then simply generated from the matrix-vector product $\vec{\xi} = \bar{k}^T \vec{\zeta}$.

While straightforward, the Cholesky approach encounters difficulties in simulations where many time steps are needed so that the kernel matrix becomes large. In this case, an alternative approach to generate the noise can be used that relies on a discretization of the correlation function in terms of its independent frequency components.[13,34] Samples of the noise can be generated from the Fourier sum

$$\xi(t) = \sqrt{\frac{2}{\pi}} \sum_{n=1}^{N} \left[J(\omega_n) \coth\left(\frac{\beta\omega_n}{2}\right) \Delta\omega \right]^{1/2} \cos(\omega_n t + \phi_n), \tag{A3}$$

where $\Delta\omega = \omega_{\max}/N$ and $\omega_n = n\Delta\omega$ with $\omega_{\max}$ chosen to be sufficiently large such that the spectral density has decayed to zero, $J(\omega_{\max}) = 0$. The phases, ϕ_n, are uniform random numbers generated on the interval $(0, 2\pi)$. The noise

sequence generated in this manner is periodic with a period of $2\pi/\Delta\omega$ which, obviously, must be at least twice the simulation time. The summation in Eq. (A3) is most efficiently performed through the use of a fast Fourier transform.

[1] Y. Tanimura, J. Phys. Soc. Jpn. **75**, 082001 (2006).
[2] A. Ishizaki and Y. Tanimura, J. Phys. Soc. Jpn. **74**, 3131 (2005).
[3] N. Makri and D. E. Makarov, J. Chem. Phys. **102**, 4611 (1995).
[4] N. Makri, J. Chem. Phys. **109**, 2994 (1998).
[5] R. Egger, L. Mühlbacher, and C. H. Mak, Phys. Rev. E **61**, 5961 (2000).
[6] L. Mühlbacher and R. Egger, J. Chem. Phys. **118**, 179 (2003).
[7] R. P. Feynman and F. L. Vernon, Ann. Phys. **24**, 118 (1963).
[8] H. Grabert, P. Schramm, and G. Ingold, Phys. Rep. **168**, 115 (1988).
[9] J. T. Stockburger and H. Grabert, Phys. Rev. Lett. **88**, 170407 (2002).
[10] J. Shao, J. Chem. Phys. **120**, 5053 (2004).
[11] J. Cao, L. W. Ungar, and G. A. Voth, J. Chem. Phys. **104**, 4189 (1996).
[12] J. M. Moix, Y. Zhao, and J. Cao, Phys. Rev. B **85**, 115412 (2012).
[13] X. Zhong and Y. Zhao, J. Chem. Phys. **135**, 134110 (2011).
[14] X. Chen, J. Cao, and R. J. Silbey, J. Chem. Phys. **138**, 224104 (2013).
[15] W. T. Strunz, Phys. Lett. A **224**, 25 (1996).
[16] J. T. Stockburger and C. H. Mak, Phys. Rev. Lett. **80**, 2657 (1998).
[17] J. T. Stockburger and C. H. Mak, J. Chem. Phys. **110**, 4983 (1999).
[18] Y. Zhou, Y. Yan, and J. Shao, EPL **72**, 334 (2005).
[19] A. Novikov, U. Kleinekathöfer, and M. Schreiber, Chem. Phys. **296**, 149 (2004).
[20] L. S. Schulman, *Techniques and Applications of Path Integration* (Wiley, New York, 1986).
[21] J. Hubbard, Phys. Rev. Lett. **3**, 77 (1959).
[22] C. K. Lee, J. Moix, and J. Cao, J. Chem. Phys. **136**, 204120 (2012).
[23] C. K. Lee, J. Cao, and J. Gong, Phys. Rev. E **86**, 021109 (2012).
[24] J. Ding, J. Xu, J. Hu, R. Xu, and Y. Yan, J. Chem. Phys. **135**, 164107 (2011).
[25] J. Strümpfer and K. Schulten, J. Chem. Theory Comput. **8**, 2808 (2012).
[26] A. Ishizaki and G. R. Fleming, J. Chem. Phys. **130**, 234111 (2009).
[27] I. R. Craig, M. Thoss, and H. Wang, J. Chem. Phys. **127**, 144503 (2007).
[28] J. Cao, J. Chem. Phys. **112**, 6719 (2000).
[29] J. Wu and J. Cao, J. Chem. Phys. **139**, 044102 (2013).
[30] A. G. Dijkstra and Y. Tanimura, Phys. Rev. Lett. **104**, 250401 (2010).
[31] T. Yu and J. H. Eberly, Science **323**, 598 (2009).
[32] M. Tanaka and Y. Tanimura, J. Phys. Soc. Jpn. **78**, 073802 (2009).
[33] C. Kreisbeck and T. Kramer, J. Phys. Chem. Lett. **3**, 2828 (2012).
[34] K. Y. R. Billah and M. Shinozuka, Phys. Rev. A **42**, 7492 (1990).
[35] Q. Shi, L. Chen, G. Nan, R. X. Xu, and Y. Yan, J. Chem. Phys. **130**, 084105 (2009).