

# Accelerating the computation of bath spectral densities with super-resolution

Thomas Markovich<sup>1</sup> · Samuel M. Blau<sup>1</sup> · John Parkhill<sup>2</sup> · Christoph Kreisbeck<sup>1</sup> · Jacob N. Sanders<sup>1</sup> · Xavier Andrade<sup>1</sup> · Alán Aspuru-Guzik<sup>1</sup>

Received: 9 February 2016 / Accepted: 17 July 2016 / Published online: 11 August 2016  
© Springer-Verlag Berlin Heidelberg 2016

**Abstract** Quantum transport and other phenomena are typically modeled by coupling the system of interest to an environment, or bath, held at thermal equilibrium. Realistic bath models are at least as challenging to construct as models for the quantum systems themselves, since they must incorporate many degrees of freedom that interact with the system on a wide range of timescales. Owing to computational limitations, the environment is often modeled with simple functional forms, with a few parameters fit to experiment to yield semi-quantitative results. Growing computational resources have enabled the construction of more realistic bath models from molecular dynamics (MD) simulations. In this paper, we develop a numerical technique to construct these atomistic bath models with better accuracy and decreased cost. We apply a novel signal processing technique, known as super-resolution, combined with a dictionary of physically motivated bath modes to derive spectral densities from MD simulations. Our approach reduces the required simulation time and provides a more accurate spectral density than can be obtained via standard Fourier transform methods. Moreover, the spectral density is provided as a convenient closed-form expression which yields an analytic time-dependent bath kernel. Exciton dynamics of the Fenna–Matthews–Olson light-harvesting complex are simulated with a second-order time-convolutionless

master equation, and spectral densities constructed via super-resolution are shown to reproduce the dynamics using only a quarter of the amount of MD data.

**Keywords** Spectral densities · Super-resolution · Hierarchical equations of motion

## 1 Introduction

Irreversible processes such as solvation, energy transfer, and chemical binding have received renewed interest in recent years. Because these processes involve large systems with many degrees of freedom, the typical approach to studying these processes is the open quantum systems formalism, in which the degrees of freedom are partitioned into a system of interest and a bath held at thermal equilibrium [1, 2]. It is commonly assumed that the system only couples weakly to the bath, making the precise nature of the bath a secondary concern in the physical theory. For example, in studying the energy transfer dynamics in a system of chromophores embedded in a protein framework, each chromophore is individually coupled to many thousands of atoms in the protein, but the system–bath formalism dramatically simplifies all of these couplings in order to make the dynamics tractable [3–5]. Renewed interest in the strong and intermediate coupling region, relevant for energy transfer in the exciton dynamics of light-harvesting complexes, has led to various studies [5–28] on the precise influence of the bath on the higher systems. Higher-order phonon processes, non-Markovian effects, and structures in the exciton–phonon coupling change the energy transfer [29–31]. Thus, details in the bath are relevant and need to be taken into account in realistic simulations. Accordingly, our goal in this paper is to apply a recent signal processing

Published as part of the special collection of articles “Festschrift in honour of A. Vela”.

✉ Alán Aspuru-Guzik  
aspuru@chemistry.harvard.edu

<sup>1</sup> Department of Chemistry and Chemical Biology, Harvard University, Cambridge, MA 02138, USA

<sup>2</sup> Department of Chemistry, The University of Notre Dame, South Bend, IN 46556, USA

technique known as super-resolution to obtain realistic atomistic models of environments containing thousands of atoms at feasible computational expense. With these atomistic bath models in hand, one can begin to evaluate the importance of a realistic bath model in a physical theory.

In the approach to open quantum systems employed in this work, we model the bath by an ensemble of non-interacting harmonic oscillators. The central mathematical object of such a model is the *spectral density*,  $J(\omega)$ , which gives the frequency-dependent strength of system–bath coupling. The spectral density can be understood as the density of bath oscillator states at each frequency. Owing to computational limitations, most studies of open quantum systems assume an extremely simple functional form for the spectral density, such as a single broad peak covering all relevant excitonic transitions of the system. With the goal of providing more physically accurate bath models and dynamics, Valleau et al. have previously obtained atomistic spectral densities for the Fenna–Matthews–Olson (FMO) complex from combined molecular dynamics (MD) [32, 33] and time-dependent density functional theory (TDDFT) [34] simulations. However, the difficulty of this more realistic approach is the high computational cost of running expensive TDDFT calculations at every step in an MD simulation. In order to obtain a spectral density of sufficient resolution, the MD–TDDFT simulation must be run for over 40 picoseconds (ps) [35], which may become computationally intractable for larger systems.

To make progress, we first observe that a typical vibrational bath is not an arbitrary function, but rather a relatively sparse collection of damped harmonic oscillators. Sparsity enables us to apply a novel numerical technique known as super-resolution in order to reconstruct the spectral density from much shorter MD–TDDFT simulations. Many methods such as the filter diagonalization method [36] or the maximum entropy [37] method have been developed to achieve similar undersampling goals. Super-resolution provides us with the ability to recover the spectral parameters of our signal with a bias toward sparsity. Super-resolution has been applied to a broad range of scientific problems, including image [38] and video compression [39], image denoising [40], astronomy [41], microscopy [42], and medical imaging [43]. To our knowledge, this paper is the first application of super-resolution to quantum dynamics. Super-resolution provides a provably convergent algorithm for the reconstruction of signals from limited time-domain measurements using a total variation minimization procedure. Super-resolution is related to compressed sensing [44–55]. Compressed sensing is a technique designed to recover sparse signals from randomly sampled data by minimizing the  $\mathcal{L}_1$ -norm of an underdetermined system of linear equations. Compressed sensing works by finding the sparsest signals consistent

with the underdetermined system of equations. This usually involves an optimization problem. Despite its success in many applications, the  $\mathcal{L}_1$ -norm minimization of compressed sensing can result in spurious signals as it emphasizes the sparsity of the solution only. Super-resolution is a numerical method that shares the spirit of compressed sensing. The difference between super-resolution and compressed sensing stems from both the choice of objective function and sampling technique. It was developed to recover sparse signals from nonrandomly undersampled data. By minimizing the  $\mathcal{L}_1$ -norm of the gradient of the function in addition to the  $\mathcal{L}_1$ -norm of the function itself, super-resolution allows for smoother solutions to the sampling problem [56–58].

Because of the ample experimental and theoretical data to compare against [5, 8–11, 13, 16–19, 21, 22, 24, 27, 31, 32, 59–66], we apply super-resolution to the FMO light-harvesting complex of *C. tepidum* but emphasize that this technique is broadly applicable. While this paper focuses on a vibrational bath which perturbs the energies of molecular electronic states, the techniques we introduce are generic for any model of a bath which is based on time correlation functions.

## 2 Super-resolution of spectral densities

In this section, we briefly review the procedure for simulating the dynamics of open quantum systems and computing spectral densities from combined MD–TDDFT simulations. We then apply the theory of super-resolution to accelerate and improve the accuracy of these computations. Computing spectral densities from atomistic calculations, rather than from semi-empirical functional forms, enables the inclusion of molecular vibrations and other physical effects (such as solvation effects) to produce a more realistic bath model [32]. Super-resolution, in turn, brings the construction of these atomistic bath models into the realm of computational feasibility.

Armed with our more realistic bath model, we will employ a second-order time-convolutionless master equation (TCL-2) to simulate the dynamics of the FMO monomer, allowing us to evaluate the physical impact of different approximations to the spectral density. TCL-2 includes non-Markovian effects up to second order in the system–bath coupling. By comparing TCL-2 with exact methods like the hierarchical equations of motion (HEOM) [67], we show that most of the relevant effects of the structured spectral density of the FMO complex are captured by TCL-2. Here we use TCL-2, since it is numerically more treatable than HEOM, in particular for structured spectral densities where HEOM becomes cumbersome and requires a high-performance GPU implementation [17,

31, 68]. We employ the equation of motion [1, 7, 9, 20, 24, 63, 69–74]:

$$\frac{d\rho_I(t)}{dt} = -\frac{i}{\hbar}[H_I, \rho_I] - \frac{1}{\hbar^2} \sum_n \int_0^t d\tau D_n(t-\tau)[H_{In}(t), [H_{In}(\tau), \rho_I(t)]] \quad (2.1)$$

$$D_n(t) = \int_0^\infty d\omega J_n(\omega) \left[ \coth\left(\frac{\hbar\omega\beta}{2}\right) \cos(\omega t) - i \sin(\omega t) \right] \quad (2.2)$$

where  $H$  is the system Hamiltonian,  $\rho$  is the system density matrix,  $D$  is our bath kernel, the subscript  $I$  indicates that we are in the interaction picture, the summation runs over all sites, and  $J(\omega)$  is the spectral density computed via super-resolution [17, 64, 65, 75]. The bath kernel is heavily dependent on our spectral density, causing it to play a central role in our dynamics. Therefore, a more physical bath picture should provide more physically intuitive dynamics.

In our atomistic bath model, molecular vibrations in the environment (e.g., a protein framework or solvation effects) create fluctuations in the energy gaps between the ground and excited states of the system (e.g., a set of chromophores). These time-dependent energy gaps are computed from TDDFT calculations run on each of the chromophores at each step of the MD simulation. The key object in the computation of spectral densities is the correlation function of the energy gap time series,

$$C(t) = \text{Tr}_b[\hat{\Delta}(t)\hat{\Delta}(0)\hat{\rho}_b], \quad (2.3)$$

where  $\hat{\Delta}(t)$  is the time-dependent energy gap between the ground and the first excited state of the system (as calculated with TDDFT),  $\hat{\rho}_b$  is the density matrix of the bath at thermal equilibrium, and  $C(t)$  is the correlation function obtained after tracing over all the modes of the bath. We discretize this equation by using an unbiased autocorrelation function,

$$C_k = \frac{1}{N-k} \sum_{i=1}^{N-k} (\Delta_i - \bar{\Delta})(\Delta_{i+k} - \bar{\Delta}), \quad (2.4)$$

where  $\bar{\Delta}$  is the mean energy gap and  $i$  and  $k$  denote discrete time indices. Note that  $C_k$  involves comparing energy gaps that are  $k$  time steps apart ( $\Delta_i$  and  $\Delta_{i+k}$ ), and  $N-k$  is the total number of included comparisons.

The frequency-dependent spectral density,  $J(\omega)$ , is typically obtained by computing the Fourier transform of the correlation function [32]. From the definition of  $C_k$  above, it is easy to check that the correlation function is real and symmetric ( $C_k = C_{N-k}$ ), which implies that the Fourier transform should be real and symmetric as well. Because

quantum mechanical spectral densities must instead be antisymmetric and obey detailed balance, it is necessary to introduce a prefactor that enforces these two properties. Many choices are possible [76], but Valleau et al. have previously shown that a harmonic prefactor,  $\beta\hbar\omega/2$ , produces the most physical temperature dependence [32]. With this choice, the spectral density becomes the cosine transform

$$J(\omega) = \frac{\beta\hbar\omega}{2} \int_{-\infty}^{\infty} \cos(\omega t) C(t) dt, \quad (2.5)$$

which characterizes the frequency-dependent coupling strength of the system to all of the nuclear vibrational modes.

The standard approach to performing this integral is the fast Fourier transform. Unfortunately, the fast Fourier transform requires sampling on a uniform grid at the Shannon sampling rate. This means that a relatively long time series,  $C(t)$ , must be computed in order to obtain good resolution of the spectral density in the frequency domain [35, 77]. Given the computational cost of MD simulations, and the even greater expense of running TDDFT calculations on top of these simulations, any method which can reduce the required length of the time series  $C(t)$  unplugs the computational bottleneck in deriving physically accurate atomistic spectral densities. That is our main goal in this paper.

While reducing the amount of time required to reproduce  $J(\omega)$ , we also choose a basis of functions which has a convenient physical form. When decomposed into a basis of damped cosines,

$$g_{ij}(t) = e^{-\gamma_i t} \cos(\Omega_j t), \quad (2.6)$$

the function  $C(t)$  is smooth and sparse. This allows for the use of the machinery of super-resolution.

To apply the super-resolution method, we discretize in time and cast our task as an inversion problem

$$C_k = \lambda_{ij} e^{-\gamma_i t} \cos(\Omega_j t_k), \quad (2.7)$$

where we seek the basis expansion coefficients  $\lambda_{ij}$  and have assumed Einstein summation convention over repeated indices. This can be rewritten as

$$C_k = A_{ijk} \lambda_{ij}, \quad (2.8)$$

where

$$A_{ijk} = e^{-\gamma_i t} \cos(\Omega_j t_k) \quad (2.9)$$

is a matrix of damped cosines and  $\lambda_{ij}$  is the set of basis coefficients we seek to recover.

The central idea of super-resolution is that the sparsity of  $\lambda_{ij}$  enables its full recovery even when the system  $C_k = A_{ijk} \lambda_{ij}$  is underdetermined, which is to say the number of time samples  $C_k$  is significantly smaller than the number of total expansion coefficients  $\lambda_{ij}$  we seek to recover. Hence, we can recover the expansion coefficients

on a dense grid of frequencies  $\Omega_j$  and damping coefficients  $\gamma_i$  from fewer time samples  $C_k$ . Of the many possible solutions to our underdetermined system, super-resolution simply selects a balance between the smoothest and sparsest (with an emphasis on smoothness) set of basis expansion coefficients. Formally, this is done by finding the vector  $\lambda_{ij}$  that minimizes

$$\underset{\lambda_{ij}}{\operatorname{argmin}} \{ \|\nabla \lambda_{ij}\|_1 + \mu \|\lambda_{ij}\|_1 \} \quad (2.10)$$

subject to  $\|A_{ijk}\lambda_{ij} - C_k\|_2 < \eta$ ,

where the subscript 1 represents the  $\mathcal{L}_1$ -norm (sum of absolute values),  $\mu$  represents a sparsity penalty,  $\nabla \lambda_{ij}$  represents the total variation norm, and  $\eta$  represents the solution tolerance. By minimizing  $\|\nabla \lambda_{ij}\|_1$ , or the total variation term, we are enforcing smoothness in the time domain on the reconstructed signal. This throws out the “peaky” solutions that can appear with compressed sensing [78, 79]. The total variation norm also provides us with a provably exact technique for recovering peak position at the expense of peak amplitude [56], which solves one of the issues seen previously with compressed sensing [78].

Recovering the expansion coefficients  $\lambda_{ij}$  in this manner by solving an underdetermined matrix inversion problem takes advantage of the natural sparsity of the problem and, as we will see in the next section, enables the construction of a well-resolved spectral density with far less time-domain data. Even more attractive, with the  $\lambda_{ij}$  coefficients in hand, it is possible to construct an analytical representation of the spectral density by taking the cosine transform of the basis functions  $g_{ij}(t)$  and applying the appropriate prefactors:

$$J(\omega) = \frac{\lambda_{ij}}{\sqrt{\pi}} \left( \frac{\beta \hbar \omega \gamma_i}{\gamma_i^2 + (\omega - \Omega_j)^2} + \frac{\beta \hbar \omega \gamma_i}{\gamma_i^2 + (\omega + \Omega_j)^2} \right), \quad (2.11)$$

where the Einstein summation convention has again been assumed. This is an analytical representation of the spectral density in Drude–Lorentz form, and it explicitly provides the oscillation frequencies which characterize the system–bath coupling. We note that the Drude–Lorentz basis naturally provides us with a width parameter,  $\gamma$ , that can be understood as the lifetime of oscillations in the bath. This is seen by examining the time-dependent formula, Eq. (2.7), where this  $\gamma$  parameter determines the strength of damping. It is important to note that in the limit as  $\gamma \rightarrow 0$ , we recover the cosine basis in the time domain and a Dirac delta distribution in the frequency domain. By using this super-resolution technique in concert with the Drude–Lorentz basis, we see that we can recover a small set of peaks with physically relevant information. Additionally, the parameters

that characterize the Drude–Lorentz spectral densities can be input directly into both TCL-2 and HEOM without any additional parameter fitting or numerical integration.

### 3 Numerical methods

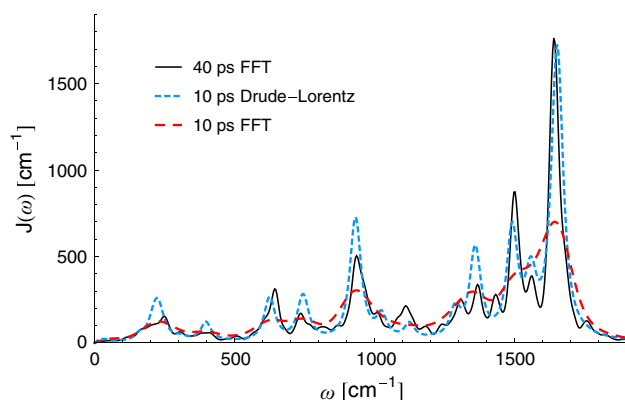
We employ the proposed Drude–Lorentz super-resolution method described above and apply it to a monomer of the Fenna–Matthews–Olson (FMO) photosynthetic energy transfer complex of the green sulfur bacterium *C. tepidum*. The FMO monomer is a system of seven chlorophyll molecules which are excitonically coupled to each other, as well as to the vibrations of the atoms in the protein framework. It functions as a molecular excitonic wire, passing excitons from the light-harvesting antenna complex to the reaction center, where a biochemical cascade is initiated. To create spectral densities for the FMO complex, we use the MD–TDDFT results of Shim et al. [63]. The calculations were done in an isothermal–isobaric ensemble at 77 K using the AMBER force field [80, 81]. These calculations began with a 2-ns equilibration before performing the production computations. The production steps ran for a total of 40 ps with a 2-fs time step, and the optical gap was calculated for each fragment every 4 fs using TDDFT with the BLYP [82–84] functional in the 3-21G basis set in Q-Chem [85].

To perform super-resolution numerically, we require an algorithm which minimizes the total variation norm to solve the minimization problem described by Eq. (2.10). In our implementation, we use the two-step iterative shrinkage thresholding (TwIST) algorithm [57, 58], which combines computational efficiency with strong convergence. To construct the measurement matrix  $A$  described in Eq. (2.9), we must select a grid of possible frequencies ( $\{\Omega_j\}$ ) and linewidths ( $\{\gamma_i\}$ ). In our implementation, we use a grid of frequencies ranging from 0 to 2000  $\text{cm}^{-1}$  in 2  $\text{cm}^{-1}$  intervals, and a grid of linewidths ranging from 0 to 160  $\text{cm}^{-1}$  in 6  $\text{cm}^{-1}$  intervals. We assume that our calculations are converged when  $\eta < 10^{-7}$  [in Eq. (2.10)], or the solution vector remains constant for 100 iterations. Finally, we perform an  $\mathcal{L}_2$  minimization of  $A_{ijk}\lambda_{ij} - C_k$  while freezing the recovered nonzero basis functions, allowing us to further minimize the error. We refer to this procedure as debiasing because it partly removes the bias toward sparsity and smoothness introduced by the  $\mathcal{L}_1$  minimization. This debiasing procedure reduces our solution tolerance to  $\eta < 10^{-9}$ , allowing convergence to a better solution. It is important to note that, in general, the super-resolution technique is robust to an over-complete basis.

## 4 Results

Figure 1 shows the results of employing the Drude–Lorentz super-resolution method to recover the spectral density for site 1 of FMO. The figure compares Drude–Lorentz super-resolution with 10 ps of MD to a standard fast Fourier transform approach with both 10 and 40 ps of MD. We take the fast Fourier transform with 40 ps of MD as our standard for comparison. By comparing the two methods with 10 ps of MD, it is clear that super-resolution resolves more features of the spectral density than the standard fast Fourier transform from the same amount of time-domain data. Moreover, super-resolution captures most of the features of the fast Fourier transform with the full 40 ps of MD: We see the expected CO stretch at  $1600\text{ cm}^{-1}$ , which we attribute to the amides in the protein scaffold, as well as all of the other major peaks in the spectral density. We attribute a significant amount of the error in our spectral density reconstruction to the fact that the truncated MD series does not explore the phase space as thoroughly in only 10 ps.

The Drude–Lorentz basis also provides significant sparsity gains in comparison with the cosine basis: We require only 56 Drude–Lorentz peaks to create the spectral density given in Fig. 1. This sparsity provides a significant computational advantage for excitonic propagation in both hierarchical equations of motion (HEOM) [67] and second-order time-convolutionless master equation [1] (TCL-2) approaches because the propagations scale factorially and linearly, respectively, as a function of the number of peaks included. In the excitonically accessible regime of  $0\text{--}540\text{ cm}^{-1}$ , we recover only 20 Drude–Lorentz peaks, and six of them have amplitudes that are two orders of

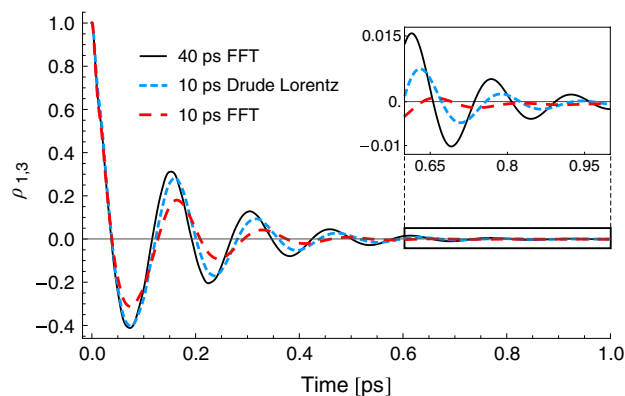


**Fig. 1** Comparison of the spectral density for site 1 of the FMO complex as a function of time and technique for spectral density recovery. Compared to the fast Fourier transform at 40 ps, much of the fine structure is easily recovered by super-resolution in the Drude–Lorentz basis, even with significant undersampling by a factor of four

magnitude smaller than the rest. These Drude–Lorentz peaks can be entered directly into master equation simulations, including HEOM codes, without the need to perform any intermediate fitting [68]. In summary, super-resolution yields a well-resolved spectral density using less time-domain data than is required by the standard fast Fourier transform approach and precludes the need for additional fitting.

As mentioned above, the TCL-2 propagation of the exciton dynamics of the FMO complex, with the Hamiltonian coming from [60], was carried out using the Drude–Lorentz spectral densities obtained from super-resolution. We propagated 1 ps of dynamics and obtained the populations of sites 1–3, as well as the coherence between sites 1 and 3.

Figure 2 shows the coherence between excitonic eigenstates 1 and 3 as a function of time. Compared to the 40-ps fast Fourier transform, we see that the 10-ps Drude–Lorentz super-resolution more faithfully reproduces the coherence dynamics than the 10-ps fast Fourier transform, in terms of both the oscillation frequency and the overall damping. The fast Fourier transform with 10 ps of MD data introduces serious overdamping as well as a significant shift in oscillation frequency. In contrast, the Drude–Lorentz expansion with 10 ps of MD data introduces only a small shift in oscillation frequency, resulting in more accurate coherence dynamics overall. We attribute most of the discrepancies to slight relative differences in the reorganization of each site between spectral densities constructed with 10 and 40 ps of MD data. It appears that while the oscillations are extremely sensitive to the relative reorganization energies between the sites, the damping is more dependent on the fine structure of



**Fig. 2** Comparison of the coherences between excitonic eigenstates 1 and 3 as a function of time and technique for spectral density recovery. Compared to the fast Fourier transform at 40 ps, the 10-ps Drude–Lorentz decomposition introduces a slight shift in oscillation frequency, but nevertheless yields more accurate dynamics than the equivalently sampled fast Fourier transform at 10 ps

the spectral densities. The Drude–Lorentz super-resolution (10-ps MD–TDDFT data) reproduces the coherence lifetimes obtained by fast Fourier transform recovered using all 40 ps of MD–TDDFT data—representing a factor of four improvement.

The contrast between the two approximation techniques becomes even more significant when we simulate dynamics beginning with an exciton fully localized on site 1. In Fig. 3, we have plotted the populations of the first three sites as a function of time. The Drude–Lorentz expansion with 10 ps of MD yields good qualitative agreement with our standard of comparison. The fast Fourier transform on 10 ps overestimates population transfer to site 3 at short times and grows much more quickly from there, whereas the Drude–Lorentz expansion slightly underpredicts the population transfer at long times. We attribute these errors in the asymptotic behavior to slight differences in the reorganization energies for the spectral densities of each of the sites: Since each site is embedded in a different

environment, the reorganization process of the individual pigments is different. This sensitivity affects overall dissipation, and even small changes in the spectral density of the Drude–Lorentz expansion (10 ps) when compared to the standard of comparison affects energy relaxation. Beyond that, the Drude–Lorentz expansion is capable of reproducing the oscillations at 0.2 and 0.4 ps in the data for sites 1 and 2, whereas the fast Fourier transform reproduces them less faithfully. In summary, the Drude–Lorentz super-resolution technique provides us with much more physical behavior.

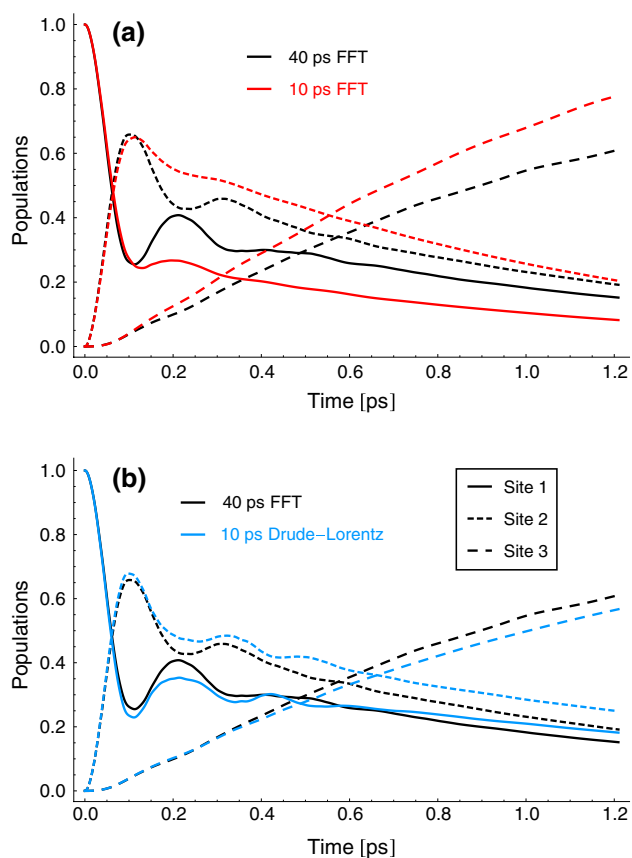
## 5 Conclusions

We have shown that the Drude–Lorentz super-resolution method provides significant computational advantages for the construction of atomistic bath models. In particular, the super-resolution calculations require only 10 ps of MD–TDDFT simulations to obtain reasonable atomistic spectral densities and system dynamics; this is one quarter the amount of data needed in standard fast Fourier transform-based calculations. Ultimately, this will permit the use of more physically accurate calculations or larger systems. Given the computational expense of running TDDFT calculations at every MD simulation step, we believe that the super-resolution method will enable the treatment of larger systems than previously possible.

One of the most significant advantages of our super-resolution method is the decomposition of these atomistic spectral densities into a naturally sparse basis of Drude–Lorentz oscillators. This makes it easy to perform fast master equation simulations within either the TCL-2 or HEOM formalisms by exploiting analytic integrals of the spectral density. Beyond this, we also directly extract physically important parameters such as the coherence lifetimes of all the oscillators in the bath. In the future, it is easy to imagine turning this technique on its head to create new spectral densities in a constructive fashion from a set of Drude–Lorentz oscillators.

While the application that we have presented here focuses on the construction of bath spectral densities, it is easy to imagine applying our presented technique to any correlation function with a damped oscillatory functional form. Such applications could include the construction of colored noise kernels, force autocorrelation functions, and real-time Raman spectroscopy.

**Acknowledgments** We acknowledge S. Valleau for useful discussions and computer code. We acknowledge the financial support of Defense Advanced Research Projects Agency Grant N66001-10-1-4063 and the Defense Threat Reduction Agency under Contract No. HDTRA1-10-1-0046. T.M. acknowledges support from the National Science Foundation (NSF) through the Graduate Research Fellowship



**Fig. 3** Comparison of the populations for sites 1–3 as a function of time and technique for spectral density recovery. Compared to the fast Fourier transform at 40 ps, the 10-ps Drude–Lorentz decomposition recovers the overall shape and provides much more faithful dynamics than the equivalently sampled fast Fourier transform at 10 ps

Program (GRFP). S.B. acknowledges support from the Department of Energy (DoE) through the Computational Sciences Graduate Fellowship (CSGF). J.N.S. acknowledges support from the Department of Defense (DoD) through the National Defense Science & Engineering Graduate Fellowship (NDSEG) Program. A.A.G. thanks the Corning Foundation.

## References

- Breuer H-P, Petruccione F (2007) The theory of open quantum systems. Oxford University Press, Oxford
- Shenvi N, Schmidt JR, Edwards ST, Tully JC (2008) Phys Rev A 78:022502
- Berkelbach TC, Hybertsen MS, Reichman DR (2013) J Chem Phys 138:4102
- Berkelbach TC, Hybertsen MS, Reichman DR (2013) J Chem Phys 138:114103
- Singh N, Brumer P (2011) Faraday Discuss 153:41
- Jang S, Newton MD, Silbey RJ (2007) J Phys Chem B 111:6807
- Mohseni M, Rebentrost P, Lloyd S, Aspuru-Guzik A (2008) J Chem Phys 129:174106
- Plenio MB, Huelga SF (2008) New J Phys 10:113019
- Rebentrost P, Mohseni M, Kassal I, Lloyd S, Aspuru-Guzik A (2009) New J Phys 11:033003
- Cao J, Silbey RJ (2009) J Phys Chem A 113:13825
- Ishizaki A, Fleming GR (2009) Proc Natl Acad Sci 106:17255
- Ishizaki A, Fleming GR (2011) J Phys Chem B 115:6227
- Sarovar M, Ishizaki A, Fleming GR, Whaley KB (2010) Nat Phys 6:462
- Abramavicius D, Mukamel S (2010) J Chem Phys 133:064510
- Wu J, Liu F, Shen Y, Cao J, Silbey RJ (2010) New J Phys 12:105012
- Moix J, Wu J, Huo P, Coker D, Cao J (2011) J Phys Chem Lett 2:3045
- Kreisbeck C, Kramer T, Rodríguez M, Hein B (2011) J Chem Theory Comput 7:2166
- Skochdopole N, Mazzotti DA (2011) J Phys Chem Lett 2:2989
- Ritschel G, Roden J, Strunz WT, Aspuru-Guzik A, Eisfeld A (2011) J Phys Chem Lett 2:2912
- Rebentrost P, Aspuru-Guzik A (2011) J Chem Phys 134:101103
- Pachón LA, Brumer P (2012) Phys Chem Chem Phys 14:10094
- Vlaming SM, Silbey RJ (2012) J Chem Phys 136:055102
- Caruso F, Saikin SK, Solano E, Huelga SF, Aspuru-Guzik A, Plenio MB (2012) Phys Rev B 85:125424
- Zhu J, Kais S, Rebentrost P, Aspuru-Guzik A (2011) J Phys Chem B 115:1531
- Roden J, Eisfeld A, Wolff W, Strunz W (2009) Phys Rev Lett 103:058301
- Olbrich C, Kleinekathöfer U (2010) J Phys Chem B 114:12427
- Olbrich C, Jansen TLC, Liebers J, Aghtar M, Strümpfer J, Schulten K, Knoester J, Kleinekathöfer U (2011) J Phys Chem B 115:8609
- Hein B, Kreisbeck C, Kramer T, Rodríguez M (2012) New J Phys 14:023018
- Christensson N, Kauffmann HF, Pullerits T, Mančal T (2012) J Phys Chem B 116:7449
- Chin AW, Prior J, Rosenbach R, Caycedo-Soler F, Huelga SF, Plenio MB (2013) Nat Phys 9:113
- Kreisbeck C, Kramer T (2012) J Phys Chem Lett 3:2828
- Valleau S, Eisfeld A, Aspuru-Guzik A (2012) J Chem Phys 137:224103
- Tuckerman M (2010) Statistical mechanics: theory and molecular simulation. OUP, Oxford
- Runge E, Gross EK (1984) Phys Rev Lett 52:997
- Mallat S (2008) A wavelet tour of signal processing, 3rd edn: the sparse way. Academic Press, London
- Hu H, Van QN, Mandelshtam VA, Shaka AJ (1998) Reference deconvolution, phase correction, and line listing of NMR spectra by the 1D filter diagonalization method. J Magn Reson 134(1):76–87
- Hoch JC, Maciejewski MW, Mobli M, Schuyler AD, Stern AS (2014) Nonuniform sampling and maximum entropy reconstruction in multidimensional NMR. Acc Chem Res 47(2):708–717
- Freeman WT, Jones TR, Pasztor EC (2002) IEEE Comput Graph Appl 22:56
- Patti AJ, Sezan MI, Murat Tekalp A (1997) IEEE Trans Image Process 6:1064
- Elad M, Feuer A (1997) IEEE Trans Image Process 6:1646
- Puschmann KG, Kneer F (2005) Astron Astrophys 436:373
- Mccutchen CW (1967) J Opt Soc Am 57:1190
- Kouame D, Ploquin M (2009) Super-resolution in medical imaging: an illustrative approach through ultrasound. In: IEEE international symposium on biomedical imaging: from nano to macro, 2009. ISBI '09, pp 249–252
- Ma J (2010) IEEE Trans Instrum Meas 59:1600
- Donoho DL (2006) IEEE Trans Inf Theory 52:1289
- Oka A, Lampe L (2009) A compressed sensing receiver for bursty communication with UWB impulse radio. In: IEEE international conference on ultra-wideband, 2009. ICUWB 2009, pp 279–284
- Herman MA, Strohmer T (2009) IEEE Trans Signal Process 57:2275
- Qiu C, Lu W, Vaswani N (2009) Real-time dynamic MR image reconstruction using Kalman filtered compressed sensing. In: IEEE international conference on acoustics, speech and signal processing, 2009. ICASSP 2009, pp 393–396
- Lustig M, Donoho D, Pauly JM (2007) Magn Reson Med 58:1182
- Nagahara M, Quevedo DE, Ostergaard J (2012) Sparse representations for packetized predictive networked control. In: 2012 IEEE 51st annual conference on decision and control (CDC), pp 1362–1367
- Tuma T, Rooney S, Hurley P (2009) On the applicability of compressive sampling in fine grained processor performance monitoring. In: 2009 14th IEEE international conference on engineering of complex computer systems. IEEE, pp 210–219
- Shabani A, Kosut RL, Mohseni M, Rabitz H, Broome MA, Almeida MP, Fedrizzi A, White AG (2011) Efficient measurement of quantum dynamics via compressive sensing. Phys Rev Lett 106(10):100401
- Mishali M, Eldar YC (2010) IEEE J Sel Top Signal Process 4:375
- Duarte MF, Davenport MA, Takhar D, Laska JN, Sun T, Kelly KF, Baraniuk RG (2008) IEEE Signal Process Mag 25:83
- Coulter WK, Hillar CJ, Isley G, Sommer FT (2010) Adaptive compressed sensing — A new class of self-organizing coding models for neuroscience. 2010 IEEE International conference on acoustics, speech and signal processing, Dallas, TX, p 5494–5497
- Candès EJ, Fernandez-Granda C (2014) Commun Pure Appl Math 67:906–956
- Bioucas-Dias JM, Figueiredo MA (2007) IEEE Trans Image Process 16:2992
- Bioucas-Dias JM, Figueiredo MAT (2007) Two-step algorithms for linear inverse problems with non-quadratic regularization. In: IEEE international conference on image processing, 2007. ICIP 2007, pp I–105–I–108
- Vulto SIE, de Baat MA, Louwe RJW, Permentier HP, Neef T, Miller M, van Amerongen H, Aartsma TJ (1998) J Phys Chem B 102:9577

60. Adolphs J, Renger T (2006) *Biophys J* 91:2778
61. Chen X, Cao J, Silbey RJ (2013) *J Chem Phys* 138:224104
62. Yuen-Zhou J, Krich JJ, Aspuru-Guzik A (2012) *J Chem Phys* 136:234501
63. Shim S, Rebentrost P, Valleau SP, Aspuru-Guzik AN (2012) *Biophys J* 102:649
64. Kolli A, Nazir A, Olaya-Castro A (2011) *J Chem Phys* 135:154112
65. Jang S (2011) *J Chem Phys* 135:034105
66. Ishizaki A, Fleming GR (2009) *J Chem Phys* 130:234111
67. Tanimura Y, Kubo R (1989) *J Phys Soc Jpn* 58:101
68. Kreisbeck C, Kramer T, Rodríguez TM, Hein B (2011) *J Chem Theory Comput* 7(7):2166–2174
69. Pereverzev A, Bittner ER (2006) *J Chem Phys* 125:4906
70. Timm C (2011) *Phys Rev B* 83:115416
71. Ahn D (1994) *Phys Rev B* 50:8310
72. Heinz-Peter B, Kappler B, Petruccione F (2000) Decoherence: theoretical, experimental, and conceptual problems. 233
73. Shabani A, Lidar DA (2005) *Phys Rev A* 71:020101
74. Smirne A, Vacchini B (2010) *Phys Rev A* 82:022110
75. Kleinekathöfer U (2004) *J Chem Phys* 121:2505
76. Berens PH, White SR, Wilson KR (1981) *J Chem Phys* 75:515
77. Candes EJ, Romberg J, Tao T (2006) *IEEE Trans Inf Theory* 52:489
78. Andrade X, Sanders JN, Aspuru-Guzik A (2012) *Proc Natl Acad Sci* 109:13928
79. Sanders JN, Saikin SK, Mostame S, Andrade X, Widom JR, Marcus AH, Aspuru-Guzik A (2012) *J Phys Chem Lett* 3:2697
80. Cornell WD, Cieplak P, Bayly CI, Gould IR, Merz KM, Ferguson DM, Spellmeyer DC, Fox T, Caldwell JW, Kollman PA (1995) *J Am Chem Soc* 117:5179
81. Ceccarelli M, Procacci P, Marchi M (2003) *J Comput Chem* 24:129
82. Becke AD (1988) *Phys Rev A* 38:3098
83. Miehl B, Savin A, Stoll H, Preuss H (1989) *Chem Phys Lett* 157:200
84. Lee C, Yang W, Parr RG (1988) *Phys Rev B* 37:785
85. Shao Y, Molnar LF, Jung Y, Kussmann J, Ochsenfeld C, Brown ST, Gilbert AT, Slipchenko LV, Levchenko SV, O'Neill DP (2006) *Phys Chem Chem Phys* 8:3172