



Published in final edited form as:

J Chem Theory Comput. 2012 August 14; 8(8): 2808–2816. doi:10.1021/ct3003833.

Open Quantum Dynamics Calculations with the Hierarchy Equations of Motion on Parallel Computers

Johan Strümpfer[†] and Klaus Schulten^{*,†,‡}

[†]Center for Biophysics and Computational Biology, University of Illinois at Urbana-Champaign

[‡]Department of Physics and Beckman Institute, University of Illinois at Urbana-Champaign

Abstract

Calculating the evolution of an open quantum system, i.e., a system in contact with a thermal environment, has presented a theoretical and computational challenge for many years. With the advent of supercomputers containing large amounts of memory and many processors, the computational challenge posed by the previously intractable theoretical models can now be addressed. The hierarchy equations of motion present one such model and offer a powerful method that remained under-utilized so far due to its considerable computational expense. By exploiting concurrent processing on parallel computers the hierarchy equations of motion can be applied to biological-scale systems. Herein we introduce the quantum dynamics software *PHI*, that solves the hierarchical equations of motion. We describe the integrator employed by *PHI* and demonstrate *PHI*'s scaling and efficiency running on large parallel computers by applying the software to the calculation of inter-complex excitation transfer between the light harvesting complexes 1 and 2 of purple photosynthetic bacteria, a 50 pigment system.

1 Introduction

Quantum dynamics plays a major role in many biological processes.^{1–4} To model such processes one has to account for the environment in which they occur, that is, one where thermal noise at an ambient temperature is present. Although many methods have been developed that account to some degree for the interaction between a quantum system and a thermal environment,^{4–16} they typically involve limitations in regard to the relative interaction strengths found within a system and between the system and environment, or make approximations regarding the quantum nature of the system, or neglect non-Markovian effects. Such limitations and approximations are often permissible for a particular system and then the respective methods yield valuable insight, but they do not apply in general.

Purple bacteria are one of the simplest photosynthetic life forms, capable, through efficient light harvesting, of living in photon-poor environments.¹⁷ The primary pigment-protein complexes found in purple bacteria, light harvesting complex 2 (LH2), light harvesting complex 1 (LH1) and the reaction center (RC), have evolved to rapidly transfer excitation energy across many tens of nanometers and convert it, before the energy is lost through fluorescence or internal conversion, to a more stable form resulting in over 90% light-harvesting efficiency.^{18,19} This remarkable efficiency due to rapid excitation transfer has been the subject of many experimental^{20–29} and theoretical^{30–37} investigations. It has been shown that excitation transfer benefits not only from quantum coherence,^{38–40} but also from environmental noise.^{40–47} In order to model excitation dynamics in purple bacteria, one has

*To whom correspondence should be addressed kschulte@ks.uiuc.edu.

to take into account the strong interactions not only between the excited states themselves, but also with their surrounding environment.^{34,48–54}

There are only a few quantum dynamics methods for computation of open quantum system density matrices that take into account arbitrary intra-system and system-environment interactions, include non-Markovian effects, and make no assumptions about the quantum nature of the dynamics.^{55–61} The hierarchy equations of motion (HEOM) of Tanimura and Kubo,⁵⁸ which employ a hierarchy of auxiliary density matrices to account for non-Markovian dynamics without resorting to perturbative approximations,^{58–71} is one of the earliest open quantum system descriptions and remains the only one applicable to large systems.^{34–36,71} Although the HEOM are computationally expensive, typically requiring many gigabytes (GBs) of computer memory and many days of computation time, their application to large systems has been possible by exploiting GPUs⁷² and highly parallel computers.^{34–36} The GPU-HEOM implementation⁷² yields tremendous improvement in computation time, but is severely limited to small (less than 10 pigments) systems by the small amount of memory available with current generation GPUs. By performing the HEOM integration on multi-processor CPUs, the system size is limited by the much larger main computer memory instead of GPU memory, allowing much larger systems to be modeled as shown here.

Implementing the HEOM on a multi-processor computer is, however, difficult due to the highly connected nature of the auxiliary density matrices. To achieve peak performance, the communication between processing elements needs to be minimized through a partitioning scheme that assigns each matrix in the hierarchy to a different processor. Additional improvements to calculation time can also be achieved by employing adaptive integration methods: either by adaptively filtering the number of auxiliary matrixes or by adaptively adjusting the integration step size.

Besides an application to large photosynthetic systems, such as those found in cyanobacteria, plants and artificial light harvesting,^{73–77} the HEOM method is applicable to a wide range of other quantum systems. Indeed, any process that can be described by a set of N discrete quantum states in contact with a bosonic environment can be modeled.⁶¹

In the following section we present the partitioning scheme and integration methods used to implement the HEOM, as well as performance benchmarks of our parallel hierarchy integrator *PHI*.

2 Methods

Many systems can be modeled with an effective Hamiltonian^{78–80} that describes the subset of relevant states $|n\rangle$, with energies E_n and interactions V_{nm} written as

$$H_0 = \sum_n E_n |n\rangle \langle n| + \sum_{n,m} V_{nm} |n\rangle \langle m|. \quad (1)$$

At present *PHI* accounts only for Hermitian effective Hamiltonian operators, and thereby assumes a single excitation subspace. Systems such as those in a biological context are embedded in a $T = 300$ K thermal environment. Such environment can be modeled as an infinite set of harmonic oscillators¹¹

¹*PHI* is presently available at <http://www.ks.uiuc.edu/Research/phi>

$$H_B = \sum_{\xi} \left[\frac{p_{\xi}^2}{2m_{\xi}} + \frac{m_{\xi}\omega_{\xi}^2 q_{\xi}^2}{2} \right]. \quad (2)$$

The thermal bath is coupled to the states $|n\rangle$ through the system-bath interaction term

$$H_{SB} = \sum_a^M \widehat{F}_a u_a, \quad (3)$$

where $u_a = \sum_{\xi} c_{a\xi} q_{\xi}$ with $c_{a\xi}$ being the coupling strength of vibrational mode ξ to the system via operator \widehat{F}_a . In the case that each state $|n\rangle$ is independently coupled to the environment holds $M = N$ and $\widehat{F}_a = |a\rangle\langle a|$, where $|a\rangle$ denotes one of the states $|n\rangle$ arising in Eq. (1). The system-bath coupling introduces a shift in the minimum energy positions of the bath coordinates q_{ξ} , which can be countered by the renormalization term

$$H_{REN} = \sum_{ab}^M \widehat{F}_a \widehat{F}_b \sum_{\xi} \frac{c_{a\xi} c_{b\xi}}{2m_{\xi}\omega_{\xi}^2}. \quad (4)$$

In the case of system-bath coupling that is independent for each state $|n\rangle$, the

renormalization term reduces to $H_{REN} = \sum_a^N |a\rangle\langle a| \lambda_a$, where $\lambda_a = \sum_{\xi} c_{a\xi}^2 / 2m_{\xi}\omega_{\xi}^2$ is the bath reorganization energy. The total Hamiltonian is thus given by $H_T = H_S + H_B + H_{SB}$, where $H_S = H_0 + H_{REN}$ is the renormalized system Hamiltonian.

The bath degrees of freedom are not of interest and, thus, are traced out of the dynamics to calculate the bath averaged density matrix evolution as

$$\widehat{\rho}(t) = \text{tr}_B \left\{ \int e^{-i\mathcal{L}_T t / \hbar} \widehat{\rho}(0) \otimes e^{-\beta H_B} dt \right\} / \text{tr}_B \{ e^{-\beta H_B} \} = \langle \mathcal{U}(t) \rangle_B \widehat{\rho}(0), \quad (5)$$

where $\beta = 1/k_B T$. Eq. (5) corresponds to the so-called factorized initial condition that does not include initial correlations between system and environment. Although initial correlations between system and environment would not be present in the case of instantaneous photo-excitation,⁶¹ they can be included when modeling other processes using the HEOM.^{59,65} Once the thermal average is performed, the influence of the environment enters only through the bath correlation functions given by,⁴

$$C_{ab}(t) = \langle u_a(t) u_b(0) \rangle_B = \frac{1}{\pi} \int_{-\infty}^{\infty} d\omega J_{ab}(\omega) \frac{e^{-i\omega t}}{1 - e^{-\beta \hbar \omega}}, \quad (6)$$

where the spectral density $J_{ab}(\omega)$ is given by

$$J_{ab}(\omega) = \frac{\pi}{2} \sum_{\xi} \frac{c_{a\xi} c_{b\xi}}{m_{\xi} \omega_{\xi}} \delta(\omega - \omega_{\xi}). \quad (7)$$

As the coupling of environmental fluctuations to off-diagonal components in the system Hamiltonian are typically smaller than the coupling to diagonal components, i.e. $|C_{ab}| \ll |C_{aa}|$, they are usually ignored.⁵⁴ In the present work only diagonal coupling to the environment is accounted for.

Currently only a few methods^{55–58,60,61} offer numerically tractable equations of motion for the solution, Eq. (5), for system sizes as they arise in the case of pigment systems of photosynthetic light harvesting. The hierarchy equations of motion^{58–63,66,67,69,71,81} (HEOM) is computationally expensive, but nevertheless furnishes an attractive method due to its applicability across a broad range of interaction scenarios.^{34–36,66,69,72,80–86}

Correlation functions

The HEOM arise by exploiting bath correlation functions $C_{ab}(t)$ described by sums of exponentially decaying terms.⁵⁸ The HEOM work particularly well when the spectral density arising in Eq. (6) and Eq. (7) can be expressed as a so-called Drude spectral density,⁴ given by

$$J_a(\omega) = 2\lambda_a \frac{\omega\gamma_a}{(\omega^2 + \gamma_a^2)}, \quad (8)$$

where $\lambda_a = \int (J_a(\omega)/\pi\omega) d\omega$ is the bath reorganization energy that determines the system-bath interaction strength, and $1/\gamma_a$ is the bath response time. The Drude spectral density describes over-damped harmonic oscillator modes ξ . The correlation function associated with the Drude spectral density is⁶⁹

$$C_a(t) = \sum_{k=0}^{\infty} c_{ak} \exp(-\nu_{ak}t), \quad (9)$$

where $\nu_{a0} = \gamma_a$, $\nu_{ak-1} = 2\pi k/\beta\hbar$ are the Matsubara frequencies⁶⁴ and the coefficients c_{ak} are

$$c_{a0} = \gamma_a \lambda_a [\cot(\beta\hbar\gamma_a/2) - i] \quad (10)$$

$$c_{ak \geq 1} = \frac{4\lambda_a \gamma_a}{\beta\hbar} \frac{\nu_{ak}}{\nu_{ak}^2 - \gamma_a^2}. \quad (11)$$

The summation to infinity in Eq. (9) needs to be truncated at a finite level, which is achieved by exploiting that beyond some finite K , $\nu_{aK} \exp(-\nu_{aK}t) \approx \delta(t)$.⁶⁰

The expansion of the correlation function using Matsubara frequencies, Eq. (9) and (11), is not the only possible expansion, and recent work has shown that alternative expansions can yield better results for low cut-offs K .^{87,88}

In principle one can also employ correlation functions for under-damped harmonic oscillators, namely ones of the form $C(t) = c \exp(-\gamma t) \cos(\Omega t)$, to derive the HEOM,^{89–91} thus enabling the use of arbitrary, multi-peaked spectral densities through, e.g., a Meier-Tannor decomposition.⁹² The Lorentzian peaks in such spectral densities, however, lead to many more terms in the HEOM which are only valid to second-order in system-bath coupling, such that in practice only Drude spectral densities are employed.^{90,91} At present, *PHI* treats spectral densities consisting of single or sums of Drude spectral densities.

Hierarchy equations of motion

In deriving the HEOM each exponential term in Eq. (9) leads to the introduction of a set of operators, $\hat{\rho}_{\mathbf{n}}$, called auxiliary density matrices (ADMs).⁶¹ The $\hat{\rho}_{\mathbf{n}}$ are indexed by a vector $\mathbf{n} = (n_{10}, \dots, n_{1K}, \dots, n_{M0}, \dots, n_{MK})$. The equations of motion, which couple the ADMs in a hierarchical structure, are^{34,35,61,71}

$$\begin{aligned} \dot{\hat{\rho}}_{\mathbf{n}} = & -\frac{i}{\hbar} [\widehat{H}_s, \widehat{\rho}_{\mathbf{n}}] \\ & - \sum_{a=1}^M \sum_{k=0}^K n_{ak} \nu_{ak} \widehat{\rho}_{\mathbf{n}} - \sum_{a=1}^M \left(\frac{2\lambda_a}{\beta\hbar^2 \gamma_a} - \sum_{k=0}^K \frac{c_{ak}}{\hbar \nu_{ak}} \right) [\widehat{F}_a, [\widehat{F}_a, \widehat{\rho}_{\mathbf{n}}]] - i \sum_{a=1}^M \left[\widehat{F}_a, \sum_{k=0}^K \widehat{\rho}_{\mathbf{n}_{ak}^+} \right] - \frac{i}{\hbar} \sum_{a=1}^M \sum_{k=0}^K n_{ak} (c_{ak} \widehat{F}_a \widehat{\rho}_{\mathbf{n}_{ak}^-} - \widehat{\rho}_{\mathbf{n}_{ak}^-} \widehat{F}_a c_{ak}^*). \end{aligned} \quad (12)$$

Each ADM with index \mathbf{n} is coupled to ADMs with indices $\mathbf{n}_{ak}^{\pm} = (n_{10}, \dots, n_{ak} \pm 1, \dots, n_{MK})$. As there are no negative indices n_{ak} arising in the derivation of the HEOM,⁶⁷ any ADM with a negative value for any n_{ak} is set to zero. The system density matrix $\hat{\rho}$ is the density matrix with index vector $\mathbf{n} = 0$. Each density matrix in the hierarchy is assigned to a hierarchy level $L = \sum_{a=1}^M \sum_{k=0}^K n_{ak}$. The coupling of the HEOM leads to a structure that can be visualized as a Pascal's d -simplex, where $d = M(K+1)$. Example hierarchies for $d = 2, 3$ and 4 are shown Figure 1.

The number of matrices rapidly increase with hierarchy level L (Figure 2) and, in principle, is infinite. The ADMs account for the non-Markovian extent of the dynamics, meaning that more ADMs are needed to capture dynamics with greater non-Markovian character. Due to this behavior, however, a natural truncation L_T can be chosen such that $L_T \min_{ak} (\nu_{ak}) \gg \omega_{\max}$, where ω_{\max} is the maximum oscillation frequency in the system, without compromising the dynamics.⁶¹ Employing such truncation is often unfeasible due to computational limitations and, thus, convergence needs to be checked for each system to arrive at feasible values of L_T .

Having chosen a truncation level L_T , the manner in which the HEOM is cut-off must also be specified. There are two ways to truncate the HEOM, either with so-called time non-local (TNL) truncation that sets all ADMs in levels L_T or greater to zero,⁶⁰ or with so-called time-local (TL) truncation that employs the Markovian approximation for ADMs in level L_T .⁹³ In case of the Markovian approximation it is assumed that for all ADMs with $L = L_T - 1$

$$\sum_{k=0}^K \widehat{\rho}_{\mathbf{n}_{ak}^+} \approx -i (\widehat{Q}_a^K(t) \rho_{\mathbf{n}} - \rho_{\mathbf{n}} \widehat{Q}_a^K(t)^\dagger), \quad (13)$$

where

$$\widehat{Q}_a^K(t) = \int_0^t \left(\sum_{k=0}^K c_{ak} \exp(-\nu_{ak}\tau) \right) \exp\left(-\frac{i}{\hbar} H_s \tau\right) \widehat{F}_a \exp\left(\frac{i}{\hbar} H_s \tau\right) d\tau. \quad (14)$$

It has been shown that TNL truncation produces spurious peaks in absorption spectra^{66,71} and, thus, TL truncation should be employed whenever possible.

3 Numerical integration of HEOM

The HEOM present a significant computational challenge for all systems but one, a two state system. Within single pigment-protein complexes in photosynthetic systems, pigment numbers lie typically between 6 and 90, depending on species and growth conditions.¹⁷ To overcome the issues of poor scaling with system size and truncation (Figure 2), and because of the highly coupled nature of the equations, the HEOM need to be implemented computationally in the most efficient manner. This section describes the implementation of the HEOM in the parallel hierarchy integration software *PHI* capable of calculating excitation dynamics for systems containing up to 60 pigments at permissible truncation levels.³⁶

Adaptive integration

Since the HEOM describes thermally averaged time evolution, the ADMs will exponentially approach, in the absence of an external driving field, a steady-state. Any initial state of the system will relax to a steady state, leading during the relaxation to increasingly smaller changes in time of the elements of each ADM.

Shi et al.⁶⁹ proposed a rescaling of the ADMs that permits an automatic removal of the ADMs while integrating the HEOM. During integration, the ADMs are adaptively removed when their maximum absolute value is less than a defined filter accuracy δ_F . This filtering greatly reduces the memory requirements for integrating the HEOM for moderate-sized systems,^{69,72} but can lead to artifacts with choices of δ_F values that are too large.

By employing the same rescaling of the ADMs as employed by Shi et al.⁶⁹ one can exploit the damped dynamics by integrating the HEOM with an adaptive timestep. Among the many adaptive integration methods, the Runga-Kutta-Fehlberg (RKF45) method⁹⁴ stands out as most practical as it is a 5th order integration algorithm that does not require the re-calculation of a timestep with a different step size in order to evaluate the integration error. The method, instead, re-uses the interim Runga-Kutta steps to also compute a 4th order integration step; at the end of a single timestep, the 4th order timestep calculation is compared with the 5th order timestep calculation and the step is kept or discarded according to a tolerance δ_I and the stepsize is increased or reduced accordingly.

Adaptive integration by the RKF45 algorithm requires 6 evaluations of the equations of motion, compared to 4 evaluations when using Runga-Kutta 4 (RK4) fixed timestep integration. The additional computation per RK45 step compared to a RK4 step, however, typically results in an order of magnitude decrease in the number of integration steps compared to RK4 with an overall reduction in total computation time. Rapid initial dynamics as well as long-time relaxation to the steady-state are both covered efficiently, with an overall short calculation time and a pre-defined integration error tolerance.

Adaptive timestep integration, however, cannot typically be used in conjunction with the adaptive ADM filtering method of Shi et al. as the latter introduces artifacts in the time evolution of the ADMs that result in the adaptive integration method taking smaller time-steps than without the automatic filtering. As will be shown in Results below, suitable choices of δ_F and δ_I values still result in speed-up within an acceptable error.

Partitioning

Central to efficient integration of the HEOM is the use of parallel computing. Unfortunately, large numbers of processors linked in computer clusters (i.e., in distributed memory computers) with a low network bandwidth cannot be used since the HEOM, particularly for systems of many pigments, are extremely coupled, and require a high bandwidth for inter-processor communication. This communication cost can, however, be managed by employing multiple threads running on a multi-processor, shared memory computer.

In principle it is relatively straightforward to assign P processing threads (where each thread is an independently executing process able to communicate with other threads and access the same shared memory) to a set of ADMs. Integration then proceeds by computing and applying the integration updates from the HEOM in parallel, pausing between each step to synchronize data between the P threads. The communication overhead due to the highly coupled nature of the HEOM means that even with modern symmetric multi-processing (SMP) hardware, however, the ADMs need to be well partitioned to avoid memory contention and maximize use of memory bandwidth available to each processor.

An algorithm for generating the index vectors \mathbf{n} of the ADMs and, thus, the vertices of the associated Pascal's d -simplex, is given in Listing 1. The partitioning algorithm (Listing 2) employs the numbering scheme in Listing 1 to assign each vertex to a set such that the number of inter-set edges is minimal whilst ensuring that each set contains almost the same number of vertices. An example of the partitioning of a hierarchy of ADMs associated with a Pascal's 3-simplex is shown in Figure 3. The partitioning algorithm presented here reduces the inter-set connections compared to either block- or cyclic-partitioning schemes often employed in parallel programming. Assigning each set to a separate thread and, crucially, instructing each thread to allocate the memory required to store the ADMs in its set, reduces inter-processor memory contention and dramatically improves the performance of the parallel HEOM integration software *PHI*.

4 Results

Excitation dynamics is a key process in photosynthetic light harvesting. Excitation needs to be transferred over many nanometers for light energy to be utilized by the cell. Since the excitation life time is rather short, namely only 1 ns, transfer needs to be fast to be efficient. In purple photosynthetic bacteria, such as *Rps. molischianum*, the needed efficiency is achieved by placing pigment molecules into ring-like structures as shown in Figure 4 and transferring excitation between the ring-like structures. Light energy typically absorbed by light harvesting complex 2 (LH2, B850 ring)³² is passed by inter-complex excitation transfer to light-harvesting complex 1 (LH1, B875 ring)⁹⁵ prior to charge separation at the photosynthetic reaction center.¹⁷ The close packing arrangements and pigment orientations in the ring-like structures result in quantum coherence that enhances excitation transfer between the rings.³⁹

We present here the effects of employing different HEOM integration parameters and methods. First, we examine the effect of adaptive timestep integration for a pigment dimer and for an 18-pigment ring, the so-called B850 ring in LH2.³² Second, we employ the B850 ring as a model system to investigate the parallel performance of *PHI*. Finally, the transfer of excitation between the two ring-structures, the B850 ring and the 32-pigment B875 ring in LH1⁹⁵ (shown in Figure 4), is computed.

Effect of integration timestep and method

Below, results are presented to demonstrate the effect of fixed and adaptive timesteps on integration speed and accuracy. First, a model dimer system is employed to highlight the importance of choosing an appropriate timestep. The dimer is specified by $HS = \Delta E \hat{\sigma}_z + V \hat{\sigma}_x$, where $\Delta E = 100 \text{ cm}^{-1}$, $V = 200 \text{ cm}^{-1}$, reorganization energies are $\lambda_1 = \lambda_2 = 50 \text{ cm}^{-1}$, and response frequencies are $\gamma_1 = \gamma_2 = 1/(0.1 \text{ ps})$; $\hat{\sigma}_z$ and $\hat{\sigma}_x$ are Pauli matrices.¹⁵ As can be seen from Figure 5, using too large a timestep can rapidly cause integration errors. Fortunately, in such a simple system such gross errors are easy to detect and rectify. For larger systems, however, finding the optimum timestep can be time-consuming. By employing the RKF45 adaptive timestep integration method (see Methods) such effort can be avoided and it becomes much easier to achieve greatly improved performance.

Effect of integration timestep and method

To examine the accuracy and speed of different integration options we consider the case of light excitation of one of the primary absorbing states of the B850 ring of pigments in light harvesting complex 2 (LH2). The excitation dynamics of the 18-pigment B850 ring was previously characterized³⁴ and calculations are repeated here to investigate the effect of different integration methods. The resulting exciton dynamics is shown in Figure 6 for the reference calculation; Table 1 presents the accuracy and computation time of adaptive and

non-adaptive integration methods. One can clearly discern that the Shi et al. adaptive filtering dramatically reduces computation time in the case of Runga-Kutta 4 (RK4) integration. Excellent performance can also be achieved using RKF45 adaptive timestep integration, with typically shorter calculation times and improved accuracy than needed for the other methods listed. Both adaptive methods can be combined for optimal performance as long as the RKF45 tolerance is larger than the ADM filter accuracy, i.e., for $\delta_T > \delta_F$; otherwise longer computation times result. In general, optimal RK4 timestep and ADM filter accuracy both depend on system parameters and, thus, require time-consuming benchmark calculations. Employing RKF45 adaptive timestep integration yields significant performance gains without prior benchmarks.

Parallel performance

Key to employing the HEOM to calculate excitation dynamics in large pigment systems is scalability to many processors. With efficient scaling, the size of problem that can be treated depends only on the size of the computer available and is not limited by the speed of a single processor. To endow *PHI* with optimal scaling capabilities, three steps were taken: (1) memory assignment was changed from being initially assigned to a single large block to being assigned to separate smaller blocks, one for each integration thread; (2) the partitioning of the ADMs was done satisfying minimum inter-thread communication; (3) each thread was assigned to a specific processor. The effect of the three steps, shown in Figure 7, is a greater than 3-fold reduction in calculation time when employing 48-cores, and a significantly better scaling compared to cyclic decomposition, a simple and obvious choice, but not an optimal one.

A significant boost in parallel performance is thus gained from the hierarchical partitioning scheme detailed in Listings 1 and 2. The performance saturation seen around 48 cores in Figure 7 indicates that inter-processor communication is still a bottleneck and any scheme to reduce this bottleneck will result in further performance improvement. The rapid bandwidth saturation, arising even with the high-bandwidth processor-memory bus (42 GB/sec for the AMD Magny-Cours CPU), suggests that parallel performance on a distributed memory computer is relatively poor.

Excitation transfer between LH2 and LH1

With the improvements in scalability and calculation time for integration of the HEOM described above, one can model excitation transfer between two large pigment-protein complexes, LH2 and LH1 (Figure 4).

In the photosynthetic light harvesting apparatus of purple bacteria most excitation transfer steps occur between the 18-pigment B850 ring in LH2 and the 32-pigment B875 ring in LH1.^{18,19,32} In order to describe this transfer one needs to account for $N = 50$ pigments. Modeling a system of this size is not possible with a single-processor implementation of the HEOM, but requires *PHI* running on a parallel computer.

The parameters describing the B850 and B875 Hamiltonians are taken from Ref. 48, with bath parameters $\lambda_n = \lambda = 180 \text{ cm}^{-1}$ and $1/\gamma_n = 1/\gamma = 100 \text{ fs}$, a choice that had been shown to reproduce experimental spectra.⁵⁴ The population of an initially excited LH2 complex is shown in Figure 8 for different truncations and inter-complex separation distances. The results show that TL truncation with a cut-off of $L_T = 3$ yields a population decay close to that of an $L_T = 5$ cut-off. The lower truncation requires 50 MB of memory to store the hierarchy of ADMs and 415 MB of memory for RKF45 integration; the higher truncation requires, however, 12 GB of memory to store the hierarchy of ADMs and 96 GB of memory for RKF45 integration.

The excitation transfer time can be calculated by fitting a simple kinetic model

$$\frac{dP_{B850}}{dt} = -\frac{P_{B850}}{\tau_{B850 \rightarrow B875}} + \frac{1 - P_{B850}}{\tau_{B875 \rightarrow B850}} \quad (15)$$

describing excitation transfer between B850 and B875 BChl rings (Figure 8). In our calculations we assumed three separations between the proteins LH1 and LH2 holding the B875 and B850 BChl rings, respectively, namely, edge-edge separations of 2 nm, 1 nm, and 0 nm (contact). In case of contact, transfer times (summarized in Table 2) are $\tau_{B850 \rightarrow B875} = 5.3$ ps and $\tau_{B875 \rightarrow B850} = 9.2$ ps; inter-protein separations of 1 nm and 2 nm result in 260% and 590% longer transfer times, respectively.

5 Conclusion

The hierarchy equations of motion (HEOM) employ a highly coupled set of density matrices to calculate the noise-averaged time evolution of a quantum system coupled to a thermal environment. The large number of matrices required to model multi-pigment systems using the HEOM has restricted the method to systems with only very few pigments. By employing adaptive numerical integration methods (hierarchy filtering of Shi et al.⁶⁹ and adaptive timestep integration⁹⁴) and a novel hierarchy partitioning scheme, efficient scaling has been realized. The software implementation available in the program *PHI* can be run on parallel computers allowing the HEOM to be solved for pigment numbers as large as found in biological systems. The performance of *PHI* has been demonstrated here by investigating the accuracy of the adaptive integration methods and by applying *PHI* to model excitation transfer in a fifty pigment system, namely the 18-BChl B850 ring in light harvesting complex 2 (LH2) and 32-BChl B875 ring in light harvesting complex 1 (LH1) of purple photosynthetic bacteria. It was shown that excitation transfer between LH2 and LH1 practically occurs as a Poisson process, i.e., an incoherent hopping process, described by a single exponential decay of excited state population. The excitation transfer times for LH2→LH1 and LH1→LH2 were calculated for direct protein-protein contact, as well as for protein-protein separations of 1 nm and 2 nm. Resulting transfer rates compare well with observations²³ and generalized Förster theory results;⁷⁸ the latter agreement should be considered an a posteriori justification of generalized Förster theory.

The efficient parallel implementation of the HEOM makes it possible to model even larger light harvesting systems than studied here such as photosystem 1 and photosystem 2 in cyanobacteria and plants,^{73,74} as well as conjugated polymers⁷⁶ and pigment dendrimers⁷⁷ used in artificial light harvesting. The current implementation of the HEOM would also benefit the calculation of exact quantum dynamics in Wigner space, recently shown to also be computationally challenging due to large memory requirements,⁹⁶ and is a promising future feature of *PHI*. Additionally, the general applicability of the HEOM lets *PHI* to be used for investigations of other open quantum systems such as charge transfer reactions,⁹⁰ e.g., in DNA⁹⁷ and in the photosynthetic reaction center,⁹⁸ or coherently coupled spin systems such as those employed in quantum computing.^{99–101}

Acknowledgments

The authors are grateful for insightful discussions with John Stone. Funding for J.S. and K.S. was provided by NSF grant PHY0822613 and NIH grants MCB-0744057 and P41-RR05969.


References

1. Hopfield JJ. Proc. Natl. Acad. Sci. USA. 1974; 71:3640–3644. [PubMed: 16592178]

2. Knox, RS. Primary Processes of Photosynthesis. Barber, J., editor. Amsterdam: Elsevier; 1977. p. 55-97.
3. Marcus RA, Sutin N. *Biochim. Biophys. Acta.* 1985; 811:265–322.
4. May, V.; Kühn, O. Charge and Energy Transfer Dynamics in Molecular Systems. Berlin: WILEYVCH; 2000.
5. Förster T. *Ann. Phys. (Leipzig).* 1948; 2:55–75.
6. Redfield A. *IBM Journal of Research and Development.* 1957; 1:19–31.
7. McLachlan AD. *Mol. Phys.* 1964; 8:39.
8. Micha DA. *J. Chem. Phys.* 1983; 78:7138–7145.
9. Haken H, Reineker P. *Z. Physik A.* 1972; 249:253–268.
10. Haken H, Strobl G. *Z. Physik A.* 1973; 262:135–148.
11. Caldeira AO, Leggett AJ. *J. Ann. Phys.(N.Y.).* 1983; 149:374–456.
12. Yang M, Fleming GR. *Chem. Phys.* 2002; 282:163–180.
13. Renger T, Marcus R. *J. Chem. Phys.* 2002; 116:9997.
14. Mukamel, S. Principles of Nonlinear Optical Spectroscopy. New York: Oxford University Press; 1995.
15. Weiss, U. Quantum dissipative systems. World Scientific Publishing Company; 2008.
16. Breuer, HP.; Petruccione, F. The theory of open quantum systems. USA: Oxford University Press; 2002.
17. Blankenship, RE. Molecular Mechanisms of Photosynthesis. Malden, MA: Blackwell Science; 2002.
18. Sener MK, Olsen JD, Hunter CN, Schulten K. *Proc. Natl. Acad. Sci. USA.* 2007; 104:15723–15728. [PubMed: 17895378]
19. Sener M, Strumpfer J, Timney JA, Freiberg A, Hunter CN, Schulten K. *Biophys. J.* 2010; 99:67–75. [PubMed: 20655834]
20. Freiberg A, Godik VI, Pullerits T, Timpmann K. *Biochim. Biophys. Acta.* 1989; 973:93–104.
21. Visscher KJ, Bergstrom H, Sundström V, Hunter C, van Grondelle R. *Photosyn. Res.* 1989; 22:211–217.
22. Pullerits T, Visscher KJ, Hess S, Sundström V, Freiberg A, Timpmann K. *Biophys. J.* 1994; 66:236–248. [PubMed: 8130341]
23. Hess S, Chachisvilis M, Timpmann K, Jones M, Fowler G, Hunter C, Sundström V. *Proc. Natl. Acad. Sci. USA.* 1995; 92:12333–12337. [PubMed: 11607622]
24. Leegwater JA. *J. Phys. Chem.* 1996; 100:14403–14409.
25. Sundström, V.; Grondelle, R. Anoxygenic Photosynthetic Bacteria. Blankenship, R.; Madigan, M.; Bauer, C., editors. Vol. Vol. 2. Springer Netherlands: Advances in Photosynthesis and Respiration; 2004. p. 349-372.
26. van Grondelle R, Novoderezhkin VI. *Phys. Chem. Chem. Phys.* 2006; 8:793–807. [PubMed: 16482320]
27. Lee H, Cheng YC, Fleming GR. *Science.* 2007; 316:1462–1465. [PubMed: 17556580]
28. Pullerits T, Polivka T, Sundström V. *Ultrashort Laser Pulses in Biology and Medicine.* 2008
29. Caruso F, Saikin SK, Solano E, Huelga SF, Aspuru-Guzik A, Plenio MB. *Phys. Rev. B.* 2012; 85:125424.
30. Linnanto J, Korppi-Tommola J. *Chem. Phys.* 2009; 357:171–180.
31. Caycedo-Soler F, Rodríguez FJ, Quiroga L, Johnson NF. *New J. Phys.* 2010; 12:095008.
32. Hu X, Ritz T, Damjanovi A, Schulten K. *J. Phys. Chem. B.* 1997; 101:3854–3871.
33. Damjanovi A, Kosztin I, Kleinekathoefer U, Schulten K. *Phys. Rev. E.* 2002; 65:031919.
34. Strümpfer J, Schulten K. *J. Chem. Phys.* 2009; 131:225101. [PubMed: 20001083]
35. Strümpfer J, Schulten K. *J. Chem. Phys.* 2011; 134:095102. [PubMed: 21385000]
36. Hsin J, Strümpfer J, Sener M, Qian P, Hunter CN, Schulten K. *New J. Phys.* 2010; 12:085005.
37. Huelga SF, Rivas A, Plenio MB. *Phys. Rev. Lett.* 2012; 108:160402. [PubMed: 22680702]
38. Ishizaki A, Fleming G. *J. Phys. Chem. B.* 2011; 115:6227–6233. [PubMed: 21488648]

39. Strümpfer J, Sener M, Schulten K. *J. Phys. Chem. Lett.* 2012; 3:536–542. [PubMed: 22844553]
40. Rebentrost P, Mohseni M, Aspuru-Guzik A. *J. Phys. Chem. B.* 2009; 113:9942–9947. [PubMed: 19603843]
41. Plenio M, Huelga S. *New J. Phys.* 2008; 10:113019.
42. Caruso F, Chin A, Datta A, Huelga S, Plenio M. *J. Chem. Phys.* 2009; 131:105106.
43. Chin A, Datta A, Caruso F, Huelga S, Plenio M. *New J. Phys.* 2010; 12:065002.
44. Caruso F, Chin A, Datta A, Huelga S, Plenio M. *Phys. Rev. A.* 2010; 81:062346.
45. Mohseni M, Rebentrost P, Lloyd S, Aspuru-Guzik A. *J. Chem. Phys.* 2008; 129:174106. [PubMed: 19045332]
46. Rebentrost P, Mohseni M, Kassal I, Lloyd S, Aspuru-Guzik A. *New J. Phys.* 2009; 11:033003.
47. Rebentrost P, Chakraborty R, Aspuru-Guzik A. *J. Chem. Phys.* 2009; 131:184102. [PubMed: 19916593]
48. Hu X, Damjanovi A, Ritz T, Schulten K. *Proc. Natl. Acad. Sci. USA.* 1998; 95:5935–5941. [PubMed: 9600895]
49. Timpmann K, Trinkunas G, Qian P, Hunter CN, Freiberg A. *Chem. Phys. Lett.* 2005; 414:359–363.
50. Trinkunas G, Freiberg A. *J. Luminesc.* 2006; 119–120:105–110.
51. Valkunas L, Janusonis J, Rutkauskas D, van Grondelle R. *J. Luminesc.* 2007; 127:269–275.
52. Janusonis J, Valkunas L, Rutkauskas D, van Grondelle R. *Biophys. J.* 2008; 94:1348–1358. [PubMed: 17921215]
53. Olaya-Castro A, Lee CF, Olsen FF, Johnson NF. *Phys. Rev. B.* 2008; 78:085115.
54. Freiberg A, Ratsep M, Timpmann K, Trinkunas G. *Chem. Phys.* 2009; 357:102–112.
55. Prior J, Chin A, Huelga S, Plenio M. *Phys. Rev. Lett.* 2010; 105:50404.
56. Chin A, Rivas Á, Huelga S, Plenio M. *J. Math. Phys.* 2010; 51:092109.
57. Topaler M, Makri N. *Chem. Phys. Lett.* 1993; 210:448–457.
58. Tanimura Y, Kubo R. *J. Phys. Soc. Jpn.* 1989; 58:1199–1206.
59. Tanimura Y. *Phys. Rev. A.* 1990; 41:6676–6687. [PubMed: 9903081]
60. Ishizaki A, Tanimura Y. *J. Phys. Soc. Jpn.* 2005; 74:3131–3134.
61. Ishizaki A, Fleming GR. *J. Chem. Phys.* 2009; 130:234111–234110. [PubMed: 19548715]
62. Yan Y, Yang F, Liu Y, Shao J. *Chem. Phys. Lett.* 2004; 395:216–221.
63. Xu RX, Cui P, Li XQ, Mo Y, Yan YJ. *J. Phys. Chem.* 2005; 122:041103.
64. Yan YJ, Xu RX. *Annu. Rev. Phys. Chem.* 2005; 56:187. [PubMed: 15796700]
65. Tanimura Y. *J. Phys. Soc. Jpn.* 2006; 75:082001.
66. Schröder M, Kleinekathöfer U, Schreiber M. *J. Chem. Phys.* 2006; 124:084903. [PubMed: 16512738]
67. Xu RX, Yan YJ. *Phys. Rev. E.* 2007; 75:031107–031111.
68. Ishizaki A, Tanimura Y. *J. Phys. Chem. A.* 2007; 111:9269–9276. [PubMed: 17880172]
69. Shi Q, Chen L, Nan G, Xu R, Yan Y. *J. Chem. Phys.* 2009; 130:084105-4. [PubMed: 19256595]
70. Xu RX, Tian BL, Xu J, Yan YJ. *J. Chem. Phys.* 2009; 130:074107-8. [PubMed: 19239284]
71. Chen L, Zheng R, Shi Q, Yan Y. *J. Chem. Phys.* 2009; 131:094502. [PubMed: 19739856]
72. Kreisbeck C, Kramer T, Rodriguez M, Hein B. *J. Chem. Theor. Comp.* 2011; 7:2166–2174.
73. Wientjes E, Van Stokkum I, Van Amerongen H, Croce R. *Biophys. J.* 2011; 100:1372–1380. [PubMed: 21354411]
74. Croce R, Van Amerongen H. *J. Photochem. Photobiol. B.* 2011; 104:142–153. [PubMed: 21402480]
75. Scholes G, Fleming G, Olaya-Castro A, van Grondelle R. *Nat. Chemistry.* 2011; 3:763–774.
76. Bolinger JC, Traub MC, Adachi T, Barbara PF. *Science.* 2011; 331:565–567. [PubMed: 21292973]
77. Imahori H. *J. Phys. Chem. B.* 2004; 108:6130–6143. [PubMed: 18950092]
78. Hu X, Ritz T, Damjanovi A, Autenrieth F, Schulten K. *Quart. Rev. Biophys.* 2002; 35:1–62.

79. Sener, M.; Schulten, K. *Energy Harvesting Materials*. Andrews, DL., editor. Singapore: World Scientific; 2005. p. 1-26.
80. Ishizaki A, Fleming GR. *Proc. Natl. Acad. Sci. USA*. 2009; 106:17255. [PubMed: 19815512]
81. Zhu K, Xu R, Zhang H, Hu J, Yan Y. *J. Phys. Chem. B*. 2011; 115:5678–5684. [PubMed: 21452824]
82. Kleinekathöfer U. *J. Chem. Phys.* 2004; 121:2505. [PubMed: 15281847]
83. Schröder M, Schreiber M, Kleinekathöfer U. *J. Chem. Phys.* 2007; 126:114102. [PubMed: 17381191]
84. Ishizaki A, Fleming GR. *J. Chem. Phys.* 2009; 130:234110-8. [PubMed: 19548714]
85. Chen X, Silbey RJ. *J. Phys. Chem. B*. 2011; 115:5499–5509. [PubMed: 21384851]
86. Hein B, Kreisbeck C, Kramer T, Rodríguez M. *New J. Phys.* 2012; 14:023018.
87. Tian B, Ding J, Xu R, Yan Y. *J. Chem. Phys.* 2010; 133:114112. [PubMed: 20866131]
88. Hu J, Xu R, Yan Y. *J. Chem. Phys.* 2010; 133:101106. [PubMed: 20849157]
89. Tanimura Y, Mukamel S. *J. Phys. Soc. Jpn.* 1994; 63:66–77.
90. Tanaka M, Tanimura Y. *J. Phys. Soc. Jpn.* 2009; 78:73802.
91. Tanaka M, Tanimura Y. *J. Chem. Phys.* 2010; 132:214502. [PubMed: 20528026]
92. Meier C, Tannor D. *J. Chem. Phys.* 1999; 111:3365.
93. Tanimura Y, Wolyne P. *Phys. Rev. A*. 1991; 43:4131–4142. [PubMed: 9905511]
94. Fehlbeg E. *Computing*. 1985; 34:265–270.
95. Hu X, Schulten K. *Biophys. J.* 1998; 75:683–694. [PubMed: 9675170]
96. Sakurai A, Tanimura Y. *J. Phys. Chem. A*. 2011
97. Dijkstra A, Tanimura Y. *New J. Phys.* 2010; 12:055005.
98. Xu D, Schulten K. *Chem. Phys.* 1994; 182:91–117.
99. Palma GM, Suominen KA, Ekert AK. *Proc. R. Soc. A*. 1996; 452:567–584.
100. Stamp P, Gaita-Ariño A. *J. Mater. Chem.* 2009; 19:1718–1730.
101. Caruso F, Huelga SF, Plenio MB. *Phys. Rev. Lett.* 2010; 105:190501. [PubMed: 21231154]

**Figure 1.**

The first 4 hierarchy levels of Pascal's d -simplices for $d = M(K+1) = 2, 3,$ and 4 . Each circle corresponds to a single ADM with the corresponding index vectors \mathbf{n} shown in blue text. The numbering inside the circles corresponds to the scheme described in the hierarchy generation algorithm in Listing 1. The shaded regions indicate which ADMs are at the same level. Connectivity of the ADMs to those in the levels above and below is not shown.

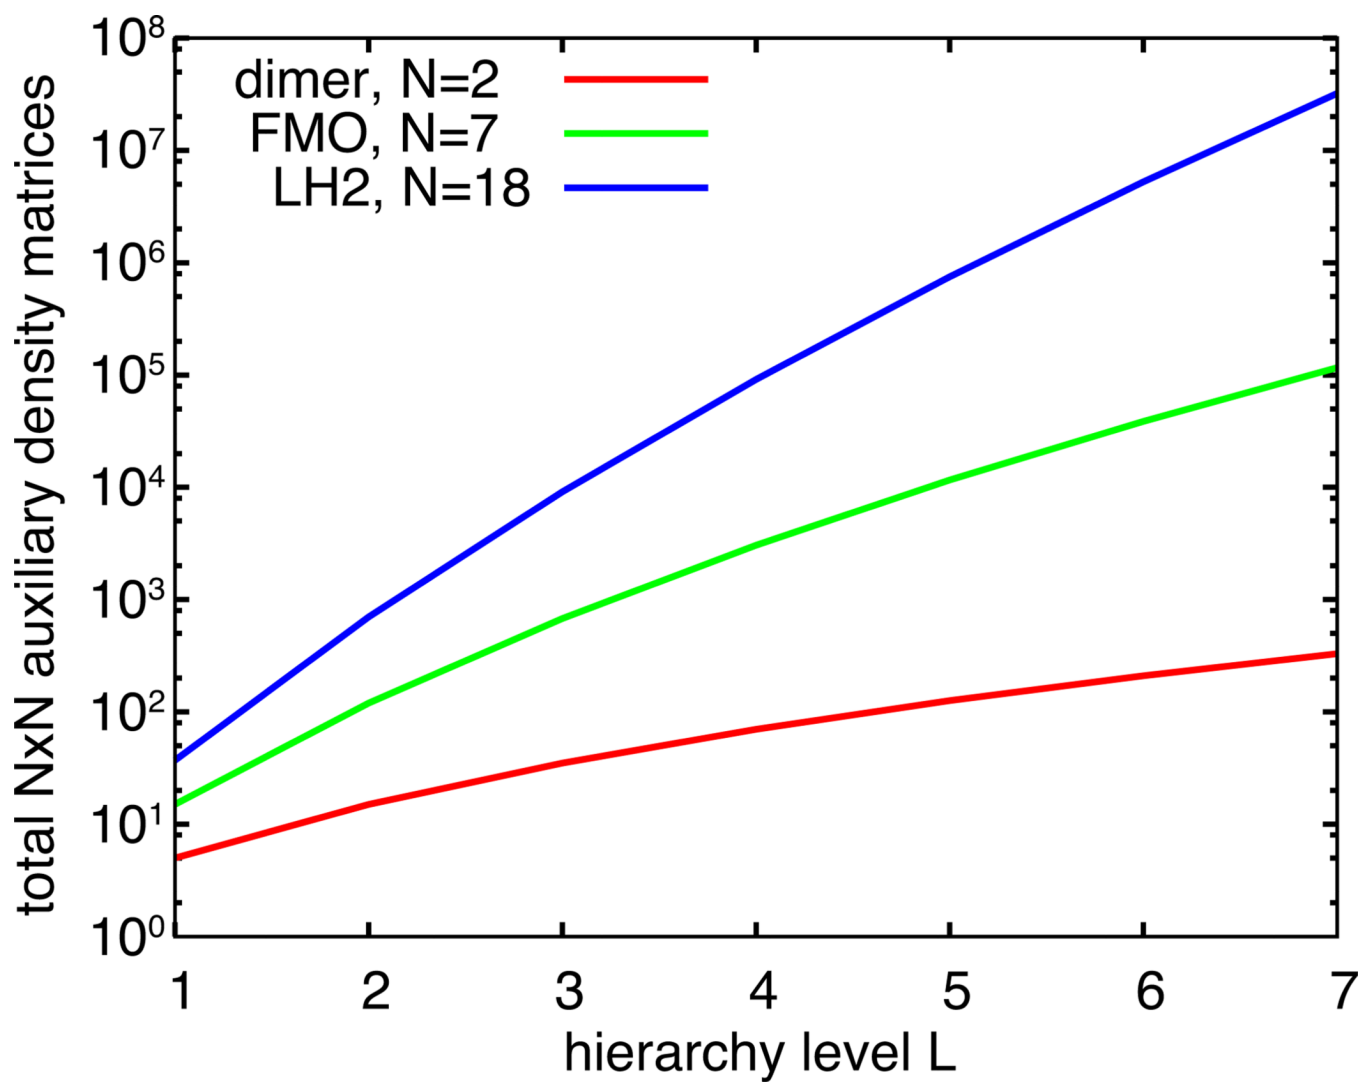


Figure 2. Total number of ADMs with increasing hierarchy level L for a pigment dimer, for the Fenna-Mathews-Olson (FMO) complex⁸⁰ and for light-harvesting complex 2 (LH2).³⁴

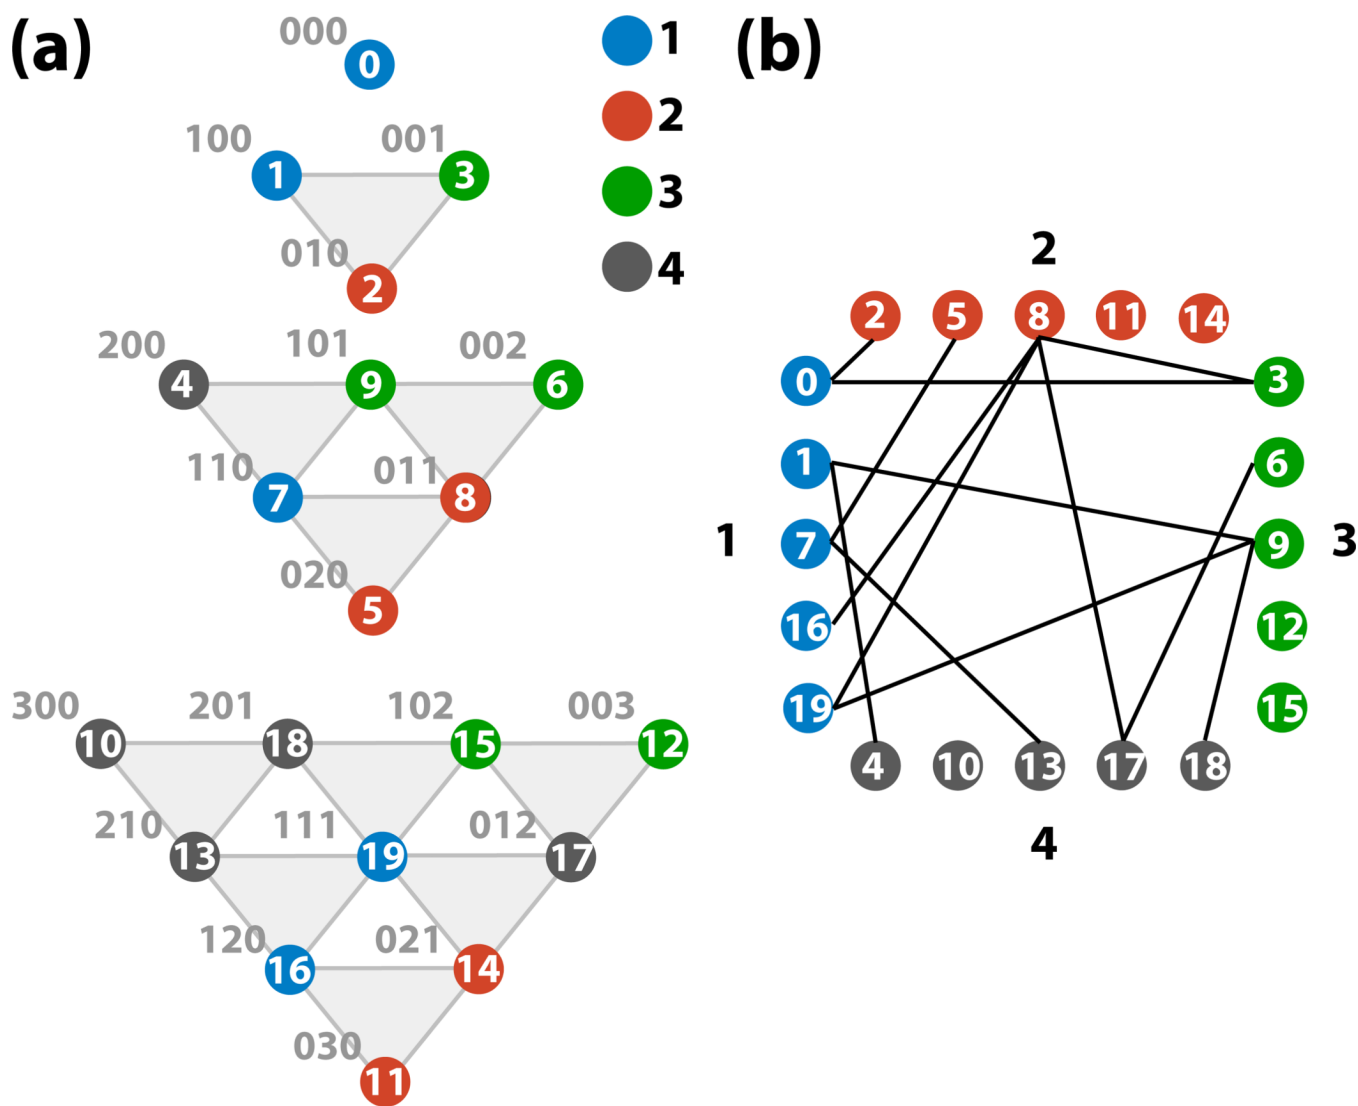


Figure 3. Partitioning of a Pascal's 3-simplex into four sets. The partitioning (a) results in each subset containing 5 vertices, with (b) 13 edges crossing between subsets.




Figure 4. Pigment-protein complexes light harvesting complex 2 (LH2) and light harvesting complex 1 (LH1). Shown are the 18 pigments forming the B850 ring of LH2 and the 32 pigments forming the B875 ring of LH1.^{32,95}




Figure 5. Comparison of HEOM integration schemes. Integration on 4 processors (CPU cores) took 0.23 minutes for Runge-Kutta 4 integration using 0.1 fs timesteps, and 0.03 minutes using Runge-Kutta-Fehlberg 4/5 adaptive integration with a minimum timestep of 0.1 fs and a tolerance of 0.0001.

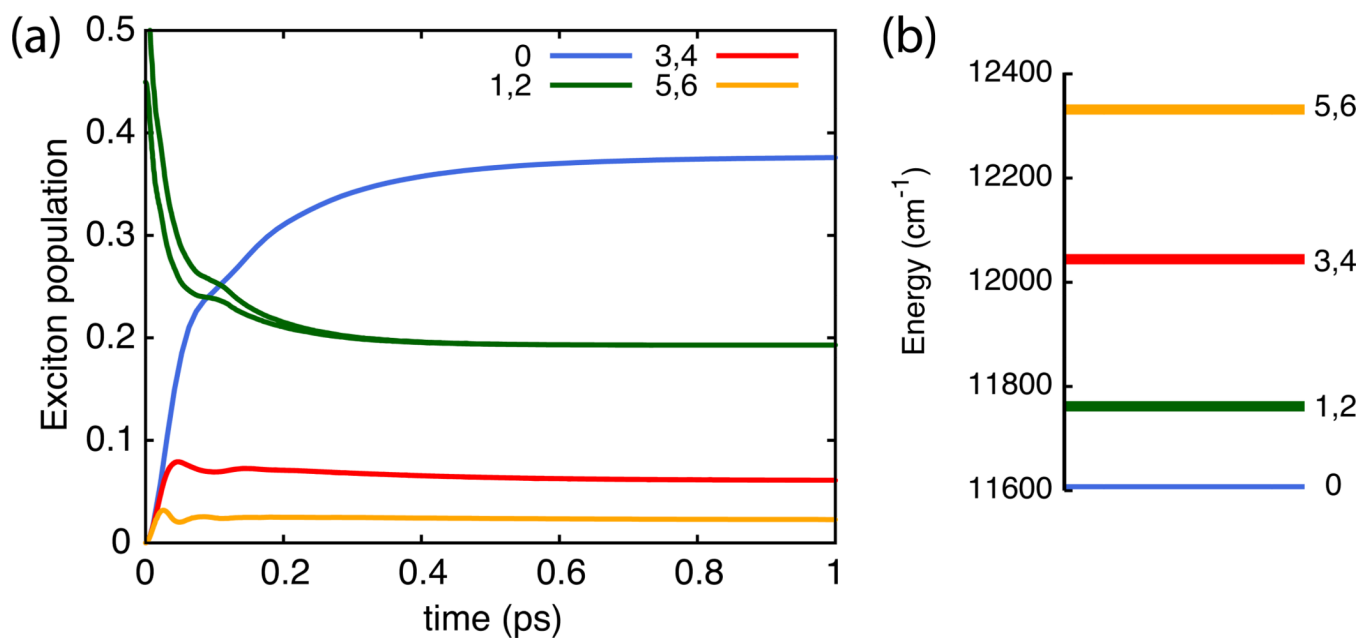


Figure 6. Exciton dynamics of the seven lowest-energy exciton states of the B850 ring in LH2 with the initial state given by a super-position of the degenerate 850 nm exciton states 1 and 2. Populations are shown in (a) and energies of associated states in (b).

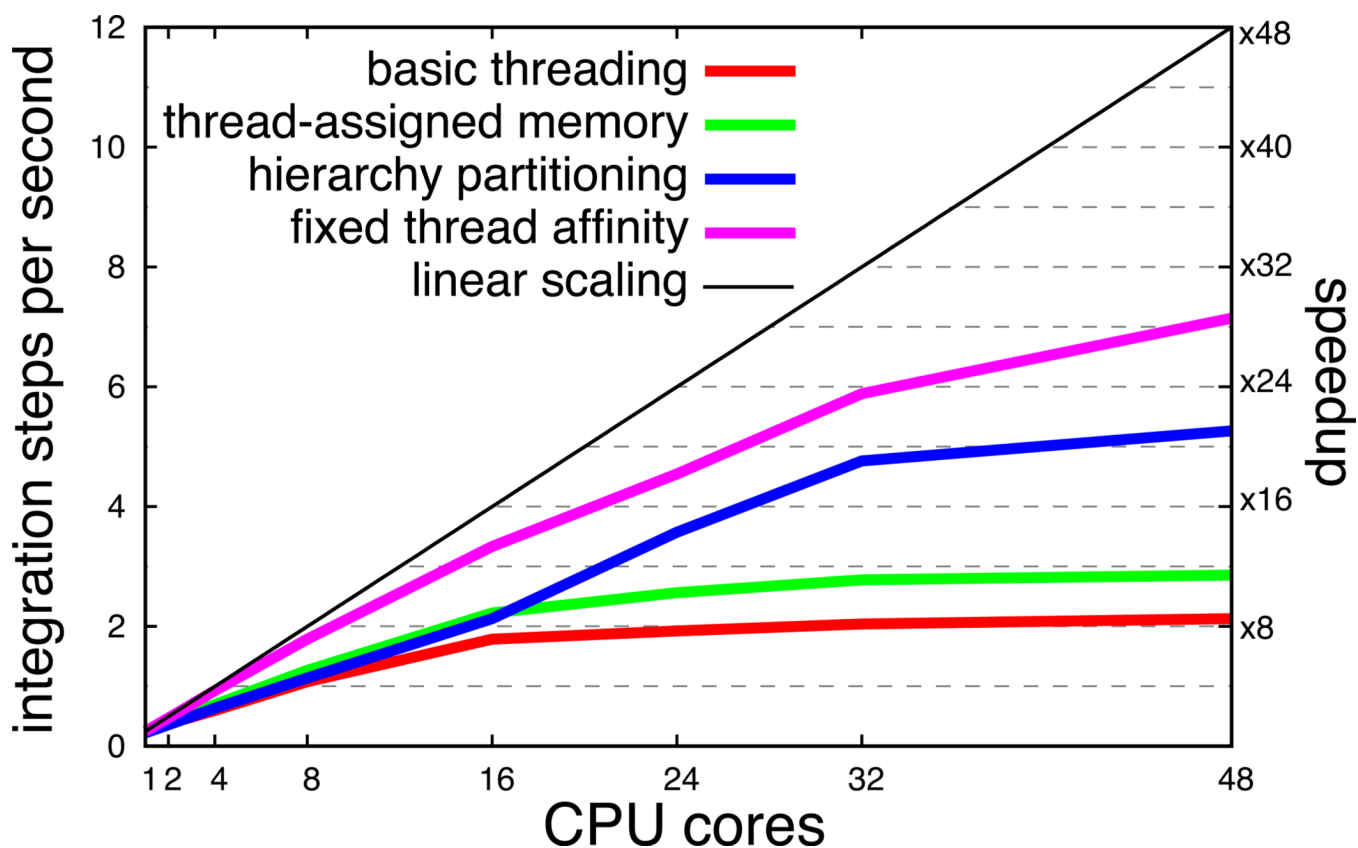


Figure 7.

Parallel scaling of *PHI* on a 48-core AMD Magny-Cours shared-memory computer running Linux. The system employed for benchmarking is the set of 18 pigments in the B850 ring of LH2 with $L_T=4$ and $K=1$. Note that increasing the number of Matsubara frequencies K results in a dramatic reduction in performance of *PHI*. “Basic partitioning” refers to a cyclic partitioning of ADMs and the memory for all ADMs being assigned to a consecutive block; “thread-assigned memory” refers to instructing each thread to perform memory assignment for its set of ADMs in the cycling partitioning scheme; “hierarchy partitioning” refers to employing the partitioning described in Listings 1 and 2 and using thread-assigned memory; “fixed thread affinity” refers to employing the partitioning scheme in Listings 1 and 2, using thread-assigned memory and instructing the operating system’s thread scheduler to affix each thread to a particular processor such that no thread migration occurs.




Figure 8. Total excited state population of LH2 reflecting inter-complex excitation transfer from LH2 to LH1 at different protein (LH2) - protein (LH1) edge-edge separations.

```

1  d = M*(K+1);
2  vertexID = 0;
3  indexVectorList[vertexID][] := vector of d zeros;
4  verticesInLevel[0][0] = vertexID;
5  numVertsInLevel[0] = 1
6  vertexID += 1;
7  totalNumVertices = 1;
8  for k = 1 to truncationLevel-1
9      vertexCount = 0;
10     for n = 0 to d-1
11         for i = 0 to numVerticesInLevel[k-1]-1
12             prevVertexID = verticesInLevel[k-1][i]
13             tmpIndexVector[] := indexVectorList[prevVertexID][];
14             m = (i+n) modulo d;
15             tmpIndexVector[m] += 1
16             if tmpIndexVector not in indexVectorList[]
17                 indexVectorList[vertexID][] := tmpIndexVector[];
18                 verticesInLevel[k][vertexCount] = vertexID;
19                 vertexID += 1;
20                 vertexCount += 1;
21         numVertsInLevel[k] = vertexCount;
22         totalNumVertices += vertexCount;

```

Listing 1.

Pseudo-code to generate index vectors \mathbf{n} for any Pascal's d -simplex, where $d = M(K+1)$ for the HEOM, up to a specified truncation level. Scalar assignment operations are indicated by “=”; vector assignment operations are indicated by “:=”.

```

1  d=M*(K+1)
2  numVertsInSet[] := vector of P zeros;
3  verticesInSet[0][0] = 0;
4  numVertsInSet[0] = 1;
5  for vertexID = 1 to d
6    setID = (vertexID - 1) modulo P;
7    verticesInSet[setID][numVertsInSet[setID]] = vertexID;
8    numVertsInSet[setID] += 1;
9    nextCornerVertexID = vertexID + d;
10  for k = 2 to truncationLevel-1
11    verticesInSet[setID][numVertsInSet[setID]] = nextCornerVertexID;
12    numVertsInSet[setID] += 1;
13    nextCornerVertexID += numVertsInLevel[k]
14  for vertexID = 2*d+1 to totalNumVertices-1
15    if vertexID not in verticesInSet[][]
16      if P > d and vertexID < P+d+1
17        setID = vertexID modulo P;
18        verticesInSet[setID][numVertsInSet[setID]] = vertexID;
19        numVertsInSet[setID] += 1;
20      else
21        connectionsToSet[] := vector of P zeros;
22        for j = 0 to d-1
23          tmpIndexVector[] := indexVectorList[vertexID][];
24          tmpIndexVector[j] -= 1;
25          tmpVertexID = vertex associated with tmpIndexVector
26          setID = id of set containing tmpVertexID
27          connectionsToSet[setID] += 1;
28        connMax[] := all p where connectionsToSet[p] == max(connectionsToSet[]);
29        setID = p in connMax[] where numVertsInSet[p] == min(numVertsInSet[]);
30        if numVertsInSet[setID] >= ceil(totalNumVertices / P)
31          setID = id where numVertsInSet[id] == min(numVertsInSet[]);
32        verticesInSet[setID][numVertsInSet[setID]] = vertexID;
33        numVertsInSet[setID] += 1;

```

Listing 2.

Algorithm to partition the vertices of a Pascal's d -simplex into P sets, minimizing inter-set connections while ensuring the numbers of vertices in each set differ by at most 1. The algorithm employs the results of the index vector generation code in Listing 1. Scalar assignment operations indicated using "="; vector assignment operations indicated using ":="; the "==" operator indicates a test for equality.

Table 1

Accuracy and performance of different integration methods for calculating excitation dynamics of the B850 pigments of LH2 with parameters and truncation as given in Ref. 34.

integration method	ADM filter [†] δ_f	integration tolerance [‡] δ_t	integration stepsize (fs)	average error	calculation time (min)
RK4			0.1	0	920
RK4			1.0	2.0×10^{-7}	91.8
RK4			2.0	5.4×10^{-7}	46.0
RK4	10^{-6}		1.0	2.0×10^{-7}	83.9
RK4	10^{-6}		2.0	5.5×10^{-7}	42.1
RK4	10^{-4}		1.0	7.8×10^{-6}	43.4
RK4	10^{-4}		2.0	1.1×10^{-5}	25.7
RK4	10^{-2}		1.0	2.7×10^{-3}	34.5
RK4	10^{-2}		2.0	2.9×10^{-3}	21.9
RKF45		10^{-6}	$1.0 \rightarrow 7.5$	6.1×10^{-7}	35.6
RKF45		10^{-4}	$1.0 \rightarrow 7.4$	8.2×10^{-7}	33.3
RKF45		10^{-2}	$1.0 \rightarrow 8.1$	1.2×10^{-6}	33.1
RKF45	10^{-6}		$1.0 \rightarrow 2.0$	3.3×10^{-7}	66.1
RKF45	10^{-6}	10^{-4}	$1.0 \rightarrow 4.6$	6.7×10^{-7}	34.1
RKF45	10^{-6}	10^{-2}	$1.0 \rightarrow 8.1$	9.9×10^{-7}	33.0
RKF45	10^{-4}	10^{-6}	$0.6 \rightarrow 1.2$	7.8×10^{-6}	71.0
RKF45	10^{-4}	10^{-4}	$1.0 \rightarrow 2.4$	1.1×10^{-5}	34.4
RKF45	10^{-4}	10^{-2}	$1.0 \rightarrow 4.6$	1.5×10^{-5}	24.1
RKF45	10^{-2}	10^{-6}	$0.3 \rightarrow 1.2$	2.0×10^{-3}	93.2
RKF45	10^{-2}	10^{-4}	$0.7 \rightarrow 2.4$	2.8×10^{-3}	50.9
RKF45	10^{-2}	10^{-2}	$1.0 \rightarrow 3.8$	3.3×10^{-3}	27.9

[†] filter constant for Shi et al. adaptive ADM filtering.⁶⁹

[‡] integration tolerance for Runge-Kutta-Fehlberg 4/5 adaptive integration.⁹⁴

$\|error\| = (N^2 N_t)^{-1} \sum_n \sum_i^N |\rho_{ij}(t_n) - \rho_{ij}^{ref}(t_n)|$ is the average absolute error in each element of the system density matrix per timestep.

Table 2

Excitation transfer times between the B850 pigments in LH2 and the B875 pigments in LH1.

Hierarchy Truncation	Inter-protein distance (nm)	$\tau_{\text{B850} \rightarrow \text{B875}}$ (ps)	$\tau_{\text{B875} \rightarrow \text{B850}}$ (ps)
$L_T = 5$	0.0	4.1	18.4
$L_T = 3$	0.0	4.0	18.1
$L_T = 3$	1.0	10.4	46.7
$L_T = 3$	2.0	23.6	105