ACCOUNTS of chemical research

Article

pubs.acs.org/accounts

Ab Initio Nonadiabatic Dynamics of Multichromophore Complexes: A Scalable Graphical-Processing-Unit-Accelerated Exciton Framework

Aaron Sisto,[†,‡] David R. Glowacki,*[,†,‡,§,∥] and Todd J. Martinez*[,†,‡]

[†]PULSE Institute and Department of Chemistry, Stanford University, Stanford, California 94305, United States

[‡]SLAC National Accelerator Laboratory, Menlo Park, California 94025, United States

[§]School of Chemistry, University of Bristol, Bristol BS8 1TS, U.K.

[∥]Department of Computer Science, University of Bristol, Bristol BS8 1UB, U.K.

CONSPECTUS: Although advances in computer hardware and algorithms tuned for novel computer architectures are leading to significant increases in the size and time scale for molecular simulations, it remains true that new methods and algorithms will be needed to address some of the problems in complex chemical systems, such as electrochemistry, excitation energy transport, proton transport, and condensed phase reactivity. Ideally, these new methods would exploit the strengths of emerging architectures. Fragment based approaches for electronic structure theory decompose the problem of solving the electronic Schrodinger equation into a series of much smaller problems. Because each of these smaller problems is largely independent, this strategy is particularly well-suited to parallel architectures. It appears that the most significant advances in computer architectures will be toward increased parallelism, and therefore fragment-based approaches are an ideal match to these trends. When the computational effort involved scales with the third (or higher) power of the molecular size, there is a large benefit to fragment-based approaches even on serial architectures. This is the case for many of the well-known methods for solving the electronic structure theory problem, especially when wave function-based approaches including electron correlation are considered. A major issue in fragment-based approaches is determining or improving their accuracy. Since the Achilles' heel of any such method lies in the approximations used to stitch the smaller problems back together (i.e., in the treatment of the cross-fragment interactions), it can often be important to ensure that the size of the smaller problems is "large enough." Thus, there are two frontiers that need to be extended in order to enable molecular simulations for large systems and long times: the strongly coupled problem of medium sized molecules (100−500 atoms) and the more weakly coupled problem of decomposing ("fragmenting") a molecular system and then stitching it back together. In this Account, we address both of these problems, the first by using graphical processing units (GPUs) and electronic structure algorithms tuned for these architectures and the second by using an exciton model as a framework in which to stitch together the solutions of the smaller problems.

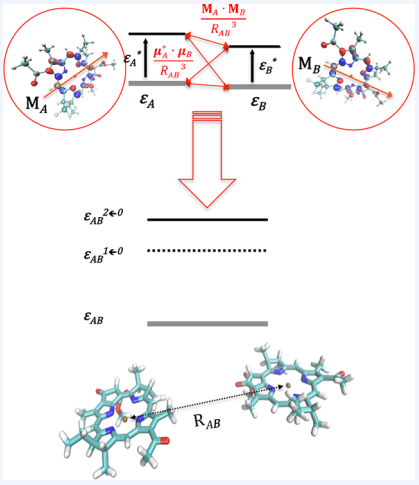

The multitiered parallel framework outlined here is aimed at nonadiabatic dynamics simulations on large supramolecular multichromophoric complexes in full atomistic detail. In this framework, the lowest tier of parallelism involves GPU-accelerated electronic structure theory calculations, for which we summarize recent progress in parallelizing the computation and use of electron repulsion integrals (ERIs), which are the major computational bottleneck in both density functional theory (DFT) and time-dependent density functional theory (TDDFT). The topmost tier of parallelism relies on a distributed memory framework, in which we build an exciton model that couples chromophoric units. Combining these multiple levels of parallelism allows access to ground and excited state dynamics for large multichromophoric assemblies. The parallel excitonic framework is in good agreement with much more computationally demanding TDDFT calculations of the full assembly.

1. INTRODUCTION

In 1929, Dirac famously wrote: "The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble. It therefore becomes desirable that approximate practical methods of applying quantum mechanics should be developed, which can lead to the main features of complex atomic systems without too much computation."[1] Dirac's statement is arguably as accurate now as it was then, although increases in central processing unit (CPU) speed and the recent introduction of graphical processing units (GPUs) has led to continual revision of what constitutes "too much computation". While CPUs excel

Special Issue: Beyond QM/MM: Fragment Quantum Mechanical Methods

Received: July 11, 2014
Published: September 4, 2014

 © 2014 American Chemical Society dx.doi.org/10.1021/ar500229p | *Acc. Chem. Res.* 2014, 47, 2857−2866

at sequential execution, coarse-grained parallelism, and control-flow dominated code, GPUs are extremely efficient for data-parallel floating point operations.[2,3] GPUs emphasize massive parallelism, and their effective use requires integrating software design with hardware architecture. The potential of GPU-based algorithms has been recognized in chemistry, most notably in molecular dynamics (MD)[4−6] and quantum chemistry.[7−17]

The quest for efficient and accurate electronic structure theory methods is driven by the desire to gain microscopic insight into increasingly large nanoscale molecular systems including biomolecular systems, condensed phases, and nanomaterials.[17−19] Empirical force field based MD simulations are applicable to large systems for long times, and new computing platforms, including specialized hardware[20,21] and cloud-based architectures,[22] have enabled millisecond simulations of protein dynamics. Unfortunately, these are limited by the difficulty in describing polarization, charge transfer, covalent bond formation/rupture, and electronic excitation with empirical force fields.

Quantum mechanics/molecular mechanics (QM/MM) approaches[23,24] combine the benefits of electronic structure and molecular mechanics by splitting the molecular system into a QM part where electrons are modeled explicitly and an MM part that is described with an empirical force field.[17,23−25] Difficulties in practical implementations arise because of ambiguities at the boundary between the QM and MM regions. Results can depend both on the placement of the boundary and also on the specific scheme used to calculate the interactions between QM and MM regions. Assessing how these uncertainties affect phenomenological observables is difficult *a priori* and requires analyzing whether an observable is converged with respect to the size of the QM region. Such convergence tests are rarely undertaken because of their cost; however, in cases where they have been undertaken (for solvation energies,[26] hydrogen bonding,[27] NMR shifts,[28,29] absorption spectra,[17] and enzymatic reaction barriers[30]), rather large QM regions with hundreds of atoms have been found necessary.

The need to treat large QM regions motivates the development of hardware-adapted software frameworks for electronic structure theory. Efficiency is particularly important for complex systems with thousands of atoms, where the potential energy landscape includes a host of thermally accessible local minima. In these cases, small-molecule concepts like that of a stationary point are not very useful; one must dynamically sample a statistically meaningful portion of the relevant phase space. In the first part of this Account, we outline recent work aimed at building efficient GPU-based algorithms for molecular electronic structure theory. These acceleration approaches are broadly applicable, but we focus on their application in both density functional theory (DFT) and time dependent density functional theory (TDDFT), enabling efficient characterization of both ground and excited states of systems with up to 1000 atoms.

An open question is whether these GPU-accelerated methods can be extended to even larger systems. We are especially interested in electronic excited states and thus focus here on multichromophoric assemblies. As an example, consider the LH2 photosynthetic complex found in purple bacteria (shown in Figure 1), which is composed of 27 bacteriochlorophyll-a (Bchla) chromophores (137 atoms each). A single LH2 complex contains 3699 atoms, in addition to the protein scaffold surrounding it. Excited state molecular dynamics simulations on systems of this size surpass the capabilities of the most efficient TDDFT implementations and require additional levels of parallelism adapted to exploit modern architectures.

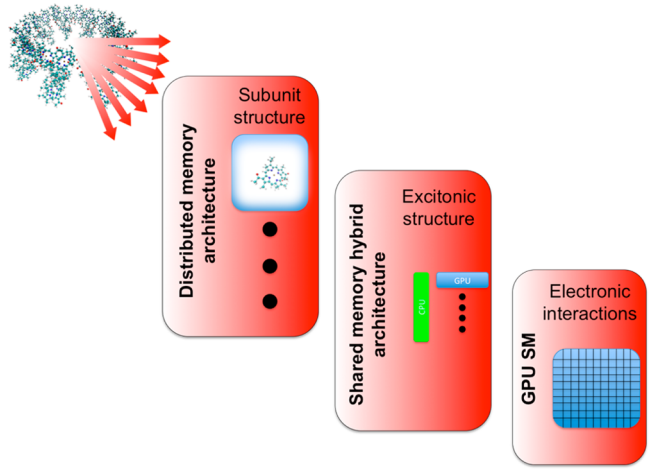

Figure 1. Parallel implementation on hybrid distributed/shared memory architectures. The structure of the LH2 chromophores is shown in the upper left hand corner.

In the latter part of this Account, we outline recent work developing parallelizable distributed-memory exciton-type approaches to calculate ground and excited state energies and gradients for large multichromophoric systems like LH2. Figure 1 schematically identifies the three different tiers of parallelism that may be exploited in carrying out excited-state calculations on a large supramolecular complex such as LH2 following its partitioning into n smaller subunits (e.g., individual Bchla's). The top tier of parallelism corresponds to each of the n subunit calculations split across nodes linked through a distributed memory parallel framework. At the middle tier of parallelism, multiple CPUs manage calculations across several GPUs on a particular node. The bottom tier of parallelism is the many-core shared-memory architecture of the individual GPUs on a particular node, used to calculate electronic interactions.

This multitiered parallel framework is well-adapted to large-scale nonadiabatic dynamics simulations with atomic-level resolution. A critical aspect of any approach that splits a larger calculation into smaller constituent parts concerns the treatment of coupling between the n different subunits. A many-body expansion is the most accurate way of calculating coupling between subunits;[31] however it can also lead to considerable expense. In the latter half of this Account, we introduce some approximations for these couplings that effectively allow us access to a computational scaling on the order of n·QM/MM. Ultimately, our aim is to build computational frameworks that extend quantum chemistry from the domain of teraFLOPs to the domain of petaFLOPs, allowing us to model nanometer-sized supramolecular structures in atomistic detail.

2. TERAFLOP MOLECULAR QUANTUM CHEMISTRY ON GPUS

2.1. GPU Architectures and Quantum Chemistry

GPU architectures allow users to (1) exploit multiple levels of parallelism and (2) map algorithmic kernels directly onto the computing architecture using flexible programming models. The many-core architecture of GPUs allows multithreaded, fine-grained, data-parallel computing. The parallel capabilities

of the GPU are exposed to the user through a parallel programming language such as Open Compute Language (OpenCL) or the Compute Unified Device Architecture (CUDA). The calculations described herein (implemented within the TeraChem program[32,33]) utilize CUDA, in which GPU programmatic control is invoked through an explicitly defined "kernel", or C-type function. Kernel execution runs simultaneously on thousands of GPU threads. Although logically autonomous, each GPU thread is associated at varying levels with all others to form a parallel hierarchy. The simplest association between threads is the grouping of 32 threads to form a warp, which executes identical instruction streams using the single-instruction−multiple-thread (SIMT) scheme. Warps are encapsulated in collective groups of up to 1024 threads, known as thread blocks. Arrays of thread blocks, in up to three dimensions, form the fundamental interface between the host CPU and the GPU. When a CUDA kernel is invoked by the host, the thread blocks are distributed onto any available streaming multiprocessors (SM), where each thread of a particular thread block executes simultaneously. The physical GPU is itself composed of several multithreaded SM units, providing a direct mapping of the thread block grid onto the SM array. Each individual SM consists of a set of streaming processor cores, on-chip shared memory, and a multithreaded instruction unit, which allows each SM to manage multiple thread blocks and concurrently execute hundreds of threads. The GPU architecture includes a memory hierarchy to support the structure of the SM units and the interface with the host CPU.

Execution of a GPU kernel by the host initiates a memory transaction between the host memory and the global GPU memory via the PCIe interface. Global memory allows simultaneous access by each thread in a warp to contiguous memory locations, but incurs high latencies (on the order of hundreds of GPU clock cycles). This feature leads to a paradigm shift in GPU programming philosophy vs CPU programming philosophy: whereas CPU-optimized code strives to minimize clock cycles on the premise that memory access is cheap, GPU-optimized code tends to minimize global memory access on the premise that clock cycles are cheap. Accordingly, whenever possible, GPU-accelerated programs restrict memory access to shared on-board GPU-memory, whose low latency and high bandwidth make rapid intrablock communication possible.

GPU threads are extremely lightweight: each thread is allocated a small memory space to store intermediate results before returning data to shared global arrays. Efficient scheduling prevents latency in the concurrent execution of thousands of SM threads. For the applications outlined in this Account, the number of blocks and threads in the CUDA grid generally surpasses the number of physical execution units on the GPU by many orders of magnitude. This is an advantage because saturation of the GPU with many more threads than cores hides issues related to GPU DRAM access latency and allows cores to remain continuously active throughout kernel execution. An additional efficiency consideration relates to data parallelism of the algorithm itself. Synchronization between threads and blocks reduces the effective concurrency. Thus, it is preferable to recompute intermediate quantities when needed, rather than rely on shared memory, interthread communication, and synchronization.

2.2. Ground and Excited State Quantum Chemistry

The fine-grained parallelism of GPUs is well-suited to quantum chemistry.[9,13,32,34−36] The most commonly used methods for describing the single-reference, electronic ground state of an arbitrary chemical system are self-consistent field (SCF) methods, for example, Hartree−Fock (HF) and DFT, effective one-electron theories describing each electron in a mean field of all other electrons. Time-dependent Hartree−Fock (TDHF) theory[37] and time-dependent density functional theory (TDDFT)[38] are extensions of these methods that describe the response of the ground state electron density to time-dependent external radiation fields. SCF methods such as HF and DFT rely on solving

$$\mathbf{FC} = \mathbf{ESC} \tag{1}$$

where **F** is the Fock operator, **C** contains the molecular orbital (MO) coefficients, and **S** is the atomic orbital (AO) overlap matrix. For molecules with less than 1000 atoms, construction of the Fock matrix elements is usually the rate-limiting step, because it requires calculating electron repulsion integrals (ERIs). The two-electron part of the Fock matrix consists of Coulomb (**J**) and exchange (**K**) components:

$$J_{\mu\nu} = \sum_{\lambda\sigma} (\mu\nu|\lambda\sigma) P_{\lambda\sigma} \tag{2}$$

$$K_{\mu\nu} = \sum_{\lambda\sigma} (\mu\lambda|\sigma\nu) P_{\lambda\sigma} \tag{3}$$

where (...|...) indicates an ERI of the form:

$$(\mu\nu|\lambda\sigma) = \iint \frac{\chi_\mu(\vec{r}_1)\chi_\nu(\vec{r}_1)\chi_\lambda(\vec{r}_2)\chi_\sigma(\vec{r}_2)}{|\vec{r}_1 - \vec{r}_2|}\, \mathrm{d}\vec{r}_1\, \mathrm{d}\vec{r}_2 \tag{4}$$

where $\chi(\vec{r})$ are atom-centered Gaussian basis functions and the density matrix elements are defined in terms of the MO coefficients (shown for the simplified case of a closed shell determinant with N_{el} electrons) as

$$P_{\mu\nu} = 2 \sum_{i}^{N_{el}/2} C_{\mu i} C_{\nu i}^{*} \tag{5}$$

Construction of the **J** and **K** matrices in both HF and DFT theory represents a significant computational bottleneck: it formally requires evaluating N^4 integrals, and ERI parallelization can lead to significant efficiency gains. The general algorithmic strategy to achieve high performance in SCF calculations involves reformulation of the integral evaluation structure to exploit the GPU architecture and take advantage of sparsity as much as possible. In a number of cases, we also avoid storing intermediate quantities and instead duplicate computations. This can be more efficient on the GPU because storage is at a premium as discussed above. For example, the integrals $(\mu\nu|\lambda\sigma)$ and $(\lambda\sigma|\mu\nu)$ are computed independently during the formation of **J**, even though they are identical (see eq 4).

In large chemical systems, the **J** and **K** matrices are sparse because many of the ERIs are numerically insignificant.[39−45] High performance gains come from efficient recognition of insignificant ERI matrix elements that do not need to be considered. For ERIs of the form $(\mu\nu|\lambda\sigma)$, the Cauchy−Schwarz inequality provides a bound on the magnitude of the ERI:

$$|(\mu\nu|\lambda\sigma)| \leq (\mu\nu|\mu\nu)^{1/2} (\lambda\sigma|\lambda\sigma)^{1/2} \tag{6}$$

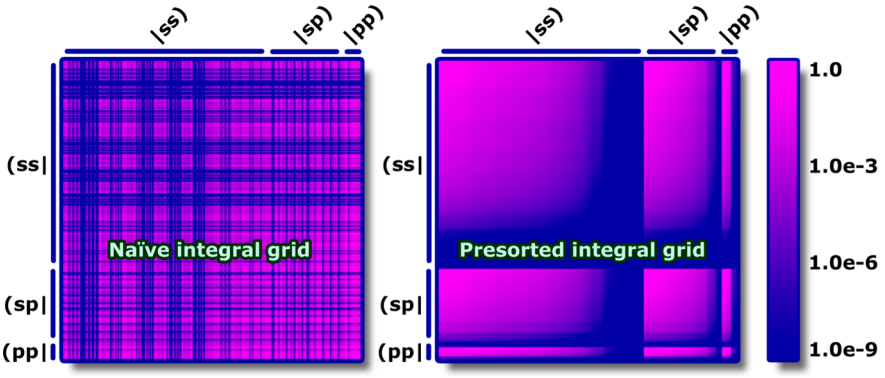

Figure 2. Organization of ERIs for Coulomb formation. Rows and columns correspond to primitive bra and ket pairs, respectively. Each ERI is colored according to the magnitude of its Schwarz bound. Pink/red indicates a large value of the ERI integral, and blue a small value.

The most efficient schemes for exploiting these bounds involve robust ordering; otherwise insignificant ERIs are interspersed among the important ones (Figure 2), and several cores are occupied computing non-negligible matrix elements while others remain idle. To resolve this, the primitive bra and ket arrays are sorted by decreasing Cauchy–Schwarz bounds prior to computing **J**, as shown in Figure 2. A similar sorting is carried out before computing **K**. These sorting procedures concentrate all the significant integrals in the upper left corner of a set of grids. In this way, groups of threads can "walk" through the integral matrix and the entire group will be occupied with significant integrals. This approach avoids both warp divergence and logic associated with testing the Schwarz bound. Because the significance of an integral is guaranteed to decrease as a thread moves across rows or down columns, all threads can stop processing as soon as the cutoff threshold has been reached (within a grid).

The algorithms outlined above, which include s, p, and d functions,[37] can be extended to excited state methods such as single-excitation configuration interaction[46] (CIS) and Tamm–Dancoff TDDFT,[47] where one of the key steps is again the construction of a Fock matrix (based on transition density matrices in place of the density matrix shown in eq 5). The same principles apply, but the gains are much larger since the transition density matrices are usually much sparser than density matrices. Essentially, this simplification arises because a transition density matrix describes the change due to a single electron while the density matrix describes the distribution of all electrons.

As an example of the kinds of calculations that are possible with GPU-accelerated DFT, we show explicit timings for a series of water clusters (with up to 1000 water molecules) in Figure 3. Scaling of the overall SCF calculation is quadratic with molecular size; the cubic scaling of the diagonalization algorithm only becomes problematic for systems with more than 2000 atoms (7000 basis functions).

3. EXCITONIC MOLECULAR QUANTUM CHEMISTRY ON GPUS

Although GPU-based quantum chemistry represents a significant advance over previous CPU-based algorithms, molecular size remains limited by the cubic scaling of matrix diagonalization and physical constraints on the number of GPUs that can coexist in a single node. Approaches for further tiers of parallelization are needed. In the case of multichromophoric assemblies, an excitonic framework is well suited

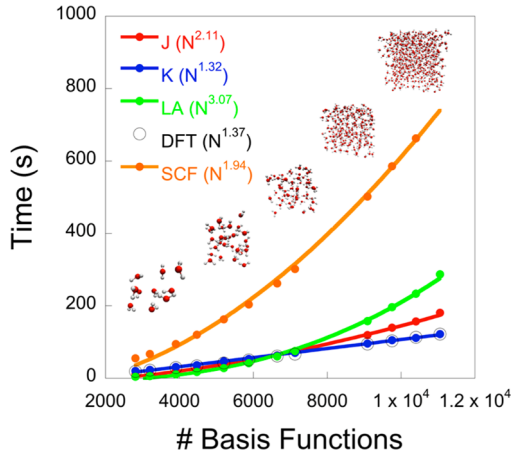

Figure 3. First SCF iteration timings in seconds for cubic water clusters. The total SCF times are broken down into **J**, **K**, linear algebra (LA), and DFT exchange-correlation contributions. All calculations were performed using a single Tesla M2090 GPU and the 6-31g basis set. Power fits show scaling with increasing system size, and the exponent for each fit is provided in the legend.

to this purpose. An exciton is an excited electronic state delocalized over several spatially separated molecular units, described as a superposition of coupled excitations on the individual molecular units. We illustrate the following discussion with a dimer and hexamer of Bchla chromophores, shown in Figures 4 and 5.

In the weak coupling limit, excitonic states are largely localized on each individual chromophore. In contrast, strong coupling leads to delocalized excitonic states. The exciton Hamiltonian is typically written as a matrix **H** in a basis of single excitations over N chromophores as[48]

$$\mathbf{H} = \sum_{i=1}^{N} E_i |i\rangle\langle i| + \sum_{j\neq i} V_{ij} |i\rangle\langle j| \tag{7}$$

where chromophore i is excited in $|i\rangle$. The diagonal site energies and the off-diagonal couplings between site-specific excitations are given as E_i and V_{ij}, respectively. When the interactions are not too strong, the coupling can be approximated using a Forster-type expression,[49,50] which accounts for Coulombic interaction between the different exciton basis states. In cases where the coupling is strong, generally because the chromophores are closely spaced, Dexter-type integrals (which describe through-bond electronic

dx.doi.org/10.1021/ar500229p | Acc. Chem. Res. 2014, 47, 2857–2866

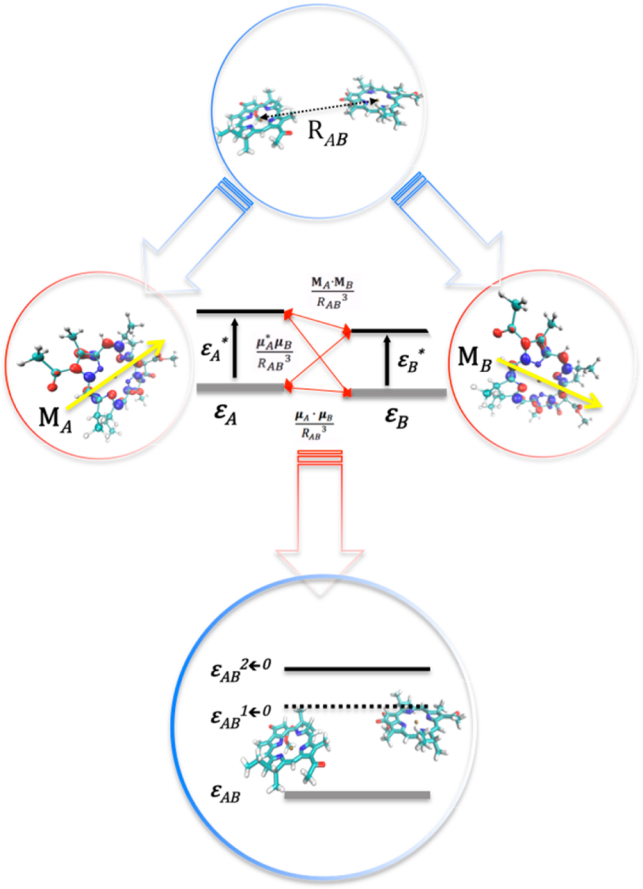

Figure 4. Schematic of the exciton model construction. Initially, the dimer is decomposed into exciton bases, corresponding to two Bchla molecules. Next, the ground and excited states of each chromophore are characterized using an excited state electronic structure method such as TDDFT. Ground and excited state electronic properties such as dipole and transition dipole moments are computed for each chromophore separately. Finally, the exciton Hamiltonian is constructed and diagonalized to obtain the ground and excited states of the full dimer complex.

exchange)[51] and so-called "penetration interactions" (which describe the interactions between electrons and screened nuclei)[52] are important,[53,54] although these generally fall off exponentially with interchromophore distance.[52]

In the exciton procedure outlined below, TDDFT provides the ground and electronic excited state energies of each chromophore, as well as interchromophoric couplings. Despite its shortcomings, TDDFT offers an efficient and practical way forward in calculating ground and excited state properties in many-electron systems.[38]

Observable quantities such as the oscillator strength are influenced by the magnitude of the coupling between ground and excited states. In the exciton model, we adopt the following notation for the excitonic basis functions: $|\varphi_0\rangle$ denotes the ground-state wavefunction of the full system and $|\varphi_{i,k}\rangle$ denotes the wavefunction of the k^{th} excitation of chromophore i, with all other chromophores in their respective ground states. The matrix element corresponding to the ground state energy of the coupled supramolecular system is given by:

$$E_0 = \langle\varphi_0|\hat{H}|\varphi_0\rangle = \sum_i \varepsilon_{(i,0)} + \sum_{j>i} \frac{\vec{\mu}_{(i,0)}\cdot\vec{\mu}_{(j,0)} - 3(\vec{n}_{ij}\cdot\vec{\mu}_{i,0})(\vec{n}_{ij}\cdot\vec{\mu}_{(j,0)})}{R_{ij}^3} \tag{8}$$

where $\varepsilon_{i,0}$ is the ground state (i.e., state 0) energy of chromophore i, and $\vec{\mu}_{(i,0)}$ is the ground state dipole moment of the ith chromophore. The individual chromophores are coupled through dipole−dipole interactions between the ground state wave functions that fall off as the cube of the interchromophore distance, R_{ij}, which we take as the center-of-mass (COM) separation. The neglect of overlap in eq 8 means that it is most accurate when the chromophores are not too strongly interacting. Equation 8 could be made more accurate using a many-body expansion as typical in fragment-based molecular calculations.[55−58] For the cases explored below, the separation between the individual chromophores is large enough that the dipolar Coulombic ground state coupling is sufficient.

Diagonal matrix elements corresponding to excited states are calculated as

$$E_{(i,k)} = \sum_{j\neq i} \varepsilon_{(j,0)} + \varepsilon_{(i,k)} + \delta \tag{9}$$

where $\varepsilon_{(i,k)}$ is the energy of the k^{th} electronic state on the i^{th} chromophore. The optional δ (which we do not use here) can correct for discrepancies between TDDFT and experimental excitation energies.

Considering N_x excited states for each of N chromophores, there are $N_x(N+1)$ basis states. Off-diagonal Hamiltonian elements are evaluated with the dipolar Forster expression:

$$V_{(i,k)(j,l)} = \left\langle \varphi_{(i,k)}\middle|\hat{H}\middle|\varphi_{(j,l)}\right\rangle = \frac{1}{\varepsilon_r}\frac{\vec{M}_{(i,k)\leftarrow(i,0)}\cdot\vec{M}_{(j,l)\leftarrow(j,0)} - 3(\vec{n}_{ij}\cdot\vec{M}_{(i,k)\leftarrow(i,0)})(\vec{n}_{ij}\cdot\vec{M}_{(j,l)\leftarrow(j,0)})}{R_{ij}^3} \tag{10}$$

where $\vec{M}_{(i,k)\leftarrow(i,0)}$ is the transition dipole for the $k \leftarrow 0$ transition on chromophore i, $\vec{n}_{ij}$ is a unit vector along the COM separation $\vec{R}_{ij}$, and ε_r is an effective dielectric constant. Equation 10 is very accurate in the limit of large R. Similarly, the off-diagonal matrix elements, which couple the ground state to an electronic state with an excited chromophore, are

$$V_{(0)(j,l)} = \left\langle \varphi_0\middle|\hat{H}\middle|\varphi_{(j,l)}\right\rangle = \frac{1}{\varepsilon_r}\sum_{i\neq j}\frac{\vec{\mu}_{(i,0)}\cdot\vec{\mu}_{(j,l)} - 3(\vec{n}_{ij}\cdot\vec{\mu}_{(i,0)})(\vec{n}_{ij}\cdot\vec{\mu}_{(j,l)})}{R_{ij}^3} \tag{11}$$

where $\vec{\mu}_{(i,k)}$ is the dipole moment of the kth excited state of chromophore i.

The general framework of the exciton model is illustrated schematically in Figure 4 for a Bchla dimer extracted from the LH2 complex, with COM separation R_{AB}: the ground states of each chromophore couple through dipole−dipole interactions as $(\vec{\mu}_A\cdot\vec{\mu}_B)/\mathbf{R}_{AB}^3$, coupling between ground and excited states occurs through $(\vec{\mu}_{A^*}\cdot\vec{\mu}_B)/\mathbf{R}_{AB}^3$, and $(\vec{\mu}_A\cdot\vec{\mu}_{B^*})/\mathbf{R}_{AB}^3$ terms, and coupling between excited states occurs through the $\vec{M}_A\cdot\vec{M}_B/\mathbf{R}_{AB}^3$ term (where * denotes an excited state). The resulting eigenstates describe the full N-chromophore complex. For example, treating a six-chromophore system with a single

dx.doi.org/10.1021/ar500229p | Acc. Chem. Res. 2014, 47, 2857−2866

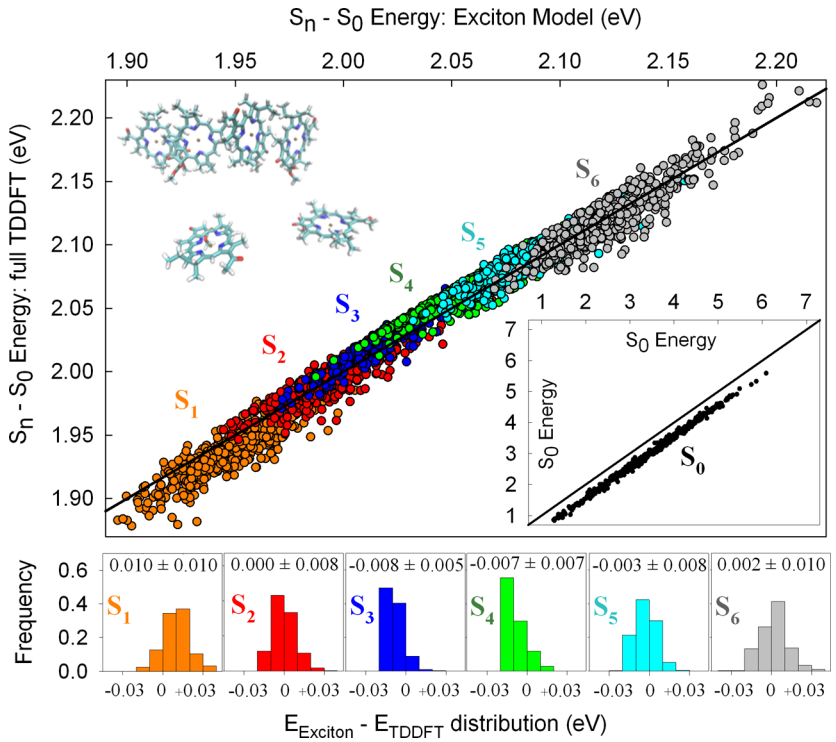

Figure 5. Top panel compares energies for the first six excited singlet states obtained using (1) the excitonic framework and (2) full TDDFT calculations on the six-chromophore system. The diagonal line represents perfect agreement between TDDFT and the exciton model. The bottom panel shows histograms of the error in the exciton model (compared with full TDDFT), along with the average error ± the standard deviation for each state.

excitation on each chromophore yields a symmetric excitonic Hamiltonian of the form:

$$\mathbf{H} = \begin{bmatrix} E_0 & V_{(0)(1,1)} & V_{(0)(2,1)} & \cdots & V_{(0)(6,1)} \\ V_{(0)(1,1)} & E_{(1,1)} & V_{(1,1)(2,1)} & \cdots & V_{(1,1)(6,1)} \\ V_{(0)(2,1)} & V_{(1,1)(2,1)} & E_{(2,1)} & \cdots & V_{(2,1)(6,1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ V_{(0)(6,1)} & V_{(1,1)(6,1)} & V_{(2,1)(6,1)} & \cdots & E_{(6,1)} \end{bmatrix} \quad (12)$$

where E_0 is given by eq 8, the remaining diagonal matrix elements are given by eq 9, first row and column off-diagonal elements are defined in eq 11, and all other off-diagonal matrix elements are defined in eq 10.

The inset of Figure 5 shows a set of Bchla's forming a subset of LH2. Within this subset, two of the Bchlas participate in LH2's B800 ring, and four in the B850 ring. This model system (450 atoms, 2700 basis functions) is a challenging target for our exciton method. We test the accuracy of the exciton framework by direct comparison to TDDFT calculations of the entire system, using a range-corrected exchange-correlation functional, ωPBEh (ω = 0.2), in the 6-31g basis set. Dynamical sampling effects are included through 1000 structures obtained from ground-state empirical force field MD simulations of the full LH2 complex.

For a set of 1000 MD snapshots, Figure 5 shows correlation between the energies of the S_0, S_1, S_2, S_3, S_4, S_5, and S_6 states obtained from (1) diagonalization of the exciton Hamiltonian in eq 12 and (2) full TDDFT. Perfect correlation between the TDDFT and excitonic energies is indicated by a diagonal line. As shown in the bottom panel of Fig 5, the agreement between the excitonic approach and the TDDFT results is good, with an average error over all six excited states of 0.008 eV.

Using the eigenvectors obtained from diagonalization of the exciton Hamiltonian, one can calculate transition dipoles, oscillator strengths, nonadiabatic coupling vectors, and absorption spectra. In terms of the excitonic basis functions, the I^{th} eigenstate of the excitonic system is

$$|\psi_I\rangle = c_0^I|\varphi_0\rangle + \sum_{i,k} c_{(i,k)}^I|\varphi_{(i,k)}\rangle \quad (13)$$

where $c_{(i,k)}^I$ are the coefficients specifying the contribution of the $|\varphi_{(i,k)}\rangle$ basis function to the I^{th} excitonic eigenstate, and c_0^I are the coefficients specifying the contribution of the ground state basis function, $|\varphi_0\rangle$ to the I^{th} excitonic eigenstate. The transition dipole moment between the ground ψ_0 and excited eigenstate ψ_J is

$$\begin{aligned} \vec{T}_J &= \langle\psi_0|\hat{\mu}|\psi_J\rangle \\ &= c_0^0 c_0^J\langle\varphi_0|\hat{\mu}|\varphi_0\rangle + c_0^0\sum_{j,l} c_{(j,l)}^J\langle\varphi_0|\hat{\mu}|\varphi_{(j,l)}\rangle \\ &\quad + c_0^J\sum_{i,k} c_{(i,k)}^0\langle\varphi_{(i,k)}|\hat{\mu}|\varphi_0\rangle + \sum_{i,j,k,l} c_{(i,k)}^0 c_{(j,l)}^J\langle\varphi_{(i,k)}|\hat{\mu}|\varphi_{(j,l)}\rangle \end{aligned} \quad (14)$$

where $\hat{\mu}$ is the dipole moment operator. Considering only transitions between electronic states on a single chromophore i allows us to simplify the final term in eq 14, giving

dx.doi.org/10.1021/ar500229p | Acc. Chem. Res. 2014, 47, 2857−2866

$$\vec{T}_J = \langle\psi_0|\hat{\mu}|\psi_J\rangle \approx c_0^0 c_0^J \langle\varphi_0|\hat{\mu}|\varphi_0\rangle + c_0^0 \sum_{j,l} c_{(j,l)}^J \vec{M}_{(j,l)\leftarrow(j,0)} - c_0^J \sum_{i,k} c_{(i,k)}^0 \vec{M}_{(i,k)\leftarrow(i,0)} + \sum_{i,k\neq l} c_{(i,k)}^0 c_{(i,l)}^J \vec{M}_{(i,l)\leftarrow(i,k)} \quad (15)$$

The first term describes the contribution of the ground state dipole moments of each chromophore. The second term accounts for the transition dipole moments between the ground and l^{th} excited state on chromophore j. The third term treats interactions between the ground and kth excited state on chromophore i. The final term describes transitions between excited states k and l on chromophore i, where $\vec{M}_{(i,l)\leftarrow(i,k)}$ is the corresponding transition dipole moment vector. From Fermi's Golden Rule, the oscillator strength is

$$f_I = \frac{2}{3}\omega_I |\vec{T}_I|^2 \quad (16)$$

where $\vec{T}_I$ is the transition dipole vector of the Ith excited state of the full system and ω_I is the corresponding excitation energy. Within the Franck−Condon approximation, the absorption spectrum can be calculated from single-point excitation energies and oscillator strengths sampled from MD snapshots:

$$I(\omega) = \sum_{I,\alpha} f_{I,\alpha} \frac{1}{\Gamma\sqrt{2\pi}} \exp\left(\frac{-(\Omega_{I,\alpha} - \omega)^2}{\Gamma^2}\right) \quad (17)$$

where $f_{I,\alpha}$ is the oscillator strength of the Ith eigenstate for the αth dynamics snapshot, $\Omega_{I,\alpha}$ is the corresponding excitation energy, and Γ is the homogeneous broadening width.

Figure 6 compares the absorption spectrum obtained using full TDDFT with that obtained using the excitonic framework

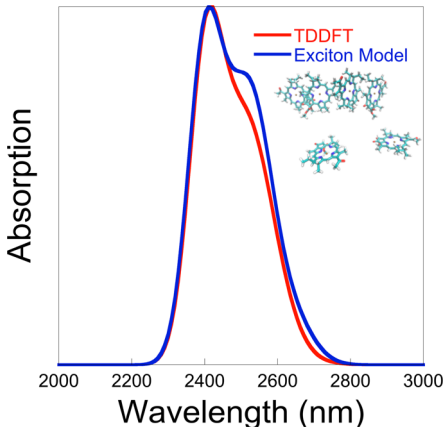

Figure 6. Absorption spectrum of a six-chromophore system computed with the exciton model and TDDFT. The inset shows the six-chromophore arrangement of Bchla molecules, with orientations extracted from MD simulations of LH2. Both spectra have been scaled to a maximum amplitude of one for comparison.

(eqs 16 and 17) with Γ = 30 meV sampled over 1000 MD geometries obtained from constant temperature classical molecular dynamics simulations of the full LH2 complex (chromophores plus environment) using a GPU-accelerated version of the AMBER package. Overlap between the TDDFT spectra and the excitonic spectra is maximum with ε_r in eqs 10 and 11 set to a value 1.05, suggesting that dielectric screening effects are rather weak. The only noticeable difference between the excitonic and TDDFT spectra concerns the shoulder on the blue side of the band, linked to the two B800 chromophores.

The encouraging results of Figures 5 and 6 led us to investigate dynamics with the exciton Hamiltonian. This requires ground- and excited-state gradients and differentiation of the excitonic Hamiltonian matrix elements in eqs 8−11. The derivative of the ground state energy, E_0, may be written as

$$\frac{\partial E_0}{\partial R} = \sum_i \frac{\partial \varepsilon_{(i,0)}}{\partial R} + \sum_{j>i} \frac{\partial}{\partial R}\left[\frac{\vec{\mu}_{(i,0)}\cdot\vec{\mu}_{(j,0)} - 3\left(\vec{n}_{ij}\cdot\vec{\mu}_{(i,0)}\right)\left(\vec{n}_{ij}\cdot\vec{\mu}_{(j,0)}\right)}{R_{ij}^3}\right] \quad (18)$$

The derivatives of the diagonal matrix elements corresponding to excited chromophores are given as

$$\frac{\partial E_{(i,k)}}{\partial R} = \frac{\partial \varepsilon_{(i,k)}}{\partial R} + \sum_{j\neq i} \frac{\partial \varepsilon_{(j,0)}}{\partial R} \quad (19)$$

Ground and excited state gradients on individual chromophores are readily computed from TDDFT. Obtaining derivatives for the off-diagonal matrix elements is more complicated. Differentiation of eq 10 in particular requires gradients of the transition dipole moment vectors, $\partial\mathbf{M}/\partial\mathbf{R}$. Here, we avoid this by using a simple approximation to the transition dipole derivatives: $\partial\mathbf{M}/\partial\mathbf{R}$ between a particular set of states for a given atom in a particular Cartesian direction is set equal to its average value sampled over a series of MD snapshots from dynamically significant regions of the configuration space. This effectively amounts to assuming a fixed atomic point charge model in both the ground state and the excited states over the dynamically relevant portions of the supramolecular configuration space. Relating the transition dipole moments to the molecular geometry in this fashion allows efficient differentiation of the excitonic state energies by application of the Hellman−Feynman theorem:

$$\vec{F}_I = -\langle\psi_I|\frac{\partial H}{\partial R}|\psi_I\rangle \quad (20)$$

where $|\psi_I\rangle$ is the wave function of the Ith eigenstate of the excitonic Hamiltonian. Similarly efficient expressions may be used to obtain nonadiabatic coupling vectors between excitonic states required in nonadiabatic MD simulations, that is,

$$\vec{d}_{IJ} = \frac{\langle\psi_I|\frac{\partial H}{\partial R}|\psi_J\rangle}{\varepsilon_J - \varepsilon_I} \quad (21)$$

where $\vec{d}_{IJ}$ is the nonadiabatic coupling vector between electronic states I and J, and ε_I and ε_J are the corresponding energies.

Figure 7 shows significant performance gains using this exciton model compared with full TDDFT. As a further test of the accuracy of the results in Figure 7, we selected two adjacent Bchla molecules from the LH2 subset in Figure 6 and compared excited-state dynamics run using the exciton framework with those obtained using full TDDFT. Simulations were initialized from a random conformation from the aforementioned MD snapshots, with velocities initialized using a 300 K Boltzmann distribution, and assuming a vertical

dx.doi.org/10.1021/ar500229p | Acc. Chem. Res. 2014, 47, 2857−2866

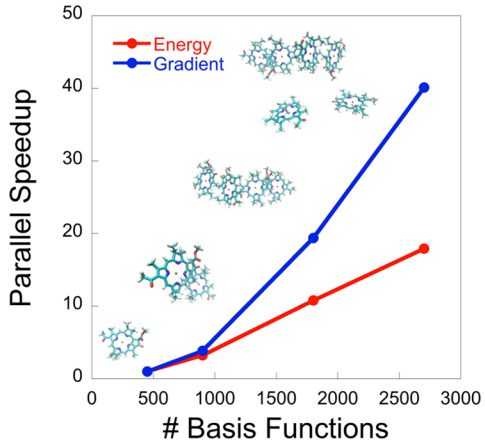

Figure 7. Speedup of the exciton model relative to full TDDFT. All calculations were performed using GTX Titan hardware with 12 CPU cores and 4 GPUs for the TDDFT calculation and the exciton model parallelized over 12 CPU cores and 4 GPUs per exciton basis function.

Franck−Condon excitation from S_0 to S_1. Propagation on S_1 used a velocity Verlet algorithm with a 0.2 fs time step. Figure 8

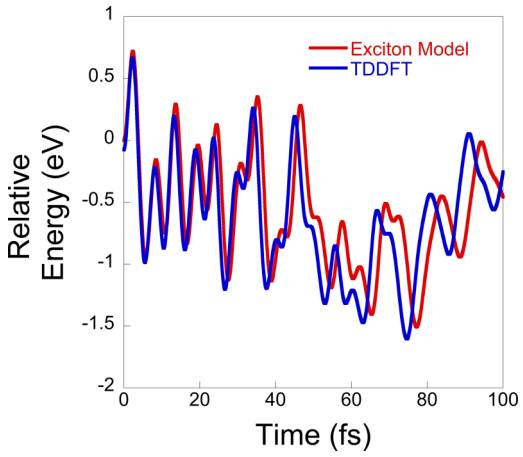

Figure 8. Time evolution of the energy of the first singlet state, S_1, of the Bchla dimer, computed independently with the exciton model and full TDDFT. Each trajectory was initialized with the same set of coordinates and velocities. Excitation energies are shown relative to the absolute energy of S_1 from the exciton model at time zero.

compares the time-dependent S_1 energy from the exciton and TDDFT calculations. Good agreement is observed for the first 50 fs, and some drift is evident thereafter due to accumulation of error in the respective gradients. After 100 fs, the deviation in the S_1 energy is small, on the order of ∼0.05 eV. Over the subpicosecond time scales relevant to electronic relaxation and energy transfer, Figure 8 suggests that the exciton model is quite accurate.

4. CONCLUSIONS

Using shared memory hybrid CPU/GPU architectures, we are able to carry out ground and excited state electronic structure and dynamics for large molecular systems. These acceleration approaches are applicable to a range of electronic structure theory methods, but we focused here on density functional theory and time dependent density functional theory (TDDFT), demonstrating that excited state *ab initio* molecular dynamics is now feasible for molecules with O(1000) atoms.

Beyond CPU/GPU shared memory parallelism, we outlined an additional "top-level" layer of parallelism well-suited to distributed memory frameworks, which we have demonstrated using an excitonic framework. This enables *ab initio* molecular dynamics calculations on large supramolecular systems. Comparison between excitonic and full TDDFT simulations is quite good. In fact, one significant advantage of efficient TDDFT algorithms is that they offer standards against which to test and compare various levels of approximation. A particularly attractive feature of the excitonic framework described herein is the fact that each of the building blocks from which it is composed can be systematically improved. Therefore, it can accommodate increasingly accurate methods for calculating both diagonal site energies and off-diagonal couplings.

Understanding electronic excitation transfer in condensed phase supramolecular systems remains a significant challenge. Future work will adapt the frameworks described herein to carrying out dynamics simulations on larger systems. This presents a number of challenges, in terms of both theoretical and algorithmic development. For example, to study the nonadiabatic dynamics of systems with very strong coupling, we will need to investigate efficient ways in which to go beyond Coulombic dipole couplings, ideally using functional forms that are able to efficiently treat both strong and weak coupling regimes. For condensed phase systems (for example, in liquids or biomolecular systems), we will need to implement efficient strategies that treat the dynamics of the environment and its interaction with the chromophores. Because the hybrid/multitiered parallel approach outlined herein scales extremely well over parallel distributed-memory architectures built from large arrays of GPU nodes, it should allow us to carry out fully atomistic nonadiabatic dynamics simulations of massive supramolecular complexes, taking full advantage of the combined shared memory/distributed memory architectures of the world's largest supercomputers. Exascale computing architectures are forecast to emerge in the next 5−10 years, and it is likely that these machines will rely on massive multitiered parallelism of the sort discussed herein, fusing low level shared memory parallelism with top level distributed memory parallelism. Theoretical frameworks like those outlined herein will allow full exploitation of these architectures, enabling us to gain microscopic insight into some of nature's most challenging molecular energy transfer systems.

■ AUTHOR INFORMATION

Corresponding Authors

*E-mail: drglowacki@gmail.com.
*E-mail: toddjmartinez@gmail.com.

Notes

The authors declare no competing financial interest.

Biographies

Aaron Sisto is a Ph.D. candidate at Stanford University working with Todd Martinez on GPU-accelerated quantum chemistry. Previously, he has worked at Lawrence Berkeley, Lawrence Livermore, and Sandia National Laboratories on topics ranging from surface catalysis to energy policy. A few of his current research interests include nonadiabatic quantum dynamics, high-performance computing, metamaterial design, and machine learning for scientific data analysis.

Dr. David Glowacki is a Royal Society Research fellow with broad interests across theoretical chemistry and computer science. One of his main research emphases is focused on the use of high-performance

computing to tackle problems in molecular dynamics and kinetics across both gas and condensed phases. He obtained his undergraduate degree at the University of Pennsylvania in 2003, his M.A at Manchester University in 2004, and his Ph.D. at the University of Leeds in 2008. He moved to Bristol in 2009 to carry out his postdoctoral research. At present, he is based across the University of Bristol and Stanford University. His work to understand energy transfer in chemical reaction dynamics was recently recognized by the Royal Society of Chemistry's 2014 Harrison-Meldola Memorial Prize.

Todd Martínez holds a B.A. in Chemistry from Calvin College (1989) and a Ph.D. in Chemistry from UCLA (1994). He was a Fulbright Junior Researcher and UC President's Postdoctoral Fellow at UCLA and Hebrew University before joining the faculty at the University of Illinois, Urbana–Champaign, in 1996. He has been the Ehrsam and Franklin Professor in the departments of Chemistry and Photon Sciences at Stanford University and SLAC National Accelerator Laboratory since 2009.

■ ACKNOWLEDGMENTS

This work was supported by the AMOS program within the Office of Basic Energy Sciences, U.S. Department of Energy, under Contract No. DE-AC02-76SF00515. D.R.G. is a Royal Society research fellow. A.S. is a DOE Computational Science Graduate Fellow and NSF Graduate Research Fellow. T.J.M. is a National Security Science and Engineering Fellow (Office of the Assistant Secretary of Defense for Research and Engineering).

■ REFERENCES

(1) Dirac, P. A. M. Quantum Mechanics of Many-Electron Systems. *Proc. R. Soc. London, Ser. A* **1929**, *123*, 714–733.
(2) Horn, D. R.; Houston, M.; Hanrahan, P. Clawhmmer: A Streaming hmmer-Search Implementation. In *Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*; IEEE Computer Society: Washington, D.C., 2005; p 11.
(3) McIntosh-Smith, S.; Gillan, C.; Sanna, N.; Scott, S.; Steinke, T. Special issue of the Journal of Parallel and Distributed Computing (JDPC) on novel architectures for high-performance computing. *J. Parallel Distrib. Comput.* **2013**, *73*, 1415–1416.
(4) Friedrichs, M. S.; Eastman, P.; Vaidyanathan, V.; Houston, M.; LeGrand, S.; Beberg, A. L.; Ensign, D. L.; Bruns, C. M.; Pande, V. S. Accelerating Molecular Dynamic Simulation on Graphical Processing Units. *J. Comput. Chem.* **2009**, *30*, 864–872.
(5) Levine, B. G.; LeBard, D. N.; DeVane, R.; Shinoda, W.; Kohlmeyer, A.; Klein, M. L. Micellization Studied by GPU-Accelerated Coarse-Grained Molecular Dynamics. *J. Chem. Theor. Comput.* **2011**, *7*, 4135–4145.
(6) Glowacki, D. R.; O'Connor, M.; Calabro, G.; Price, J.; Tew, P.; Mitchell, T.; Hyde, J.; Tew, D.; Coughtrie, D. J.; McIntosh-Smith, S. A GPU-accelerated immersive audiovisual framework for interaction with molecular dynamics using consumer depth sensors. *Faraday Disc.* **2014**, DOI: 10.1039/C4FD00008K.
(7) Wu, X.; Koslowski, A.; Thiel, W. Semiempirical Quantum Chemical Calculations Accelerated on a Hybrid Multicore CPU-GPU Computing Platform. *J. Chem. Theor. Comput.* **2012**, *8*, 2272–2281.
(8) Asadchev, A.; Gordon, M. S. New Multithreaded Hybrid CPU/GPU Approach to Hartree-Fock. *J. Chem. Theor. Comput.* **2012**, *8*, 4166–4176.
(9) Ufimtsev, I. S.; Martínez, T. J. Quantum Chemistry on Graphical Processing Units. 1. Strategies for Two-Electron Integral Evaluation. *J. Chem. Theor. Comput.* **2008**, *4*, 222–231.
(10) Yasuda, K. Accelerating Density Functional Calculations with Graphics Processing Unit. *J. Chem. Theor. Comput.* **2008**, *4*, 1230–1236.
(11) Genovese, L.; Ospici, M.; Deutsch, T.; Mehaut, J.-F.; Neelov, A.; Goedecker, S. Density Functional Theory Calculation on Many-Cores Hybrid Central Processing Unit-Graphic Processing Unit Architectures. *J. Chem. Phys.* **2009**, *131*, No. 034103.
(12) Andrade, X.; Aspuru-Guzik, A. Real-Space Density Functional Theory on Graphical Processing Units: Computational Approach and Comparison to Gaussian Basis Set Methods. *J. Chem. Theor. Comput.* **2013**, *9*, 4360–4373.
(13) Luehr, N.; Ufimtsev, I. S.; Martínez, T. J. Dynamic Precision for Electron Repulsion Integral Evaluation on Graphical Processing Units (GPUs). *J. Chem. Theor. Comput.* **2011**, *7*, 949–954.
(14) Vogt, L.; Olivares-Amaya, R.; Kermes, S.; Shao, Y.; Amador-Bedolla, C.; Aspuru-Guzik, A. Accelerating Resolution-of-the-Identity Second-Order Møller–Plesset Quantum Chemistry Calculations with Graphical Processing Units. *J. Phys. Chem. A* **2008**, *112*, 2049–2057.
(15) Ma, W.; Krishnamoorthy, S.; Villa, O.; Kowalski, K. GPU-Based Implementations of the Noniterative Regularized-CCSD(T) Corrections: Applications to Strongly Correlated Systems. *J. Chem. Theor. Comput.* **2011**, *7*, 1316–1327.
(16) DePrince, A. E., III; Hammond, J. R. Coupled Cluster Theory on Graphics Processing Units. I. The Coupled Cluster Doubles Method. *J. Chem. Theor. Comput.* **2011**, *7*, 1287–1295.
(17) Isborn, C. M.; Goetz, A. W.; Clark, M. A.; Walker, R. C.; Martinez, T. J. Electronic Absorption Spectra from MM and ab Initio QM/MM Molecular Dynamics: Environmental Effects on the Absorption Spectrum of Photoactive Yellow Protein. *J. Chem. Theor. Comput.* **2012**, *8*, 5092–5106.
(18) Ufimtsev, I. S.; Luehr, N.; Martinez, T. J. Charge Transfer and Polarization in Solvated Proteins from Ab Initio Molecular Dynamics. *J. Phys. Chem. Lett.* **2011**, *2*, 1789–1793.
(19) Kulik, H. J.; Luehr, N.; Ufimtsev, I. S.; Martinez, T. J. Ab Initio Quantum Chemistry for Protein Structures. *J. Phys. Chem. B* **2012**, *116*, 12501–12509.
(20) Shaw, D. E.; Chao, J. C.; Eastwood, M. P.; Gagliardo, J.; Grossman, J. P.; Ho, C. R.; Lerardi, D. J.; Kolossváry, I.; Klepeis, J. L.; Layman, T.; McLeavey, C.; Deneroff, M. M.; Moraes, M. A.; Mueller, R.; Priest, E. C.; Shan, Y.; Spengler, J.; Theobald, M.; Towles, B.; Wang, S. C.; Dror, R. O.; Kuskin, J. S.; Larson, R. H.; Salmon, J. K.; Young, C.; Batson, B.; Bowers, K. J. Anton, a Special-Purpose Machine for Molecular Dynamics Simulation. *Commun. ACM* **2008**, *51*, 91–95.
(21) Shaw, D. E.; Maragakis, P.; Lindorff-Larsen, K.; Piana, S.; Dror, R. O.; Eastwood, M. P.; Bank, J. A.; Jumper, J. M.; Salmon, J. K.; Shan, Y.; Wriggers, W. Atomic-Level Characterization of trhe Structural Dynamics of Proteins. *Science* **2010**, *330*, 341–346.
(22) Kohlhoff, K. J.; Shukla, D.; Lawrenz, M.; Bowman, G. R.; Konerding, D. E.; Belov, D.; Altman, R. B.; Pande, V. S. Cloud-Based Simulations on Google Exacycle Reveal Ligand Modulation of GPCR Activation Pathways. *Nat. Chem.* **2014**, *6*, 15–21.
(23) Warshel, A.; Levitt, M. Theoretical Studies of Enzymic Reactions: Dielectric, Electrostatic and Steric Stabilization of the Carbonium Ion in the Reaction of Lysozyme. *J. Mol. Biol.* **1976**, *103*, 227–249.
(24) Senn, H. M.; Thiel, W. QM/MM Studies of Enzymes. *Curr. Opin. Chem. Biol.* **2007**, *11*, 182–187.
(25) van der Kamp, M. W.; Mulholland, A. J. Combined Quantum Mechanics/Molecular Mechanics (QM/MM) Methods in Computational Enzymology. *Biochemistry* **2013**, *52*, 2708–2728.
(26) Fox, S. J.; Pittock, C.; Fox, T.; Tautermann, C. S.; Malcolm, N.; Skylaris, C.-K. Electrostatic Embedding in Large-Scale First Principles Quantum Mechanical Calculations on Biomolecules. *J. Chem. Phys.* **2011**, *135*, No. 224107.
(27) Meier, K.; Thiel, W.; van Gunsteren, W. F. On the Effect of a Variation of the Force Field, Spatial Boundary Condition and Size of the QM Region in QM/MM MD Simulations. *J. Comput. Chem.* **2011**, *33*, 363–378.
(28) Flaig, D.; Beer, M.; Ochsenfeld, C. Convergence of Electronic Structure with the Size of the QM Region: Example of QM/MM NMR Shieldings. *J. Chem. Theor. Comput.* **2012**, *8*, 2260–2271.

(29) Johnson, E. R.; DiLabio, G. A. Convergence of Calculated Nuclear Magnetic Resonance Chemical Shifts in a Protein with Respect to Quantum Mechanical Model Size. *J. Mol. Struct. THEOCHEM* **2009**, *898*, 56−61.
(30) Liao, R.-Z.; Thiel, W. Convergence in the QM-only and QM/MM Modeling of Enzymatic Reactions: A Case Study for Acetylene Hydratase. *J. Comput. Chem.* **2013**, *34*, 2389−2397.
(31) Murrell, J. N. The Many-Body Expansion of the Potential Energy Function for Elemental Clusters. *Int. J. Quantum Chem.* **1990**, *37*, 95−102.
(32) Ufimtsev, I. S.; Martínez, T. J. Quantum Chemistry on Graphical Processing Units. 3. Analytical Energy Gradients, Geometry Optimization, and First Principles Molecular Dynamics. *J. Chem. Theor. Comput.* **2009**, *5*, 2619−2628.
(33) http://www.petachem.com, accessed 7/1/14.
(34) Ufimtsev, I. S.; Martínez, T. J. Graphical Processing Units for Quantum Chemistry. *Comput. Sci. Eng.* **2008**, *10*, 26−34.
(35) Ufimtsev, I. S.; Martínez, T. J. Quantum Chemistry on Graphical Processing Units. 2. Direct Self-Consistent-Field Implementation. *J. Chem. Theor. Comput.* **2009**, *5*, 1004−1015.
(36) Titov, A. V.; Ufimtsev, I. S.; Luehr, N.; Martinez, T. J. Generating Efficient Quantum Chemistry Codes for Novel Architectures. *J. Chem. Theor. Comput.* **2013**, *9*, 213−221.
(37) McLachlan, A. D.; Ball, M. A. Time-Dependent Hartree−Fock Theory for Molecules. *Rev. Mod. Phys.* **1964**, *36*, 844−855.
(38) Adamo, C.; Jacquemin, D. The Calculations of Excited-State Properties with Time-Dependent Density Functional Theory. *Chem. Soc. Rev.* **2013**, *42*, 845−856.
(39) Whitten, J. L. Coulombic Potential Energy Integrals and Approximations. *J. Chem. Phys.* **1973**, *58*, 4496−4501.
(40) Challacombe, M.; Schwegler, E. Linear Scaling Computation of the Fock Matrix. *J. Chem. Phys.* **1997**, *106*, 5526−5536.
(41) Almlof, J.; Faegri, K.; Korsell, K. Principles for a Direct SCF Approach to LCAO-MO ab Initio Calculations. *J. Comput. Chem.* **1982**, *3*, 385−399.
(42) Kussman, J.; Beer, M.; Ochsenfeld, C. Linear-Scaling Self-Consistent Field Methods for Large Molecules. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.* **2013**, *3*, 614−636.
(43) Scuseria, G. E. Linear Sclaing Density Functional Calculations with Gaussian Orbitals. *J. Phys. Chem. A* **1999**, *103*, 4782−4790.
(44) Rudberg, E.; Rubensson, E. H.; Salek, P. Kohn-Sham Density Functional Theory Electronic Structure Calculations with Linearly Scaling Computational Time and Memory Usage. *J. Chem. Theor. Comput.* **2011**, *7*, 340−350.
(45) Ahlrichs, R.; Bär, M.; Häser, M.; Horn, H.; Kölmel, C. Electronic structure calculations on workstation computers: The program system turbomole. *Chem. Phys. Lett.* **1989**, *162*, 165−169.
(46) Foresman, J. B.; Head-Gordon, M.; Pople, J. A.; Frisch, M. J. Toward a Systematic Molecular Orbital Theory for Excited States. *J. Phys. Chem.* **1992**, *96*, 135−148.
(47) Hirata, S.; Head-Gordon, M. Time-Dependent Density Functional theory within the Tamm−Dancoff Approximation. *Chem. Phys. Lett.* **1999**, *314*, 291−299.
(48) May, V.; Kuhn, O. *Charge and Energy Transfer Dynamics in Molecular Systems*; Wiley-VCH: Weinheim, Germany, 2011.
(49) Förster, T. Delocalized Excitation and Excitation Transfer. In *Modern Quantum Chemistry*; Sinanoglu, O., Ed.; Academic Press: New York, 1965; pp 93−137.
(50) Kreuger, B. P.; Schioles, G. D.; Fleming, G. R. Calculation of Couplings and Energy Transfer Pathways between the Pigments of LH2 by the ab Initio Transition Density Cube Method. *J. Phys. Chem. B* **1998**, *102*, 5378−5386.
(51) Dexter, D. L. A Theory of Sensitized Luminescence in Solids. *J. Chem. Phys.* **1953**, *21*, 836−850.
(52) Harcourt, R. D.; Scholes, G. D.; Ghiggino, K. P. Rate Expressions for Excitation Transfer. II. Electronic Considerations of Direct and Through-Configuration Exciton Resonance Interactions. *J. Chem. Phys.* **1994**, *101*, 10521−10525.
(53) Thompson, A. L.; Gaab, K. M.; Xu, J.; Bardeen, C. J.; Martínez, T. J. Variable Electronic Coupling in Phenylacetylene Dendrimers: The Role of Förster, Dexter, and Charge-Transfer Interactions. *J. Phys. Chem. A* **2004**, *108*, 671−682.
(54) Scholes, G. D.; Harcourt, R. D.; Ghiggino, K. P. Rate Expressions for Excitation Transfer. III. An ab Initio Study of Electronic Factors in Excitation Transfer and Exciton Resonance Interactions. *J. Chem. Phys.* **1995**, *102*, 9574−9581.
(55) Deev, V.; Collins, M. A. Approximate ab Initio Energies by Systematic Molecular Fragmentation. *J. Chem. Phys.* **2005**, *122*, No. 154102.
(56) He, X.; Zhang, J. Z. H. A New Method for Direct Calculation of Total Energy of Protein. *J. Chem. Phys.* **2005**, *122*, No. 031103.
(57) Bettens, R. P. A.; Lee, A. M. A New Algorithm for Molecular Fragmentation in Quantum Chemical Calculations. *J. Phys. Chem. A* **2006**, *110*, 8777−8785.
(58) Fedorov, D. G.; Kitaura, K. *The Fragment Molecular Orbital Method*; CRC Press: Boca Raton, FL, 2009.