Chapter 4:

Empirical Valence Bond Methods for Exploring Reaction Dynamics in the Gas Phase and in Solution

Jeremy N. Harvey, Michael O'Connor, and David R. Glowacki

4.1 Introduction

In this Chapter, we will review the use of empirical valence bond (EVB)¹ methods for studying the dynamics of chemical reactions. The focus will be on recent work that we have done in our groups studying reactions in solution, but we will also include examples of other work on gas-phase reactions. By *reaction dynamics*, we mean the detailed way in which atoms move during, prior to, and after a bond-making or breaking event.

Dynamics can be studied experimentally, *e.g.* through the determination of state-dependent cross-sections and angular distribution functions. Huge progress in this field was obtained using the crossed-molecular beam approach, which was instrumental in the award of the 1986 Nobel prize in Chemistry to Dudley Herschbach, Yuan T. Lee and John Polanyi.² Indirect information about the atomic motions during the chemical reaction is obtained from the observed distribution among energy levels for the products, as a function of the collision energy and of the initial rotational and vibrational states of the reactants. Also important is the spatial distribution of the products for a particular initial collision geometry.

Because the information provided by such experiments is somewhat *indirect*, theory has always played an important role in understanding reaction dynamics, by providing simulated models for the chemical reaction events. Such simulations are of various types. In the simplest case, with which we will mostly be concerned here, the atomic motions are assumed to occur classically on a potential energy surface as defined within the Born-Oppenheimer approximation. Multiple sets of initial conditions (atomic positions and velocities) are chosen based on some classical interpretation of the reaction conditions, and Newton's equations of motion are then integrated numerically. Reactive trajectories are monitored, the final positions and velocities mapped onto the experimental observables, and comparison is made to experiment. In more sophisticated methods, quantum effects on the nuclear motions are introduced, *e.g.* through a classical representation of the nuclear zero-point energy³ or through use of quantum mechanical laws of motion in place of Newton's laws.⁴ Also, the existence of multiple electronic states and motion of the system from one state to another may be treated. In all these cases, it is necessary

to be able to evaluate the potential energy (and its gradient) for many different arrangements of the atoms making up the reacting system.

One obvious way to do this is to solve the electronic Schrödinger equation for the corresponding nuclear coordinates. In the context of classical simulations of nuclear motions, this leads to the direct dynamics approach, or to the related Car-Parrinello simulation approach.⁵ However, these approaches are rather computationally demanding, especially for extended systems, since electronic structure methods typically scale badly with the number of atoms. In this respect, it should be noted that development of new computer architectures and of efficient algorithms is increasingly making such methods applicable to treat dynamics in solution. However, another challenge with such techniques is that the electronic structure method used always involves approximations, and thereby one is not **necessarily** guaranteed a potential energy surface of the required accuracy for the chemical question at hand. It is not straightforward to tune electronic structure methods to correctly predict important quantities such as barrier heights, reaction energy, and so on.

It is against this background – *i.e.*, accuracy balanced against efficiency – that we elected to use EVB methods in our recent work investigating the dynamics of solution-phase reactions. Over the last few years, ultrafast experimental pump-probe techniques, where a chemical reaction is initiated with one short laser pulse, followed by a spectroscopic analysis of the reaction mixture using a second fast pulse have yielded new insights into the dynamics of a number of reactions (for a recent overview, see ⁶). Like the crossed-beam and other gas-phase dynamical studies mentioned above, these approaches yield valuable but indirect information about the elementary reaction behaviour – usually in the form of a transient spectral band. Hence theory can play an important role in helping to interpret the experimental results. This requires accurate simulation of the reactions in solution, and we elected to use EVB for this given the accuracy that it can achieve with a relatively low computational cost.

4.2 EVB and Related Methods for Describing Potential Energy Surfaces

There are many different varieties of the EVB method, as discussed in **Chapter 1**. They have found widespread use in studying a number of problems. In the present chapter, we want to focus only on those studies where the EVB approach has been used to study reaction dynamics, taken in the narrow sense of quantum or classical studies of nuclear motion during chemical reactions, in both non-equilibrium and equilibrium regimes. We thus do not include studies of *rate constants* in the framework of various versions of transition state theory, or the related problem of characterizing potentials of mean force. We do note however that the borderline between advanced versions of transition state theory and ‘reaction dynamics’ is ill-defined.⁷ Also, given the Table of Contents of the present volume, we explicitly do not

set out to discuss the work that has been performed by others on the use of EVB to study the kinetics of biochemical reactions, or the dynamics of proton transport in aqueous solutions. The first of these topics will be covered by several other chapters, and the second will be described in **Chapter 3** by Gregory Voth.

Using this narrow definition of ‘reaction dynamics’ – *i.e.*, the detailed atomic motions that take place as part of a chemical reaction – it turns out that (to our knowledge) there have been rather few studies using EVB. This is an interesting observation in itself, and we start by considering the methods that *have* been used for this purpose.

Classical or quantum studies of reaction dynamics for reactions involving small numbers of atoms have tended to use other methods to represent the potential energy surface, especially various forms of global fit of the energy surface to multiple *ab initio* energies.⁸ For systems with a small number of atoms, this is a reasonable approach, since it is possible to calculate accurate *ab initio* energies, to do so at the many points needed to survey the surface, and to fit the resulting large number of points using a flexible and accurate (but complex) set of analytical basis functions. However, for larger systems, all three aspects become more challenging, and fitting global *ab initio* potential energy surfaces is no longer a viable option. As a result, many studies of reaction dynamics tend toward so-called ‘direct dynamics’, where the potential energy is calculated at each time step using electronic structure theory – usually involving a computationally inexpensive method like semiempirical theory, but also Hartree-Fock or density functional theory (DFT). Apart from choice of the atomic orbital basis set, these methods have the benefit of needing no fitting, and are able to represent a ‘model chemistry’ that encompasses a range of possible reactions within the system under study. This includes reaction pathways which are known in advance, and also those which are unknown. However, electronic structure theory is computationally demanding, especially for more accurate methods, and affordable approaches can be very inaccurate indeed, as discussed for some examples below. For the purposes of running dynamics, one has the liberty to select among a number of related methods that which best reproduces some known features of the potential energy surface. Semiempirical methods afford some flexibility insofar as it is possible to modify their parameterization in order to obtain a better agreement with benchmark values, giving rise to the so-called specific reaction parameter (SRP) semiempirical methods.⁹ There have similarly been attempts to formulate versions of DFT which allowing one to enact a diabaticization into separate reactant and product states, with tuneable coupling interactions to bring barrier heights into better agreement with more accurate approaches.¹⁰

The EVB approach represents a good alternative choice in constructing reactive potentials for systems that are too large to be exhaustively mapped out using electronic structure methods. It effectively relies on molecular mechanics force fields as the ‘basis functions’ for representing chemical reactions, and such

force fields are known to be able to capture the broad features of the potential energy surface around a given energy minimum. Since EVB is a fitted method, it can also be adjusted to reproduce exactly (or near-exactly) the energies of key stationary points on the potential energy surface. The only major disadvantage from the point of view of dynamics is that EVB potentials typically only include reactive channels that have been explicitly included; it is not possible (usually) to ‘discover’ new reactive channels when running dynamics, as happens with direct dynamics or in some cases with fitted potentials.

One elegant study of reaction dynamics was carried out by Chang, Minichino and Miller in 1992.¹¹ The EVB method was used to construct a reactive potential for the CH₂O system, describing the minimum energy configuration, the transition state (TS), and the CO + H₂ dissociation region (neither the HCOH hydroxycarbene secondary minimum nor the TS leading to it was described). The surface was parameterized by requiring the local Hessian at the TS to match the quantum chemical Hessian. The quantum mechanical method used was either second-order Møller-Plesset perturbation (MP2) or coupled-cluster (CCSD), together with a medium-sized basis set. These electronic structure approaches yield fair agreement with the best experimental estimates of barrier heights and reaction energy, but the EVB surfaces were ultimately tuned so as to fit the experimental values.

The fitted EVB surfaces were used to propagate quasiclassical trajectories for dissociation. Based on the target experimental observations, the initial conditions were chosen to represent CH₂O formed by internal conversion of specific S1 photoexcited states – *i.e.*, with high levels of initial rovibrational excitation. The trajectories were initiated close to the dissociation TS, in order to increase the efficiency of obtaining dissociated products, thereby making it possible to have meaningful comparisons with experimentally measured translational, vibrational and rotational energy of the H₂ and CO products. Especially with the CCSD-fitted surface, the agreement with experiment is overall rather good.

Beyond Chang *et al.*’s approach to modelling this reaction, several other quasi-classical trajectory studies have been performed.¹² For example, direct dynamics studies have been performed using Hartree-Fock,^{12a} MP2,^{12c} DFT^{12c} and SRP Hamiltonians.^{12b} The semiempirical methods do not yield very good agreement with experiment; arguably they are even less good than the EVB approach, prompting the authors of that study to write that even reparameterized semiempirical Hamiltonians “might be well suited for locally interpolating *ab initio* data, but not for quantitatively describing global potential energy surfaces”.^{12b} The MP2 and DFT studies perform reasonably well, more or less on a par with EVB. With Hartree-Fock, the relative energy of the TS is incorrect (overestimated by 26 kcal/mol), so it is necessary to express the average translational energy release as a proportion of the available energy in order to make a meaningful comparison with experiment. The EVB approach does not require this scaling, since by construction, the energetics on the surface are very good.

The benchmark potential energy surface for this system is one generated by Bowman and co-workers, which has formed the basis for a number of quasiclassical trajectory studies.^{12d, 12e} Along with careful experimental analysis, these studies have led to the discovery of a novel mode of dissociation for formaldehyde, in which one C-H bond nearly breaks, and the corresponding H atom travels (or ‘rooms’) around the molecule prior to ‘abstracting’ the other H atom en route to H₂ formation. This minor pathway, which is characterized by nascent H₂ with significant internal vibrational energy excitation, was not discovered in the EVB or direct dynamics studies, in part no doubt because the corresponding regions of the potential energy surfaces were not accurate enough, and also because dissociation was initiated in those studies at the standard TS, rather than at the minimum. It is nevertheless notable that the much earlier EVB study is qualitatively reasonably accurate for describing the dynamic behaviour along the main pathway. Presumably, re-parametrization of the EVB surface could be carried out if desired to include a description of those regions of the PES which are important for the ‘roaming’ pathway.

Another dynamical study of a gas-phase chemical reaction using an EVB-type surface was performed by Tishchenko and Truhlar,¹³ for the OH + H₂ → H₂O + H reaction. This too is a good test, since the dynamical behaviour of this reaction has been extensively studied both experimentally and theoretically. In fact, there has been even more work on this reaction than on the formaldehyde dissociation. In this case, the ‘EVB’ potential energy surface used was in fact generated using a modification of the multi-configuration molecular mechanics (MCMM) procedure.¹⁴ It was noted in the paper that the study of dynamics, as opposed to rate constants with transition state theory, provides a more stringent test of potential energy surface accuracy – an observation which is certainly in line with our own work and discussed in further detail below. Tishchenko and Truhlar were more focussed on testing the accuracy of the potential energy representation scheme rather than on comparisons to experiment, and it was found that the MCMM approach is in fact very accurate.

To close this section, we again reiterate that there seem to be relatively few studies of reaction dynamics using EVB potential energy surfaces. For systems with small numbers of atoms, this is easy to understand, since more accurate methods exist which are able to describe several different regions of the energy surface in a balanced way, enabling very subtle aspects of dynamics to be discovered. The example of formaldehyde dissociation and roaming is a case in point. To a certain extent, this is also understandable for larger systems, in that ‘direct dynamics’ methods have the huge advantage of being ‘off-the-shelf’ techniques that can be applied immediately, without the need for precious co-worker time to be invested in fitting of surfaces beyond that which is already baked into the basis sets.

However, we believe that the many advantages of EVB methods mean that they should be used more frequently for dynamical studies, especially for larger systems, where only qualitative or semi-

quantitative accuracy is needed. A significant strength of EVB models of reactive potentials in complex systems is the fact that they build on the significant work (*i.e.*, tens of thousands of man-hours) within the chemical molecular dynamics community that has been invested into the parameterization of molecular mechanics force fields. In what follows, we outline recent work we have recently undertaken in applying the EVB methodology to describe non-equilibrium reaction dynamics in the gas phase, and also in solution – *i.e.*, for solutes which are embedded in both weakly and strongly coupled solvents. Much of this work has been motivated by close collaboration with experimental colleagues carrying out ultrafast time-resolved infrared spectroscopy. As a result, careful comparisons with time-resolved experimental data have been an important aspect in our use of the EVB method, providing good insight into the performance of EVB methods in non-equilibrium regions of the potential energy surface in condensed phases. As our work with the EVB method has expanded to treat increasingly complex systems (particularly the solution phase dynamics results in strongly coupled solvents) by a growing user base, we have been forced to think carefully about the computational efficiency, transparency, and usability of our EVB software implementations and their ability to exploit the power of modern parallel computer architectures. For this reason, we have included within this chapter a section discussing the EVB algorithms that we have implemented as well as those which we are testing, along with measures of their performance. Readers who are interested only in the chemical applications should feel free to skip the section on algorithms; however, more expert readers may be interested in the lessons that we have learned as we have undertaken this work.

4.3 Methodology

The most common approach to EVB reactive dynamics involves a pseudo-Hamiltonian matrix $\mathbf{H}(\mathbf{q})$, which is constructed from two diabatic basis functions – a reactant (R) and product (P) function – yielding a 2×2 matrix with the following structure:

$$\mathbf{H}(\mathbf{q}) = \begin{bmatrix} V_R + \epsilon_R & H_{12} \\ H_{12} & V_P + \epsilon_P \end{bmatrix} \quad (4.1)$$

where V_R and V_P are expressions for the potential energies of the reactant and product diabatic states at a given structure, obtained from fairly standard MM force fields. The ϵ_R and ϵ_P values in **Eq. (4.1)** are constant diagonal energy shifts, usually chosen so as to reproduce the known exo or endo-thermicity of the reaction in question. H_{12} is an off-diagonal element that couples the reactant and product basis functions. The off-diagonal coupling elements are usually taken to be simple functions of the nuclear

coordinates – in our work, typically a Gaussian function of one or two internuclear distances. $\mathbf{H}$ is then diagonalized, *i.e.*, one finds the matrix $\mathbf{U}$ such that:

$$\mathbf{D} = \mathbf{U}^T \mathbf{H} \mathbf{U} \quad (4.2)$$

where $\mathbf{D}$ is a diagonal matrix, whose elements l_i are the eigenvalues of $\mathbf{H}$. $\mathbf{U}$ is a matrix containing the corresponding eigenvectors, $\mathbf{U}_i$. The adiabatic ground state energy is taken as l_0 , the lowest eigenvalue of $\mathbf{H}$; the corresponding eigenvector $\mathbf{U}_0$ contains the coefficients describing the extent to which each diabatic basis state (or force field) contributes to the state with energy l_i . Application of the Hellman-Feynman relation then gives a matrix of Cartesian atomic forces $\mathbf{F}$:

$$\mathbf{F} = -\frac{d\mathbf{D}}{d\mathbf{q}} = -\mathbf{U}^T \frac{d\mathbf{H}}{d\mathbf{q}} \mathbf{U} \quad (4.3)$$

which contains vectors $\mathbf{F}_i$ for the gradients of those adiabatic states which have corresponding eigenvalues in $\mathbf{D}$. The $\mathbf{F}_0$ vector contains those forces which correspond to the lowest eigenvalue l_0 , and is used for dynamics propagation on the adiabatic ground state.

To obtain an EVB potential, the diagonal forcefields are first compared to *ab initio* results (or experiment) for the corresponding stable species, and if necessary, small adjustments are made in order to reproduce the structures or perhaps vibrational frequencies. Then, the behavior of the diagonal forcefields (shifted appropriately to take reaction energy into account) along the reaction coordinate is assessed, and compared to *ab initio* data. As shown in **Figure 1**, the effect of diagonalization of the EVB pseudo-Hamiltonian can only be to *lower* the energy relative to the lowest of the diagonal state energies. Hence if one or more of the diagonal states is significantly lower in relative energy than the *ab initio* data in one section of the reaction coordinate, as for the case of V_2' shown in the schematic **Figure 1**, then it will be impossible to get a good fit of the potential energy surface. Note also that in **Figure 1**, V_1 and V_2 are shown to cross each other rather close (both in terms of energy and in terms of the reaction coordinate) to the maximum of the *ab initio* energy profile. **Figure 1** highlights an ideal scenario for obtaining a good fit to the *ab initio* data in the EVB model. Appropriate choice of the dependence of the off-diagonal coupling matrix element on coordinates can however in principle lead to a good fit even when the crossing occurs away from the maximum.

Figure 1: Schematic representation of an EVB reactive potential energy surface, V , obtained from two diabatic states V_1 and V_2 corresponding to reactant and product states. V is very similar to V_1 for values of the coordinates lying in the reactant region, and to V_2 in the product region. An example of an inappropriate product diabatic state function, V_2' , that lies lower than the target potential energy for some values of q , is shown.

In this respect, it should be noted that along a reaction coordinate with bond-making and/or breaking, one or both of the diagonal states becomes severely anti-bonding in the region corresponding to the minimum of the other state. Consider a case where the states involve atom or group transfer, with V_1 corresponding to an $X-Y\cdots Z$ connectivity and V_2 corresponding to $X\cdots Y-Z$ connectivity. For a structure at which V_1 is low in energy, because $r(X-Y)$ is small and $r(Y-Z)$ is large, V_2 will be high in energy due to unfavourable bond-stretching (large $r(Y-Z)$, far from equilibrium), and due to unfavourable close approach of two atoms that are considered to be non-bonded in this state (X and Y). The very steep behavior of typical non-bonded potentials (such as the Lennard-Jones 12-6 potential used in many forcefields) with respect to interatomic distance means that this second factor is typically dominant, accounting in the case considered for a very high relative energy for V_2 . We have found that it is often necessary to adjust non-bonded energy terms (of the type $Y-Z$ for V_1 and $X-Y$ for V_2) to obtain accurate EVB potentials.

Once it has been established that the diagonal forcefields produce an acceptable description, the off-diagonal terms can be fitted to the target *ab initio* data, usually through application of a non-linear least-squares protocol. In many cases, even a constant off-diagonal element can produce qualitatively reasonable results. In order to avoid persistent coupling away from the intended coupling region, the off-

diagonal element can be obtained as a Gaussian function (or a sum of two or more Gaussian functions) of one or more of the key interatomic distances, and this is the approach we have mostly used to date. For a typical “one-dimensional” Gaussian function (*i.e.*, one which depends on one interatomic distance, or on the difference between two such distances), there are three parameters to be fitted: the amplitude, centre, and width. This means that only a small amount of *ab initio* data is needed in order to obtain a numerically meaningful fit. Even with a sum of two Gaussian functions, there are only six parameters to be fitted. Typically, we use data along the minimum energy reaction path, as well perhaps as a few points displaced from the reaction path (*e.g.* for an atom abstraction reaction $X-Y\dots Z \rightarrow X\dots Y-Z$, a grid of energies for different values of $r(X-Y)$ and $r(Y-Z)$ could be used). A typical fitting set in our work might contain of the order of 50-200 structures and relative energies, though of course more data can always be incorporated into the fitting set. We note that using higher-dimensional Gaussian functions for the coupling matrix element is possible,¹¹ but this requires caution as the corresponding element can remain quite large even quite far from the reference structure.

Some of the applications discussed in the remainder of this chapter concern reaction dynamics in the gas phase; in this case, the *ab initio* calculations used to fit the EVB potential are of course carried out in vacuum. However, our primary focus – and the applications which initially led us to explore the EVB methodology – concern reactions in solution. Some of the reactions which we have studied involve relatively non-polar solutes and relatively weakly-coupled solvents.¹⁵ In such cases, it has seemed reasonable to us to fit the EVB potential in the gas phase, based on gas-phase *ab initio* data. The solute-solvent terms are then simply added in the appropriate way to the diagonal potential energy terms within the framework set up in the corresponding force-field, with the overall diagonal element obtained as a sum of ‘solute’ (meaning the reacting sub-system) terms, solvent terms, and solute-solvent terms. The system was then treated with periodic boundary conditions, and a cut-off for interactions at long range. In a typical forcefield approach, the intra-solvent terms will be the same for all diagonal elements, though obviously the solute terms are very different. Even the solute-solvent terms may be somewhat different, due to different non-bonding interactions (electrostatic, where used, and van der Waals). We have assumed that the off-diagonal terms are not modified by the solvent, an approximation that has been shown elsewhere to be accurate.¹⁶

We have also studied systems with stronger solute/solvent coupling. One of the systems, which is described in further detail below, features a more polar solvent (acetonitrile), together with a fairly polar product (hydrofluoric acid, obtained from fluorine atoms by hydrogen atom abstraction from solvent). In this system, we faced the problem of needing to describe the solvent-product interaction, and we found

that one convenient way to do so was by significantly expanding the EVB matrix to include states for which proton transfer has occurred between the hydrofluoric acid and the solvent.

4.4 Recent Applications

In this section, we describe some recent applications of EVB to generate reactive potential energy surfaces for the study of reaction dynamics in both solution phase and in the gas phase. The work discussed is from our groups in Bristol and Leuven.

Cl + CH₄ in the gas phase

A first simple example is the gas-phase reaction of chlorine atoms with methane to yield HCl and methyl radical. This reaction has been extensively studied experimentally, focussing on such aspects as the distribution of HCl rotational states formed. It has also been studied computationally in several groups. From the point of view of this chapter, this is not a particularly challenging system for which to construct an EVB potential: at the energy regime we are interested in, there is only one reaction possible, and the number of atoms involved is very small. Nevertheless, this system provides a good test of the accuracy of the EVB approach as we use it, since results can be compared both to experiment and to more accurate theoretical approaches. In the latter respect, we emphasize in particular the very accurate coupled-cluster based potential developed by Czako and Bowman.¹⁷ This is based on over 16,000 CCSD(T) energies extrapolated to the infinite basis set limit, and also includes corrections for spin-orbit coupling. The *ab initio* data are fitted with a mean accuracy of 0.2 kcal/mol in the low-energy regime where reaction occurs. The *ab initio* calculations and the dynamical studies can be used to compare experiment and theory for this system. For example, it is useful to compare the predicted distribution of rotational states of the product HCl with the experimentally observed distribution.

To model this system, we constructed a 5x5 EVB Hamiltonian, allowing for abstraction of each of the four hydrogen atoms of the methane.¹⁸ Because the potential energy expressions are identical by symmetry, and only two (at most) of the diagonal potential energies are at reasonable energy at any given structure, the effort involved in fitting is similar to that used for a simpler 2x2 case. The calculations were performed using an in-house code specifically designed to carry out EVB calculations, together with the VENUS code¹⁹ to drive the quasiclassical trajectories. Because the potential energy routines were written specifically for this and related gas-phase problems, it was possible to write efficient serial code, *i.e.* we did not use the MPI or other parallelization frameworks mentioned in later sections. Individual bond stretching, bending, and torsional terms, as well as non-bonded van der Waals terms, were computed, and assembled as needed for the different diagonal energy states. The off-diagonal terms were only introduced

between the reactant state and each of the product states. In each case, these were expressed as Gaussian functions of the C-H and H-Cl distances for the respective hydrogen atom.

Only a small amount of *ab initio* data was necessary in order to obtain a reasonably accurate EVB surface. In fact, two surfaces were constructed, with the first using only 300 *ab initio* calculations for the reactive surface within the minimum energy path region, *i.e.* with C_{3v} symmetry corresponding to collinear approach of the chlorine atom to one of the C-H bonds. In the second approach, a set of 500 points were sampled from reactive trajectories, and *ab initio* energies computed at these points. In this case, **no** structures with C_{3v} symmetry were considered, since such points are not visited during classical trajectories using random initial coordinates and velocities. For both surfaces, the reactant and product state molecular mechanics expressions were fitted to further *ab initio* data for the CH₄, CH₃ and HCl molecules. Both surfaces yielded good fits to the corresponding *ab initio* data, with mean unsigned fitting errors of 0.3 and 0.6 kcal/mol respectively.

We also computed the structure of the transition state saddle-point, and found that it matched that obtained with MP2 and CCSD(T) very well. Similarly, the vibrational frequencies at the transition state were computed and compared to those obtained with CCSD(T). The agreement was again found to be very good, with frequencies on both EVB fitted surfaces differing by about 6 % from the *ab initio* values on average, and by less than 25 % for all modes. This is impressive, given that the EVB parameterization based on collinear input has no *ab initio* input concerning the distorted structures lying away from the minimum energy, collinear, path. It is to be noted that others¹¹ have used more complex expressions for the off-diagonal coupling terms that can *guarantee* exact matching of the *ab initio* Hessian at a number of points – but this has not been used here. This means that the bending frequencies at the TS are described reasonably well only thanks to a reasonable admixture of terms such as H-C-H bending terms in the reactant state. The fact that these provide an adequate description of the TS region is a sign that the pseudo-Hamiltonian formulation of EVB incorporates the correct ‘physics’ and hence has some predictive power.

The EVB potentials do not return perfect results for the reaction dynamics, but they do provide very good results. The HCl product is predicted to be rotationally cold, albeit slightly less so than observed either in experiment or in trajectories with the best fully *ab initio* potential.¹⁷ The slight discrepancy appears to be due to slight inaccuracies in the treatment of the potential away from C_{3v} symmetry in the transitional region between the TS and products. The predicted reaction cross-section is in good agreement with the best theory.

CN + *c*-C₆H₁₂ (CH₂Cl₂ solvent)

Our first EVB potential was developed for the reaction of the cyano radical with cyclohexane, CN + *c*-C₆H₁₂ → HCN + *c*-C₆H₁₁ in CH₂Cl₂ solvent. This reaction was carried out by experimental colleagues using an ultra-fast pump-probe set-up, wherein the nascent HCN was detected using time-resolved infrared spectroscopy, which could discriminate between different HCN vibrational states. Remarkably, a large fraction of the nascent HCN product appeared to form vibrationally hot (*i.e.*, in a $\nu=1$ excited vibrational quantum state), and then undergo subsequent cooling by energy transfer to solvent, both of which could be monitored using time-resolved spectroscopy.^{15b} Surprised by the observations of relatively long-lived excited vibrational quantum states in the products, our experimental colleagues asked us whether the conclusions they had reached based on their spectral assignments could be reproduced by a molecular dynamics model.

Initially, we planned to use direct dynamics, and a gas-phase model system. This was based on previous good experience with direct dynamics methods. However, this reaction is not well suited to this approach: test calculations using high-level *ab initio* methods for the related abstraction of a secondary hydrogen atom from propane showed a highly exothermic reaction, with essentially no barrier, but instead a submerged plateau in the reactant region. Van der Waals attraction on first approach of reactants leads to a drop in energy, and then chemical reaction occurs with essentially no barrier. Using a variety of cheaper *ab initio*, DFT and semiempirical methods, we obtained a potential energy barrier that was much too high with some methods; with other methods, we failed to obtain any barrier, or indeed the plateau indicated by the benchmarks, giving energies along the reaction path that were too low. There were also problems with the calculated reaction energy, and some difficulties due to multi-reference character of the wavefunction were found in simple correlated *ab initio* methods. Overall, even considering the wide range of DFT functionals we tested, no suitable method could be found. The combination of wanting to get a low (or no) barrier with a plateau, a long-range attractive interaction, and a reasonable account of the exothermicity proved too demanding. The best functional considered, B3LYP, was only of marginally acceptable accuracy.

Essentially, it was these difficulties with on-the-fly direct dynamics approaches that prompted us to investigate the suitability of the EVB method. Using the same benchmark CCSD(T) *ab initio* data, the MMFF forcefield was modified to provide an acceptable description of the reactant and product states. Then, a simple 2x2 EVB model [as in **Eq. (4.1)**] was constructed and fitted to the *ab initio* data. A good fit was obtained, and this potential was subsequently used to model the dynamics of reactive events both in the gas phase and in dichloromethane solution. The detail of the simulations is of less importance in the present context, and the reader is referred to the original papers for this aspect.^{15, 20} The key aspect here is

that the efficiency of the EVB potential allowed us to examine the reaction dynamics in a solvent box in a well-converged way. Two hundred and fifty reactive trajectories were sampled, based on different equilibrated initial conditions, and using a simulation system comprised of the CN radical, the cyclohexane co-reactant, and 125 dichloromethane solvent molecules, in a periodically-repeated cubic box of length 23.7 Å. Each simulation involved an equilibration phase (200 ps, 0.5 fs timestep, 4×10^5 steps) and a production phase (200 ps, 0.1 fs timestep, 2×10^6 steps). In total, this makes 6×10^8 propagation steps – which is what is required in order to obtain a converged model for the reactive dynamics.

The availability of an accurate and efficient reactive potential enabled us to carry out extensive sampling and also enabled us to perform various numerical ‘experiments’, in which we varied the initial conditions in order to explore the effect on the dynamics. One of the most fascinating microscopic insights that came out of the MD modelling had a direct impact on the interpretation of the experimental results: we noted that the rate of vibrational cooling of the vibrationally hot HCN seemed to be nearly an order of magnitude faster at very short times (*i.e.*, in the immediate wake of reaction) compared to longer times. We hypothesized that this arose because transfer occurs faster to the cyclohexyl radical co-product than it did within the bulk CH_2Cl_2 solvent itself. Additional simulations were run in which the HCN and cyclohexyl radical were constrained to lie close to one another, and indeed the rate of energy transfer was extremely fast – much faster than the energy transfer that occurred for HCN in bulk CH_2Cl_2 solvent. In interpreting their data, our experimental colleagues had reasonably assumed that energy transfer from the nascent HCN had a rate identical to that which occurs in bulk CH_2Cl_2 solvent. Our EVB MD simulations showed that the energy transfer rate out of the hot HCN is in fact time-dependent – very fast at short times, and significantly slower at long times. We used our conclusions to re-analyze the experimental data, and we found that a time-dependent energy transfer rate led to significantly better fits of the experimental data at short times, providing strong evidence for the nascent HCN product emerging from the reaction even hotter than our experimental colleagues had initially guessed.²⁰ These extensive simulations, which furnished considerable microscopic insight into the measurements that could not have otherwise been obtained experimentally, were largely possible using the EVB approach. In retrospect, it is actually serendipitous that our initial investigations using direct dynamics were partly unsuccessful, since it encouraged us to explore alternative approaches in the first place.

CN + Tetrahydrofuran (Tetrahydrofuran solvent)

In a closely related study, our experimental colleagues in the group of Andrew Orr-Ewing also examined the reaction of CN radical with tetrahydrofuran (THF) solvent (and its perdeuterated form), using the same UV pump/IR probe approach.²¹ In this reaction, experimental measurements of the

vibrational state of the product HCN and DCN was more difficult due to overlap of vibrational bands, but it was still possible to get insight into the dynamics and subsequent vibrational cooling. Similar to the cyclohexane reaction, *ab initio* calculations were carried out, and used as a basis to construct an EVB potential describing the hydrogen atom transfer. Both abstraction from the 2- and 3- positions of THF were considered in the *ab initio* calculations, and the more exothermic reaction was found to be the one adjacent to the ether oxygen in the 2 position. However, both reactions are very exothermic, and have very small or no potential energy barriers. The EVB model was constructed for abstraction in the 3-position.

In the case of HCN formation in cyclohexane discussed above, the timescale for vibrational cooling of HCN in the simulations was in fair agreement with the experimentally measured one, with the initially developed, MMFF-based EVB potential. With CN in THF, this was not the case, apparently due to poor description of the spectral overlap of the C-H stretch with the vibrational spectrum of the THF. Using an adjusted second-generation potential, tailored to attempt to describe the effect of solvent interaction on the C-H stretch, much better results were obtained.

F + CD₃CN (CD₃CN solvent)

This case has been mentioned briefly above in the context of our description of methodology. Again, the study was motivated by collaboration with experimental colleagues studying the reaction using pump-probe experimental methods.²² Whereas our previous investigations were aimed at understanding thermal reaction dynamics of solutes in relatively weakly coupled solvents, this system was chosen to investigate solute dynamics where the solvent coupling was considerably stronger. UV photolysis of xenon difluoride gives rise to fluorine atoms, which then react with the solvent to produce DF. The D—F bond is much stronger than the reactant C—D bond, and therefore the reaction is considerably exothermic. Similar to gas-phase reactions of fluorine atoms, the nascent DF is therefore very vibrationally excited, with essentially zero product of $\nu=0$ DF, and the bulk of the population in either $\nu=2$ or $\nu=3$. The fate of the DF was then probed over the ps time scale using a probe IR laser pulse. The difference with respect to the cyano radical reactions (already mentioned) is that DF is quite polar, and interacts quite strongly with acetonitrile solvent molecules, as well as the CD₂CN by-product, forming a strong hydrogen bond. It was anticipated that these features would lead to interesting dynamics both for the atom abstraction reaction and the post-reaction structural adjustments and energy redistribution.

Accordingly, we constructed an EVB model for the reaction.²²⁻²³ As described previously, this proved more challenging than in the earlier cases mentioned here, as we required a potential that would describe both the atom abstraction (this proved relatively easy to achieve) and the hydrogen-bonding in product. This second feature is described reasonably well at least with respect to the structure and energy of the

equilibrium structure through the electrostatics and van der Waals terms in the product state forcefield. Also, since DF is such a weak acid, no deuteron transfer is expected in the process, and indeed no evidence for it was observed experimentally. In fact, our CCSD(T)-F12 *ab initio* calculations show that there is no secondary potential energy surface minimum corresponding to a $\text{CD}_3\text{CND}^+ \cdots \text{F}^-$ arrangement of the nuclei. However, for high vibrational excitation of the DF, partial transfer to the solvent can occur. The molecular mechanics force-field does not describe this feature very well, yielding energies which are much too high for the corresponding structures compared to the *ab initio* results, by failing to describe the significant anharmonicity in the potential energy surface, coupling the D—F bond stretch to the CD_3CN —DF hydrogen bond.

In principle, description of this feature in the product region does not require an EVB treatment. A carefully constructed single-state molecular mechanics potential energy surfaces could instead be used, likely requiring a polarizable forcefield to describe the required potential energy surface features. However, we elected to use an EVB approach as described earlier, adding extra EVB states corresponding to $\text{CD}_3\text{CND}^+ \cdots \text{F}^-$ ‘states’ of the system. This is effectively an alternative strategy for describing the significant polarizability of the DF and the acetonitrile upon close approach to one another. Specifically for this system, we treated a fluorine radical embedded in n solvent molecules with a matrix having a dimension of $(n + 2) \times (n + 2)$, with the following structure:

$$\mathbf{H} = \begin{bmatrix} V_1 + \epsilon_1 & H_{12} & 0 & 0 & \cdots & 0 \\ H_{12} & V_2 + \epsilon_2 & H_{23} & H_{24} & \cdots & H_{2n+2} \\ 0 & H_{23} & V_3 + \epsilon_3 & 0 & \cdots & 0 \\ 0 & H_{24} & 0 & V_4 + \epsilon_4 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & H_{2n+2} & 0 & 0 & \cdots & V_{n+2} + \epsilon_{n+2} \end{bmatrix} \quad (4.4)$$

where the diagonal elements V_1 and V_2 correspond to the respective energies of the reactant ($\text{F} + n \text{CD}_3\text{CN}$) and product ($\text{DF} + \text{CD}_2\text{CN} + (n-1) \text{CD}_3\text{CN}$) states, respectively. Diagonal element V_3 treats a state involving proton transfer from the product DF to the CD_2CN co-product (*i.e.*, $\text{F}^- + \text{CD}_2\text{CND}^+ + (n-1) \text{CD}_3\text{CN}$), while elements $V_4 \dots V_{n+2}$ correspond to the states in which the DF transfers its proton to all possible solvent molecules (*i.e.*, $\text{F}^- + \text{CD}_3\text{CND}^+ + \text{CD}_2\text{CN} + (n-2) \text{CD}_3\text{CN}$). The off-diagonal matrix element H_{12} (modelled using a two-dimensional ellipsoidal Gaussian function of the C—D and D—F distances for the transferring deuterium atom) is responsible for coupling together reactant and product diabatic states i and j . The matrix elements between state 2 and states 3 to $n+2$ are one-dimensional Gaussian functions, and describe the mixing in of deuteron-transfer character to the product state.

The abstraction and post-reaction potentials, shown in **Figure 2** with comparisons to the CCSD(T)-F12 results, reproduce very well the solvatochromatic shift of the infra-red stretching frequency observed between gas-phase DF, and DF within an equilibrium CD₃CN solvation environment. It also provides an excellent description of the experimentally observed abstraction dynamics, and again enables efficient reaction simulation. MD simulations using this EVB model predict a DF product formed with a very high vibrational excitation, on the order of $v=2-3$, in essentially quantitative agreement with experimental observations. The subsequent vibrational cooling that takes place as DF's excess energy leaks into the solvent is also observed to occur on the same timescale as that measured experimentally. Once again, the EVB approach is key to being able to sample enough dynamics to obtain well-converged statistical treatment of the reaction. **Many hundreds** of simulations were performed, each involving **many hundreds** of thousands of timesteps, which is affordable given the short time required to evaluate the EVB potential, despite its complexity in this case.

Figure 2: Comparison between CCSD(T)-F12 and multi-state EVB (MS-EVB) potential energy surfaces for the $F + CD_3CN$ reaction, along the reaction profile (left) and for the solvent-DF hydrogen bonding region (right).

As in the case of the cyano radical reaction, one of the other benefits of the EVB approach is that it is affordable enough to enable simulation with modified physics, which provide insight into the origin of the phenomena observed. In the present case, two sets of modified simulations were key to unravelling two competing effects that influence the appearance of the DF's transient IR spectrum. The first effect is vibrational cooling, which brings the DF from a very anharmonic region of the potential energy curve into a harmonic region of its potential energy curve, and thereby leads to an increase in the vibrational frequency (manifest as a spectral blue-shift in the transient IR spectrum of DF).

The other effect is hydrogen bond formation. In the reactants, the F atom is neutral, and interacts only weakly with solvent, perhaps through two-centre three-electron bonding with the nitrile nitrogen lone pair. Upon reaction, DF is formed, but the D is initially oriented towards the carbon atom from which it has just been abstracted. There is no possibility for hydrogen bonding to appear during the abstraction step – the nitrogen of the CD₂CN co-product radical is not in an orientation that allows it to participate in H-bonding, and neither are the other solvent molecules. Motions of the DF rapidly allow it to form H bonds, on the sub-ps timescale, and this in turn leads to a significant loosening of the D—F bond, and a associated red shift in the transient IR spectra.

Simulations with modified physics furnished the microscopic insight that allowed us to understand these red and blue shifts separately. The first key part of this analysis was the ability to carry out reactive trajectory simulations of a rather special variety – *i.e.*, modified so as to remove the excess vibrational energy of the nascent DF by suitable damping of the atomic velocities at the moment of DF formation. Hence the nascent DF is formed in the correct post-reaction solvation environment (*i.e.*, without hydrogen bonding), but is vibrationally cold. These trajectories (with DF in post-reaction solvation environment, but with its vibrational excitation damped) were critical in examining non-equilibrium post-reaction solvation dynamics. They showed a significant *red* shift that accompanies relaxation of DF's post-reaction solvation environment in the picoseconds following reaction – providing a clear link between the observed transient DF spectral profile and the ultrafast onset of the hydrogen bond following reaction.

The second key part of our spectral analysis for this reaction involved simulations in which a DF molecule was first allowed to reach an equilibrium solvation environment, and then underwent a modification of its atomic velocities so as to inject an amount of vibrational energy similar to that yielded by the reactive trajectories. These simulations effectively amounted to a Franck-Condon excitation of DF within an equilibrium solvation environment. These simulations of non-equilibrium vibrational relaxation dynamics showed that the DF – which was both hydrogen bonded *and* vibrationally excited – did indeed cool down, at roughly the same rate as observed in the full non-equilibrium reactive dynamics simulations, and that relaxation was linked to a significant blue-shift in the calculated IR stretching band.

Using these two sets of additional simulations to analyse the full non-equilibrium reactive dynamics simulations was possible owing to the efficiency of the EVB potential. They also provided valuable microscopic insight into various reactive and solvation dynamics phenomena (and their associated timescales). Experimentally, and in the full simulations, only a modest shift in the DF infra-red absorption spectrum is observed for the nascent product, which the simulations allow us to attribute to a cancellation of the blue-shift and red-shift observed following DF formation.

Diazocyclopropane ring opening

Our final example concerns a much more complex reaction: ring-opening of a diazo-cyclopropane, through loss of nitrogen to form a carbene, which undergoes C—C bond cleavage to yield an allene (**Figure 3**). This reaction was chosen by our collaborator, Barry Carpenter, as a possible candidate for solvent-induced enantioselective amplification. The reactant has a plane of symmetry, and is hence non-chiral, but the product allene can exist as either of two enantiomers (both shown in **Figure 3**). The hypothesis (at this stage only a hypothesis, since the reaction has not been carried out experimentally) is that carrying out the ring-opening in a chiral solvent might lead to observation of the product in a chiral form, enriched in one or other of the enantiomers. This would be a dynamical effect,²⁴ since the minimum energy path from reactants, over a transition state for N₂ loss, and then formation of product, starts out completely symmetric, then ‘splits’ at a valley-ridge inflexion point²⁵ into two equivalent but chiral paths leading to the two enantiomers of product.

Figure 3: Ring-opening of a diazocyclopropane, leading to enantiomeric products.

This suggestion was initially tested through some simulations using a simple model potential,²⁶ then through direct dynamics simulations with DFT.²⁷ In the latter calculations, a chirally-arranged set of point charges were used to simulate the effect of a chiral solvent.²⁷ However, it was clearly desirable to be able to model the reaction using a full solute + solvent model. Accordingly, an EVB model was constructed, and used to explore the dynamics in explicit solvent.²⁸ Intriguingly, it does appear that with some chiral solvents, a very significant excess of one enantiomer of product over the other should be obtained – and this is being further investigated.

This final example shows that the EVB approach is not limited to atom-abstraction reactions – the fact that the earlier examples we have studied have been of this type is due to the nature of the projects that our experimental collaborators have worked on as we have endeavoured to push the limits of the

EVB methodology for non-equilibrium reaction dynamics, rather than an intrinsic limitation of the EVB protocol.

4.5 Software Implementation Aspects

Throughout the course of our EVB studies of fundamental reaction dynamics, increasing numbers of colleagues, collaborators, post-docs, and students have gotten involved, and there has been a steady increase in the efforts to extend the methodology to systems of increasing complexity (*e.g.*, solution phase). Therefore, we have found that it is in our interest to build EVB frameworks which are able to benefit from the substantial research effort that has been invested in popular MD software packages and force fields developed over the last few decades, with which many workers across chemistry already have some familiarity. Some established simulation codes have EVB functionality (*e.g.*, AMBER), but it has not been of the type that we wished to use, nor has it been sufficient to handle cases where one wishes to include large numbers of states. In our groups, we have developed three different types of code. The first – perhaps the most ‘bespoke’ – is for the simple gas-phase $\text{Cl} + \text{CH}_4$ system, in which we have simply coded all aspects of the calculation of the diagonal and off-diagonal matrix elements of the EVB matrix, together with the diagonalization and gradient computation.¹⁸ Second, for the more complex solution-phase simulations, we have instead chosen to adapt existing codes to carry out EVB calculations. Third, we have recently added EVB functionality to accelerated force field routines that we have been developing over the last few years, and which are designed to run interactively for real-time steering by human users on multi-core mobile and laptop architectures.²⁹ Whereas the first code uses fairly traditional serial programming, the latter two codes have relied heavily on hardware-adapted parallelization algorithms, both as a strategy that enables existing force field codes to carry out EVB calculations in a fashion that is relatively transparent, and also to ensure computational efficiency. In the rest of this section, we give an overview of the parallelization strategies that we have implemented to accelerate our EVB MD simulations on modern computer architectures – *i.e.*, the latter two classes of code mentioned in the previous paragraph.

It needs pointing out that (by the standards of modern biomolecular simulations) the solution-phase chemical dynamics applications discussed in this chapter which we have carried out to date are rather small in terms of the number of atoms involved in each simulation: typically a few hundred. Nevertheless, they incur non-trivial computational effort, for the following reasons: (1) because we have undertaken detailed comparisons with time-resolved infrared spectroscopy measurements, we need to obtain results over a wide range of initial conditions and ensure that we have converged the simulation statistics, meaning that many trajectories are required, and (2) when large numbers of EVB states are involved – *e.g.*, in the case of F in CD_3CN solvent, the computational cost increases significantly. In an effort to

speed up our EVB molecular dynamics simulations and increase their efficiency for future studies, we have experimented with different parallelization strategies over the past couple years. The first strategy, which we outline below, is based on the message-passing interface (MPI) and exploits multi-core computer processing unit (CPU) architectures. We have also investigated strategies designed to run efficiently on single-instruction-multiple-data (SIMD) stream processing architectures like graphical processing units (GPUs).

CPU parallelization using MPI

As discussed above, our production simulations for $F + CD_3CN$ in CD_3CN solvent, which described the reactive dynamics of an F radical embedded in 62 solvent molecules, yielded a 64×64 $\mathbf{H}(\mathbf{q})$ matrix. The decision to utilize a 64-state matrix was determined through consideration of the number of CPU cores which we could reasonably exploit on the architectures available to us, the minimum size of the simulation required to quench DF without unduly heating the bath, and the fact that ~~the~~ our computational resources consisted of 8-core CPU nodes. These simulations relied on the parallelized dynamics propagation strategy schematically illustrated in **Figure 4**. The propagation algorithm works by instructing each diabatic state to calculate its energy and forces in parallel as a separate thread within an MPI framework. The master thread then: (i) gathers together the results for each state to construct the diagonal matrix elements for $\mathbf{H}(\mathbf{q})$; (ii) calculates the off-diagonal matrix elements of $\mathbf{H}$; (iii) undertakes the diagonalization required to solve **Eq. (4.2)**, yielding the l_0 eigenvalue of the $\mathbf{D}$ matrix and its corresponding eigenvector $\mathbf{U}_0$; and (iv) solves **Eq. (4.3)** to yield the Hellman-Feynman force vector $\mathbf{F}_0$. Once this is done, $\mathbf{F}_0$ and l_0 are then dispatched to each MPI thread, overwriting the force and energy data held on each thread from the previous timestep. Each thread then propagates forward a single dynamical timestep, with the identical forces and energies ensuring that the new geometry on each thread is identical. Each thread then carries out its own energy and force calculations, the results of which are specific to the connectivity of the particular diabatic state. This parallelized propagation strategy, which we henceforth refer to as *state-parallelism*, scales nearly linearly in the number of EVB states so long as: (1) a large enough multi-core architecture is available across which to parallelize the MPI threads, and (2) the cost for any given thread to carry out a force field calculation is significantly larger than what it costs the master thread to diagonalize the Hamiltonian matrix and calculate the Hellman-Feynmann forces at each time step. **Figure 5** shows the scaling that we see for our MPI-parallelized CHARMM implementation up to 64 cores (*i.e.*, eight 8-core nodes), tested on the 64-state $F + CD_3CN$ systems described above, along with analysis of the timings spent on particular computational tasks. For very large matrices, diagonalization will eventually emerge as the computational bottleneck; however, we have yet to arrive at this point for any of the examples discussed above.

Figure 4: MPI state-parallelized EVB propagation scheme

Figure 5: Left hand panel shows the scaling of the state parallelized EVB scheme in **Figure 4** (tested for F + CD₃CN) vs. the number of 8-core CPU nodes utilized, up to a maximum of 64 CPUs. The right hand panel shows a decomposition of the relative timings required to carry out various computational tasks during the MS-EVB runs for the F + CD₃CN.

Since beginning the work described in this chapter, we have implemented the state-parallelized MPI scheme shown in **Figure 4** (and benchmarked in **Figure 5**) in both CHARMM and TINKER, allowing us access to the wide range of flexible force field methods and related molecular modelling algorithms available within both of these programs. **Figure 5** shows that we get scaling which is close to linear as

one increases the number of cores used to construct and diagonalize the 64-state $\mathbf{H}(\mathbf{q})$. The fact that we have been able to implement the **Figure 4** scheme in well-known simulation programs highlights another key aspect of the MPI parallelization described above: because it essentially involves several very similar replicas of the same executable, relatively little modification is required to the underlying simulation code itself. The resulting MPI-parallelized EVB code rather requires only top level modifications, and ends up looking rather like a sort of parallel wrapper, in which the routines for calculating energy and forces are modified so as to include the appropriate MPI_GATHER commands, a matrix diagonalization, and a Hellman-Feynmann multiplication routine. To run an EVB calculation using this framework, the user has only to run the appropriate MPI executable, specifying the number of MPI threads to launch, the state topology for each thread, the functional form and parameters for the coupling elements, and the diagonal offsets.

GPU parallelization

Modern molecular mechanics simulation codes are increasingly making use of GPUs for force field evaluation. For example, AMBER, GROMACS, and NAMD all have GPU implementations of their force-field routines, which allow the fast evaluation of force fields for very large systems, and which we henceforth refer to as *force-field parallelism*. Molecular simulation code parallelized on GPUs often ends up looking rather different than CPU-parallelized code like that discussed above, as a result of the differences in their underlying hardware architectures. CPUs are suited to heavy or complex processing no more than a few parallel tasks: they excel at fast sequential execution of programs with complex logic. GPUs, on the other hand are designed to exploit massive parallelism in which (typically) thousands of lightweight threads (each with a small memory space) are run on thousands of cores designed to carry out efficient parallel floating-point arithmetic. They excel at multi-threaded, fine-grained, data parallel computing.

The observation that data-parallelism (*i.e.*, performing the same operation on a large amount of data) is prevalent in many scientific algorithms including molecular dynamics, is in fact what has led to the widespread adoption of GPU-accelerated routines in scientific applications. GPUs have their origins in graphical rendering applications, where a single thread is often tasked with evaluating a relatively simple operation on a specific pixel or small subset of pixels (in an image which is potentially composed of millions of pixels). Scientific software applications often have a structure which is not too different – *e.g.*, force field evaluation potentially requires millions of identical calculations, each of which is independent of any another. The current trend in CPU architectures is also one of increased parallelism: this includes increasing the number of cores within a CPU, and also hyper-threading and vector parallelism within a

core. Devoting time to optimally exploit parallelism in modern algorithms will be beneficial both now and in the future, as parallel computer architectures continue to evolve.

The simplest way to exploit the parallel capabilities of GPU accelerated molecular mechanics codes in efficiently evaluating EVB forces is essentially to stick with the state-parallel framework in **Figure 4**, organizing tasks such that a given state-specific force evaluation (each of which occurs on a separate MPI thread in **Figure 4**) is executed on a GPU rather than a CPU. For most force fields, the terms contributing to a given diagonal element V_i are expressed as follows:

$$V_i = V_{nonbond} + V_{bond} + V_{angle} + V_{torsion}$$

For a single state that includes N atoms, $V_{nonbond}$ is a formally $O(N^2)$ operation that involves computing non-bonded interaction between all pairs of atoms, while other terms typically scale linearly. The overall scaling of the evaluation of energy and gradients, in big-O notation, for a system with M states, is then $O(MN^2)$. Assuming that GPU-accelerated routines for fast force calculations within a particular state are available, then the most obvious strategy for GPU-accelerated parallelism involves sequentially evaluating V_i for each state, but with the internal force field terms for each V_i computed in parallel on the GPU.

The two parallel approaches discussed so far – *state parallelism* and *force field parallelism* – highlight the two factors at play in achieving high performance. State parallelism performs best where there are a large number of states, each with relatively few atoms. An example of this case is the F + CD₃CN work described above, where less than a thousand atoms were simulated, and sequential evaluation of the force fields was not a massively limiting factor. The force field parallelism approach would perform poorly in this case, because the number of atoms N needs to be sufficiently large to fully saturate the large number of stream processors typically available on the GPU, and sequentially computing a large number of states becomes the limiting factor on performance. Conversely, when the number of states is small, and N is sufficiently large, then the force field parallelism approach takes advantage of the large degree of parallelism offered for the computation of V_i on the GPU. The latter case is more typical of most EVB applications one finds within the literature.

Below we outline a hybrid approach incorporating aspects of both *state parallelism* and *force field parallelism*, which we have implemented on the GPU in a custom package designed to carry out real-time user guided reactive MD.²⁹ While there are a number of programming languages designed to expose the parallelism of the GPU, we have used the Open Computing Language (OpenCL), which provides a hardware portable framework for writing high performance programs on GPUs and multi-core CPUs. In OpenCL, the basic unit of work is referred to as a work-item, which may be thought of as a lightweight

thread performing an operation on a unique portion of data. The code for a specific work-item (*i.e.*, a C function) is referred to as a *kernel*. Work-items are grouped into *work-groups*, each of which is computed independently of every other *work-group*. Expressing the task in this way allows one to exploit significant parallelism in accomplishing the overall task, because each work-group can be enqueued independently. Work-items with a specific work-group will typically perform the same operation on a different piece of data; this structure allows for further parallelization in the form of vector operations, which perform the same task on a contiguous array of data, and lead to considerable computational acceleration. This method for expressing parallelism works well on GPUs, where large groups of threads or vector operations are computed simultaneously.

In our hybrid approach, we constructed an OpenCL kernel for each of the force field terms included in V_i . As described above, the force field for a given state may be computed in parallel, and each term contributing to V_i may be also be computed in parallel. This provides an ample amount of parallelism which can be used to accelerate EVB force field evaluation for several coupled states; however, the parallelism is fundamentally *two-dimensional*, involving both the state index and the force field term index. OpenCL is particularly well suited to this task, because it contains within it the notion of multi-dimensional work-items – *i.e.*, a work-item may be expressed in terms of a particular state i and also to a particular term within state V_i . For example, in the calculation of $V_{nonbond}$ a work-item will compute the Lennard-Jones force for a particular atom in a particular state. The parallelism within a state is performed on a per-atom basis – *i.e.*, a work-item for a particular atom in a particular state includes a loop for Lennard-Jones interactions between that atom and the other $N-1$ atoms in the system. This parallelization strategy involves a certain amount of redundant computation (*e.g.*, there are $N(N-1)$ Lennard-Jones calculations rather than $N(N-1)/2$ calculations). This may seem wasteful, but for GPU architectures this actually turns out to be the more efficient strategy owing to the fact that it allows for coalesced memory access to atomic data, which is the greatest concern in achieving high-speed access to device memory. With this parallel strategy, the number of work-items is then $M \times N$, which for large N and/or M generates a significant number of lightweight threads that can be executed in parallel. This feature is what makes the approach well suited to exploiting massively parallel architectures such as GPUs. For efficiency in both computation and memory access, the forces acting on an atom for a particular state are also accumulated during the evaluation of each term in V_i .

Once the state energies and forces are computed, the state energies are copied from the GPU to the host and $\mathbf{H}$ is constructed and diagonalized using LAPACK routines. The eigenvalues and eigenvectors of $\mathbf{H}$ are then obtained, and the forces acting on the system may be computed. Off-diagonal terms in the $\mathbf{H}$ matrix are presently calculated sequentially on the CPU; parallelizing their computation on the GPU would require a separate kernel for each type of off-diagonal term, and this work is in progress.

Figure 6: Illustration of the GPU accelerated EVB hybrid parallel framework

Figure 6 provides a schematic that illustrates the overall hybrid parallel approach that we have implemented to accelerate the EVB method on GPUs, highlighting how the notion of a multi-dimensional work-item allows one to create sufficient lightweight threads to saturate the GPU cores, and thereby most effectively exploit the parallelism available on the architecture. We benchmarked our implementation on an interactive MD simulation²⁹ of OH embedded in a bath of methane (CH₄) molecules, which is able to undergo one of atmospheric chemistry's most prototypical reaction sequences: CH₄ + OH → CH₃ + H₂O. Because the hydrogen atoms on each methane molecule are able to undergo abstraction by the OH radical, an increase in the number of methane molecules in the system increased the number of EVB states. The benchmarking was run on a 2014 Macbook Pro with a relatively modest NVIDIA GTX 570M GPU (336 CUDA cores). **Figure 7** shows results obtained for two different GPU acceleration approaches

– a force-field only parallelization, in which the force fields for different states are calculated sequentially, and the hybrid strategy in which parallelization is implemented across both the states and the force fields. **Figure 7** shows timings for the construction of H using the two different approaches, and also for the diagonalization of H . The force-field only parallelization strategy exhibits poor performance for even just one OH and one CH₄ molecule, and the hybrid parallelization strategy gives a significant speed-up. Up to approximately 41 states, the hybrid implementation does not saturate the parallel cores of the GPU, and therefore scales sub-linearly. After the saturation point, the quadratic scaling of the nonbonded term is evident (with over an order of magnitude speedup compared to the force-field only parallel implementation!). Beyond 50 states, diagonalization of H , which scales as $O(N^3)$, begins to take longer than the force evaluation. The **Figure 7** benchmark-tests utilized a serial LAPACK diagonalization algorithm; a parallelized implementation should exhibit better performance. Because condensed-phase EVB systems typically consist of many atoms and relatively few states, the cubic scaling of matrix diagonalization does not typically become a bottleneck.

Figure 7: Computational scaling of the time required to construct the H matrix, comparing the force-field only parallelization approach with the hybrid parallelization approach. The time required for matrix diagonalization is also included. Benchmarked on CH₄+OH system on a 2014 Macbook Pro with an NVIDIA GTX 750M GPU.

4.6 Conclusions and Perspectives

In this Chapter, we have provided an overview of the use and development of EVB potentials for exploring reaction dynamics. We first compared the EVB approach to other methods to describe reactive potential energy surfaces, especially global fits based on extensive *ab initio* data, and ‘direct’ or ‘on the

fly' dynamics where the potential is calculated using DFT or other cheap electronic structure methods at each time step. Globally fitted potentials are able to be highly accurate, and for small systems, EVB will not generally be competitive with them for accuracy - though we discuss one example, the reaction between Cl and methane, where an EVB potential is at least similarly accurate as a global fit. For larger systems, especially for describing dynamics in condensed phases, EVB will likely carry on as an attractive option, and is much more efficient than direct dynamics. This enables a more systematic exploration of reaction dynamics, including modifying the physics of the system to attempt to unravel the impact of different effects on the observed behaviour. We also reviewed some other applications of EVB to reaction dynamics studies - though there are relatively few such studies.

Next, we described briefly the way in which we have developed EVB potentials for studies of reaction dynamics, based on *ab initio* calculations, and fitting of relatively simple functional forms for off-diagonal coupling matrix elements. We then reviewed a number of recent applications we have made of the EVB approach to describing reaction dynamics. A simple system, abstraction of hydrogen atom from methane by chlorine, is an example of a gas-phase reaction that has been studied by many other theoretical approaches. The EVB potential we developed is of good accuracy, in terms both of its ability to reproduce the reference *ab initio* data, and its ability to predict dynamical observables that agree with experiment. We next covered three other atom abstraction reactions, for abstraction of hydrogen atoms by the cyano radical from cyclohexane or tetrahydrofuran, and by the fluorine atom from acetonitrile. In each case, we carried out simulations in solution, using periodic boundary conditions within the EVB framework. The high efficiency of the EVB framework allowed us to carry out extensive simulations, well converged in terms of the observables that were monitored by our experimental collaborators. For the reactions of cyano radical, relatively non-polar species are involved and simple 2x2 EVB models were adequate to describe the dynamics. For fluorine atom reaction in acetonitrile, the product hydrogen fluoride forms strong hydrogen bonds with the solvent, and we found that a more complex EVB model was needed, in which additional EVB states were included to describe partial proton transfer. Finally, we discuss an example of a more complex reaction, ring-opening of a diazocyclopropane in a chiral solvent.

We also discussed the algorithmic progress that we have made in developing efficient EVB codes which are adapted to modern parallel architectures. The most obvious parallel strategy (*e.g.*, that which we have implemented in CHARMM and TINKER) utilizes *state parallelism*, and involves using MPI to parallelize force fields on different states across different CPU cores. Another approach involves *hybrid parallelism*: this is a parallel two-dimensional strategy across both the state index and force field term index. Our efforts in this algorithmic and software domain are ongoing. For example, we are investigating strategies for parallelizing across multiple GPUs on separate compute nodes. The most straightforward

way to accomplish this is to use a hybrid MPI/GPU parallel framework – *e.g.*, in the case where one has access to multiple GPU compute nodes, each node is assigned the task of calculating GPU-accelerated force fields for a particular state. This will require a distributed memory approach, involving data reductions of the energies and forces across nodes in order to achieve scalable performance. We are also experimenting with identifying matrix elements of H with extremely high energies (and hence contribute nearly nothing to the final eigenvector), so that they can be screened from inclusion in H , thereby reducing the total cost of force evaluation, and the cost of diagonalization.]

Overall, the EVB method, which has had much success in modelling other aspects of chemical reactivity, seems very well suited to exploring reaction dynamics of quite complex systems in solution. It is efficient, so that extensive classical (or quantum, though this has not been done here) simulation is possible. Yet it is also fairly easy to parameterize EVB potentials based on ~~quite~~ limited *ab initio* data.

Molecular mechanics force fields continue to be improved – techniques for fitting to non-standard systems are becoming more available, accuracy is being improved, and important effects such as polarizability are now being treated increasingly frequently. As EVB is a method that builds upon molecular mechanics, this also means that EVB potentials will continue to improve. The approach we have used is but one member of the family of empirical methods based on pseudo-Hamiltonian valence-bond techniques, and the many other formalisms discussed in this volume will also continue to be successful.

In closing, it is also worth pointing out that the study of chemical reaction dynamics has traditionally focused on simple gas-phase reactions. This is not so much because only these reactions display interesting dynamical effects – but rather because these are the systems that are the most amenable to study using spectroscopic methods that are sensitive to the timescales involved. Ultrafast spectroscopy is increasingly used to study dynamics in more complex systems, and theory is often an essential partner for such studies, as the confounding effect of solvent often means that less precise experimental is available than for simpler gas-phase systems. We believe that dynamical effects will start to become investigated for systems that are of genuine relevance to synthetic chemistry – and EVB-based theoretical methods will be ideally placed to form a strong partnership with experiment in such investigations.

4.7 References

1. Warshel, A.; Weiss, R. M., An empirical valence bond approach for comparing reactions in solutions and in enzymes. *Journal of the American Chemical Society* **1980**, *102* (20), 6218-6226.
2. *The Nobel Prize in Chemistry 1986*. *Nobelprize.org*. Nobel Media AB 2014. Web. 27 Jul 2015. http://www.nobelprize.org/nobel_prizes/chemistry/laureates/1986/.

3. Bonnet, L.; Rayez, J. C., Quasiclassical trajectory method for molecular scattering processes: necessity of a weighted binning approach. *Chemical Physics Letters* **1997**, *277* (1-3), 183-190.
4. (a) Althorpe, S. C., The plane wave packet approach to quantum scattering theory. *International Reviews in Physical Chemistry* **2004**, *23* (2), 219-251; (b) Habershon, S.; Manolopoulos, D. E.; Markland, T. E.; Miller, T. F., Ring-Polymer Molecular Dynamics: Quantum Effects in Chemical Dynamics from Classical Trajectories in an Extended Phase Space. *Annual Review of Physical Chemistry*, Vol 64 **2013**, *64*, 387-413.
5. Car, R.; Parrinello, M., UNIFIED APPROACH FOR MOLECULAR-DYNAMICS AND DENSITY-FUNCTIONAL THEORY. *Physical Review Letters* **1985**, *55* (22), 2471-2474.
6. Orr-Ewing, A. J., Dynamics of Bimolecular Reactions in Solution. *Annual Review of Physical Chemistry*, Vol 66 **2015**, *66*, 119-141.
7. Liu, Y. P.; Lu, D. H.; Gonzalez-Lafont, A.; Truhlar, D. G.; Garrett, B. C., DIRECT DYNAMICS CALCULATION OF THE KINETIC ISOTOPE EFFECT FOR AN ORGANIC HYDROGEN-TRANSFER REACTION, INCLUDING CORNER-CUTTING TUNNELING IN 21 DIMENSIONS. *Journal of the American Chemical Society* **1993**, *115* (17), 7806-7817.
8. Braams, B. J.; Bowman, J. M., Permutationally invariant potential energy surfaces in high dimensionality. *International Reviews in Physical Chemistry* **2009**, *28* (4), 577-606.
9. (a) Greaves, S.; Rose, R.; Abou-Chahine, F.; Glowacki, D.; Troya, D.; Orr-Ewing, A., Quasi-classical trajectory study of the dynamics of the $\text{Cl} + \text{CH}_4 \rightarrow \text{HCl} + \text{CH}_3$ reaction. *Physical Chemistry Chemical Physics* **2011**, *13* (23), 11438-11445; (b) Troya, D.; Weiss, P. J., Ab initio and direct quasiclassical-trajectory study of the $\text{Cl} + \text{CH}_4 \rightarrow \text{HCl} + \text{CH}_3$ reaction. *The Journal of chemical physics* **2006**, *124* (7), 074313.
10. Cembran, A.; Song, L.; Mo, Y.; Gao, J., Block-localized density functional theory (BLDFT), diabatic coupling, and their use in valence bond theory for representing reactive potential energy surfaces. *Journal of chemical theory and computation* **2009**, *5* (10), 2702-2716.
11. Chang, Y. T.; Minichino, C.; Miller, W. H., Classical trajectory studies of the molecular dissociation dynamics of formaldehyde: $\text{H}_2\text{CO} \rightarrow \text{H}_2 + \text{CO}$. *The Journal of chemical physics* **1992**, *96* (6), 4341-4355.
12. (a) Chen, W.; Hase, W. L.; Schlegel, H. B., Ab initio classical trajectory study of $\text{H}_2\text{CO} \rightarrow \text{H}_2 + \text{CO}$ dissociation. *Chemical physics letters* **1994**, *228* (4), 436-442; (b) Peslherbe, G. H.; Hase, W. L., Semiempirical MNDO, AM1, and PM3 direct dynamics trajectory studies of formaldehyde unimolecular dissociation. *The Journal of chemical physics* **1996**, *104* (20), 7882-7894; (c) Li, X.; Millam, J. M.; Schlegel, H. B., Ab initio molecular dynamics studies of the photodissociation of formaldehyde, $\text{H}_2\text{CO} \rightarrow \text{H}_2 + \text{CO}$: Direct classical trajectory calculations by MP2 and density functional theory. *The Journal of Chemical Physics* **2000**, *113* (22), 10062-10067; (d) Townsend, D.; Lahankar, S.; Lee, S.; Chambreau, S.; Suits, A.; Zhang, X.; Rheinecker, J.; Harding, L.; Bowman, J., The roaming atom: Straying from the reaction path in formaldehyde decomposition. *Science* **2004**, *306* (5699), 1158-1161; (e) Lahankar, S. A.; Chambreau, S. D.; Townsend, D.; Suits, F.; Farnum, J.; Zhang, X.; Bowman, J. M.; Suits, A. G., The roaming atom pathway in formaldehyde decomposition. *The Journal of chemical physics* **2006**, *125* (4), 044303.
13. Tishchenko, O.; Truhlar, D. G., Efficient global representations of potential energy functions: Trajectory calculations of bimolecular gas-phase reactions by multiconfiguration molecular mechanics. *The Journal of chemical physics* **2009**, *130* (2), 024105.

14. Kim, Y.; Corchado, J. C.; Villa, J.; Xing, J.; Truhlar, D. G., Multiconfiguration molecular mechanics algorithm for potential energy surfaces of chemical reactions. *The Journal of Chemical Physics* **2000**, *112* (6), 2718-2735.
15. (a) Glowacki, D. R.; Orr-Ewing, A. J.; Harvey, J. N., Product energy deposition of CN+ alkane H abstraction reactions in gas and solution phases. *The Journal of chemical physics* **2011**, *134* (21), 214508; (b) Greaves, S. J.; Rose, R. A.; Oliver, T. A.; Glowacki, D. R.; Ashfold, M. N.; Harvey, J. N.; Clark, I. P.; Greetham, G. M.; Parker, A. W.; Towrie, M., Vibrationally Quantum-State-Specific Reaction Dynamics of H Atom Abstraction by CN Radical in Solution. *Science* **2011**, *331* (6023), 1423-1426.
16. Hong, G.; Rosta, E.; Warshel, A., Using the constrained DFT approach in generating diabatic surfaces and off diagonal empirical valence bond terms for modeling reactions in condensed phases. *The Journal of Physical Chemistry B* **2006**, *110* (39), 19570-19574.
17. Czakó, G.; Bowman, J. M., Dynamics of the reaction of methane with chlorine atom on an accurate potential energy surface. *Science* **2011**, *334* (6054), 343-346.
18. Hornung, B.; Harvey, J. N.; Preston, T. J.; Dunning, G. T.; Orr-Ewing, A. J., Empirical Valence Bond Theory Studies of the $\text{CH}_4 + \text{Cl} \rightarrow \text{CH}_3 + \text{HCl}$ Reaction. *submitted*.
19. Lourderaj, U.; Sun, R.; Kohale, S. C.; Barnes, G. L.; de Jong, W. A.; Windus, T. L.; Hase, W. L., The VENUS/NWChem software package. Tight coupling between chemical dynamics simulations and electronic structure theory. *Computer Physics Communications* **2014**, *185* (3), 1074-1080.
20. Glowacki, D. R.; Rose, R. A.; Greaves, S. J.; Orr-Ewing, A. J.; Harvey, J. N., Ultrafast energy flow in the wake of solution-phase bimolecular reactions. *Nature Chemistry* **2011**, *3* (11), 850-855.
21. Rose, R.; Greaves, S.; Abou-Chahine, F.; Glowacki, D.; Oliver, T.; Ashfold, M.; Clark, I.; Greetham, G.; Towrie, M.; Orr-Ewing, A., Reaction dynamics of CN radicals with tetrahydrofuran in liquid solutions. *Physical Chemistry Chemical Physics* **2012**, *14* (30), 10424-10437.
22. Dunning, G.; Glowacki, D.; Preston, T.; Greaves, S.; Greetham, G.; Clark, I.; Towrie, M.; Harvey, J.; Orr-Ewing, A., Vibrational relaxation and microsolvation of DF after F-atom reactions in polar solvents. *Science* **2015**, *347* (6221), 530-533.
23. Glowacki, D. R.; Orr-Ewing, A. J.; Harvey, J. N., Non-equilibrium reaction and relaxation dynamics in a strongly interacting explicit solvent: F + CD₃CN treated with a parallel multi-state EVB model. *Journal of Chemical Physics* **2015**, [/http://dx.doi.org/10.1063/1.4926996](http://dx.doi.org/10.1063/1.4926996).
24. Rehbein, J.; Carpenter, B. K., Do we fully understand what controls chemical selectivity? *Physical Chemistry Chemical Physics* **2011**, *13* (47), 20906-20922.
25. Quapp, W.; Hirsch, M.; Heidrich, D., Bifurcation of reaction pathways: the set of valley ridge inflection points of a simple three-dimensional potential energy surface. *Theoretical Chemistry Accounts* **1998**, *100* (5-6), 285-299.
26. Collins, P.; Carpenter, B. K.; Ezra, G. S.; Wiggins, S., Nonstatistical dynamics on potentials exhibiting reaction path bifurcations and valley-ridge inflection points. *The Journal of chemical physics* **2013**, *139* (15), 154108.
27. Carpenter, B. K., Effect of a chiral electrostatic cavity on product selection in a reaction with a bifurcating reaction path. *Theoretical Chemistry Accounts* **2014**, *133* (8), 1-8.

28. Carpenter, B. K.; Harvey, J. N.; Glowacki, D. R., Prediction of enhanced solvent-induced enantioselectivity for a ring opening with a bifurcating reaction path. *Physical Chemistry Chemical Physics* **2015**, *17* (13), 8372-8381.
29. Glowacki, D. R.; O'Connor, M.; Calabró, G.; Price, J.; Tew, P.; Mitchell, T.; Hyde, J.; Tew, D. P.; Coughtrie, D. J.; McIntosh-Smith, S., A GPU-accelerated immersive audio-visual framework for interaction with molecular dynamics using consumer depth sensors. *Faraday discussions* **2014**.