Development of Metrics for Predicting the Onset of Reactions in Temporal QM/MM Simulations

Supervisor: Dr. Nicholas Mosey
Examiner: Dr. Tucker Carrington
Department of Chemistry
Queen's University

5/4/2010

I do hereby verify that this written report is my individual work and contains my own original ideas, concepts and designs. No portion of this report has been copied in whole or in part from another source, with the possible exception of properly referenced material.

Faisal Abdulla – 5445670

Abstract

Molecular dynamics (MD) is a powerful tool for studying the atomic-level details of chemical processes. To model reactions in MD simulations, it is necessary to use quantum chemical methods, which significantly limit the time scales accessible in these simulations. One way to extend the time scales accessible in MD simulations is to switch between force-fields and quantum chemical methods as needed to describe reactions. In this work, several techniques for determining when to switch between force-fields and quantum chemical methods are developed, implemented and tested. The results shed light on potential useful techniques for identifying when switches should occur and highlight the features of effective metrics for this purpose.

Table of Contents

Table of Figures

1.0 Introduction

1.1 Motivation

With computer processors becoming more powerful every day, research and industry have turned to computer simulations to model chemical processes and reactions. One of the most common simulation methods used is known as molecular dynamics (MD, see section 1.2). Briefly, an MD simulation aims to study the atomic-level motion of a chemical system by treating the nuclei as classical particles that can be propagated according to Newton's equations. To do this, it is necessary to calculate the forces acting on the atoms (see section 1.2). Typically, this is achieved through either force field (FF) or quantum mechanical calculations (QM). FF methods use empirical models to describe the interactions between atoms in the system. The simple functional forms used in FFs facilitate rapid calculations, which allow MD simulations to simulate microsecond time scales. Unfortunately, FFs are generally limited in their ability to describe the formation and dissociation of bonds, and are thus unable to model chemical reactions. QM methods, on the other hand, explicitly consider the electronic structure of the system, and can thus describe the changes in bonding that occur during reactions. Unfortunately, solving for the electronic structure is computationally demanding, and QM-based MD simulations are limited to exploring sub-nanosecond time scales, which are far too short to observe most chemical reactions.

The time scale limitations inherent to QM-based MD simulations are usually overcome in either of two ways. The first involves the use of techniques that either apply geometric constraints or modify the potential energy surface such that the system preferentially samples reactive regions of phase space[1-8]. These methods are well-established and effective, yet require prior knowledge of the reaction one would like to observe. The second approach involves the use of functional forms in FFs that are able to describe changes in bonding[9-12]. These reactive FFs have drawbacks in the sense that they require a significant degree of parameterization and are not generally transferable across large classes of molecular systems.

Recently, an alternative strategy in which both FF and QM methods are used in MD simulations has been proposed as a means of extending the time scales accessible in MD simulations of chemical reactions[13]. This method is based on the notion that QM methods are only needed to describe the changes in bonding that occur when a system undergoes a reaction, while an FF is suitable otherwise. Specifically, the MD simulation is performed using a FF while monitoring the behaviour of the system for indications that a reaction may occur. When a possible reaction is detected, the simulation switches to a model based on QM methods, and then switches back to FFs when the reaction is complete. Since reactions are rare events in MD simulations, and proceed quickly when they do occur, QM methods are actually used sparingly in this technique. However, the targeted use of QM methods during reactions ensures that changes in bonding are described properly. Recent applications of this method have demonstrated that it can extend the time scales accessible in MD simulations by several orders of magnitude, while still allowing reactions to be observed. This method is analogous to a simulation strategy called QM/MM in which QM and FF (MM) methods are applied to different regions of a molecular system to examine large lengths scales. In the method described here, the use of QM and FF methods varies temporally instead of spatially, and hence this technique will be called temporal QM/MM. This method is discussed further in section 1.3.

1.2 MD, FF, and QM Calculations

In an MD simulation, Newton's second law is applied repeatedly to calculate accelerations of nuclei in a structure. The expression in the context of MD simulations is shown below:

$$\mathbf{F}_I = m_I \mathbf{a}_I = -\frac{\partial U}{\partial \mathbf{r}_I} \tag{1}$$

where $\mathbf{F}_I$ is the force, m_I is the mass, $\mathbf{a}_I$ is the acceleration, U is the potential energy, $\mathbf{r}_I$ is the position, all of an atom indexed I.

Once these accelerations are calculated, an algorithm called the "Verlet algorithm" is applied in order to calculate new nuclei positions and velocities. This cycle is repeated throughout the course of the simulation. The equations used for the Verlet algorithm are shown below:

$$\mathbf{r}_I(t+\Delta t) = \mathbf{r}_I(t) + \mathbf{v}_I(t)\Delta t + \frac{\mathbf{a}_I(t)}{2}\Delta t^2 \tag{2}$$

$$\mathbf{v}_I(t) = \frac{\mathbf{r}_I(t+\Delta t) - \mathbf{r}_I(t-\Delta t)}{2\Delta t} \tag{3}$$

where $\mathbf{v}_I$ is the velocity, t is time, Δt is the time step, and all other variables having the same definitions for an atom indexed I. The times steps used in these simulations must be sufficiently short to capture the fastest processes in the system. These are typically bond vibrations with periods of ~10^{-14} s. As such, the time steps are on the order of 10^{-15} s. In other words, with Δt = 10^{-15} s, it would require 1 million MD time steps, with a full potential energy and force evaluation at each step, to simulate 1 nanosecond of time.

These MD simulations are based on the potential energy function that governs the interactions between nuclei. These potential energy functions can be calculated in two different ways: using simple FF or using QM calculations. A detailed discussion of FF and QM methods is beyond the scope of this thesis. Briefly, FF calculations typically use a simple harmonic potential to describe the potential energy function. This approach has the advantage of being easy to treat computationally, thus minimizing the computational effort associated with each MD time step. Unfortunately, the use of harmonic potentials to describe bonds prevents a description of bond dissociation because the energy becomes infinite as the bond is stretched. As such, FFs are suited to efficiently describing the behaviour of chemical systems near potential energy minima, where the harmonic approximation is valid, but are not suited to describing reactions.

On the other hand, QM calculations provide approximate solutions to the electronic Schrodinger equation (or equivalent) and can thus accurately describe the changes in electronic structure that occur during changes in bonding. As such, these methods are better suited to describing reactions that FF approaches, which is evident through the comparison of potential energy curves for bond stretching obtained with QM and FF methods provided in Figure 9.1. Unfortunately, QM methods are computationally demanding and their use limits MD simulations to studying sub-nanosecond time scales.

1.3 Temporal QM/MM Calculations

Temporal QM/MM calculations[13] attempt to lengthen feasible simulation times by switching between FF and QM calculations throughout the course of an MD simulation as needed to describe chemical reactions. Specifically, the MD simulation is performed using FF potential models and the behaviour of the system is monitored for the onset of reactions. Should a reaction be predicted, the MD simulations switches to using a QM potential to describe the reactive processes, and then returns to the more efficient FF method once the reaction is complete. By relying primarily on FFs, temporal QM/MM

simulations can study much longer time scales than those accessible using QM-based MD simulations. Meanwhile, the targeted use of QM methods ensures that reactive events are described properly.

The critical aspects of a temporal QM/MM is determining when to switch between FF and QC methods and ensure switching between these two potentials does not significantly alter the trajectory. Determining when to switch is difficult because reactions do not occur during FF-based MD simulations. As such, in temporal QM/MM, one identifies events which should have led to a reaction if the correct potential was used to determine when to switch between FF and QC methods. When such an event is detected, the system is reverted to a slightly earlier point along the trajectory and that portion of the simulation is rerun while switching to QM methods to allow the reaction to occur. After a brief period of time the simulation returns to using FF methods. A time-dependent switching function[14-15] is used during the switching period to ensure continuity of the forces acting on the atoms, which is sufficient for the system to exhibit meaningful behaviour[16-18]. This switching function smoothly takes the system from moving on a purely FF potential to a purely QM potential over a brief period of time. A schematic illustrating the series of events that occurs during a temporal QM/MM calculation are outlined in Figure 9.2.

1.4 Project Goals

In order for TQMMM calculations to be effective, it is necessary to have metrics that accurately detect the onset of reactive events during FF-based MD simulations. Ideally, one would like to have a metric that identifies all reactive events that should occur without detecting any additional processes during the FF-based MD simulations that do not lead to reactions upon further study with QM methods. This is likely an unattainable goal; however, one can endeavour to develop a form of metric that accurately detects all events that correspond to reactions (needed to ensure accuracy), while minimizing the number of events detected that do not lead to reactions (which are still investigated with QM methods, thus increasing computational cost).

The goal of this project is to design, implement, and test a number of different models, and then evaluate them on their successes and failures. This will be achieved through an iterative process in which a metric is developed to meet certain criteria, studied to determine its performance, and analyzed to identify the origins of any deficiencies. This information is then used to develop a new model and these steps are repeated.

In order to test these models, they were implemented into simulation software and temporal QM/MM simulations were performed. Details are discussed in section 3. Performing these tests highlighted deficiencies within existing metrics and allowed for the development of new ones. The metrics explored during this project are reported in sections 2.3 through 6.0. Conclusions and opportunities for future work are presented in section 7.0.

2.0 Implementation and Simulation Details

2.1 Computation and Programming

The temporal QM/MM calculations are performed using a combination of the Tinker[19] and NWChem[20] software packages that have been modified and interfaced by the Mosey research group. In this code, Tinker performs all dynamics, detects possible reactive events, and controls the switching between FF and QM methods. NWChem only provides the QM potential energies and forces. As such, the implementation of all new metrics involved making modifications to the Tinker source codes, which are written using FORTRAN.

In addition to changes directly into the Tinker software, a number of Perl scripts were also written in order to speed up repeated processes such as preparing each event sample for quantum calculations (submitevents.pl, see Appendix A) and identifying which potential reactive events needed to be explored with QM calculations (filterevents.pl, see Appendix B).

2.2 Procedures Involved

A good metric for detecting reactions should have a number of properties. First, it should detect all reactive events that would be observed if the entire temporal QM/MM simulation was performed using only QM methods. Second, it should not identify too many events that look like they may be reactions, but that do not actually make it to the products after the switch to QM methods is made. Third, the metric should be able to be calculated quickly and use information that is available during simulations performed using only an FF.

To test the various metrics considered in this work, temporal QM/MM simulations were performed on 1-CN-semibullvalene, which can undergo a Cope rearrangement. This particular reaction is outlined in Figure 9.3 along with relevant geometric details, and was chosen for study because it is reversible and has a low barrier (6.3 kcal/mol at PBE/6-31G(d,p) level of theory), which will allow several reactive events to occur on short time scales, providing a lot of data for testing the temporal QM/MM method and the metric developed in this study. The simulations of this system were performed using MM3[21-23] as an FF and the PBE/6-31G(d,p)[24] level of theory for the QM calculations.

The test simulations were performed as follows. First, a 1 ns MD simulation was performed using only an FF while monitoring for potential reactive events using the developed metrics. Whenever a possible reactive event was detected, the data needed to restart the simulation from that point were saved and the FF-based simulation was continued while looking for the next possible reaction. Once this stage of the test was completed, MD simulations were started from each of the saved data points while switching to QM methods to allow for possible reactions to occur. The simulations were each performed for 0.1 ps, which is a sufficient amount of time for the reaction to occur if successfully detected. At the end of each of these simulations, the structure was analyzed to determine whether the system ended up on the product state or returned to the reactants. The numbers of successfully detected events and erroneously detected events were then determined to assess the abilities of the metric employed. To facilitate a comparison between different metrics, the initial 1 ns FF-based MD simulations performed for all metrics tested were given exactly the same starting conditions.

2.3 Defining the Parameters of the Models

All the models that were tested were based on the structure of the system. Using preliminary optimization calculations run elsewhere, the optimal structure for the reactant and transition state were know. From this information, the change in each of the bond distances, $d_{1\text{-}2}$ and $d_{10\text{-}11}$ on Figure 9.3 could be determined. Using this data alone, the models that were implemented were supposed to predict the onset of reactions based on current bond distances as calculated through MD simulations.

3.0 The Linear OP Model

3.1 Mathematics and Physical Significance

This model was the starting basis for this project and was developed by Mosey in 2009[13]. The order parameter is based on the structure of the system.

The first step in the calculation is the generation of a vector that corresponds to the change in bond distances being monitored ($d_{1\text{-}2}$ and $d_{10\text{-}11}$on Figure 9.3) as the system moves from reactant to the transition state. This vector is as follows:

$$\mathrm{v}_{ideal} = \left[\left(d_{1-2}^{\mathrm{TS}} - d_{1-2}^{\mathrm{reactant}} \right), \left(d_{10-11}^{\mathrm{TS}} - d_{10-11}^{\mathrm{reactant}} \right) \right] \quad (4)$$

where "TS" and "reactant" superscripts correspond to the transition state and reactant structures respectively. Note that this vector is constant throughout the simulation since these values have been pre-calculated from optimized structures.

At every step of the MD simulation, another vector $v_{current}$ is then calculated which corresponds to the current values of the distances. This vector is as follows:

$$\mathrm{v}_{current} = \left[\left(d_{1-2}^{\mathrm{current}} - d_{1-2}^{\mathrm{reactant}} \right), \left(d_{10-11}^{\mathrm{current}} - d_{10-11}^{\mathrm{reactant}} \right) \right] \quad (5)$$

where the "current" superscript corresponds to the current structure and all other variables have their standard definitions.

The order parameter is then defined as the ratio of dot products as follows:

$$Order\ Parameter = \frac{\mathrm{v}_{current} \bullet \mathrm{v}_{ideal}}{\mathrm{v}_{ideal} \bullet \mathrm{v}_{ideal}} \quad (6)$$

This results in an order parameter that changes smoothly from 0 to 1 as the system moves from the reactant's structure to the transition state's structure. Figure 9.4 displays this behaviour for an arbitrary structure. This metric could then be used to determine if the system is moving towards the transition state and how close the system is to the transition state.

In an FF-based MD simulation, the system is never expected to reach the transition state (due to the limitations imposed by the FF-based harmonic potential), and the order parameter could never be 1. Instead, to detect reactions in temporal QM/MM simulations, a cutoff for the order parameter is defined and if the order parameter exceeds this cutoff, a reaction is predicted. At this point, the simulation would be stopped and restarted while switching to QM methods. If lower values for the cutoff are chosen, more reactions would be predicted, ensuring reactions are observed. However, with more predicted reactions also comes more erroneously detect events that do not lead to reactions leading to wasted time using QM calculations.

3.2 Results

Figure 9.5 shows how many total events were detected and how many of those events led to a reaction. A key point to notice is that for lower cutoff values, more events were detected. This is simply because a less strict condition was imposed on the program, and hence a larger portion of the molecular geometries were able to meet these conditions.

The results show that the number of detected events increases exponentially as the OP cutoff is decreased, which is expected. The number of actual events converges for a cutoff of less than 0.16, at about 30 events. On the basis of the values obtained with the OP cutoff set at 0.16, one predicts a ratio of 25 successful events/500 detected possible events, or a success rate of 5%. Of course, the percentage of correctly identified events is higher with higher cutoffs, but some events are missed, which is undesirable.

For approx. 80 total detections (which will be used as a baseline for comparing future models), the model's accuracy lies at 10%. That is, a total of eight of those 80 events led to reactions. This was taking into consideration that it was important to detect most of the reactions in the simulation. Without this constraint, a cutoff of 0.29 could have been picked, which led to 100% of detected events leading to reactions. However, only one event was detected, making this cutoff non-feasible for use.

3.3 Model Limitations

The results in the previous section indicate that the linear OP can successfully detect reactive events, but also detects many events that do not lead to reactions but still waste resources by using QM methods. To develop an improved OP, it is necessary to determine the limitations of this one. To do this, a closer examination of events that did not lead to reactions was required. One of the events that was examined during the evaluation stage of this model showed that bond $D_{10\text{-}11}$ underwent a large drop in length. This overwhelmed the order parameter, causing it to exceed the cutoff value, and hence, a reaction was predicted, despite the small change in the $d_{1\text{-}2}$ bond length. Figure 9.6 displays this result for the event that was analyzed. Overall, this indicates that the linear OP model is not sufficiently selective for the concerted motion of the bonds.

3.4 Improving the Linear Model

In order to ensure that large changes of bond length of a single bond did not overwhelm the order parameter, a new, more complex model was required. A number of possible candidates for the new model were considered. This included arctan functions, logarithmic functions and Gaussian curves.

All the models mentioned above are capable of dealing with large changes in their independent variables without large changes in their dependant variables. In a mathematical sense, for y=f(x), where x is an interatomic distance, being the model being implemented, the limit of f'(x) (the derivative) for large values of x tends to 0. Notice how parabolic functions were not included in the above list. This is because their derivates do not follow the above limit. While the possible bond distances are physically bound in FF calculations through the harmonic potential approaching infinity, they are not in a purely mathematical sense, and hence, this model type was avoided.

The final model that was chosen was the Gaussian model. The main reason this model was picked was because of the ease of modifying the curve's shape, such as the width and steepness.

4.0 The Gaussian OP Model

4.1 Mathematics and Physical Significance

The new model first calculates an expected bond distance, which corresponds to a distance along the reaction coordinate to the transition state. It is calculated as follows:

$$d_i^{\text{expected}} = d_i^{\text{reactant}} + C\left(d_i^{\text{TS}} - d_i^{\text{reactant}}\right) \tag{7}$$

One this value has been calculated for each bond distance being examined, the order parameter is calculated as follows:

$$Order\ Parameter = e^{-A_{1-2}\left(d_{1-2}^{\text{current}} - d_{1-2}^{\text{expected}}\right)^2} + e^{-A_{10-11}\left(d_{10-11}^{\text{current}} - d_{10-11}^{\text{expected}}\right)^2} \tag{8}$$

where A is a constant that determines how wide the Gaussian curve is for each bond examined and C is a fraction of bond length change to transition state and all other variables have their standard definitions.

In a physical sense, the model treats each bond being examined separately. It also gives each bond a limit of contribution to the total value of the parameter which means that should a bond exceed its expected bond distance to the transition state, the parameter would not be affected. Furthermore, the metric is only large when *both* distances have reached their ideal values.

This model is slightly more complex than the previous one, requiring a total of four parameters (an A-parameter and a C-parameter for each bond). However, to simplify this model, A-parameters were set equal, and likewise for C-parameters. This corresponds to treating the bond formation and dissociation associated with the Cope rearrangement as a concerted processes, which is reasonable. The values of the parameters were set at A = 100 and C = 0.3. These values were selected to ensure that all events detected with the linear model would also be detected with this model.

It is important to note that highest value this parameter can take is 2.0 at the transition state, while the linear model takes a value of 1.0. This parameter could have been normalized (but this was not done) by dividing by the total number of bonds being examined.

4.2 Results

The event that was run through in Figure 9.6 was rerun through the Gaussian model to ensure that the new model developed was not sensitive to large changes in bond length. Figure 9.7 shows how the new model fared against the event that caused problems with the linear model. Specifically, the linear OP originally detected an event at frame 56, while this is actually a minimum value for the Gaussian OP and well beyond the cutoff that was assigned (cutoffs of 1.3 or greater were used in testing the new model). As such, it seems that the Gaussian OP is performing as designed in the sense that a large change in an individual bond does not overwhelm the OP.

In order to compare this new model to the linear model, different order parameter cutoffs were chosen, and the ratio between events detected and events reacted were examined. Figure 3.2 shows the results of this model. Using the same baseline as was used for the linear model (80 predicted reactions), 16 of them had led to reactions. This model's accuracy lies at roughly 20%, a two-fold increase over the initial linear model's accuracy.

4.3 Model Limitations

The Gaussian model, while being able to take into account concertedness of a reaction (or lack thereof), does not take into account any elements of time. That is, it does not consider whether or not atoms are moving towards or away from the transition state. For example, the trajectories as calculated by FF would show the program that the molecule's geometry lies relatively close to the transition state. However, the dynamics would show that the molecule is moving from the transition state to the reactant state. An example of this is shown on Figure 9.9. Intuitively, one would assume that this would not lead to a reaction, whereas the program would predict that a reaction would occur. Events that are detected under such circumstances will lead to the unnecessary use of QM methods, and thus it is desirable to avoid this behaviour.

In addition to direction of the system along the reaction coordinate, it is also important to considering how well aligned the motion is with the reaction coordinate. While the point at which these events were detected is not important, it is important to note that the event that led to a reaction had a section of its trajectory that lay close to the trajectory that would be followed if the reaction followed an exactly concerted pathway. This behaviour was not observed for the trajectory that did not lead to a reaction. As such, the Gaussian model needs to be improved to consider the progression of the system with respect to the reaction coordinate. This was a key element that needed to be exploited in a new model that was developed.

4.4 Improving the Gaussian Model

In order to improve the Gaussian model, the new model would have to take into account a number of factors to ensure that the two events shown in Figure 9.3 can be distinguished. These parameters are as follows:

i. The trajectory is moving towards the transition state
ii. The trajectory lies within a certain distance from that of a fully concerted reaction
iii. The trajectory is satisfying the above conditions for a certain amount of time

A proposed model that was expected to satisfy all the above requirements was based on the vector form of a parametric equation. Details are explained in the next section.

5.0 The Vector OP Model

5.1 Mathematics and Physical Significance

The new model was called a vector model because it coupled the distances between the two bonds being monitored in this reaction into a single vector. This vector was then parameterized based on initial bond lengths (a vector which represented the bond lengths of the reactant) and change of bond lengths to the transition state (also a vector).

The result was a vector model where the distance of the molecules path along the trajectory to the transition state could be monitored by a single variable, t.

The model takes the following form (note that it has been written in vector notation):

$$\left\langle d_{1-2}^{\text{current}}, d_{10-11}^{\text{current}} \right\rangle = \left\langle d_{1-2}^{\text{reactant}}, d_{10-11}^{\text{reactant}} \right\rangle + t \left\langle \Delta d_{1-2}, \Delta d_{10-11} \right\rangle \quad (9)$$

where t is the parametrized variable, Δd is the change in bond length from the reactant to the transition state and all other variables have their standard definitions.

With this new model came a new algorithm for determining whether a reaction would occur. All the previous models were based on a single parameter. Should that parameter be exceeded, a reaction would be predicted. However, for this new model, a more complex way of determining whether a reaction would occur had to be implemented. The algorithm is as follows:

1. Obtain $d_{1\text{-}2}^{\text{current}}$ from the FF calculations.
2. Assume the reaction is perfectly concerted. Calculated the associated t-value for this bond distance.
3. Calculate the expected $d_{10\text{-}11}$ value associated with this t-value.
4. Obtain $d_{10\text{-}11}^{\text{current}}$ from the FF calculations.
5. Calculate the difference between the calculated $d_{10\text{-}11}$ and the actual $d_{10\text{-}11}$ values.
6. If the difference between these lies below a certain value (this was a parameter called the *fuzz*) and if t has been increasing, increment a counter.
7. If the counter exceeds a certain threshold and if t exceeds a certain threshold, flag a predicted reaction.

Figure 9.10 shows what the vector model looks like, with the fuzz, trajectories and other elements labelled. The code that was implemented for this model can be found in Appendix C. This code was written in FORTRAN.

This new model attempts to take into account that the atoms between the bonds that are being monitored have to not only be at a certain position in space, but have had to be moving in a direction with a certain allowable degree of concertedness for a certain amount of time that would favour a reaction.

5.2 Results

When running this new model with parameters chosen such that all previous events that were known to react would be detected (parameters chosen were 21 time steps, fuzz of 0.143, t cutoff of 0.144), more than 2000 possible reactive events were detected. This meant that the while the model did detect all reactions that previous models detected, it was of such a low accuracy that it could not be compared to the prior two models.

5.3 Model Limitations

Upon examining values that were required of the parameters such that the events that led to reactions would still be detected by the model, it was found that for values of t that were negative (ie. the molecule's position lay pre-reactant), the values of required of the fuzz were higher than the values required when t were positive. Hence, by using a fixed value of the fuzz, these higher values raised the minimum requirements for the fuzz, and hence, allowed more events to satisfy the necessary conditions. Figure 9.11 displays this behaviour. The vertical dashed lines represent the region where the counter had been incrementing (the area where t was increasing). The region of concern is shaded.

5.4 Improving the Vector Model

Two ways of solving this problem could be used. The first would be to reduce the fuzz for larger values of t. However, due to the fact that the fuzz was already small for larger values of t, the other alternative of widening the fuzz for negative values of t was chosen. Figure 9.12 shows what the new varying OP would like like for t values less than 0 (this corresponds to a $d_{1\text{-}2}$ value of less than 1.56). Since implementing a form of step function to the fuzz would make it programmatically more complicated, a gradual change over the fuzz was implemented. The result was a fuzz that looked like a funnel.

6.0 The Funneling OP Model

6.1 Mathematics and Physical Significance

The mathematics behind this model were identical to that of the vector model. The only difference was that after the calculation of the difference between the expected $d_{10\text{-}11}$ and the actual $d_{10\text{-}11}$, the result was multiplied by a factor that was based on a function of t.

The factor that was used in this was as follows:

$$\frac{1}{2\pi}\arctan\left(10t\right)+0.75 \tag{10}$$

where t was the parameter in the original vector model.

The factors of 1/2π, 10t, and 0.75 involved in the prefactor were chosen just like any other number that was picked in the prior two models. That is, it was picked so as to ensure that events that were already detected would still be detected by this model.

The physical significance behind this model is that there are a number of different trajectories that could lead to a reaction starting from pre-reactant molecular geometries (see Figure 9.13). The old vector model did not take this into account, and hence, imposed a too-wide restriction on what the geometries could be during the path to the transition state. So by loosening the parameters around pre-reactant geometries, we are able to impose a larger restriction on geometries that are close to the transition state. A plot of what the new funneling model looks like along with trajectories that successfully lead to reactions is show in Figure 9.14.

6.2 Results

When running the new model, a large number of events were still detected. This time, 1500 events compared to the vector model's 2000 events. Hence, this model is still non-feasible for use. Also, in order to determine what properties that trajectories that led to reactions had that had them being detected by the model, a number of them had been picked from a random sample out of the 1500 events and plotted with those trajectories that led to reactions.

The resulting graph is shown on Figure 9.15. The problem now becomes one of finding a way of separating the red lines (trajectories that didn't react) from the green lines (trajectories that led to a reaction).

6.3 Proposed Solutions

While no work has been put into further refining this funneling model due to time constraints, a few suggestions as to what could be used to distinguish the red and green trajectories are listed below:

1. Notice how most of the red trajectories appear to have their endpoints below the dashed line, very near the fuzz. Perhaps an implementation of a check would be able to remove those events.
2. A further examination on the starting point and ending point of the trajectory, and calculating the slope between those two points may show that the slopes for trajectories that did not react differ from the slopes of those that did react.

The points listed above are merely ideas that could be drawn at first glance, however, they will remain as ideas until further testing can be done.

7.0 Conclusions and Recommendations

In this project, a series of metrics for detecting the onset of reactions in temporal QM/MM simulations were developed, implemented, tested and studied. The results suggest that the Gaussian model is best suited for the role of temporal QM/MM calculations due to its highest detection accuracy of 20%. This was an improvement over the linear model's detection accuracy of 10% for 80 detections. However, the funneling model appears to have all the necessary 'factors' that are assumed to play a role in making a reaction occur. Of course, only a limited set of parameters entering these metrics were investigated in this project, and thus it may be possible to improve on the results reported above.

This project also forms the basis for future work in the development and optimization of metrics for temporal QM/MM. An obvious starting point is to more extensively investigate the parameters that enter these models with a particular emphasis on the Gaussian and funnel models. In addition one could explore the possibility of allowing the parameters to vary during the simulation. Specifically, the parameters could be set to initial values that vary as information is gained during the simulation regarding parameters let to the successful detection of reactive events. Such adaptive models may aid in the selection of parameters without relying on significant foreknowledge of the reaction to be observed.

Overall, these models provide good potentials for increasing time frames of chemical reaction simulations. Development of these models is key to ensuring efficient use of computational power.

8.0 References

1. Fleurat-Lessard, P.; Ziegler, T. *J. Chem. Phys.* **2005**, 123, 084101.
2. Sprik, M.; Ciccotti, G. *J. Chem. Phys.* **1998**, 109, 7737-7744.
3. Carter, E. A.; Ciccotti, G.; Hynes, J. T.; Kapral, R. *Chem. Phys. Lett.* **1989**, 156, 472.
4. Kirkwood, J. G. *J. Chem. Phys.* **1935**, 3, 300.
5. Torrie, G. M.; Valleau, J. P. *J. Comput. Phys.* **1977**, 23, 187.
6. Miron, R. A.; Fichthorn, K. A. *J. Chem. Phys.* **2003**, 119, 6210.
7. Voter, A. F. *J. Chem. Phys.* **1997**, 11.
8. Ianuzzi, M.; Laio, A.; Parrinello, M. *Phys. Rev. Lett.* **2003**, 90, 238302.
9. Brenner, D. W.; Shenderova, O. A.; Harrison, J. A.; Stuart, S. J. *J. Phys.: Condens. Matter* **2002**, 14, 783-802.
10. van Duin, A. C. T.; Dasgupta, S.; Lorant, F.; Goddard III, W. A. *J. Phys. Chem. A* **2001**, 105, 9396-9409.
11. Stuart, S. J.; Tutein, A. B.; Harrison, J. A. *J. Chem. Phys.* **2000**, 112, 6472-6486.
12. Brenner, D. W. *Phys. Rev. B* **1990**, 42, 9458-9471.
13. Weyand, S.A.; McNeish, J. and Mosey, N.J. manuscript submitted for publication
14. Danielsson, J.; Meuwly, M. *J. Chem. Theory Comput.* **2008**, 4, 1083.
15. Nutt, D. R.; Meuwly, M. *Biophys. J.* **2006**, 91, 1191.
16. Bulo, R. E.; Ensing, B.; Sikkema, J.; Visscher, L. *J. Chem. Theory Comput.* **2009**, 5, 2212-2221.
17. Kim, D.; Schatz, G. C. *J. Phys. Chem. A* **2007**, 111, 5019-5031.
18. Salazar, M. R. *J. Phys. Chem. A* **2005**, 109, 11515-11520.
19. Ren, P.; Ponder, J. W. *J. Phys. Chem. B* **2003**, 107, 5933.
20. Bylaska, E. J.; de Jong, W. A.; Kowalski, K.; Straatsma, T. P.; Valiev, M.; Wang, D.; Apra, E.; Windus, T. L.; Hirata, S.; Hackler, M. T.; Zhao, Y.; Fan, P. -.; Harrison, R. J.; Dupuis, M.; Smith, D. M. A.; Nieplocha, J.; Tipparaju, V.; Krishnan, M.; Auer, A. A.; Nooijen, M.; Brown, E.; Cisneros, G.; Fann, G. I.; Fruchtl, H.; Garza, J.; Hirao, K.; Kendall, R.; Nichols, J.; Tsemekhman, K.; Wolinski, K.; Anchell, J.; Bernholdt, D.; Borowski, P.; Clark, T.; Clerc, D.; Dachsel, H.; Deegan, M.; Dyall, K.; Elwood, D.; Glendening, E.; Gutowski, M.; Hess, A.; Jaffe, J.; Johnson, B.; Ju, J.; Kobayashi, R.; Kutteh, R.; Lin, Z.; Littlefield, R.; Long, X.; Meng, B.; Nakajima, T.; Niu, S.; Rosing, M.; Sandrone, G.; Stave, M.; Taylor, H.; Thomas, G.; van Lenthe, J.; Wong, A.; Zhang, Z. **2006**.
21. Allinger, N. L.; Yuh, Y. H.; Lii, J. -. *J. Am. Chem. Soc.* **1989**, 8551.
22. Lii, J. -.; Allinger, N. L. *J. Am. Chem. Soc.* **1989**, 111, 8566.
23. Lii, J. -.; Allinger, N. L. *J. Am. Chem. Soc.* **1989**, 111, 8576.
24. Perdew, J. P.; Burke, K.; Ernzerhof, M. *Phys. Rev. Lett.* **1996**, 77, 3865-3868.

9.0 Figures

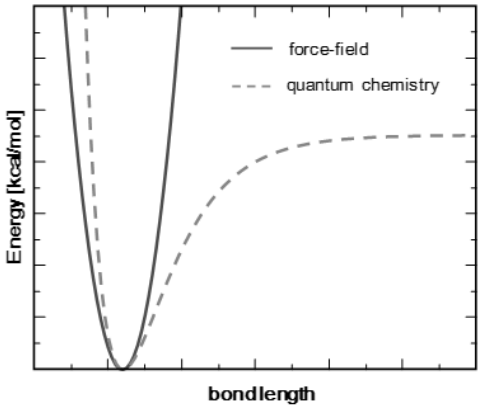

Figure 9.1: Comparison of potential energy curves between FF and QM. From the above, it can be seen that as the bond begins to stretch, FF calculations would predict that the molecule's energy tends to infinity. However, this does not accurately reflect the behavior of a molecule. QM calculations predict that the energy would increase up to a certain point, where it would remain constant.

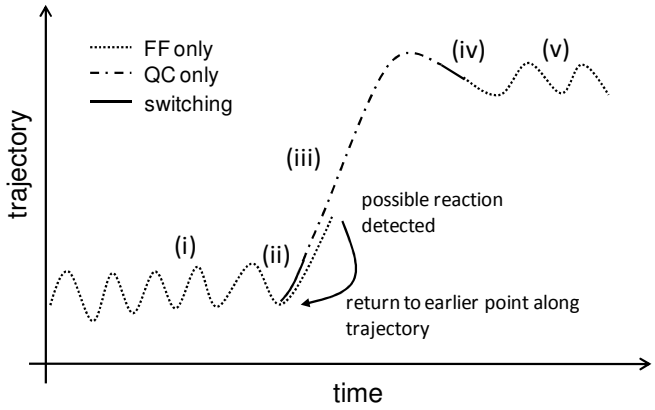

Figure 9.2: The series of events in a temporal QM/MM calculation. i) Perform a normal FF-based MD simulation and monitor for the onset of a reaction. ii) Since a reaction is detected, revert back to a previous structure. Also, enable a switching function to allow the switch from FF to QM-based models. iii) Perform MD simulation using purely QM-based potentials. iv) Check to ensure reaction has finished. Switch back to gradually to using pure FF-based potentials. v) Continue with the rest of the simulation using purely FF-based potentials.

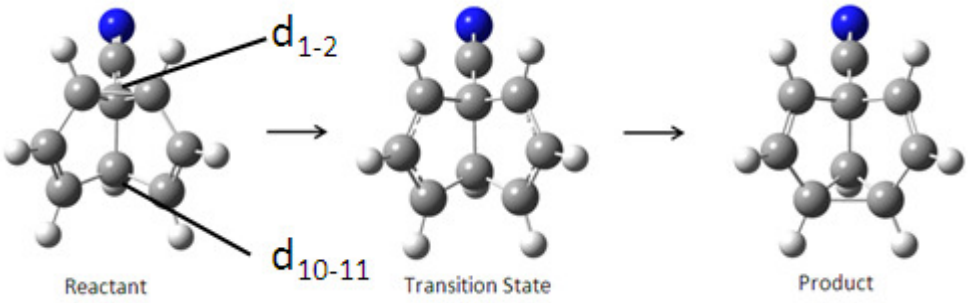

Figure 9.3: Cope rearrangement of semi-bullvalene. The images show the reactant, transition state and product structures of the molecule being examined. The two bonds that were being monitored by the metrics are also labeled.

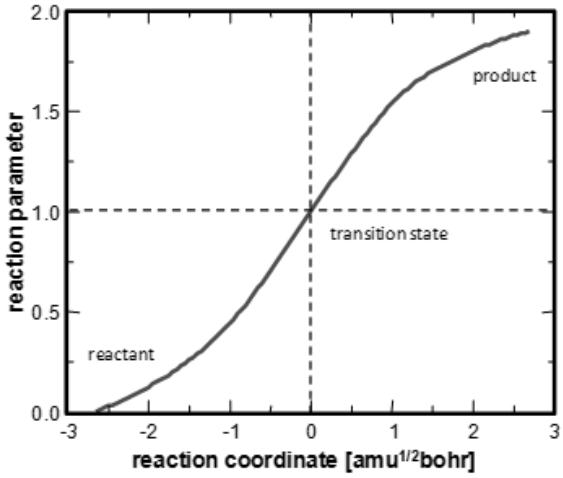

Figure 9.4: Variation of reaction parameter with reaction coordinate. The reaction parameter varies smoothly as a function of the reaction coordinate (which is based on the bond distances being measured). The parameter takes a value of 1 at the transition state structure.

Figure 9.5: Comparison of reactions detected against total events for the linear OP model. With lower cutoff values for the order parameter, more reaction predictions were made since the parameter allowed for structures further away from the transition state to be detected. However, the increase in the number of the events that did lead to reactions was relatively small.

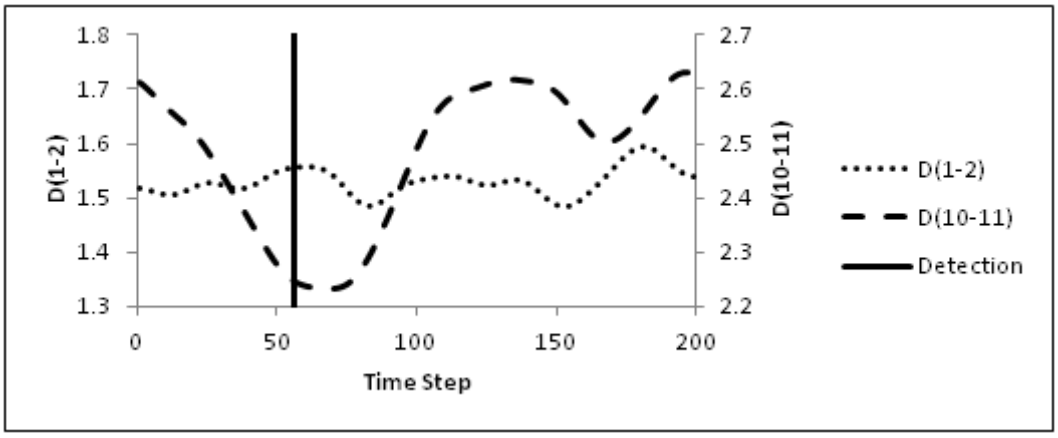

Figure 9.6: Change in the $d_{1\text{-}2}$ and $d_{10\text{-}11}$ bond lengths that led to the detection to a reaction. The vertical line indicates the point where the linear model predicts a reaction. This is caused by the extremely large change in the $d_{10\text{-}11}$ bond distance, which overwhelms the order parameter.

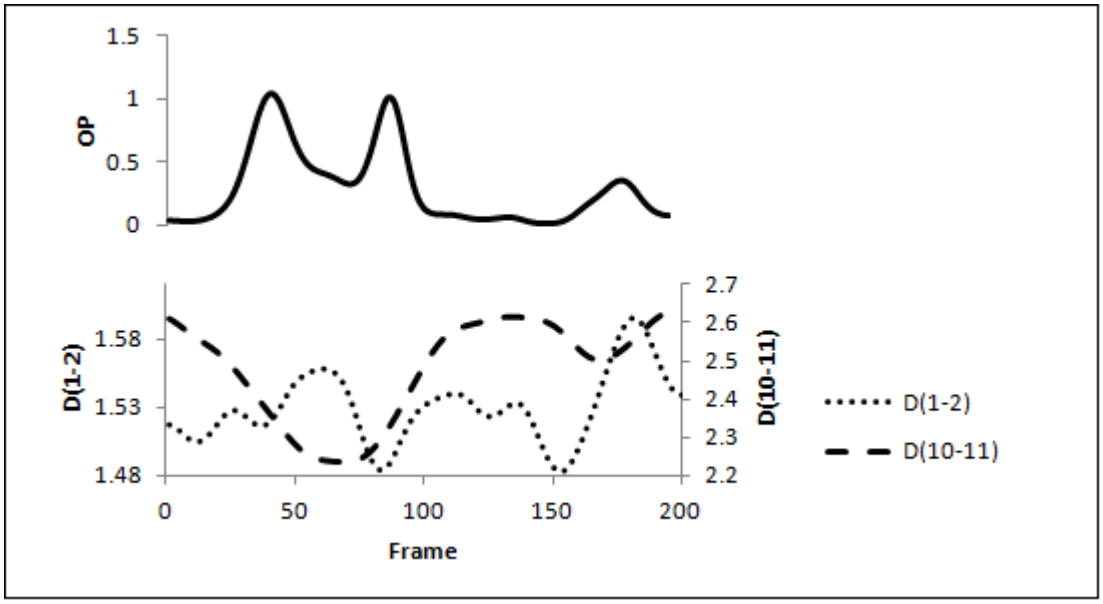

Figure 9.7: Behaviour of the new Gaussian model with large changes in bond length. While the value of d10-11 dropped significantly (and caused the initial linear model to detect a reaction), this was not the case for the Gaussian model that was implemented. The Gaussian model requires both bonds to be at their optimum values in order for the parameter to increase to a level that causes a reaction to be detected.

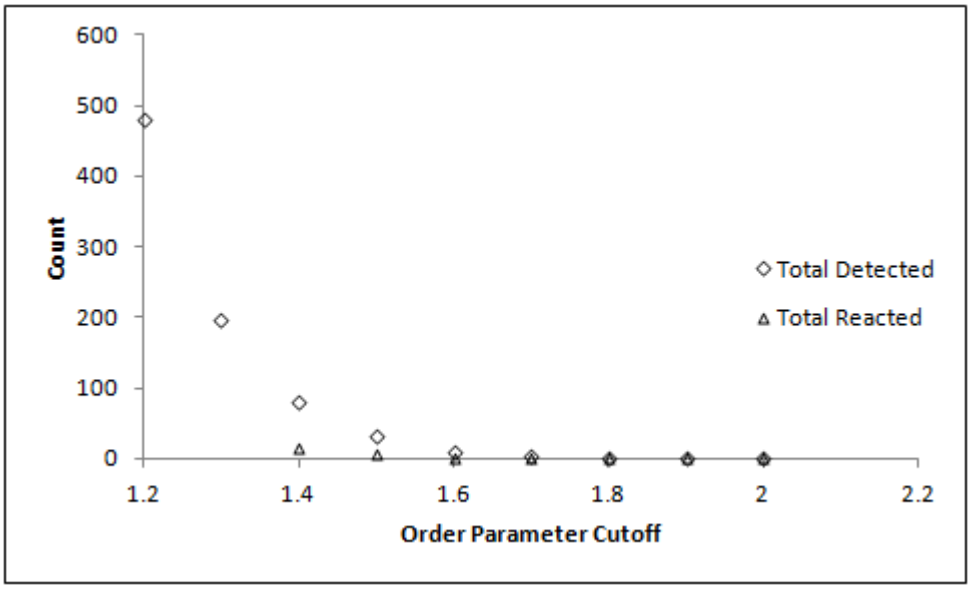

Figure 9.8: Comparison of reactions detected against total events for the Gaussian OP model. Like the previous linear model, more detections were made with lower cutoffs. However, more detections successfully led to reactions when a total of 80 predicted detections were made. About 16 detections were made with 79 predictions, leading to an accuracy of ~20%.

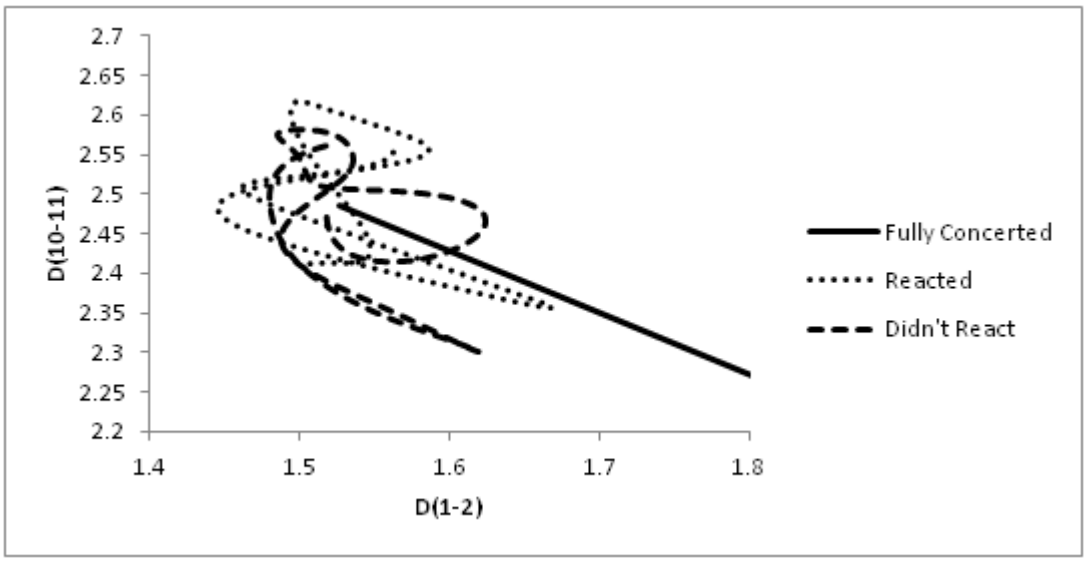

Figure 9.9: Trajectories for two events. One event reacted, whereas the other did not react. Trajectories that led to a reaction appeared to not only be moving towards the transition state, but stayed within a certain distance from the expected distances if the reaction was fully concerted.

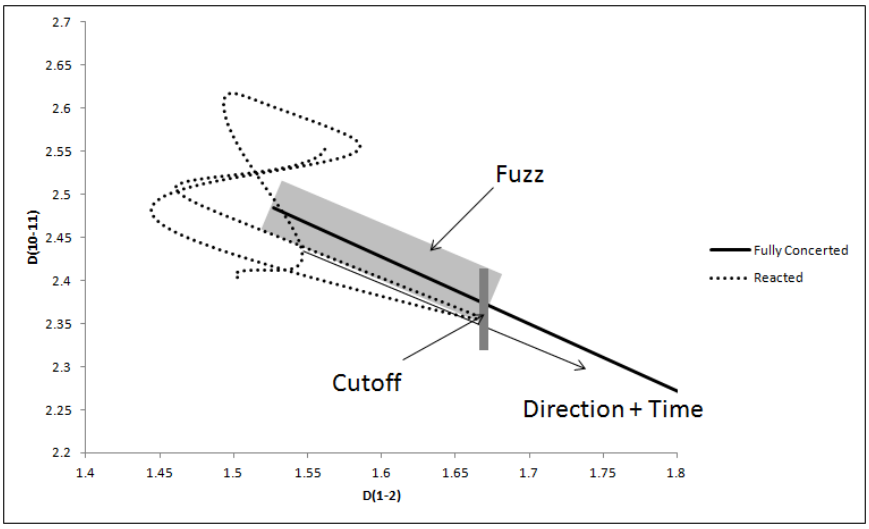

Figure 9.10: Visual representation of the vector model. A number of extra parameters were involved in this model. The trajectories had to be moving within the fuzz for a certain amount of time steps and had to reach the required cutoff in order for the reaction to be predicted.

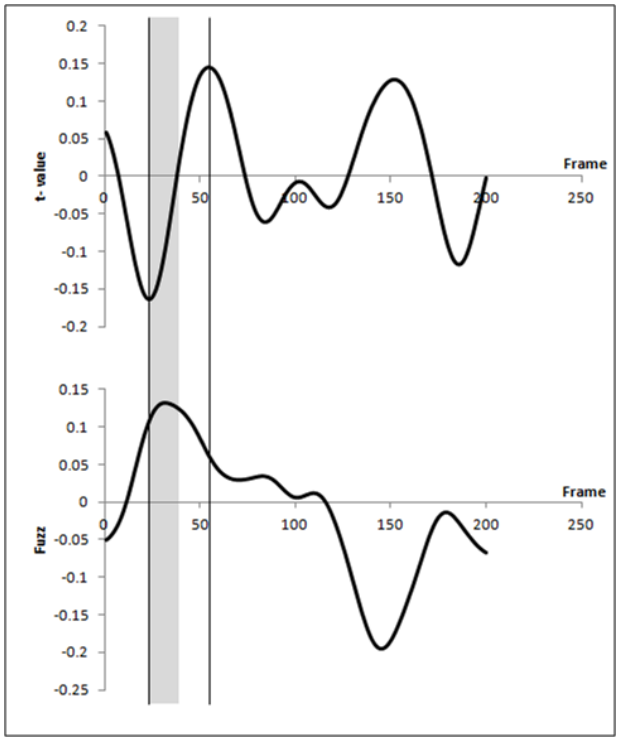

Figure 9.11: Variation of fuzz with changing values of t for an event. The region between the vertical lines indicates that where the t-value had been increasing. The shaded region indicates where the t-value was negative. Notice how the values for the fuzz were a lot higher when t-value was negative. This meant that the minimum fuzz had to be increased in order to accommodate this, hence, reducing the strictness level of the model.

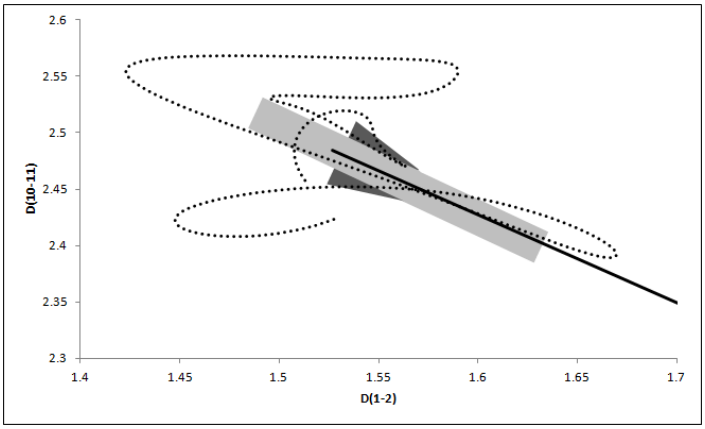

Figure 9.12: Representation of varying fuzz to capture more of the trajectory. For negative values of t, which corresponds to a molecular structure that is pre-reactant (and a d1-2) distance of about 1.56Å, the fuzz is increased. This allows the model to capture a larger portion of the trajectory of the structure, allowing for an increase in the number of time steps that the trajectory should be following in order for a reaction to be detected.

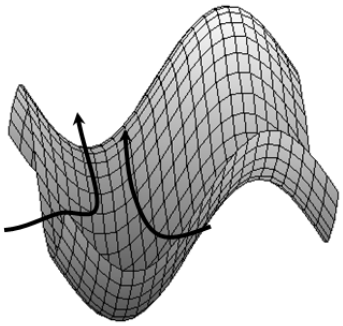

Figure 9.13: Trajectories that could lead to a reaction on a potential energy surface. While the potential energy surface may vary with structure, the concept is the same. That is, the trajectory does not have to pass 'straight through' the valley in order to result in a reaction. A number of trajectories could come in from either side and still lead to a reaction.

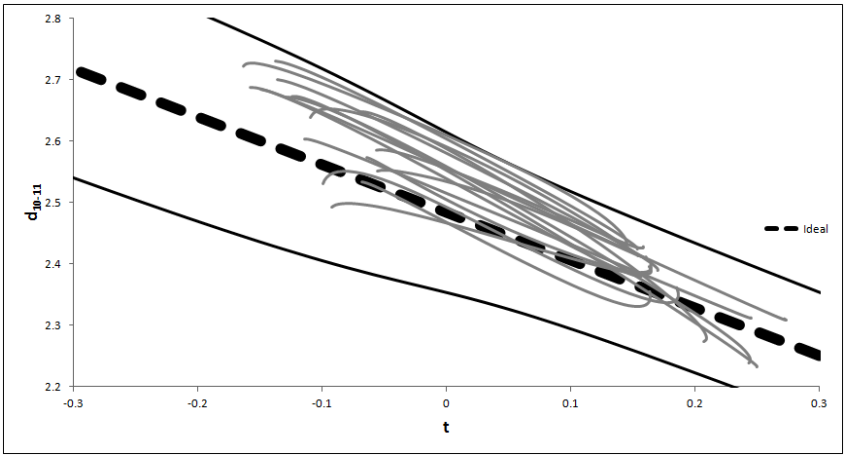

Figure 9.14: Trajectories that led to reactions drawn in the funnel. Every single grey line corresponds to a trajectory that successfully lead to a reaction. The dark lines above and below these lines correspond to the fuzz that was being implemented with this model and the funnel can be seen as the fuzz gradually changes. The dashed line in the middle corresponds to the trajectory for a fully concerted reaction.

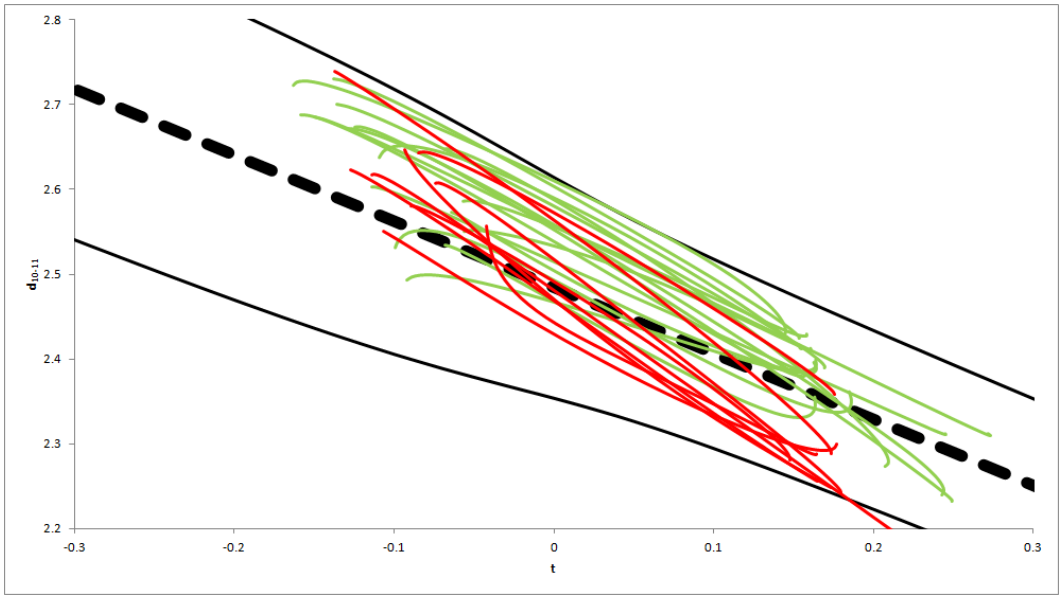

Figure 9.15: Trajectories for events that reacted (green) and those that did not react (red). While no further work was done after this point due to time constraints, it appears that trajectories that did react (green) appear to be very similar to those that didn't react (red). More work will need to be done in order to find another property of the red trajectories that can make then distinguishable from the green ones.

Appendix A
submitevents.pl source code

Source code of the script written to prepare event data generated by Tinker for submission into NWChem for QM calculations.

```
#!/usr/bin/perl

$eventfolder = $ARGV[0];
$resultsfolder =  "/work/5fjha/Test_simulations/";
$lookinfolder = $resultsfolder . $eventfolder . "/";
$copytofolder = $resultsfolder . "switching/temp/";
$basicinput = $resultsfolder . "switching/basic_input/";

opendir my($dh), $lookinfolder or die "Couldn't open dir
'$lookinfolder': $!";
my @files = readdir $dh;
closedir $dh;

foreach $file (@files) {
    if (rindex($file, "event_1") != -1)
    {
         system("mkdir " . $copytofolder . $file);
         system("cp " . $basicinput . "* " . $copytofolder . $file);
         system("cp " . $lookinfolder . $file . "/dyn.third_last " .
$copytofolder . $file);
         system("cp " . $lookinfolder . $file . "/dyn.third_last " .
$copytofolder . $file . "/semi_bull.dyn");
    }
}
```

Appendix B
filterevents.pl source code

Source code of the script written to remove events that have already been passed through NWChem's QM calculations. This removes duplicate events when an order parameter cutoff value has been reduced.

```
#!/usr/bin/perl

$cutoff1 = $ARGV[0] . "/reactive_events.dat";
$cutoff2 = $ARGV[1] . "/reactive_events.dat";
$framediff = $ARGV[2];
$count=0;

#Remove all events in cutoff2 from cutoff1 (2>1)

open(INPUT1, "< $cutoff1") || die("Cutoff1 could not be opened");
open(INPUT2, "< $cutoff2") || die("Cutoff2 could not be opened");

while (defined($line2 =  <INPUT2>))
 {
    $found=0;
    $line2 =~ m/(\d+)\s+?.*/;
    $line2 = $1;

    while ($found == 0)
     {
        $line1 =  <INPUT1>;
        $count +=1;
        $line1 =~ m/(\d+)\s+?.*/;
        $line1 = $1;
        $val = $line2-$line1;

        if ($val <= $framediff)
         {
             $found = 1;
             print "Event$count\n";
         }
     }
}

close(INPUT1);
close(INPUT2);
```

Appendix C
Vector Model source code

Source code for the vector model that was used Section 5.0.

```
dx=x(myorder(1,1))-x(myorder(2,1))
dy=y(myorder(1,1))-y(myorder(2,1))
dz=z(myorder(1,1))-z(myorder(2,1))
drx=dx*dx + dy*dy + dz*dz
drx=sqrt(drx)
dx=x(myorder(1,2))-x(myorder(2,2))
dy=y(myorder(1,2))-y(myorder(2,2))
dz=z(myorder(1,2))-z(myorder(2,2))
dry=dx*dx + dy*dy + dz*dz
dry=sqrt(dry)



newt=(drx-mydr0(1))/myop(1)
expy= mydr0(2)+(newt*myop(2))
diff=expy-dry



if (abs(diff).LE. myfuzz) then
      if (newt .GE. oldt) then
            count=count+1
      else
            count=0
      endif
else
      count=0
endif



oldt=newt



 if (count.GE.mylength .and. newt.GE.mytcut ) then
     rxn_possible=.true.
 endif
```