Algorithms for efficient calculation of nonlinear optical spectra: Ultrafast Spectroscopy Suite and its applications

by

Peter A. Rose

A thesis submitted in partial
fulfillment of the requirements for the
Doctorate in Philosophy degree in Physics

Department of Physics
Faculty of Science
University of Ottawa

Abstract

This thesis presents analytic and computational advances in the prediction of perturbative nonlinear optical spectroscopies. The contributions of this thesis are packaged together in an open source, freely available piece of software called ultrafast spectroscopy suite (UFSS). It is designed to automatically simulate nonlinear optical spectroscopies for any phase-matching or phase-cycling condition, including finite pulse effects. UFSS includes an algorithm called the diagram generator (DG) that automates the process of writing out all of the Feynman diagrams that contribute to a desired phase-matching or phase-cycling condition, and includes all pulse overlap diagrams when relevant, paving the way toward automation of perturbative calculations. Further, many diagrams can be automatically combined into composite diagrams, giving an exponential decrease in computation time of high-order calculations. Composite diagrams even allow for the efficient study of Rabi oscillations as a function of pulse amplitude, by summing many orders of perturbation theory. The perturbative calculations are done using a novel algorithm presented in this thesis called Ultrafast Ultrafast spectroscopy (UF^2). UF^2 is an efficient method for determining diagrammatic contributions to spectra including arbitrary (whether analytical or experimentally measured) pulse shapes. It uses the speed of the fast Fourier transform to be as much as 500 times faster than direct propagation techniques for small model Hamiltonians (for Hamiltonian dimension of 100 or less). UF^2 outperforms direct propagation techniques for a wide range of model systems, with the speed boost diminishing as the dimension of the model Hamiltonian increases. UF^2 can predict spectra for any model system whose relevant Hilbert space that can be described using a finite basis and that can be diagonalized numerically, and users are free to specify their own model. UFSS includes a model generator that generates Hamiltonians and Liouvillians of vibronic systems, allowing users to easily simulate NLOSs for a wide range of model system parameters. UFSS is a fully functional piece of software for simulating any NLOS, to any desired order in perturbation theory.

Acknowledgements

First and foremost I thank my advisor, Jacob J. Krich. You have been the best advisor a student could hope for, giving both remarkable scientific insight and compassion. Scientific discussions with you were always illuminating and exciting. I came out of every meeting with you excited for the next step, no matter how frustrated and stuck I felt going in. The fact that you care about your students as people is constantly clear, and a wonderful quality in an advisor. I could not possibly have made it through the winding road of a PhD without your patience, kindness and insight. It has been a great privilege to work with you.

Next I thank my wife, Larisa Rose. Your unwavering love and support were a constant source of strength, stability and joy throughout the ups and downs. I could not have finished this degree without you. Thank you to my parents for your love; for believing in me, and teaching me to believe in myself. Thank you to all of my family, near and far; you were never more than a phone call away and I felt your love every day.

Finally, I thank all of my fellow students in the Krich group and in Sunlab, both for stimulating conversations and for a good laugh. Research is so much more fun when done in good company.

Dedication

To the women in my life, both present and remembered. You have done so much for me. You are all incredible.

Statement of Original Contributions

Chapters 1-4 of this thesis provide summaries of well-known, foundational material that is helpful in understanding the original work presented in the remaining chapters. Chapter 5 was published in the *Journal of Chemical Physics* with authors Peter A. Rose and Jacob J. Krich. The project was begun in collaboration with Jacob J. Krich and I performed all analytical and numerical calculations. Jacob J. Krich provided guidance and edited the manuscript. Chapter 6 was published in the *Journal of Chemical Physics* with authors Peter A. Rose and Jacob J. Krich. The project was begun in collaboration with Jacob J. Krich and I performed all analytical and numerical calculations. Jacob J. Krich provided guidance and edited the manuscript. Chapter 7 was published in the *Journal of Chemical Physics* with authors Peter A. Rose and Jacob J. Krich. The project was begun in collaboration with Jacob J. Krich and I performed all analytical and numerical calculations. Jacob J. Krich provided guidance and edited the manuscript. Chapter 8 is not published. The project was begun in collaboration with Jacob J. Krich and I performed all analytical and numerical calculations. Jacob J. Krich provided guidance and edited the chapter. Chapter 9 is not published. The project was begun in collaboration with Jacob J. Krich and I performed all analytical and numerical calculations. Jacob J. Krich provided guidance and edited the chapter.

Contents

Chapter 1

Introduction: Nonlinear Optical Spectroscopies

Optical spectroscopy is a huge field that can perhaps be most simply described as methods using light to understand and characterize atoms, molecules and materials. Linear absorption and emission spectra provide fingerprints for atoms and small molecules, with applications ranging from chemistry to astrophysics. Linear absorption of molecules at low temperature in the gas phase can give beautifully detailed snapshots of the electronic, vibrational and rotational energy levels. Though rich with information, it is, at the end of the day, a picture. It is a static measurement. Nonlinear optical spectroscopies (NLOSs) are a wide class of techniques that allow us to probe the dynamics of molecules and materials, and gain insight into their properties, in both the gas and condensed phases. However, the experimental data from NLOSs are complex and often difficult to interpret. It is therefore important to develop efficient methods for simulating NLOSs, so that we can fit a model to the data. This thesis describes open source, freely available python code that combines all the ingredients needed for modelling NLOSs into an efficient and easy to use software package we call the ultrafast spectroscopy suite (UFSS).

In order to probe the dynamics of a system, two key ingredients are required: some way to bring the system out of equilibrium, and some way to monitor the resulting dynamics with high time resolution. Ultrafast optical pulses as short as 5-10 fs provide both of these ingredients [1, 2]. These pulses are ideally suited to excite optically sensitive molecules and materials, bringing them out of equilibrium. The same pulses can then be used to monitor the resulting dynamics via changes in the system's optical response to a subsequent pulse. The ultrashort time period of both excitation and detection makes it possible to map out the optical response in steps of 10 fs or less, which means that we can, in some sense, film the molecular dynamics with a frame rate up to the order of 100 THz.

There are many different types of NLOSs, each of which is sensitive to different system dynamics. We give more detailed examples of a few of the most popular NLOSs in Chapter 4. Briefly, each NLOS is defined by the number of

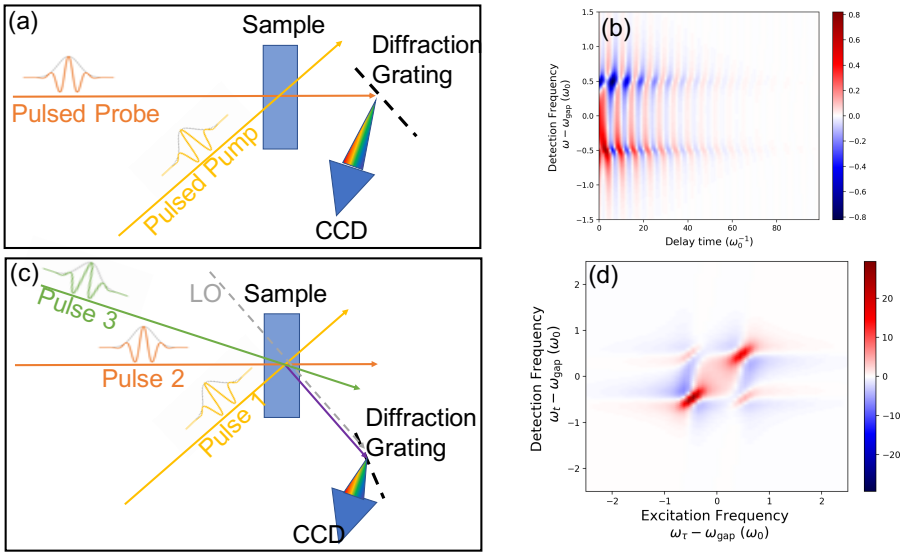

Figure 1.1: (a) Schematic of transient absorption (TA) spectroscopy. The absorption of the probe is measured for each delay time between the pump and probe pulses. (c) Simulation of the transient change in probe absorption spectra (same as Fig. 4.2(c) in Sec. 4.2). (c) Schematic of 2D photon echo spectroscopy. Pulses 1-3 interact in the sample to generate a nonlinear signal, shown in purple. The signal is typically measured using a reference pulse call the local oscillator (LO). The LO does not participate in generating the 2D photon echo signal. (d) Simulation of the 2D photon echo signal as a function of excitation frequency and detection frequency (same as Fig. 4.5(d) in Sec. 4.4). For more information about these TA and 2DPE simulations, see Chapter 4.

pulses that are used, and something called a phase-discrimination condition, which we define in Chapter 7. In many cases, the phase-discrimination condition corresponds to where a detector is placed, relative to the pulse propagation directions. For example, transient absorption (TA), shown schematically in Fig 1(a), involves measuring the change in absorption of a probe pulse due to the pump pulse, where the probe pulse arrives some time T after the pump pulse. Another way of saying this is that a detector monitors the signal generated along the direction of the probe pulse. In contrast, 2D photon echo spectroscopy (2DPE), shown schematically in Fig 1(c), involves 4 pulses, with a detector monitoring the signal generated along the direction of the 4th pulse. Changing the placement of the detector changes the phase-discrimination condition, which changes the signal that is measured.

The signals that are detected in a NLOS are, of course, not 3D movies of the electrons and nuclei of the system under study. Rather, they are complex, information rich spectra that can be quite difficult to interpret, as evidenced by the decade long debate that ensued about the interpretation of NLOS data taken for a photosynthetic complex from purple sulphur bacteria called the Fenna-Matthews-Olson complex (see Sec. 4.4 for a brief description of this issue) [3–6]. The fact that NLOSs can be difficult to interpret should not come as a surprise, because they are essentially the projection of the motion of many (10s, 100s, perhaps more) degrees of freedom onto, typically, 2 or 3 dimensional data sets.

In order to accurately interpret the data that come from NLOSs, we need an efficient way to fit a model to the experimental data. In order to simulate NLOSs, we need three key ingredients: (1) the model for the molecular system that we are attempting to fit to the data, (2) a model for the interaction between the ultrafast pulses and the

system, and (3) a method for predicting the specific signal that is determined by the phase-discrimination condition. The process of going from model system to spectra can be quite computationally intensive, and therefore fitting a model to experimental spectra, which requires repeated simulations using varying system parameters, is that much more computationally demanding than a single simulation [7–10].

Furthermore, in many cases, simulations of NLOSs often ignore the shape of the ultrafast pulses, opting to work in the so-called impulsive limit, which we discuss in more detail in Chapter 3. Quantitative prediction of spectra benefits from the inclusion of the pulse shapes used in the experiment. In this thesis, we introduce a novel algorithm that is an efficient method for simulating NLOSs that we call ultrafast ultrafast (UF^2) spectroscopy, and detail in Chapters 5 and 6. UF^2 uses the speed of the fast Fourier transform to efficiently include any pulse shape (either analytical or experimental).

Perhaps the most general contribution this thesis makes to the field of simulating NLOSs lies in codifying the concept of a phase-discrimination condition, distilling it into a simple input parameter. The algorithm we have written for this purpose is called the diagram generator (DG), which is detailed in Chapter 7, and which we further improve in Chapter 8. The DG allows us to automate the calculation of any NLOS by simply writing down the phase-discrimination condition as a short list of tuples. With the power of the DG, the difference between simulating TA and 2DPE lies in defining two additional pulses, their associated delay times, and altering a simple tuple representing the phase-discrimination condition. The DG interprets the phase-discrimination condition, automatically determines what calculations are needed, and dispatches them to UF2. A user who moves their detector to monitor a different signal (changing the phase-discrimination condition), can change the associate simulation just by changing a single input parameter.

It is worth taking a moment to discuss the philosophical approach to modeling molecular systems that we take in this thesis, which is to seek the simplest effective model that can accurately predict the data from a NLOS experiment. We demonstrate in Chapters 5 and 6 that UF^2 can be hundreds of time faster than alternative methods, and that UF^2 gains the biggest speed advantage for small model systems, which is consistent with the goal of finding the simplest model that can fit the experimental results. If a small model will not suffice (dimension of the system Hamiltonian is greater than $\approx 10^5$ for closed systems or greater than $\approx 10^3$ for open systems coupled to a Markovian bath; see Chapters 5 and 6 for detailed analysis of where UF^2 outperforms direct propagation techniques), UF^2 will not be a suitable choice.

Before moving further into this thesis, we briefly describe two types of NLOS, and direct the reader to Chapter 4 for a broader (though far from exhaustive) overview of some of the signals that can be simulated using UFSS. Returning to TA, the absorption of the probe is measured for many different delay times between the arrival of the pump and probe pulses. The absorption of the probe pulse can be measured either using a photodiode, or using a diffraction grating and a CCD, in order to obtain spectral resolution of the absorption process, as shown schematically in Fig. 1.1(a). Oscillations in the absorption of the probe as a function of pump-probe delay time

reveal the energy differences between energy levels (either electronic or vibrational) in the electronic excited and ground states of the molecules under study, as seen in Fig. 1.1(b). Decay of certain absorption bands can reveal energetic relaxation of the excited states, as well as relaxation to the nuclear equilibrium position in the electronic excited state. Decay of the entire TA signal gives insight into the relaxation rate to the electronic ground-state of the molecule. For example, TA has been used to study the rapid internal conversion process crucial to vision that takes place in rhodopsin [11, 12]. An example TA spectra is shown in Fig. 1(b). This simulation is done using UF2, and an explanation of the simulation can be found in Chapter 4.

Although a powerful technique, TA sacrifices spectral control during the excitation process in order to obtain high time resolution. The time-resolution of TA is limited by the time duration of the pulses. However, ultrashort pulses necessarily include many different wavelengths. It is therefore not possible to selectively excite at a single wavelength, and also obtain high time resolution. However, by using two pump pulses, instead of one, 2D photon echo (2DPE) spectroscopy can obtain spectrally resolved information about the excitation process, allowing for the study of 2D frequency-frequency correlation plots [13]. 2DPE involves 3 pulses that interact with the sample and create a nonlinear signal, and a fourth reference pulse that is used to detect that signal. By varying the time delay between the first two pulses, one obtains the temporal profile of the excitation process, the Fourier transform of which shows excitation as a function of frequency. The time delay between pulses 2 and 3, called the population time, serves the same role as the delay between pump and probe in TA. 2DPE allows one to study dynamics, just as with TA, but now as a function of both excitation frequency and detection frequency. An example 2DPE spectrum as a function of both excitation and detection frequency is shown in Fig. 1(d). Fig. 1(d) shows a simulation for a single population time, but typically many population times are studied. The simulation in Fig. 1(d) is done using UF2, and an explanation of the simulation can be found in Chapter 4.

From a theoretical perspective, NLOSs are analogous to nonlinear optics, but typically involve studying the resonant aspects of the nonlinear response [14]. NLOSs are often thought of as methods for measuring the nonlinear response of the material under study, due to a perturbation caused by weak optical pulses. Much of the intuition for these experiments comes from constructing the third-order response functions (introduced in Sec. 3.3), based upon a mathematical model for the material system of interest. Accurate prediction of experimental signals should involve the convolution of these response functions with the optical pulses used in experiments. However, this last step is often not done. Or if it is, it is done in an incomplete manner, for example by using approximate methods and/or by using idealized pulse shapes like Gaussians or Lorentzians [15–20]. Further, fitting experimental data ideally would involve recalculating the signal (including exact pulse shapes) many times by varying the model parameters. However, calculating the signal once, without the inclusion of pulse shapes, is often a very expensive process, and so the fitting process is often not done. The motivation of this thesis is to develop a fast and convenient numerical method for simulating NLOSs, including arbitrary pulse shapes, to make this fitting procedure fast, thus enabling the quantitative prediction of NLOSs.

1.1 Outline of Thesis

Chapters 2 and 3 represent summaries of well-known, standard foundational elements of the field of NLOSs. Chapter 2 describes models of molecular systems, introduces the Born-Oppenheimer approximation, and describes the model systems that we study. Chapter 3 describes the light-matter interaction, including time dependent perturbation theory, the electronic dipole approximation, and introduces the concepts of response functions. After establishing the formalisms of NLOSs, Chapter 4 shows simulations of a few of the most well-known NLOSs using a simple model system, with the aim of giving a brief look at what NLOSs are, and some of the effects that can be studied using NLOSs.

The remaining chapters describe our own work. Chapters 5-7 are published articles, while chapters 8 and 9 are unpublished. I am the first author on all published works, with my advisor being the sole other author, and all published work represents my own work under the direction of and in collaboration with my advisor. The unpublished work similarly represents work done by me under the supervision of my advisor. Chapter 5 introduces an efficient algorithm for calculating NLOSs for closed systems, called ultrafast ultrafast (UF^2) spectroscopy. It challenges the conventional wisdom that diagonalizing the model for the molecular Hamiltonian is computational more costly than direct propagation of the Hamiltonian dynamics. We show, by comparison to a direct-propagation algorithm we include in the software, called RKE, that UF^2 outperforms direct propagation techniques over a wide range of system sizes.

Chapter 6 extends UF^2 to treat open systems. This extension is based upon diagonalizing the equations of motion for the density matrix (the Liouvillian), which is rarely done. However, we point out that when working with the Redfield equation, one must diagonalize the system Hamiltonian. In the secular approximation, diagonalizing the Liouvillian is no more costly than diagonalizing the Hamiltonian, and that in this case, UF^2 always outperforms direct propagation methods. We also demonstrate the range of system sizes for which UF^2 outperforms direct propagation for the full Redfield equations, as well as for the diabatic Lindblad formalism.

Chapter 7 describes the Feynman diagram generator (DG), which is a component in UFSS that automatically generates all of the Feynman diagrams that give non-zero contributions to any perturbative spectroscopy that can be described by a phase-matching or phase-cycling condition. The DG generates all pulse overlap diagrams when they are relevant (when pulses overlap), and generates only the time-ordered diagrams when pulses do not overlap. The DG can generate diagrams for any order of perturbation theory, which allows us to automate spectroscopic calculations to any desired order. The DG can also output the diagrams as pdf files so that the diagrams can be visually inspected.

Chapter 8 is unpublished work that demonstrates how to combine many single Feynman diagrams together into what we call composite diagrams. This essentially replaces the DG for calculation purposes, and offers an exponential speed-up in calculation of higher-order perturbative calculations, opening the door to studying higher-order spectroscopies including pulse overlap effects, and to studying non-perturbative effects in the field-strength.

We use the composite diagram method to efficiently simulate damped Rabi oscillations in transient absorption as a function of pump pulse amplitude in a two-level and a three-level system.

Chapter 9 goes into the details of how we represent and diagonalize the model systems introduced in Chapter 2. We represent (an)harmonic vibrational modes using the ladder operators of the quantum harmonic oscillator. We compare using the ladder operators of the ground-state configuration to ladder operators that represent fictitious vibrational modes that have an equilibrium geometry halfway between the equilibria of the ground and excited states, and have vibrational frequencies that are the geometric means of the ground and excited state vibrational frequencies. We demonstrate that we can represent the model Hamiltonian more efficiently in this basis, and also discuss an efficient method for truncating the basis states used to represent the Hamiltonian.

Chapter 10 gives concluding remarks.

1.2 Software testing and validation

Since one of the primary goals of this thesis is to describe a freely available piece of software, we briefly discuss the tests that have been done to ensure that the algorithms described here and implemented in the Python package UFSS produce accurate results. As mentioned in Chapter 5, we have compared a simulation done using UFSS to a simulation done using a different, independently developed, piece of software introduced in Ref. 21. The results of each simulation agreed to within less than 1%, which is the convergence threshold we have typically used throughout our work, unless otherwise stated. This validation of the results from UFSS by an outside source give confidence in the reliability of UFSS.

In addition, as discussed in Chapters 5 and 6, we have developed two independent methods for calculating nonlinear optical spectra. One is the novel algorithm introduced in those two chapters, Ultrafast Ultrafast (UF^2), while the other is our own implementation of one of the standard methods from the literature, which we call Runge-Kutta Euler (RKE). Throughout the development of UFSS, we have tested to make sure that the results of UF^2 and RKE agree, providing a strong internal check. In section 6.4 we present runtime comparisons for how long it takes to generate transient absorption (TA) spectra using both UF^2 and RKE. We calculated TA spectra, which are 2-dimensional datasets consisting of the frequency-detected change in absorption for many different delay times between the pump and probe pulses, using both methods for a variety of different model systems. For each model system, we compared the spectra generated by UF^2 and RKE by computing the ℓ_2-norm difference between the two spectra. Specifically, for each runtime ratio presented in Figures 6.3 and 6.4, we verified that the spectra produced by UF^2 and RKE agreed within 2%.

Further, in chapter 6 we show that UF^2 converges to an analytical result from Ref. [19] as dt^2.

Each of the tests described above is included in the software repository as a test that can be run, to check (a) that UFSS gives the same result as a spectrum generated using the code from Ref. 21, (b) that UF^2 gives the same

result as RKE for a test case, and (c) that UF^2 gives the same result as the analytical solution from Ref. [19].

Chapter 2

Simple Models for Molecular Systems

In this chapter we discuss the way that we model molecular systems, and aggregates or polymers of such systems. The field of developing accurate models for specific molecules is a vast and rich field. The starting point for modeling molecules ab initio usually begins with electronic structure calculations using density functional theory [22, 23]. However, we do not seek to model any particular molecule, but rather use a model that can capture a few of the most generic features of molecules, with the hope that we can make qualitative predictions about molecular systems as a class, rather than predictions for particular molecules.

2.1 Molecular Hamiltonians and the Born-Oppenheimer Approximation

We begin by discussing some of the common approximations that are made when developing models for molecular systems. This section represents a brief summary of material drawn largely from chapter 9 of Ref. 24. Molecules consist of n electrons and N atoms. Each particle has its own mass, charge and spin. Although electrons can be thought of as point charges, nuclei are made up of many individual charges, and thus the electric fields produced by each of the nuclei are more complicated. If we neglect the effects of spin, we arrive at the rovibronic (rotational-vibrational-electronic) Schrödinger equation,

$$\left(\underbrace{-\frac{\hbar^2}{2}\sum_{i=1}^{n+N}\frac{\nabla_i^2}{m_i}}_{T} + \underbrace{\frac{\hbar^2}{2M}\sum_{i,j=1}^{N+n}\nabla_i\cdot\nabla_j}_{T'} + \underbrace{\sum_{i<j=1}^{N+n}\frac{q_i q_j}{4\pi\epsilon_0 R_{ij}}}_{V} \right)\Psi_k(\boldsymbol{r},\boldsymbol{R}) = E_k\Psi_k(\boldsymbol{r},\boldsymbol{R}),$$

where m_i are the masses of each particle, M is the total mass of the molecule, q_i are the charges of each particle, R_{ij} indicates interparticle distance, $\boldsymbol{r}$ are the positions of the electrons and $\boldsymbol{R}$ are the positions of the nuceli. These

terms can be broken down as follows

$$T = \underbrace{-\frac{\hbar^2}{2m_e}\sum_{i=1}^{n}\nabla_i^2}_{T_e} \underbrace{-\frac{\hbar^2}{2}\sum_{i=n+1}^{n+N}\frac{\nabla_i^2}{m_i}}_{T_N}$$

$$V = \underbrace{\sum_{i<j=1}^{n}\frac{e^2}{4\pi\epsilon_0|r_i - r_j|}}_{V_{ee}} + \underbrace{\sum_{i=1}^{n}\sum_{j=n+1}^{n+N}\frac{-Z_j e^2}{4\pi\epsilon_0|r_i - R_j|}}_{V_{eN}} + \underbrace{\sum_{i<j=n+1}^{n+N}\frac{Z_i Z_j e^2}{4\pi\epsilon_0|R_i - R_j|}}_{V_{NN}}.$$

By invoking the Born-Oppenheimer approximation, we separate the rovibronic Hamiltonian into a purely electronic part that depends only parametrically upon the nuclear positions, and a purely rotational-vibrational part that is independent of the electron positions. The Born-Oppenheimer approximation is intuitively a statement that the electrons, which are many orders of magnitude lighter than the nuclei, move much more quickly than the nuclei, and thus rearrange instantaneously on the scale of nuclear motion. Thus the electrons create a potential energy surface upon which the nuclei move. This is generally an accurate assumption only when the energetic separation between different electronic states is much larger than the energetic separation between different nuclear states. By taking the ansatz that $\Psi(\boldsymbol{r}, \boldsymbol{R}) = \psi(\boldsymbol{r}; \boldsymbol{R})\chi(\boldsymbol{R})$, and neglecting the term T', we separate the rovibronic Schrödinger equation into the electronic part,

$$\left(T_e + V_{ee}(\boldsymbol{r}) + V_{eN}(\boldsymbol{r}, \boldsymbol{R})\right)\psi_i(\boldsymbol{r}; \boldsymbol{R}) = V_{\text{elec},i}(\boldsymbol{R})\psi_i(\boldsymbol{r}; \boldsymbol{R}), \tag{2.1}$$

where $V_{\text{elec},i}(\boldsymbol{R})$ is the R-dependent eigenvalue of the i^{th} electronic state. Since the Hamiltonian in Eq. 2.1 commutes with $\boldsymbol{R}$, it can be solved with $\boldsymbol{R}$ fixed. By solving for $V_{\text{elec},i}(\boldsymbol{R})$ at many different values of $\boldsymbol{R}$, a potential energy surface in the i^{th} electronic state is mapped out for the rotational-vibrational Schrödinger equation, which is

$$\left(T_N + V_{NN}(\boldsymbol{R}) + V_{\text{elec},i}(\boldsymbol{R})\right)\chi_\nu^{(i)}(\boldsymbol{R}) = E_\nu^{(i)}\chi_\nu^{(i)}(\boldsymbol{R}). \tag{2.2}$$

Even after making so many approximations, the task of solving Eq. 2.1 is a very computationally challenging task, since there are still $3n$ degrees of freedom. Further, some essential effects of the neglected electron spin degrees of freedom must be taken into account (at a minimum the Pauli exclusion principle must be satisfied).

However, from the perspective of this thesis, the form of $\psi(\boldsymbol{r}; \boldsymbol{R})$ is practically unimportant. We will assume that there are two independent electronic states of interest, the ground and excited state, denoted $|g\rangle$ and $|e\rangle$, respectively. We will also assume that the combination $V_{NN} + V_{\text{elec},i}$, which defines the total potential energy surface of the nuclei, is either harmonic, or nearly so, and well approximated by a power series expansion in $\boldsymbol{R}$. By

writing $V_g = V_{NN} + V_{\text{elec},g}$ and $V_e = V_{NN} + V_{\text{elec},e}$, we can write the vibronic Hamiltonian of a two-level system as

$$H = T_N + V_g(\boldsymbol{R}) \left|g\right\rangle \left\langle g\right| + V_e(\boldsymbol{R}) \left|e\right\rangle \left\langle e\right|,$$

or

$$H = T_N + V_g(\boldsymbol{R}) + c^\dagger c \tilde{V}_e(\boldsymbol{R}), \tag{2.3}$$

where $c^\dagger$ is the electronic excitation creation operator, and $\tilde{V}_e(\boldsymbol{R}) = V_e(\boldsymbol{R}) - V_g(\boldsymbol{R})$. Eq. 2.3 represents the fundamental building block that we use to describe aggregates and polymers.

Molecules that consist of N atoms have $3N-6$ vibrational degrees of freedom, 3 translational degrees of freedom, and 3 rotational degrees of freedom. A change of basis allows one to separate out the motion of the center of mass of the molecule, and express T_N and $V(\boldsymbol{R})$ in terms of the normal modes of the molecule. In this case, the real-space coordinates $\boldsymbol{R}$ are replaced with mass-weighted coordinates $\boldsymbol{q}$ that have units of $\sqrt{m_i}R$. We will further neglect the rotational degrees of freedom. Even so, including all of the vibrational degrees of freedom in a quantum mechanical simulation would require working in a Hilbert space whose dimension grows exponentially with the number of vibrational degrees of freedom, which is computationally unfeasible even for small N. A great deal of work has been done to include the effects of these vibrational modes at various levels of approximation. One method of overcoming the exponential scaling of quantum-mechanical calculations is to treat all of the nuclei classically, using nonequilibrium molecular dynamics simulations. By treating each nucleus as a classical massive point charge, the scaling cost can be reduced to polynomial scaling in N. Hybrid methods have also been developed to simulate several nuclei quantum mechanically, while treating the rest of the system classically. Several approximate quantum mechanical methods have been developed that can simulate hundreds of vibrational modes. The multi-layer multi-configurational time-dependent Hartree (ML-MCTDH) method is an approximate quantum mechanical simulation that has been used to treat 300 degrees of freedom [25]. The Tree Tensor Network state (TTNS) algorithm has been used to simulate 252 degrees of freedom using [26].

Rather than trying to explicitly model all of the vibrational modes in a molecule, an alternative method involves identifying "important" degrees of freedom that will be studied exactly, and treating the remaining degrees of freedom as a perturbation. In many cases, a relatively small number of degrees of freedom are actually relevant to the system dynamics of interest. For example, simulations of the absorption spectrum of pyrazine showed that only 4 of the 24 degrees of freedom are important for the qualitative dynamics of the conical intersection, while the remaining 20 degrees of freedom are primarily responsible for relaxation and dephasing [27].

The idea that most degrees of freedom of complex systems serve primarily to cause relaxation and dephasing leads to the concept of the reduced density matrix (see Sec. 2.2.1 below for a brief description of open systems and density matrices). In the density matrix formalism of open systems, a small number of degrees of freedom are identified as being the "system", while the others are identified with the "bath". By performing a partial trace over

the bath, a reduced density matrix is obtained, which is small enough to be computationally tractable. In order to propagate the reduced density matrix, approximate equations of motion are derived by treating the bath degrees of freedom as a perturbation on the system.

The heierarchical equations of motion (HEOM) algorithm is a method for propagating the reduced density matrix to any level of precision in the system-bath coupling (though as with all approximate methods, higher accuracy requires more computational resources, making a practical cut-off of the accuracy) [28–30]. If the system-bath coupling is weak, less computationally expensive methods can be employed. In many cases a quantum master equation can be derived to give approximate equations of motion for the system (see for example Ref. 31 and Sec. 2.2.1 below). If the bath retains information about the system dynamics, the bath is said to have memory, and the master equation is time-dependent. If one approximates that the time-scale on which the bath relaxes is faster than any of the system dynamics, the time-dependence (and memory) of the bath is neglected, which greatly simplifies the master equation. This neglect of the bath memory is known as the Markovian approximation.

We rely on the Markovian approximation throughout this thesis. We harness two different formalisms for describing Markovian processes: Lindblad and Redfield theory, which we discuss in further detail below. The overall approach taken in this thesis is to include a small number of vibrational modes in the system, and, in the case of open systems, treat the bath using either the Redfield or the Lindblad formalism. We start by introducing the vibronic Hamiltonian, and then discuss coupling to the bath.

2.2 Polymer Model

The goal of this section is to introduce a model of s coupled two-level systems (TLS), each with an electronic ground state $|g\rangle$, and a single optical excitation $|i\rangle = c_i^\dagger |g\rangle$, interacting with K vibrational modes with generalized momentum and coordinates $\vec{p} = \{p_1, p_2, ..., p_K\}$ and $\vec{q} = \{q_1, q_2, ..., q_K\}$, respectively. The general form of this Hamiltonian that we consider is

$$H = T(\boldsymbol{p}) + V_g(\boldsymbol{q}) + \sum_{i=1}^{s} c_i^\dagger c_i \tilde{V}_i(\boldsymbol{q}) + \sum_{i \neq j} c_i^\dagger c_j V_{ij}(\boldsymbol{q}), \tag{2.4}$$

where T is the momentum operator, V_g is the potential energy hypersurface associated with ground electronic state of the system, $\tilde{V}_i = V_i - V_g$ is the change to the potential energy hypersurface due to excitation of site i, V_{ij} are the possibly coordinate-dependent couplings between excitations i and j, and $c_i^\dagger$ is the creation operator for an excitation of site i. Eq. 2.4 is a sum of s individual Hamiltonians given by Eq. 2.3, with the addition of the inter-site coupling term $V_{ij}(\boldsymbol{q})$.

By expressing V_g and $\tilde{V}_i$ as power-series in $\vec{q}$, we can solve Eq. 2.4 using the harmonic oscillator wavefunctions as a basis. To do this we can make the substitution $q_\alpha \to \sqrt{\frac{\hbar}{2\omega_{\alpha,g}}} \left(b_\alpha + b_\alpha^\dagger\right)$, where b_α are the lowering operators for each vibrational mode with ground-state frequency $\omega_{\alpha,g}$. By using a finite‘ number of basis terms, we can express

H as a matrix that can be diagonalized numerically for a small number of vibrational modes. When working with closed systems, we identify a small number of vibrational modes $k \ll K$, which are most strongly coupled to the electronic states, and partition the total Hamiltonian as

$$H = H_S \otimes I_B + I_S \otimes H_B + H_{SB},$$

with

$$H_S = T(p_1, p_2, ..., p_k) + V_g(q_1, q_2, ..., q_k) + \sum_{i=1}^{s} c_i^\dagger c_i \tilde{V}_i^{(0)}(q_1, q_2, ..., q_k) + \sum_{i \neq j} c_i^\dagger c_j V_{ij}^{(0)}(q_1, q_2, ..., q_k) \tag{2.5}$$

$$H_B = T(p_{k+1}, p_{k+2}, ..., p_K) + V_g(q_{k+1}, q_{k+2}, ..., q_K) \tag{2.6}$$

$$H_{SB} = \sum_{i=1}^{s} c_i^\dagger c_i \left(\tilde{V}_i(q_1, q_2, ..., q_k, ..., q_K) - \tilde{V}_i^{(0)}(q_1, q_2, ..., q_k) \right) + \sum_{i \neq j} c_i^\dagger c_j \left(V_{ij}(q_1, q_2, ..., q_{k+1}, ..., q_K) - V_{ij}^{(0)}(q_1, q_2, ..., q_k) \right), \tag{2.7}$$

where $\tilde{V}_i^{(0)}(q_1, q_2, ..., q_k)$ is the zeroth-order expansion of $\tilde{V}_i(q_1, q_2, ..., q_K)$ with respect to the bath modes $\{q_\alpha\}_{\alpha=k+1}^{K}$. I_B is an identity operator on the bath, while I_S is an identity operator upon the system subspace. In the case of closed systems, we ignore the bath, and work exclusively with H_S.

2.2.1 Open systems

In the case of open systems (the following brief discussion of open quantum systems is largely drawn from chapter 3 of Ref. 31), we expand H_{SB} to lowest order in the bath degrees of freedom, leading to a linear coupling to the bath

$$H_{SB} \approx \sum_r O_{r,S} \otimes \underbrace{\sum_{\alpha'=k+1}^{K} \lambda_{r,\alpha'} q_{\alpha'}}_{O_{B,r}},$$

where $O_{r,S}$ are system operators, such as $c_i^\dagger c_i$, q_α, etc. Assuming that the linear couplings $\lambda_{r,\alpha'}$ are small, master equations can be derived by treating H_{SB} perturbatively in $\lambda_{r,\alpha'}$ to arrive, for example, at the Redfield equation. By approximating the bath to be a continuum of oscillators, and invoking the Markov approximation, one derives a time-independent master equation.

When working with open systems, we define equations of motion for the density matrix, which is an operator on the Hilbert space of H and can represented as

$$\rho = \sum_n a_n |n\rangle \langle n|,$$

where $|n\rangle$ are wavefunctions. If there is a basis such that $a_0 = 1$ and $a_n = 0$ for all $n > 0$, ρ is said to be a pure

state. Otherwise, ρ is a mixed state. The time-dependent Schrödinger equation for the density matrix is

$$\dot{\rho} = -\frac{i}{\hbar}\left[H, \rho\right], \tag{2.8}$$

and the expectation value of an operator A is

$$\langle A \rangle = \mathrm{Tr}\left[A\rho\right].$$

Eq. 2.8 represents the equations of motion for a closed system, and the solutions to Eq. 2.8 are the same as the solutions to the time-dependent Schrödinger equation for wavefunctions. The power of working with density matrices lies in the formalism of open systems and quantum master equations, which allows us to write down effective equations of motion for a small number of degrees of freedom, rather than having to treat all of the degrees of freedom explicitly.

We work with a quantum master equation that defines effective equations of motion for the reduced density matrix, which is defined as

$$\rho_S(t) = \mathrm{Tr}_B \rho_{SB}(t),$$

where Tr_B indicates a partial trace over only the bath degrees of freedom. We assume that the initial density matrix is a product state:

$$\rho_{SB}(0) = \rho_S(0) \otimes \rho_B(0).$$

The state of the system at some later time is, in principle, given by

$$\begin{aligned}\rho_S(t) &= \mathrm{Tr}_B \left\{U(t,0)\rho_{SB}^{(0)}U^\dagger(t,0)\right\} \\ &= V(t)\rho_S(0),\end{aligned}$$

where $V(t)$ is called a dynamical map. In order to describe a physical process, $V(t)$ must preserve the positivity and Hermiticity of the complete density matrix ρ_{SB}. In the Markovian limit, we have time-independent equations of motion

$$\dot{\rho} = \mathscr{L}\rho,$$

where $\mathscr{L}$ is a superoperator, and is the generator of the quantum dynamical semigroup $V(t)$. It has been shown that the most general form of the generator $\mathscr{L}$ that satisfies the physicality conditions of preserving the Hermiticity and complete positivity of the system and bath, is the Lindblad form, which can be expressed as

$$\dot{\rho} = -\frac{i}{\hbar}\left[H, \rho\right] + \sum_{k=1}^{N^2-1} \gamma_k \left(A_k \rho A_k^\dagger - \frac{1}{2}\left\{A_k^\dagger A_k, \rho\right\}\right),$$

where N is the dimension of the system Hilbert space, A_k are operators on the system Hilbert space, and γ_k are coupling constants that have dimensions of inverse time [32, 33].

Redfield theory allows the derivation of a form for $V(t)$ that arises from a microscopic treatment of the bath (see Sec. 3.3 of Ref. 31 for a general derivation of the Redfield equation, and see Sec. 6.3.2.1 of this thesis for our implementation of Redfield theory in the context of vibronic systems, which follows the form used in Ref. 34). However, it cannot be cast into Lindblad form, and therefore does not guarantee total positivity of ρ_{SB} (in other words, Redfield theory can allow probabilities of the system and/or bath to become negative). Lindblad theory, in contrast, is often criticized due to the fact that it requires parametrization that is often not motivated by a microscopic treatment of the bath. However, Redfield theory has recently been shown to be mappable to Lindblad theory [35] in the range of parameters where the Redfield propagator preserves total positivity, which provides a very compelling way forward.

Chapter 3

Light-matter interaction

The previous chapter's focus was upon constructing model systems of interest, and here we turn to how to probe those systems using light (the following discussion is adapted from pages 92-103 of Ref. 14). Let us assume that we are interested in studying a system that evolves (in the absence of light) as

$$\dot{\rho} = \mathscr{L}_0 \rho,$$

where $\mathscr{L}_0$ is the Liouvillian (see Section 2.2.1) describing the time evolution of the system density operator, ρ. Let us assume that the system is in a stationary state ρ_0 of the system, such that $\mathscr{L}_0 \rho_0 = 0$ (before any optical pulses are incident on the system). Since the goal of nonlinear optical spectroscopies is often deducing the dynamics of the excited states, it is of crucial importance to understand the light-matter interaction, which is responsible for creating the initial states on each excitation manifold and probing the final states at various delay times. In general, nonlinear optical spectroscopy involves the use of L optical pulses, whose electric field we will treat classically as

$$\mathbf{E}(t) = \sum_{j=a,b,\ldots,L} \underbrace{\mathbf{e}_j \varepsilon_j(t)}_{\text{rotating}} + \underbrace{\mathbf{e}_j^* \varepsilon_j^*(t)}_{\text{counter-rotating}} \quad . \tag{3.1}$$

where $\mathbf{e}_j$ is the possibly complex polarization vector, and the complex amplitude ε_j of each pulse is defined as

$$\varepsilon_j(t) = A_j(t - t_j) e^{-i(\omega_j(t-t_j) - \mathbf{k}_j \cdot \mathbf{r} - \phi_j)}, \tag{3.2}$$

where A_j is the pulse envelope, ω_j is the central frequency, t_j is the arrival time, $\mathbf{k}_j$ is the wavevector, and ϕ_j is the phase. $\varepsilon_j(t)$ is typically referred to as the rotating term, while $\varepsilon_j^*(t)$ is typically referred to as the counter-rotating term. In the rotating-wave approximation (see Section 3.3) the rotating term excites the ket and de-excites the bra of the density matrix, while the counter-rotating term de-excites the ket and excites the bra. The rotating-wave approximation in optical spectroscopies is related to the use of rotating coordinate systems used in the context

of nuclear magnetic resonance (NMR) [36–38]. We will assume that the interaction of the pulses with the system can be well approximated using the electric dipole approximation (described below in Section 3.1), so that the interaction Hamiltonian is given by

$$H'(t) = -\boldsymbol{\mu} \cdot \mathbf{E}(t),$$

where $\boldsymbol{\mu}$ is the electric dipole operator. The full equations of motion in the presence of the electric fields are then

$$\dot{\rho} = \mathscr{L}_0 \rho + \mathscr{L}'(t)\rho, \tag{3.3}$$

where $\mathscr{L}'(t) = -\frac{i}{\hbar}\left[H'(t), \rho\right]$. After all L pulses have finished interacting with the system, it will be in a final state $\rho(t)$, which depends upon both the system dynamics $\mathscr{L}_0$, and the interaction with the pulses $H'(t)$, and can be determined by solving Eq. 3.3.

At this point the goal is to perform a measurement of the system, which will ideally give insight into the system dynamics $\mathscr{L}_0$. The most common measurement modes either look for the optical field produced by the excited system [13] or a signal proportional to the excited state population, such as fluorescence or photocurrent [39].We will briefly describe how both of these signal types are related to the phases of the optical pulses.

In the most common cases, the desired quantity will be the signal field, $\mathbf{E}_s(t)$, that is emitted by the sample. The emitted field is generated by the macroscopic dipole moment $\mathbf{P}(\mathbf{r}, t)$ of the entire sample. Up to this point we have considered only a single molecular system, but in the common experimental case there are actually many of them dispersed over a macroscopic volume. If we assume that there are many identical, non-interacting molecules that fill a volume l^3, each with a dipole operator $\boldsymbol{\mu}$, then we can calculate the macroscopic dipole moment using

$$\mathbf{P}(\mathbf{r}, t) = \sum_i \mathrm{Tr}\left[\boldsymbol{\mu}\rho_i(t)\right], \tag{3.4}$$

where ρ_i is the density matrix for a single molecule at position r_i, and the sum runs over all molecules in the volume l^3. Given the initial assumption that all of the molecules are identical, the $\mathbf{r}$ dependence comes from the $\mathbf{r}$ dependence of the optical pulses.

The bulk dipole moment $\mathbf{P}(\mathbf{r}, t)$ acts as a source term in Maxwell's equations and gives rise to far-field radiation, which is the signal that is ultimately detected. In this case, we write Maxwell's equations as

$$\nabla \times \nabla \times \mathbf{E}(\mathbf{r}, t) + \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\mathbf{E}(\mathbf{r}, t) = -\frac{4\pi}{c^2}\frac{\partial^2}{\partial t^2}\mathbf{P}(\mathbf{r}, t), \tag{3.5}$$

where Eqs. 3.3, 3.4 and 3.5 are all coupled and must be solved self-consistently throughout the volume l^3. We will, however, assume that the sample is optically thin, so that the sample does not absorb a large fraction of the incident fields, and that the generated signal fields are much weaker than the incident fields.

We can therefore solve Eq. 3.3 using only the incident electric fields from Eq. 7.3, and solve for the bulk

polarization in Eq. 3.4 using a single molecule at a single location in space. The $\mathbf{r}$ dependence of $\mathbf{P}$ is then due entirely to the accumulation of phase-factors $e^{\pm i\mathbf{k}_j\cdot\mathbf{r}}$ from the incident fields. $\mathbf{P}(\mathbf{r},t)$ is therefore proportional to $\sum_s e^{i\mathbf{k}_s\cdot\mathbf{r}}\mathbf{P}_s(r) + c.c.$, with many different signal wavevectors $\mathbf{k}_s = m_a\mathbf{k}_a + m_b\mathbf{k}_b + ... + m_L\mathbf{k}_L$, where m_j are integers. Assuming that $\frac{1}{l} \ll k_s$, constructive interference builds up for each signal E_s along the associated phase-matching direction $\mathbf{k}_s$, which allows us to pick out individual components of $\mathbf{P}$ by placing a detector in the $\mathbf{k}_s$ direction. Solutions to Eq. 3.5 (for example see pages 96-98 of Ref. 14) show that the generated signal fields obey $\mathbf{E}_s(\mathbf{r},t) \propto i\mathbf{P}_s(\mathbf{r},t)$, where the factor of i indicates that the radiated field is $\pi/2$ out of phase with respect to the bulk polarization vector. Without being concerned about the prefactors, which are generally not important for spectroscopies, we will take $i\mathbf{P}_s(t)$ to be a proxy for the emitted signal field .

That emitted field is often detected as a function of frequency by using a diffraction grating and a CCD. The field can either be detected directly, in which case the intensity, $I_s \propto E_s^2$ is measured (termed homodyne detection), or (more commonly) can be heterodyne detected using a local oscillator $\mathbf{E}_{LO} = \mathbf{e}_{LO}\varepsilon_{LO}(t) + \mathbf{e}^*_{LO}\varepsilon^*_{LO}(t)$ traveling along the $\mathbf{k}_s$ direction, giving a total measured signal

$$\begin{aligned} S_{\text{total}}(\omega) &= |\mathbf{e}_{LO}\varepsilon_{LO}(\omega) + i\mathbf{P}_s(\omega)|^2. \\ S_{\text{total}}(\omega) &= |\varepsilon_{LO}(\omega)|^2 + i\varepsilon^*_{LO}(\omega)\mathbf{e}^*_{LO}\cdot\mathbf{P}_s(\omega) - i\varepsilon_{LO}(\omega)\mathbf{e}_{LO}\cdot\mathbf{P}^*_s(\omega) + |P_s(\omega)|^2, \end{aligned} \tag{3.6}$$

where we use the convention that

$$\tilde{f}(\omega) = \int_{-\infty}^{\infty} f(t)e^{i\omega t}.$$

In perturbative spectroscopies, $|P_s(\omega)|$ is easily smaller than even a relatively weak reference pulse, such that $|\varepsilon_{LO}(\omega)| \gg |P_s(\omega)|$, so we can neglect the term $|P_s(\omega)|^2$. The intensity profile of the local oscillator can be obtained by measuring $|\varepsilon_{LO}(\omega)|^2$ alone, and then subtracting it from the total signal in Eq. 3.6 to give the desired nonlinear signal

$$S(\omega) \propto \text{Im}\left[\tilde{\varepsilon}^*_{LO}(\omega)\mathbf{e}^*_{LO}\cdot\tilde{\mathbf{P}}(\omega)\right]. \tag{3.7}$$

Eq. 3.7 shows that it is possible to determine the complex-valued $\mathbf{P}(\omega)$, if the relative phase between ε_{LO} and $\mathbf{P}$ can be controlled. Extracting the relative phase of $\mathbf{P}_s$ is not possible in a homodyne experiment.

The other common spectroscopic observable is the fluorescence or photo-current of the sample instead of the coherent emitted radiation from the molecular polarization. In such cases, we must still assume that the sample is optically thin. However, we no longer require that $\frac{1}{l} \ll k_s$, because the measured quantity does not retain the phase-matching information encoded in the pulse propagation direction $\mathbf{k}_j$. Experiments that measure fluorescence or photo-current often use collinear beam geometries where all the pulses travel along the same path. The measurements instead use phase-cycling in order to pick out only the signals that are proportional to $e^{i\phi_s}$, where $\phi_s = \sum_j m_j\phi_j$ [39–41].We describe phase-cycling in more detail in Chapter 7, where we unify phase-matching and

phase-cycling under the umbrella term phase-discrimination.

This chapter will describe the standard methods for theoretically treating the light-matter interaction, and highlight two common approximations that this thesis has the capability to move beyond: (1) the Condon approximation, and (2) the impulsive limit. We begin with a brief discussion of the electric dipole approximation. We then give a short derivation of time-dependent perturbation theory (TDPT), followed by the optical response function. This chapter will lay the groundwork for discussing nonlinear optical spectroscopies in more detail in the next chapter.

3.1 Electric Dipole Approximation

This section represents a brief summary of material drawn largely from chapter 14 of Ref. 24. All matter is composed of positively charged nuclei that are surrounded by negatively charged electrons. Neutral atoms and molecules can be thought of as complicated charge distributions that have a net zero charge. The electric field of any neutral material can be described by expanding into electric moments, with the largest moment typically being the dipole moment (followed by quadrupole, and so on). The dipole operator is defined as

$$\mu(r,R) = \underbrace{e\sum_{j=1}^{N} Z_j R_j}_{\mu_N(R)} \underbrace{-e\sum_{i=1}^{n} r_i}_{\mu_e(r)},$$

where R_j and r_i are the nuclear and electron positions, respectively, measured relative to the molecular center of mass, e is the fundamental electric charge, and Z_j are the atomic numbers of the nuclei. Most molecules do not have a permanent magnetization, and thus the magnetic moments are typically much smaller than the electric moments. Higher-order electric moments and/or magnetic moments are typically ignored unless the dipole transition strength vanishes (typically referred to as a forbidden transition). We assume that all interactions between external electric fields and molecules proceed via the dipole interaction, ignoring all higher moments.

The Born-Oppenheimer approximation (introduced in Sec 2.1) splits the total molecular wavefunction into two parts:

$$\Psi_{i,\nu}(r,R) = \psi_i(r;R)\chi_\nu^{(i)}(R),$$

where $\chi_\nu^{(i)}(R)$ is the nuclear wavefunction in the i^{th} electronic state, and $\psi_i(r;R)$ is the i^{th} electronic wavefunction, depending parametrically upon R. In the event that the Born-Oppenheimer approximation breaks down, the functions $\Psi_{i,\nu}$ can be used as a basis to represent the true wavefunctions.

We can calculate the dipole matrix elements within the Born-Oppenheimer approximation as

$$\mu_{ij,\nu\mu} = \int dR \chi_\nu^{(i)*}(R)\chi_\mu^{(j)}(R) \int dr \psi_i^*(r;R)\psi_j(r;R)\mu(r,R)$$
$$\mu_{ij,\nu\mu} = \int dR \chi_\nu^{(i)*}(R)\chi_\mu^{(j)}(R) \int dr \psi_i^*(r;R)\psi_j(r;R)\mu_e(r) + \int dR \chi_\nu^{(i)*}(R)\chi_\mu^{(j)}(R)\mu_N(R) \int dr \psi_i^*(r;R)\psi_j(r;R).$$

The term involving $\mu_N(R)$ drops out because the electronic wavefunctions are orthogonal. Thus if the Born-Oppenheimer approximation holds, only the positions of the electrons (and not the nuclei) are important for calculating the dipole matrix elements. Thus we have

$$\mu_{ij,\nu\mu} = \int dR \chi_\nu^{(i)*}(R)\chi_\mu^{(j)}(R) \underbrace{\int dr \psi_i^*(r;R)\psi_j(r;R)\mu_e(r)}_{\mu_{ij}(R)},$$

where the r-integral must be performed for each relevant nuclear configuration R. In his 1928 paper, Condon expands the dipole operator as a power series in R

$$\mu(R) = \mu_o \left(1 + \kappa(R - R_0) + O(R^2)\right),$$

where R_0 should be taken to be the value at which $\chi_\nu^{(i)*}(R)\chi_\mu^{(j)}(R)$ is maximal [42]. The optical spectroscopic literature almost universally takes only the constant term of this expansion, which has become known as the Condon approximation. [i]

There is a small body of literature that studies the implications of the breakdown of the Condon approximation within the Born-Oppenheimer approximation (see for example Ref. 43). It is worth noting that if the Born-Oppenheimer approximation breaks down, that the derivation of the object $\mu_{ij}(R)$ is no longer valid, since it relies on the separation of the electronic and nuclear wavefunctions. If this separation does not occur, then the matrix elements of the dipole operator must be calculated as

$$\mu_{\alpha,\beta} = \int dR \int dr \Psi_\alpha^*(r,R)\Psi_\beta(r,R)\mu(r,R).$$

There is a larger body of literature that investigates dipole transitions when the Born-Oppenheimer approximation fails. A coordinate dependence upon the dipole operator emerges in the BO basis when the true wavefunctions are given by a superposition of different electronic wavefunctions, with coordinate-dependent coefficients. Taking an example from page 436 of Ref. 24, if one of the true eigenstates can be approximately represented as

$$\Psi(r,R) = N\chi_\nu(R)\left(\psi_1(r;0) + CR_t\psi_2(r;0)\right), \tag{3.8}$$

[i] Note that vibrational spectroscopy typically does not involve a change in electronic state, and so $\mu_{ii,\nu\mu} = 0$ if $\mu(r,R)$ is independent of R. In the context of vibrational spectroscopy, the Condon approximation refers to the lowest-order non-vanishing term, which is $\mu(R) \approx \mu_0\kappa(R - R_0)$.

where R_t is taken to be the coordinate of the t^{th} vibration, then following the procedure outlined above, one will find that, in the BO basis, $\mu(R) = \mu_{1g} + CR_k\mu_{2g}$, where $\mu_{ig} = \langle i| \mu_e(r) |g\rangle$ are the dipole strengths associated with each BO electronic state. The types of vibronic couplings that lead to eigenstates that have the form Eq. 3.8 are typically referred to as Herzberg-Teller effects. Ref. 44 studies vibronic systems, and compares absorption spectra that arise from Herzberg-Teller couplings to absorption spectra that follow the BO approximation but include non-Condon terms in the dipole operator.

Throughout this thesis, we will assume the Condon approximation, unless stated otherwise. However, it is worth noting that the techniques presented later in this thesis are well suited to exploring the breakdown of the Condon approximation.

3.2 Time-dependent perturbation theory

In this section we outline time-dependent perturbation theory for the light-matter interaction with weak electric fields. The derivation in this section is somewhat similar to derivations in Ref. 14 on pages 26-28 for wavefunctions and pages 61-62 for density matrices (in both cases for closed systems). The most important difference is that here we do the derivation for open systems. The full Hamiltonian of the molecular system and laser pulses is given by

$$H(t) = H_0(t) + H'(t),$$

where H_0 describes the molecular system, and $H'(t) = -\boldsymbol{\mu} \cdot \mathbf{E}(t)$ is the light-matter interaction in the dipole approximation.

Since the goal of many spectroscopies is to use light as a perturbation, most spectroscopies are performed using low-power pulses. For a detailed derivation of what constitutes "low-power", see Ref. **?** , Appendix B. The basic idea is that the laser pulses should be weak enough that only a small percentage of molecules in the sample will be excited by the pulse. In this regime, we are able to take the full Hamiltonian of the molecule and optical pulses, and treat the pulses perturbatively. This has the advantage of both providing intuitive insight into the processes that are being probed, and of offering a more tractable analytical or computational problem.

The total equations of motion for the density matrix are given by

$$\dot{\rho} = -\frac{i}{\hbar}\left[H_0(t), \rho\right] - \frac{i}{\hbar}\left[H'(t), \rho\right],$$

for a closed system. The equations of motion for an open systems (see Sec. 2.2.1) include an additional term $D(t)$,

$$\dot{\rho} = -\frac{i}{\hbar}\left[H_0(t), \rho\right] + D(t)\rho - \frac{i}{\hbar}\left[H'(t), \rho\right],$$

where $D(t)$ is a possibly time-dependent superoperator that describes the dephasing and dissipation processes that result from a coupling to the bath. We define two superoperators

$$\mathscr{L}_0(t)\rho = -\frac{i}{\hbar}\left[H_0(t), \rho\right] + D(t)\rho$$

and

$$\mathscr{L}'(t) = -\frac{i}{\hbar}\left[H'(t), \rho\right],$$

which allows us to write the EOM compactly as

$$\dot{\rho} = \mathscr{L}_0(t)\rho + \mathscr{L}'(t)\rho. \tag{3.9}$$

If $\mathscr{L}'(t) = 0$, then we have

$$\dot{\rho} = \mathscr{L}_0(t)\rho, \tag{3.10}$$

which we can formally integrate, defining $\mathcal{T}_0(t, t_0)$ to be

$$\rho(t) = \mathcal{T}_0(t, t_0)\rho(t_0). \tag{3.11}$$

Now substituting Eq. 3.11 into Eq. 3.10, we have

$$\frac{\partial}{\partial t}\mathcal{T}_0(t, t_0)\rho(t_0) = \mathscr{L}_0(t)\mathcal{T}_0(t, t_0)\rho(t_0),$$

and since this must be true for any initial state $\rho(t_0)$, we have

$$\frac{\partial}{\partial t}\mathcal{T}_0(t, t_0) = \mathscr{L}_0(t)\mathcal{T}_0(t, t_0). \tag{3.12}$$

We define the interaction picture by saying that there is a state $\rho_I(t)$ such that

$$\rho(t) = \mathcal{T}_0(t, t_0)\rho_I(t). \tag{3.13}$$

Now substituting Eq. 3.13 into Eq. 3.9, we get

$$\frac{\partial}{\partial t}\left(\mathcal{T}_0(t, t_0)\rho_I(t)\right) = \left(\mathscr{L}_0(t) + \mathscr{L}'(t)\right)\mathcal{T}_0(t, t_0)\rho_I(t),$$

and by Eq. 3.12 we have

$$\begin{aligned}\mathscr{L}_0(t)\mathcal{T}_0(t,t_0)\rho_I(t) + \mathcal{T}_0(t,t_0)\frac{\partial}{\partial t}\rho_I(t) &= \mathscr{L}_0(t)\mathcal{T}_0(t,t_0)\rho_I(t) + \mathscr{L}'(t)\mathcal{T}_0(t,t_0)\rho_I(t)\\ \mathcal{T}_0(t,t_0)\frac{\partial}{\partial t}\rho_I(t) &= \mathscr{L}'(t)\mathcal{T}_0(t,t_0)\rho_I(t).\\ \frac{\partial}{\partial t}\rho_I(t) &= \mathcal{T}_0^{-1}(t,t_0)\mathscr{L}'(t)\mathcal{T}_0(t,t_0)\rho_I(t).\end{aligned}$$

For shorthand purposes, we can define

$$\mathscr{L}'_I(t) = \mathcal{T}_0^{-1}(t,t_0)\mathscr{L}'(t)\mathcal{T}_0(t,t_0) \tag{3.14}$$

which gives us

$$\frac{\partial}{\partial t}\rho_I(t) = \mathscr{L}'_I(t)\rho_I(t).$$

Integrating both sides from t_0 to t, we get

$$\int_{t_0}^{t} d\tau \frac{\partial}{\partial \tau}\rho_I(\tau) = \int_{t_0}^{t} d\tau \mathscr{L}'_I(\tau)\rho_I(\tau),$$

where by the fundamental theorem of calculus, we get

$$\begin{aligned}\rho_I(t) - \rho_I(t_0) &= \int_{t_0}^{t} d\tau \mathscr{L}'_I(\tau)\rho_I(\tau),\\ \rho_I(t) &= \rho_I(t_0) + \int_{t_0}^{t} d\tau \mathscr{L}'_I(\tau)\rho_I(\tau).\end{aligned}$$

Making repeated re-substitutions into this integral form we get

$$\rho_I(t) = \rho_I(t_0) + \sum_{n=1}^{\infty}\int_{t_0}^{t} d\tau_n \int_{t_0}^{\tau_n} d\tau_{n-1} ... \int_{t_0}^{\tau_2} d\tau_1 \mathscr{L}'_I(\tau_n)\mathscr{L}'_I(\tau_{n-1})...\mathscr{L}'_I(\tau_1)\rho_I(t_0),$$

which naturally leads to the perturbative expansion

$$\rho_I(t) = \rho_I^{(0)}(t) + \rho_I^{(1)}(t) + \rho_I^{(2)}(t) + ...$$

where $\rho_I^{(0)}(t) = \rho_I(t_0)$ and

$$\rho_I^{(n)}(t) = \int_{t_0}^{t} d\tau \mathscr{L}'_I(\tau)\rho_I^{(n-1)}(\tau),$$

for $n > 0$. Multiplying both sides by $\mathcal{T}_0(t,t_0)$, and using Eqs. 3.13 and 3.14 to move back to the Schrödinger picture we get

$$\rho^{(n)}(t) = \mathcal{T}_0(t,t_0)\int_{t_0}^{t} d\tau \mathcal{T}_0^{-1}(\tau,t_0)\mathscr{L}'(\tau)\rho^{(n-1)}(\tau).$$

We now make a change of variables $t' = t - \tau$ to arrive at

$$\rho^{(n)}(t) = \mathcal{T}_0(t, t_0) \int_0^{t-t_0} dt' \mathcal{T}_0^{-1}(t - t', t_0) \mathscr{L}'(t - t') \rho^{(n-1)}(t - t').$$

Using the composition theorem for the time-evolution operator $\mathcal{T}_0(t_3, t_2)\mathcal{T}_0(t_2, t_1) = \mathcal{T}_0(t_3, t_1)$ and the fact that the inverse time-evolution operator moves the state backwards in time, so $\mathcal{T}_0^{-1}(t_1, t_2) = \mathcal{T}_0(t_2, t_1)$, we have

$$\rho^{(n)}(t) = \int_0^{t-t_0} dt' \mathcal{T}_0(t, t - t') \mathscr{L}'(t - t') \rho^{(n-1)}(t - t').$$

Now assuming that $\rho^{(0)} = \rho(t_0)$ is a stationary state of $\mathscr{L}_0(t)$, we are free to take the limit as $t_0 \to -\infty$, which gives

$$\rho^{(n)}(t) = \int_0^{\infty} dt' \mathcal{T}_0(t, t - t') \mathscr{L}'(t - t') \rho^{(n-1)}(t - t'),$$

or equivalently

$$\rho^{(n)}(t) = \mathcal{T}_0(t, 0) \int_0^{\infty} dt' \mathcal{T}_0^{-1}(t - t', 0) \mathscr{L}'(t - t') \rho^{(n-1)}(t - t'). \tag{3.15}$$

We note here that this derivation can be followed precisely the same way for the wavefunction $|\psi\rangle$, since Eq. 3.9 has the same relevant mathematical properties as

$$\frac{d\,|\psi(t)\rangle}{dt} = -\frac{i}{\hbar} H_0(t)\,|\psi(t)\rangle - \frac{i}{\hbar} H'(t)\,|\psi(t)\rangle\,.$$

By replacing ρ with $|\psi\rangle$, $\mathscr{L}_0(t)$ with $-\frac{i}{\hbar}H_0(t)$, $\mathscr{L}'(t)$ with $-\frac{i}{\hbar}H'(t)$, and $\mathcal{T}_0(t, t_0)$ with $U_0(t, t_0)$ we arrive at

$$\left|\psi^{(n)}(t)\right\rangle = -\frac{i}{\hbar} U_0(t, 0) \int_0^{\infty} dt' U_0^{-1}(t - t', 0) H'(t - t') \left|\psi^{(n-1)}(t - t')\right\rangle, \tag{3.16}$$

and so we see that TDPT is a technique for solving time-dependent ordinary differential equations.

At this point, we have a formal solution to Eq. 3.9, which demonstrates the validity of the perturbative expansion even if H_0 and/or D are time-dependent. However, in such situations, the object $\mathcal{T}_0(t, t_0)$ cannot be expressed in a useful analytical form, and must be numerically approximated. This thesis focuses on the case where both H_0 and D are time independent, and therefore we have

$$\mathcal{T}_0(t, t_0) = e^{\mathscr{L}_0(t - t_0)},$$

which can be solved numerically by diagonalizing $\mathscr{L}_0$. This concept of diagonalizing the equations of motion lies at the heart of the numerical techniques derived in this thesis. Expressing eqs. 3.15 and 3.16 as an efficient numerical algorithm (which we call UF^2) is described in Chapter 5 for closed systems and Chapter 6 for open systems.

Briefly, we note that $\mathscr{L}_0$ can either be expressed as a superoperator acting on the $N \times N$ matrix ρ (where N is

the dimension of the Hilbert space of H_0), or as a $N^2 \times N^2$ matrix acting on the Liouville-space vector $|\rho\rangle\!\rangle$. The expectation value of A can be expressed as

$$\langle\!\langle A|\rho\rangle\!\rangle = \mathrm{Tr}[A\rho].$$

3.3 Optical Response Functions

Having derived Eq. 3.15, we are ready to discuss the optical response function (ORF). The ORF forms a conceptual foundation for NLOSs, and is a convenient way to introduce the concept of phase-discrimination, which lies at the heart of NLOSs, as well as two commonly used approximations, the rotating wave approximation (RWA) and the impulsive limit. The numerical techniques outlined in this thesis are designed around the concept of phase-discrimination, and are designed specifically to go beyond the impulsive limit. We will often invoke the RWA, but our numerical methods can do calculations beyond the RWA if desired.

The ORF is an appealing object because it is independent of the details of the electric field $\mathbf{E}(t)$, and depends only upon properties of the system, namely $\mathscr{L}_0$ and $\boldsymbol{\mu}$. We begin the discussion of ORFs by considering linear absorption of an ultrafast pulse. The total induced polarization of the molecular system can be expanded in powers of the electric field amplitude as

$$\mathbf{P}(t) = \mathbf{P}^{(1)}(t) + \mathbf{P}^{(3)}(t) + \mathbf{P}^{(5)}(t) + ...,$$

where we have assumed that the material has no permanent dipole moment (as is the case for centrosymmetric crystals and randomly oriented molecules), so that the even orders $P^{(2n)}$ cancel. Linear absorption of a macroscopic medium is related to the microscopic polarization that is linear in E, which is

$$\mathbf{P}^{(1)}(t) = \mathrm{Tr}\left[\boldsymbol{\mu}\rho^{(1)}(t)\right].$$

In order to determine the total signal that is measured in the far field, we must solve Maxwell's equations by using the linear molecular polarization $\mathbf{P}(t)$ as a source term. Following the derivation on pages 96-98 of Ref. 14, we assume that there are many identical non-interacting molecules that occupy a cube of volume l^3. We specify that the incoming electric field $\mathbf{E}(t)$ is traveling along the z-direction, and ignore the spatial variation of the field along the x and y axes.‘

From Eq. 3.15, and explicitly writing ρ as a Liouville-space vector, we have that

$$\begin{aligned}
|\rho^{(1)}(t)\rangle\!\rangle &= \mathcal{T}_0(t,0)\int_0^\infty dt'\mathcal{T}_0^{-1}(t-t',0)\mathscr{L}'(t-t')|\rho^{(0)}\rangle\!\rangle \\
&= \int_0^\infty dt'\mathcal{T}_0(t',0)\mathscr{L}'(t-t')|\rho^{(0)}\rangle\!\rangle \\
&= \frac{i}{\hbar}\int_0^\infty dt'\mathcal{T}_0(t',0)\left(\varepsilon_j(t-t')\mu^K|\rho^{(0)}\rangle\!\rangle + \varepsilon_j^\star(t-t')\mu^K|\rho^{(0)}\rangle\!\rangle - \varepsilon_j(t-t')\mu^B|\rho^{(0)}\rangle\!\rangle - \varepsilon_j^*(t-t')\mu^B|\rho^{(0)}\rangle\!\rangle\right) \\
&= \frac{i}{\hbar}\int_0^\infty dt'\varepsilon_j(t-t')\mathcal{T}_0(t',0)\left(\mu^K|\rho^{(0)}\rangle\!\rangle - \mu^B|\rho^{(0)}\rangle\!\rangle\right) - \frac{i}{\hbar}\int_0^\infty dt'\varepsilon_j^*(t-t')\mathcal{T}_0(t',0)\left(\mu^B|\rho^{(0)}\rangle\!\rangle - \mu^K|\rho^{(0)}\rangle\!\rangle\right),
\end{aligned}$$

where $\mu^K|\rho\rangle\!\rangle \equiv \mu\rho$ and $\mu^B|\rho\rangle\!\rangle \equiv \rho\mu$. Since both μ and $\rho^{(0)}$ are Hermitian operators, $\mu^K|\rho^{(0)}\rangle\!\rangle$ is the Hermitian conjugate of $\mu^B|\rho^{(0)}\rangle\!\rangle$. The time-evolution operator $\mathcal{T}_0(t,t_0)$ must preserve the Hermiticity of the density operator, and thus if A and B are Hermitian conjugates, then $\mathcal{T}_0(t,t_0)A$ and $\mathcal{T}_0(t,t_0)B$ must also be Hermitian conjugates. Thus we have

$$\rho^{(1)}(t) = \frac{i}{\hbar}\int_0^\infty dt'\varepsilon(t-t')\mathcal{T}_0(t',0)\left(\mu^K\rho^{(0)} - \mu^B\rho^{(0)}\right) + h.c.$$

At this point it is informative to expand ε_j using Eq. 3.2 to get

$$\rho^{(1)}(t) = \frac{i}{\hbar}e^{i(\mathbf{k}_j\cdot\mathbf{r}+\phi_j)}\int_0^\infty dt' A_j(t-t')e^{-i\omega_j(t-t')}\mathcal{T}_0(t',0)\left(\mu^K\rho^{(0)} - \mu^B\rho^{(0)}\right) + h.c.$$

$$\rho^{(1)}(t) = \frac{i}{\hbar}e^{i(\mathbf{k}_j\cdot\mathbf{r}+\phi_j)}\mathcal{T}_0(t,0)\int_0^\infty dt'\left(A_j(t-t')e^{-i\omega_j(t-t')}\mathcal{T}_0^{-1}(t-t',0)\mu^K\rho^{(0)} - A_j(t-t')e^{-i\omega_j(t-t')}\mathcal{T}_0^{-1}(t-t',0)\mu^B\rho^{(0)}\right) + h.c.$$

As long as the envelope $A_j(t)$ is wide enough to resolve the optical carrier frequency, the term involving $\mu^B\rho^{(0)}$ will be much smaller than the term involving $\mu^K\rho^{(0)}$. Neglecting the $\mu^B\rho^{(0)}$ term for this reason is called the rotating-wave approximation (RWA). To see why it is negligible, let us assume that $D = 0$ (so we have a closed system), and therefore

$$\mathcal{T}_0(t,t_0)|\rho\rangle\!\rangle = U_0(t,t_0)\rho U^\dagger(t,t_0).$$

Expressing $|\rho^{(0)}\rangle\!\rangle$ in the eigenbasis of H_0, we have

$$|\rho^{(0)}\rangle\!\rangle = \sum_i p_i\,|i\rangle\langle i|\,,$$

where p_i are the thermal occupation probabilities of each eigenstate of H_0, where $H_0|i\rangle = \hbar\omega_i|i\rangle$. Then

$$\mu^K|\rho^{(0)}\rangle\!\rangle = \mu\sum_i p_i\,|i\rangle\langle i|$$

$$\mu^K|\rho^{(0)}\rangle\!\rangle = \sum_{i,k} p_i\mu_{ki}\,|k\rangle\langle i|$$

and

$$\mathcal{T}_0^{-1}(t,0)\mu^K|\rho^{(0)}\rangle\!\rangle = \sum_{i,k} p_i\mu_{ki}e^{i(\omega_k-\omega_i)t}\,|k\rangle\langle i|$$

while

$$\mathcal{T}_0^{-1}(t,0)\mu^B|\rho^{(0)}\rangle\!\rangle = \sum_{i,k} p_i\mu_{ik}e^{i(\omega_i-\omega_k)t}\,|i\rangle\langle k|\,.$$

The dipole matrix elements μ_{ki} connect states that are separated by energies on the order of the optical gap. In general optical pulses will be used that have a carrier frequency similar to the optical gap of the material system, so that $\omega_j \approx \omega_k - \omega_i$ for the relevant i,k in the system under study. Thus the integrand involving $\mu^K|\rho^{(0)}\rangle\!\rangle$ has the form

$$\begin{aligned} A_j(t-t')e^{-i\omega_j(t-t')}\mathcal{T}_0^{-1}(t-t',0)\mu^K|\rho^{(0)}\rangle\!\rangle &= A_j(t-t')\sum_{i,k} p_i\mu_{ki}e^{i(\omega_k-\omega_i-\omega_j)(t-t')}\,|k\rangle\langle i| \\ &\approx A_j(t-t')\sum_{i,k} p_i\mu_{ki}e^{i\omega_{\text{slow}}(t-t')}\,|k\rangle\langle i| \end{aligned}$$

where $\omega_{\text{slow}} \ll \omega_j$. In contrast

$$\begin{aligned} A_j(t-t')e^{-i\omega_j(t-t')}\mathcal{T}_0^{-1}(t-t',0)\mu^B|\rho^{(0)}\rangle\!\rangle &= A_j(t-t')\sum_{i,k} p_i\mu_{ik}e^{i(\omega_i-\omega_k-\omega_j)(t-t')}\,|i\rangle\langle k| \\ &\approx A_j(t-t')\sum_{i,k} p_i\mu_{ij}e^{i\omega_{\text{fast}}(t-t')}\,|i\rangle\langle k|\,, \end{aligned}$$

where $\omega_{\text{fast}} \approx 2\omega_j$. Since $\omega_{\text{slow}} \ll \omega_{\text{fast}}$, the integrand involving $\mu^B|\rho^{(0)}\rangle\!\rangle$ (the counter-rotating term) will oscillate many more times in the region of integration as compared with the integrand involving $\mu^K\rho^{(0)}$ (the rotating term). Thus upon integration, the rotating term will make a much larger contribution to $\rho^{(1)}$ than will the counter-rotating term. The only circumstance under which this would not hold is if $A_j(t)$ is too short to resolve ω_{fast}, in which case it is possible for the counter-rotating term to make a significant contribution. In general, the rotating term excites the ket and de-excites the bra, whereas the counter-rotating term excites the bra and de-excites the ket. By invoking the RWA, we have

$$\rho^{(1)}(t) = \int_0^\infty dt' A_j(t-t')\frac{i}{\hbar}e^{i(\mathbf{k}_j\cdot\mathbf{r}+\phi_j)}e^{-i\omega_j(t-t')}\mathcal{T}_0(t',0)\mu^K|\rho^{(0)}\rangle\!\rangle + h.c.,$$

The linear polarization is

$$\begin{aligned} P^{(1)}(t) &= \langle \mu \rho^{(1)} \rangle \\ &= \int_{-\infty}^{\infty} dt' A_j(t-t') e^{-i\omega_j(t-t')} e^{i(\mathbf{k}_j \cdot \mathbf{r} + \phi_j)} \frac{i}{\hbar} \theta(t') \langle\!\langle \mu | \mathcal{T}_0(t',0) \mu^K | \rho^{(0)} \rangle\!\rangle + c.c. \\ &= \int_{-\infty}^{\infty} dt' A_j(t-t') e^{-i\omega_j(t-t')} e^{i(\mathbf{k}_j \cdot \mathbf{r} + \phi_j)} R^{(1)}(t') + c.c., \end{aligned}$$

where $R^{(1)}$ is the linear response function of the material. If we assume that $A_j(t)$ is short with respect to all of the relevant time-scales contained in $R^{(1)}(t)$, but is simultaneously long enough to resolve the optical carrier frequency ω_j, then we can approximate that $A_j(t) \approx \delta(t)$, and thus that

$$P^{(1)}(t) \approx e^{i(\mathbf{k}_j \cdot \mathbf{r} + \phi_j)} R^{(1)}(t).$$

This is called the impulsive limit.

The conceptual advantage of the response function is thus two-fold. First, it is in some sense the "ideal" signal. The goal of nonlinear optical spectroscopy experiments is often to use pulses to deduce the properties of the system, contained in $R^{(1)}$, not the convolution of the system properties with the pulse shape, contained in $P^{(1)}$. Given that the response functions are often complicated enough by themselves, and that connecting the results of experiments with theoretical predictions are often hard enough, without including the details of the pulse shapes, it is tempting to ignore the pulse shapes. There has been some work done trying to understand in general how finite pulse effects modify the true signal, as compared with the response function. However, when comparing experiment to a theoretical model, the pulse shapes are often ignored.

The exponential factor in front of $R^{(1)}(t)$ is the phase-discrimination condition (see the beginning of this Chapter for a discussion of phase-discrimination). In this case, the phase-discrimination condition shows that the sample will radiate the signal $E_s \propto iP^{(1)}(t)$ only along the direction $\mathbf{k}_j$. This comes as no surprise when measuring linear absorption. We expect to measure the linear absorption signal by placing a detector behind the sample, in the path of the original laser. However, for non-linear signals, this will not necessarily be the case. Instead, signals will radiate in many different directions defined by $\mathbf{k}_s = \sum_j m_j \mathbf{k}_j$, where m_j are integers.

Moving on to the nonlinear case, we write down the third-order response function as

$$R(t_1,t_2,t_3) = \left(\frac{i}{\hbar}\right)^3 \theta(t_1)\theta(t_2)\theta(t_3) \sum_{O,O',O''=K,B} \underbrace{\langle\!\langle \mu | \mathcal{T}_0(t_3,t_2) \mu^{O''} \mathcal{T}_0(t_2,t_1) \mu^{O'} \mathcal{T}_0(t_1,0) \mu^O | \rho^{(0)} \rangle\!\rangle}_{R_\alpha(t_3,t_2,t_1)},$$

where we suppress the (3) superscript, since the presence of three time arguments implies that this is a third-order response function. Since there are 2 options for each μ^O, there are $2^3 = 8$ different third-order response functions. However, these 8 can be broken down into 4 complex conjugate pairs of response functions. Following Mukamel's

numerical identification on page 120 of Ref. 14, we have

$$R_1(t_1,t_2,t_3) = \left(\frac{i}{\hbar}\right)^3 \theta(t_1)\theta(t_2)\theta(t_3)\langle\!\langle \mu|\mathcal{T}_0(t_3,t_2)\mu^B\mathcal{T}_0(t_2,t_1)\mu^B\mathcal{T}_0(t_1,0)\mu^K|\rho^{(0)}\rangle\!\rangle$$
$$R_2(t_1,t_2,t_3) = \left(\frac{i}{\hbar}\right)^3 \theta(t_1)\theta(t_2)\theta(t_3)\langle\!\langle \mu|\mathcal{T}_0(t_3,t_2)\mu^B\mathcal{T}_0(t_2,t_1)\mu^K\mathcal{T}_0(t_1,0)\mu^B|\rho^{(0)}\rangle\!\rangle$$
$$R_3(t_1,t_2,t_3) = \left(\frac{i}{\hbar}\right)^3 \theta(t_1)\theta(t_2)\theta(t_3)\langle\!\langle \mu|\mathcal{T}_0(t_3,t_2)\mu^K\mathcal{T}_0(t_2,t_1)\mu^B\mathcal{T}_0(t_1,0)\mu^B|\rho^{(0)}\rangle\!\rangle$$
$$R_4(t_1,t_2,t_3) = \left(\frac{i}{\hbar}\right)^3 \theta(t_1)\theta(t_2)\theta(t_3)\langle\!\langle \mu|\mathcal{T}_0(t_3,t_2)\mu^K\mathcal{T}_0(t_2,t_1)\mu^K\mathcal{T}_0(t_1,0)\mu^K|\rho^{(0)}\rangle\!\rangle.$$

Third-order spectroscopies are often performed using 3 pulses, so that

$$E(t) = \sum_{j=1}^{3} \varepsilon_i(t) + \varepsilon_i^*(t).$$

In this case, the third-order polarization field that builds up in the sample is

$$P^{(3)}(t) = \int dt_3 \int dt_2 \int dt_1 E(t-t_3)E(t-t_3-t_2)E(t-t_3-t_2-t_1)R(t_1,t_2,t_3).$$

Since $E(t)$ is composed of 6 terms, and there are 8 response functions, this in principle involves 8×6^3 terms. Half of these will be overall complex conjugate partners. At this point, it is of crucial importance to remember the exponential phase factors $e^{i(\mathbf{k}_j\cdot\mathbf{r}+\phi_j)}$. If we are only interested in the signal proportional to $e^{i(\mathbf{k}_s\cdot\mathbf{r}+\phi_s)}$, where $\mathbf{k}_s = \pm\mathbf{k}_1\pm\mathbf{k}_2\pm\mathbf{k}_3$ and $\phi_s = \pm\phi_1\pm\phi_2\pm\phi_3$, then there are only $8\times 3! = 48$ terms, since we are only interested in terms where, after expanding E^3, the product $\varepsilon_1^{(*)}\varepsilon_2^{(*)}\varepsilon_3^{(*)}$ is present, where ε_j contributes $+\mathbf{k}_j$, and ε_j^* contributes $-\mathbf{k}_j$. For example, if we are interested in the rephasing signal, which is proportional to $e^{i(\mathbf{k}_s\cdot r+\phi_s)}$, where $\mathbf{k}_s = -\mathbf{k}_1+\mathbf{k}_2+\mathbf{k}_3$ and $\phi_s = -\phi_1+\phi_2+\phi_3$, then we are only interested in terms where the product $\varepsilon_1^*\varepsilon_2\varepsilon_3$ appears. Note that $\varepsilon_1^*\varepsilon_2\varepsilon_3 R_1^{(3)}$ is not the complex conjugate of $\varepsilon_1^*\varepsilon_2\varepsilon_3 R_1^{(3)*}$, so all 8 response functions are needed for this counting argument. In addition, if all three pulses are coincident in time, there are 6 different unique ways of permuting the electric-field indices of the product $\varepsilon_1^*(t-t_3)\varepsilon_2(t-t_3-t_2)\varepsilon_3(t-t_3-t_2-t_1)$, that give the desired phase-discrimination condition.

By invoking the RWA, and by using the fact that the thermal occupation of electronic excitations is negligible at room temperature, we can neglect 32 of the 48 terms, leaving us with 16 well-known terms. An example of one of the 32 neglected terms is

$$\int dt_3 \int dt_2 \int dt_1 \varepsilon_3(t-t_3)\varepsilon_2(t-t_3-t_2)\varepsilon_1^*(t-t_3-t_2-t_1)\left(\frac{i}{\hbar}\right)^3 \theta(t_1)\theta(t_2)\theta(t_3)\langle\!\langle \mu|\mathcal{T}_0(t_3,t_2)\mu^K\mathcal{T}_0(t_2,t_1)\mu^B\mathcal{T}_0(t_1,0)\mu^K|\rho^{(0)}\rangle\!\rangle.$$

In the RWA the term $\varepsilon_1^*\mu^K$ de-excites the ket, and assuming that $|\rho^{(0)}\rangle\!\rangle$ has no population in an excited electronic state, that term will be zero.

When calculating nonlinear optical spectroscopy, a final assumption is often invoked, which is called time-

ordering. When the three pulses arrive in the order implied by their index, and do not overlap in time, only 3 out of 16 terms contribute to the rephasing signal. Those three terms are

$$P^{(3)}_{\text{rephasing}}(t) = \int dt_3 \int dt_2 \int dt_1 \varepsilon_3(t-t_3)\varepsilon_2(t-t_3-t_2)\varepsilon_1^*(t-t_3-t_2-t_1)\left(R_1^*(t_1,t_2,t_3)+R_2(t_1,t_2,t_3)+R_3(t_1,t_2,t_3)\right).$$

In the impulsive limit, the other 13 overlap terms only contribute when two or more pulses are perfectly coincident. Otherwise, those terms are strictly zero in the impulsive limit, and are thus ignored in impulsive calculations.

Each of the terms that contribute to a given polarization field $P^{(n)}(t)$ can be visualized using double-sided Feynman diagrams. All 16 diagrams that contribute to third-order signals are displayed in Figure 7.1, and the three that contribute to the rephasing signal when the pulses are time-ordered are surrounded by a black box in that figure. Double-sided Feynman diagrams are constructed using two vertical lines, one representing the ket-side , and the other representing the bra-side, of the density matrix. Arrows that point up and to the right are referred to as rotating terms (and are associated with ε_j), while arrows that point up and to the left are called counter-rotating terms (and are associated with ε_j^*).

Double-sided Feynman diagrams used for describing NLOSs originally developed from single-sided diagrams describing the absorption and emission of photons as pioneered by Feynman (see Ref. 45). Single-sided diagrams were used to describe resonant linear and off-resonant nonlinear light-matter interactions (for example in Ref. 46). Double-sided diagrams were introduced initially to show the evolution of the wavefunction and its complex-conjugate (in other words a density matrix) during multiple-photon absorption/emission processes, for the purpose of calculating the transition probabilities between different excited states (for example see Ref. 47). Ref. 48 uses this double-sided diagram approach to study resonant Raman processes. Perhaps the first formal use of double-sided diagrams for the organization of different NLOS signals occurs in Ref. 49, although the diagrams shown there look visually dissimilar to the modern double-sided diagrams (rather than two separate lines, the diagrams in Ref. 49 involve circles, indicating that the density matrix begins in a population, ends in a population, and in between the bra and ket sides evolve along different trajectories). Double-sided Feynman diagrams have become ubiquitous throughout the field.

Although inconvenient, performing calculations using finite-pulses at 3rd-order is certainly doable. Although some studies of finite pulse effects consider only the impact of including finite pulse shapes in the time-ordered diagrams [17, 19, 20], in order to self-consistently account for finite pulse effects, one must include all of the relevant overlap diagrams as well [13, 18]. Calculating 16 diagrams instead of 3 may be tedious, but it is well within reason. However, when one considers higher-order spectroscopies, the number of necessary diagrams grows rapidly. For example, the total fifth-order polarization field is given by

$$P^{(5)}(t) = \int dt_3 \int dt_2 \int dt_1 E(t-t_5)E(t-t_5-t_4)E(t-t_5-t_4-t_3)E(t-t_5-t_4-t_3-t_2)E(t-t_5-t_4-t_3-t_2-t_1)R(t_1,t_2,t_3,t_4,t_5),$$

where

$$R(t_1,t_2,t_3,t_4,t_5) = \left(\frac{i}{\hbar}\right)^5 \theta(t_1)\theta(t_2)\theta(t_3)\theta(t_4)\theta(t_5) \sum \underbrace{\langle\!\langle \mu | \mathcal{T}_0(t_5,t_4)\mu^{O''''}\mathcal{T}_0(t_4,t_3)\mu^{O'''}\mathcal{T}_0(t_3,t_2)\mu^{O''}\mathcal{T}_0(t_2,t_1)\mu^{O'}\mathcal{T}_0(t_1,0)\mu^{O}|\rho^{(0)}\rangle\!\rangle}_{R_\alpha(t_1,t_2,t_3,t_4,t_5)},$$

where O, O', O'', O''', O'''' can be either K or B. For the 5^{th}-order case, there are $2^5 = 32$ different combinations of μ^K/μ^B, with half being complex conjugate partners, as above in the case of the 3^{rd}-order response functions. If $E(t)$ is a sum over 5 separate pulses, we will have $5! = 120$ different pulse permutations that yield a given phase-matching condition associated with a signal that goes to first-order in each pulse. This will yield 3840 different diagrams. Invoking the RWA in combination with taking the initial state to be entirely in the ground state brings the total down to 960 (according to the automated diagram generation process, introduced in Chapter 7). Nearly 1000 diagrams really requires automation, which is what we do in Chapter 7.

Although the numerical algorithms we call UF^2 (derived in chapters 5 and 6) can be used to calculate response functions, the motivation of those algorithms, and this thesis, is to provide a computationally convenient and efficient method for moving beyond the impulsive limit by directly calculating $P^{(n)}(t)$ for any n and any phase-discrimination condition. We achieve this by writing an efficient algorithm, UF^2, to calculate $\rho^{(n)}$ given $\rho^{(n-1)}$, and by writing an algorithm to automatically generate all of the necessary diagrams associated with any phase-discrimination condition. We will discuss phase-discrimination, and how it is achieved experimentally, further in the next section.

Chapter 4

Examples of Nonlinear Optical Spectroscopies

Having developed the foundational theory needed to calculate NLOSs in Chapters 2 and 3, we pause here to give a brief (non-exhaustive) overview of optical NLOSs, to give a flavor of some of signals that are measured. This chapter gives several examples of NLOSs using a simple model system, described in Sec. 4.1. All of the simulations presented below are performed using ultrafast spectroscopy suite (UFSS), the software suite that is described in Chapters 5-9, and which is available at github. UFSS is designed to calculate any order perturbative NLOS due to any number of optical pulses, interacting resonantly with the bound states of a model system. UFSS does not treat chemical reaction dynamics or photoionization, which involve coupling to continua of unbound states, or non-resonant processes, which involve coupling to virtual states. With some very minor development, UFSS could be adapted to simulate IR vibrational spectroscopies as well.

All of the spectra presented in this chapter can be produced in the Jupyter notebook `ThesisChapter1.ipynb` (which is included in the GitHub repository). This notebook gives a guide of how to use UFSS to calculate linear and transient absorption, 2 pulse photon echo, and 2D photon echo simulations, including inhomogeneous broadening and isotropic averaging. Although the default is set to study a simple system with only 2 coupled excited electronic states (outlined in Sec. 4.1), the notebook includes options to include additional electronic states, as well as vibrational modes. This chapter, and the associated Jupyter notebook, thus serve as an introductory user's guide to UFSS. This chapter does not include any novel contributions to the field. It is the final chapter that lays the foundation for the novel results presented in all subsequent chapters of this thesis.

4.1 Example System

In order to describe NLOSs, it is helpful to discuss simulations of spectra using a simple example system. NLOSs directly probe electronic states in molecules by virtue of the fact that the optical pulses couple to the system via the electronic dipole of the system (see Sec. 3.1). In many cases they indirectly probe vibrational degrees of freedom, due to the fact that the vibrational degrees of freedom are coupled to the electronic degrees of freedom, and it is usually of critical importance to understand those couplings in order to correctly interpret spectra. However, the intuition about NLOSs is often founded on considering just the electronic states. Here we consider a system consisting of two coupled subunits, each with a ground state and a single electronic excited state. Each subunit is an identical two-level system (2LS), and their singly excited states interact via a simple Coulomb interaction term. The Hamiltonian of the system is

$$H = \hbar \left(\sum_{i=1,2} \omega_{\text{gap}} c_i^\dagger c_i + 0.5\omega_0 \left(c_1^\dagger c_2 + c_1 c_2^\dagger \right) \right), \tag{4.1}$$

where $c_i^\dagger$ creates an excitation on site i and $\hbar\omega_{\text{gap}}$ is the energetic gap between the ground and excited electronic state of each isolated subunit, and ω_0 sets the energy scale for the excited states in the model. Note that Eq. 4.1 corresponds to the case $s = 2, k = 0$ of Eq. 2.5, with $V_i(\mathbf{q}) = \hbar\omega_{\text{gap}}$ and $V_{12}(\mathbf{q}) = 0.5\omega_0$ (note the lack of $\mathbf{q}$-dependence, due to the fact that we do not include any explicit vibrational degrees of freedom in this model).

The eigenstates of H include one ground state, one doubly excited state, and two singly excited states that have energies $\omega_{\text{gap}} - 0.5\omega_0$ and $\omega_{\text{gap}} + 0.5\omega_0$. Although we wish to keep this model as simple as possible, closed systems exhibit unrealistic, perfectly sharp absorption lines. We therefore include dephasing and relaxation processes via the secular Redfield equations, which we describe in more detail in Sec. 6.3.2.1. Briefly, we model the coupling of the electronic states to a bath using an Ohmic spectral density of coupling strength $0.05\omega_0$, and a Lorentz-Drude cut-off function with cut-off frequency ω_0, and temperature $T_{th} = \hbar\omega_0/k_B$, where k_B is Boltzmann's constant (with the exception of Fig. 4.6(a-c), which shows results for $k_B T_{th} = 0.25\omega_0$). To put this into perspective, at room temperature $k_B T_{th} = 25$ meV, corresponding to a frequency $f_0 = \omega_0/2\pi$ of about 200 cm^{-1} or 6 THz, and a period of about 160 fs. The coupling to the bath implies a relaxation rate of about 250 fs, with a corresponding dephasing time of 500 fs, both of which can be seen in Fig. 4.6(e).

We model the light-matter interaction (introduced in Chapter 3) using the electric dipole approximation (see Sec. 3.1), giving us the interaction Hamiltonian

$$H'(t) = -\mu \cdot \mathbf{E}(t).$$

We treat the electric fields classically, as a sum over L pulses

$$\mathbf{E}(t) = \sum_{i=a,b,\ldots,L} \mathbf{e}_i \varepsilon_i(t) + \mathbf{e}_i \varepsilon_i(t),$$

For the simulations in this section, all of the pulses are linearly polarized in the same direction ($\mathbf{e}_i = \mathbf{x}$ for all i), the electric field amplitude is modeled as a Gaussian

$$\varepsilon_i(t) = \frac{1}{\sqrt{2\pi}\sigma} e^{-t^2/2\sigma^2 - i(\omega_c - \mathbf{k}_i \cdot \mathbf{r} + \phi_i)},$$

where the standard deviation of each pulse is taken to be $\sigma = \omega_0^{-1}$, and the pulses are centered on the optical gap of an isolated 2LS, so that $\omega_c = \omega_{\text{gap}}$ (except for Fig. 4.1(d), where the pulse is centered at $\omega_c = \omega_{\text{gap}} - 0.5\omega_0$). All pulses are taken to be identical, except that the pulses may travel along different directions ($\mathbf{k}_i$ are not the same for all pulses). If we take $f_0 = 6$ THz, then we have $\sigma = 25$ fs, or a full-width half maximum (FWHM) of about 60 fs, which is easily attainable in current experimental labs.

We model the light-matter interaction using the electric dipole approximation, where both molecules are taken to have equal dipole magnitude, but are perpendicular to one another, so that

$$\mu_1 = (c_1^\dagger + c_1)\mathbf{x}$$
$$\mu_2 = (c_2^\dagger + c_2)\mathbf{y},$$

and the total dipole operator for the system is

$$\mu = \mu_1 \otimes \mathbb{I}_2 + \mathbb{I}_1 \otimes \mu_2$$

All simulations in the following section include isotropic averaging assuming a random, static, distribution of molecules in the lab frame (see Appendix B of Ref. 50.

4.2 Linear and Transient Absorption

We begin exploring NLOSs by first considering linear absorption spectroscopy. Figure 4.1(a) shows an experimental schematic of linear absorption, which usually involves scanning a tunable light source (either a variable CW laser or a white light source with a monochromator) across many different wavelengths. The laser is directed into a sample, and then a detector on the other side of the sample. Fig 4.1(b) shows an alternative experimental schematic for acquiring part of the absorption spectrum in a single shot by using an ultrashort laser pulse, a diffraction grating, and a detector. The absorption spectrum acquired using an ultrashort pulse is equivalent to the CW absorption

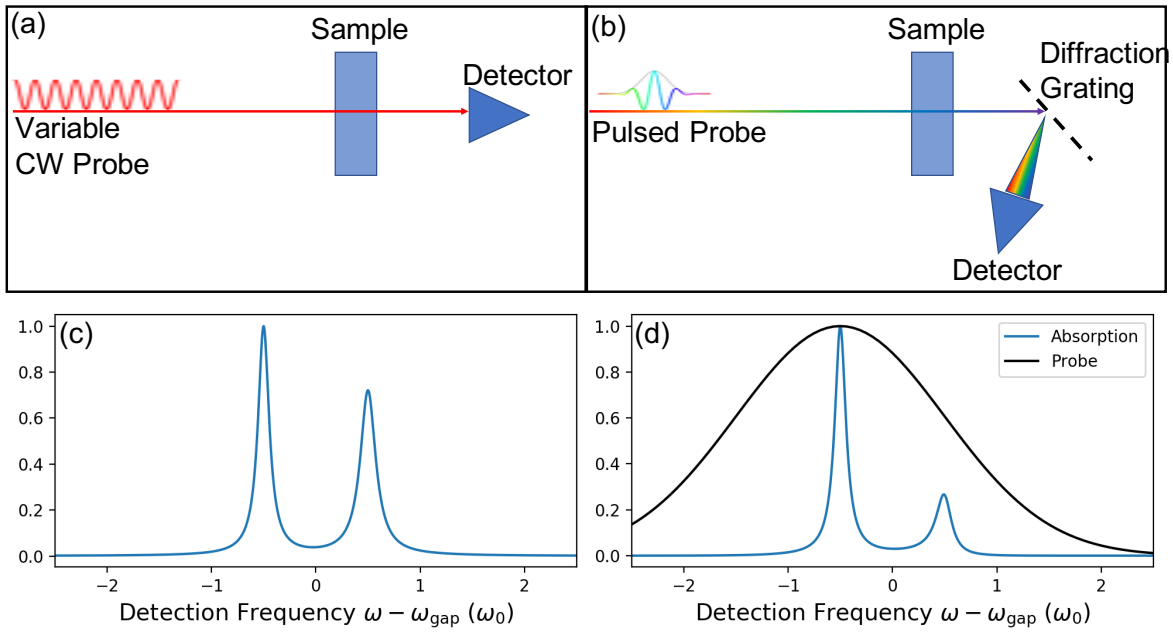

Figure 4.1: (a) Schematic of CW absorption spectrum, obtained using a tunable CW laser or a white-light source and a monochromator. (b) Schematic of ultrafast absorption spectrum. (c) Linear absorption spectrum obtained by using a variable continuous wave laser. Note the higher-energy state has a larger homogeneous line-width, due to the relaxation rate that couples it to the lower energy state. The integrated peaks are roughly equivalent, as both transitions have equal dipole strength. (d) Linear absorption spectrum obtained by using a single ultrafast probe pulse, centered on the lower energy absorption peak. The probe amplitude is displayed in black. Each peak is weighted by the probe intensity, showing how the bandwidth of ultrashort pulses can effect the spectra that are measured.

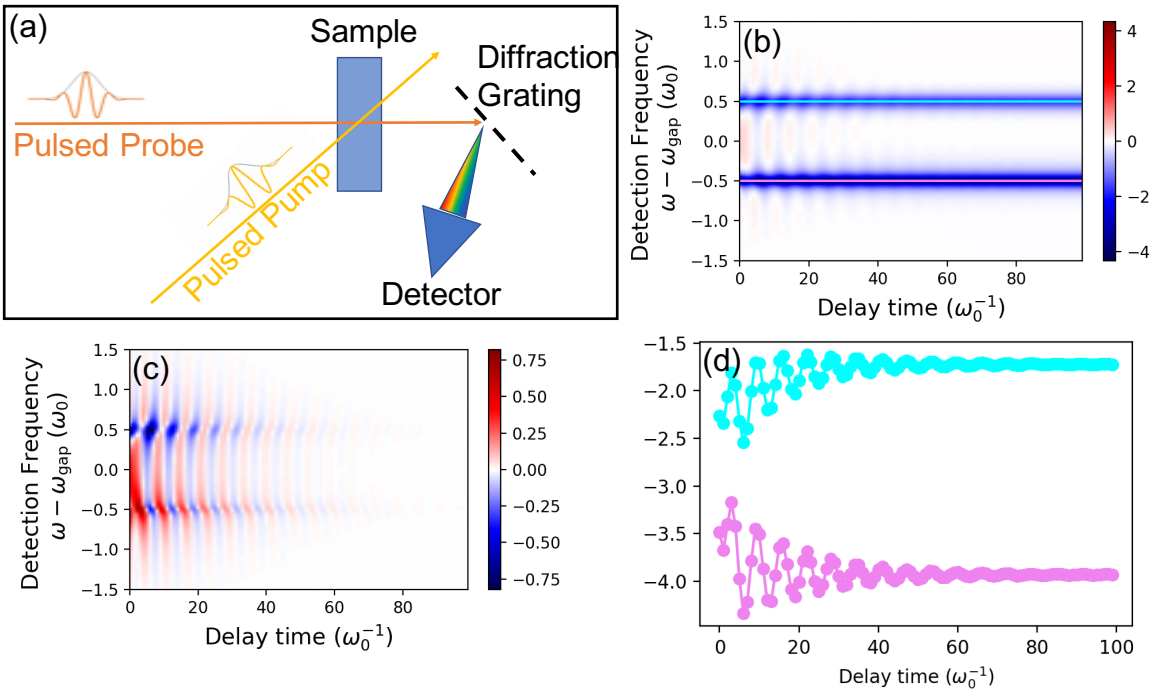

Figure 4.2: (a) Schematic of transient absorption (TA) spectroscopy. The absorption of the probe is measured for each delay time between the pump and probe pulses. (b) Change in probe absorption due to the pump pulse for a simple dimer of two identical two-level systems with coupling $0.5\omega_0$ (see Eq. 4.1). Note that the negative signal at $\omega_{\text{gap}} \pm 0.5\omega_0$ indicates that the absorption of the probe pulse is diminished at the linear absorption peaks. (c) Change in probe absorption spectra with transients only (static change in the probe absorption at a delay time of $1000\omega_0^{-1}$ has been subtracted away). (d) Line-cuts of TA spectra shown in (b) at $\omega = 0.5\omega_0 + \omega_{\text{gap}}$ (cyan) and $\omega = -0.5\omega_0 + \omega_{\text{gap}}$ (purple), which reveals that the magnitude of the signal at $\omega = 0.5\omega_0 + \omega_{\text{gap}}$ is smaller than the magnitude of the signal at $\omega = -0.5\omega_0 + \omega_{\text{gap}}$. Note that the oscillations are consistent with a dephasing time of $20\omega_0^{-1}$, and if we take $\omega_0 = 2\pi \times 6$ THz, this dephasing rate would be 500 fs. The relaxation time is somewhat difficult to see, but is also visible with a time constant of $10\omega_0^{-1}$.

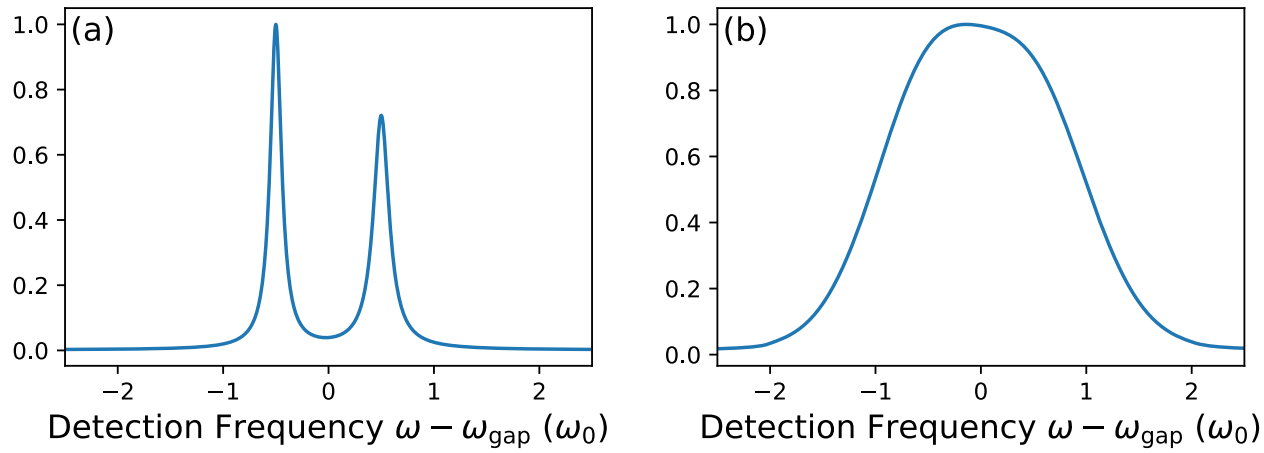

Figure 4.3: (a) Absorption spectrum for two coupled 2LS's (same as Fig. 4.1(a)). (b) Absorption spectrum for the same system as (a), but including inhomogeneous broadening of the optical gap with standard deviation $0.4\omega_0$. We do not include inhomogeneous broadening over the coupling or the individual site energies .

spectrum, multiplied by the intensity of the laser pulse as a function of frequency. Fig. 4.1 shows the CW absorption spectrum in panel (c), and the ultrafast absorption spectrum in panel (d), along with the the pulse amplitude profile (note that in this simulation only, the pulse is centered at $\omega_c = \omega_{\text{gap}} - 0.5\omega_0$, instead of $\omega_c = \omega_{\text{gap}}$, in order to highlight the fact that the absorption peaks are weighted by the laser pulse).

The ultrafast absorption spectrum is an important starting point for understanding transient absorption (TA), which uses the ultrafast absorption of a probe beam as a reference. TA begins by measuring the absorption spectrum using a probe pulse, just as in Fig. 4.1(b). Then, as shown in Fig. 4.2(a), a pump pulse is added, traveling along a different spatial path. The time delay between the pump pulse and the probe pulse is varied. For each value of the delay time, the absorption of the probe is measured. By subtracting the baseline probe absorption, measured in the absence of the pump, the change in absorption is plotted as a function of delay time. Fig. 4.2(b) shows a TA simulation. Panel (c) shows the TA simulation after subtracting out the static absorption (technically obtained at a delay time of $1000\omega_0^{-1}$), and panel (d) shows two linecuts from panel (b). The dynamics of the excited states of the system are visible as oscillations that decay due to the coupling to the bath. The average strength of signal at $\omega_{\text{gap}} - 0.5\omega_0$ increases slightly, while the strength of the signal at $\omega_{\text{gap}} + 0.5\omega_0$ decreases slightly as a function of the delay time, due to thermal relaxation of the initial excitation created by the pump pulse.

4.3 Inhomogeneous Broadening and 2 Pulse Photon Echo

The example spectra from Section 4.2 do not include inhomogeneous broadening, which is the variation of molecular properties, both in time and in space, due to interactions with the environment. In Fig. 4.3(b) we show the absorption spectrum when including a Gaussian distribution of static disorder for the optical gap of the coupled 2LS. We take the standard deviation of the disorder to be $\sigma = 0.4\omega_0$. As can be seen, the individual peaks have

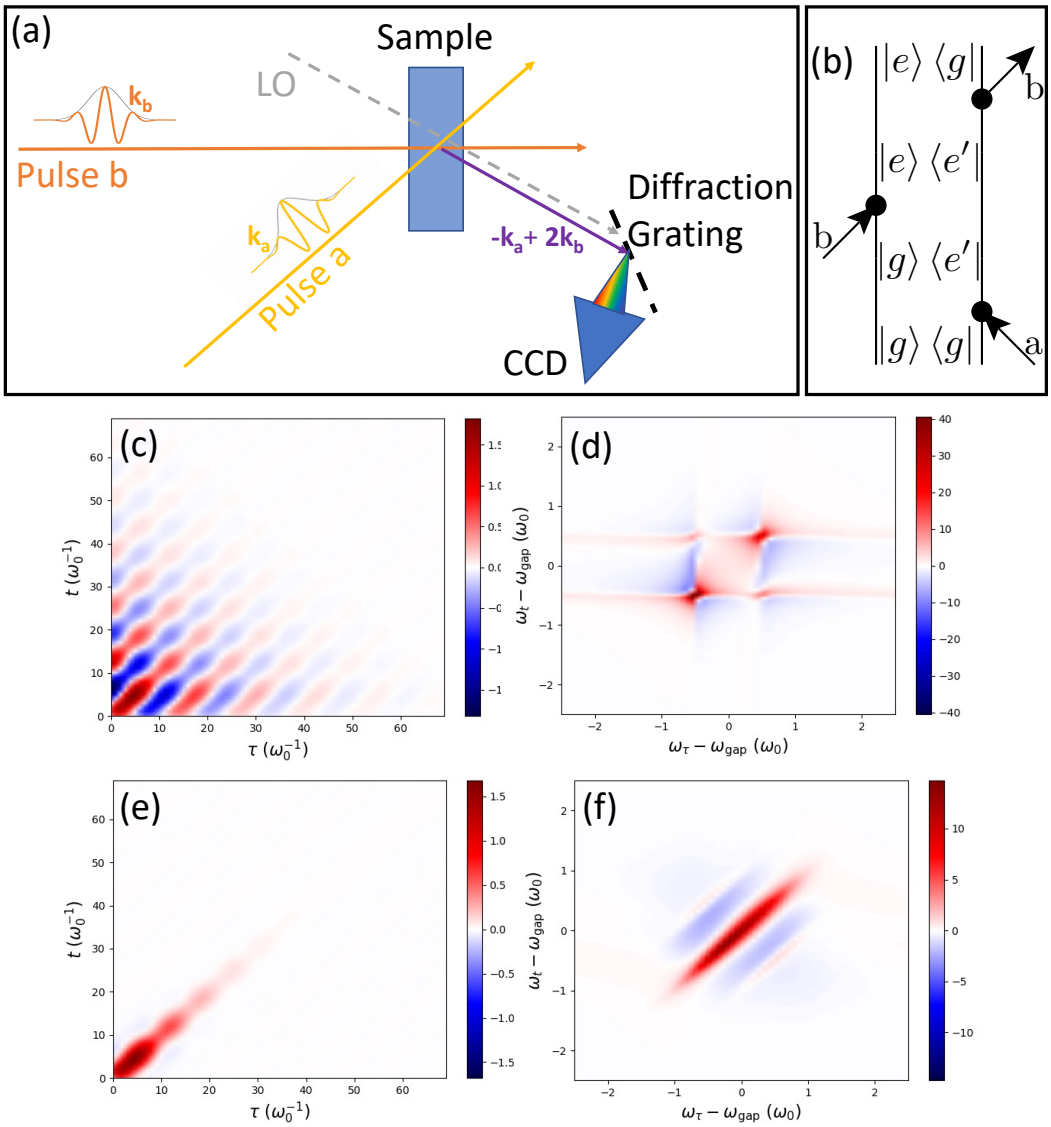

Figure 4.4: (a) Schematic of 2 pulse photon echo experiment. Signal is measured in the $-k_a + 2k_b$ direction. The delay time between pulse a and b is called the coherence time, and is usually labeled τ. The signal field is usually mixed with a reference pulse called a local oscillator. Data can be collected using a diffraction grating and CCD, as shown, in order to obtain data as a function of ω_t, the detection frequency, for each τ. Alternatively the diffraction grating may be replaced by a photodiode. In this case both the coherence time, τ, and rephasing time, t, must be varied independently. The rephasing time is defined as the delay between pulse b and the local oscillator. In some cases, the experiment is done for $t = \tau$, so as to extract the homogeneous decay time. (b) Feynman diagram describing the stimulated emission process in the 2 pulse photon echo experiment (see Chap. 7 for more details on Feynman diagrams). (c) 2 pulse photon-echo simulation without inhomogeneous broadening shown as a function of τ and t. (d) 2D Fourier transform of (c), showing the 2 pulse photon-echo simulation as a function of the Fourier conjugate variables ω_τ and ω_t. (e) 2 pulse photon-echo simulation with inhomogeneous broadening shown as a function of τ and t. Note that the decay along the line $\tau = t$ is unchanged by inhomogeneous broadening of the optical gap. (f) 2D Fourier transform of (e), showing the 2 pulse photon-echo simulation as a function of the Fourier conjugate variables ω_τ and ω_t. The diagonal shows the inhomogeneous broadening modeled as Guassian disorder with standard deviation $\sigma = 0.4\omega_0$, with a similar shape to the linear absorption spectrum in Fig. 4.3(b), while the anti-diagonal shows the pure homogeneous linewidth, with similar shape to the linewidths seen in Fig. 4.3(a), controlled by the relaxation process with a time constant of $10\omega_0^{-1}$, modeled using secular Redfield theory.

been smeared out to the point that they are no longer distinguishable.

A technique called the two pulse photon echo (2PE) can remove the inhomogenous broadening and reveal the homogeneous line shapes. The photon echo experiment is the optical analogue to the spin-echo experiment in NMR. When working with spin-1/2 systems, it is straight-forward to create $\pi/2$ and π pulses. A $\pi/2$ pulse moves the entire ensemble into the excited state. A π pulse moves the entire ensemble from ground, to excited, and back to ground (or from excited, down to ground, and back to excited). However, in many cases the 2 pulse photon echo experiment has been carried out in the perturbative regime, owing to the difficulty of creating $\pi/2$ and π pulses for optical excitation of complex molecular systems [13]. However, there is ongoing interest in using two pulse photon echoes for all-optical memory and multimode quantum memories [51] in solid-state systems. Highly localized excitations of quantum dots can be treated as 2LSs, and have been shown to exhibit Rabi oscillations [52]. Rabi oscillations have also been studied in CdTe/(Cd,Mg)Te quantum well systems [53].

For our purposes, we will focus on the perturbative regime, so that we are not considering $\pi/2$ and π pulses. Rather, we are considering the case where only a small percentage of the ensemble interacts with the incident fields. 2PE can be used to obtain the homogeneous linewidth from condensed phase samples. 2PE can be done using two pulses, just as with transient absorption. However, for the photon echo experiment, the detector is placed in a different location, as shown in Fig. 4.4a, in order to measure the signal in the $-\mathbf{k}_a + 2\mathbf{k}_b$ direction. In the case of TA, the change in absorption of the second pulse is measured, and so the nonlinear signal that is probed is generated in the same direction as the second pulse. The second pulse thus acts both as the final interaction with the system, as well as the reference pulse used when detecting the non-linear signal. The fact that the second, probe, pulse is also the reference pulse, leads to the expression "self-heterodyne detection", which is distinguished from homodyne detection because TA is sensitive to the amplitude of the nonlinear signal, rather than the intensity (see the beginning of Chapter 3 for a discussion of heterodyne and homodyne detection). In the case of 2PE, however, the detector is placed in a different location, and is therefore sensitive to a different type of nonlinear signal. Since detectors measure intensity, a detector in the $-\mathbf{k}_a + 2\mathbf{k}_b$ will only be sensitive to the intensity of the nonlinear signal, and not the amplitude. However, by using a reference pulse with a controlled phase, one can measure the complex amplitude of the 2PE signal, if desired.

We can use time-dependent perturbation theory (TDPT) to gain insight into NLOSs by visualizing each of the perturbative calculations in the form of a Feynman diagram (introduced in Chap. 3 and discussed in more detail in Chap. 7). One of the two diagrams that contribute to the 2PE signal is shown in Fig. 4.4(b) (this figure shows the stimulated emission term; the other term is called excited-state absorption, and can be seen in Figure 7.1). Briefly, when looking at a Feynman diagram, time advances moving upwards. Arrows that point towards the right are associated with the rotating term, and thus with the $e^{i\mathbf{k}_j\cdot\mathbf{r}}$ phase term, while arrows that point towards the left are associated with the counter-rotating term, and thus with the $e^{-i\mathbf{k}_j\cdot\mathbf{r}}$ phase term. The left vertical line represents the ket-side of the system density matrix, while the right vertical line represents the bra-side of the system density

matrix. In the rotating-wave approximation, arrowheads that touch the bra or ket vertical lines cause electronic excitation to the bra or ket, respectively, while arrow tails that touch the bra or ket vertical lines cause electronic de-excitation of the bra or ket, respectively. Thus we see from Fig. 4.4(b) that the stimulated emission term involves excitation of the bra-side by the first pulse. The second pulse both excites the ket-side, and de-excites the bra-side of the density matrix. This gives insight into the system evolution between pulses, as described in the following paragraph.

In the photon echo experiment, the first pulse excites the bra-side of the density matrix, causing the system to evolve in the coherence $|g\rangle\langle e'|$. Considering only the unitary evolution of the system, this means that each individual molecule evolves at a frequency $e^{i(\omega_{e'}-\omega_g)\tau} \approx e^{i\omega_{\text{gap}}\tau}$ (approximate due to the fact that there are really two states with eigenenergies of $\omega_{\text{gap}} \pm 0.5\omega_0$). We can understand the first pulse as setting up a bulk polarization field in the sample that oscillates at the same frequency as the optical transition frequency of the molecules, ω_{gap}, and decays due to the homogeneous linewidth of the molecule. In the case of dissolved molecules, each molecule in the solution is surrounded by a slightly different environment. The difference in environment introduces inhomogeneous broadening, and causes ω_{gap} to be slightly different for each nominally identical molecule in the solution. The bulk polarization field rapidly loses spatial coherence because ω_{gap} acts as a randomly distributed variable across the solution. When a second pulse, delayed from the first by a coherence time, τ, arrives, the polarization of each individual molecule has acquired a phase $e^{i\omega_{\text{gap}}\tau}$. The second pulse causes the polarization field of each molecule to "flip", because after interaction with the second pulse, the system is in a coherence $|e\rangle\langle g|$, and so picks up phase $e^{i(\omega_g-\omega_e)t} \approx e^{-i\omega_{\text{gap}}t}$. Thus the polarization field then begins accumulating a phase $e^{-i\omega_{\text{gap}}t}$, where t is the time measured from the second pulse. When $t = \tau$, the polarization field of each individual molecule has returned to a total phase of $e^0 = 1$, and the bulk polarization is once more spatially in phase (it has rephased). At this point, the signal field is maximum. Thus if the detected signal field is plotted as a function of t for any value of τ, the field shows a peak at $t = \tau$ (so at a time 2τ after the first pulse). This peak is often called an "echo". The magnitude of the echo is diminished be a factor of $e^{-2\tau\gamma}$, where γ is the homogeneous dephasing rate. This echo effect does not occur throughout the sample. The sample experiences a bulk rephasing effect only for detection along the particular direction $-\mathbf{k}_a + 2\mathbf{k}_b$. By placing a detector along this direction, the echo can be detected, typically using a reference pulse called a local oscillator.

Figure 4.4(a) shows a schematic for an experimental setup that could measure the 2PE. Fig. 4.4 shows the simulated signal field emitted along the $-\mathbf{k}_a + 2\mathbf{k}_b$ direction as a function of the interpulse delay, τ, and the time after the second pulse, t, for (b) no inhomogeneous broadening and (d) with inhomogeneous broadening. The echo effect is seen in Fig. 4.4(d) as a slowly decaying signal along the $\tau = t$ line. The decay rate along the $\tau = t$ line reveals the homogeneous dephasing rate for a single dissolved molecule, without the ensemble averaging inherent in an absorption spectrum. The Fourier transform along this line thus reveals the homogeneous lineshape, which is seen along the antidiagonal of Fig. 4.4(e), which shows the 2D Fourier transform as a function of the conjugate

variables ω_τ and ω_t. Along the diagonal ($\omega_t = \omega_\tau$) we see the inhomogeneous lineshape, analogous to the lineshape in Fig. 4.3(b). In contrast, Fig. 4.4(c) shows the 2D Fourier transform of panel (b), without any inhomogeneous broadening.

Figs. 4.4(c) and (e) are called frequency-frequency correlation plots, and are somewhat analogous to CW double resonance spectra [13]. In a CW double resonance experiment the change to the absorption spectrum is measured due to excitation at a particular frequency. By plotting the change to the absorption spectrum at many different excitation frequency, a frequency-frequency 2D spectrum is obtained. In the 2PE experiment, the ω_τ-axis is analogous to the excitation axis, and the ω_t is the detection axis. The difference, of course, is that in the 2PE experiment, all frequencies are simultaneously excited and detected, due to the use of ultrashort pulses that contain many frequencies. In the absence of inhomogeneous broadening, we see two diagonal peaks along the line $\omega_t = \omega_\tau$. The diagonal line is associated with excitation at some frequency, followed by detection at the same frequency. There are also two cross-peaks visible, which correspond to excitation at one frequency, and detection at another frequency. The cross-peaks are thus sensitive to couplings between states. If, for example, the system under study consisted of two uncoupled 2LSs (perhaps two molecules dissolved in the same solution), one with an optical gap of $\omega_{\text{gap}} - 0.5\omega_0$, the other with an optical gap $\omega_{\text{gap}} + 0.5\omega_0$, we would see only the diagonal peaks.

4.4 2D Photon Echo and 2D Electronic Spectroscopy

The addition of a third pulse to the 2PE experiment yields 2D photon echo (2DPE) spectroscopy (shown schematically in Fig. 4.5(a)), which allows for the study of the excited state dynamics. The delay between the second and third pulses, T, called the population time, reports on the electronic excited- and ground-state dynamics, just as with TA. The advantage over TA lies in the frequency resolution of the excitation, provided by taking the Fourier transform with respect to the coherence time τ, the delay between the first two pulses. 2DPE yields frequency-frequency correlation plots as a function of T, allowing for the study of frequency-dependent excitation and detection as the excited system evolves. If one integrates along ω_τ (which is equivalent to measuring at $\tau = 0$), one recovers the TA spectrum, along with additional phase information, which is provided by the reference, or local oscillator, pulse. At a population time $T = 0$, the 2DPE signal is the same as the 2PE signal. This similarity can be seen by comparing the Feynman diagrams from Fig. 4.5(b) and Fig. 4.4(b). When the third pulse arrives at the same time as the second pulse, the stimulated emission signal is the same for both the 2DPE and the 2PE.

2DPE is an information rich experiment, but can be very difficult to interpret. Intuition that comes from the impulsive limit, and a neglect of vibrational motion, leads to very simple interpretation of spectra. Fig. 4.6(a) and (b) show 2DPE simulations with $k_B T_{th} = 0.25\omega_0$, where k_B is Boltzmann's constant and T_{th} is the temperature, and for population times $T = 0$ and $T = 25\pi/\omega_0$, respectively. Both panels clearly show four peaks, two along the diagonal $\omega_t = \omega_\tau$, and two off diagonal peaks, called cross-peaks. Fig. 4.6(c) shows each of those peaks as

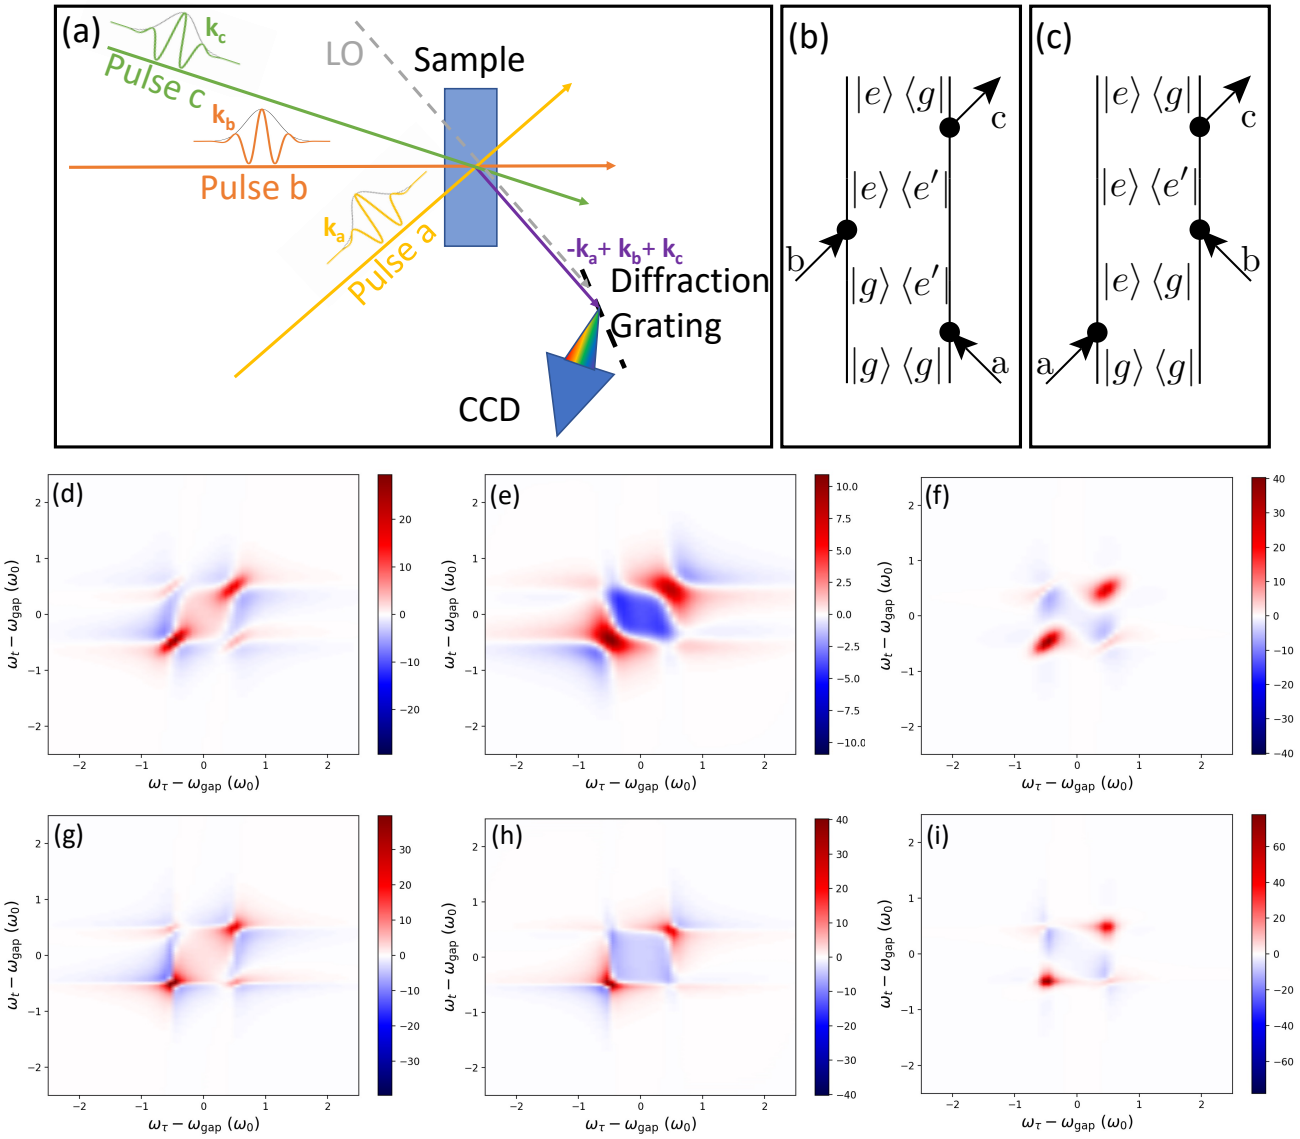

Figure 4.5: (a) Schematic of the rephasing 2D photon echo experiment, where the delay between pulses a and b is the coherence time τ, and the delay between pulses b and c is the population time T. The signal field is detected as a function of frequency ω_t, for many values of τ and T. The Fourier transform along the τ axis yields the excitation frequency ω_τ. The non-rephasing signal is measured by changing the order of the first two pulses (in which case T is the delay between pulse a and pulse c), or maintaining the pulse order and placing a detector in the $\mathbf{k}_a - \mathbf{k}_b + \mathbf{k}_c$ direction (which preserves the definition of T as the delay between pulse b and pulse c). Adding the rephasing and non-rephasing spectra yields the 2D electronic spectroscopy (2DES) signal. Feynman diagrams describing the stimulated emission process in (b) the rephasing experiment and (c) the non-rephasing experiment (see Chap. 7 for more details on Feynman diagrams). Rephasing 2DPE at population time $T = 0$ (d) with inhomogeneous broadening and (g) without inhomogeneous broadening. Non-rephasing 2DPE at population time $T = 0$ nonrephasing 2DPE, (e) with inhomogeneous broadening and (h) without inhomogeneous broadening. Sum of rephasing and non-rephasing at population time $T = 0$ (f) with inhomogeneous broadening and (i) without inhomogeneous broadening.

a function of population time T. Figs. 4.6(d-e) are directly analogous to (a-c), except that (d-e) are simulated for $k_B T_{th} = \hbar\omega_0$. In Fig. 4.6(c) we see that the lower energy diagonal peak is nearly constant as a function of T, while the higher-energy diagonal peak decays as a function of T. This is interpreted as energy relaxing from the higher-energy state to the lower-energy state, and is a common signature seen in 2DPE, displaying excitations thermalizing and flowing "downhill" in the energy landscape. We see a corresponding rise in the signal at the lower cross-peak at $\omega_\tau = \omega_{\text{gap}} + 0.5\omega_0$, $\omega_t = \omega_{\text{gap}} - 0.5\omega_0$, which corresponds to excitation at high energy and detection at low energy. Both cross-peaks oscillate with frequency ω_0 as a function of T, showing that the system was excited into a coherent superposition of two states separated by an energy $\hbar\omega_0$. In this model system, we of course know that the oscillations represent a coherent superposition of two different excitations, which only exists due to the coupling J included in our model. However, a model that includes a single 2LS linearly coupled to a vibrational mode of vibrational frequency ω_0 would also exhibit cross-peak oscillations at ω_0 as a function of T.

In fact, the difficulty of distinguishing between vibrational coherences and electronic coherences in 2DPE experiments lies at the heart of a contentious debate that was sparked when Ref. 3 claimed that oscillations with a dephasing time of about 700 fs observed in their experiment were due to electronic coherences in the Fenna-Matthews-Olson (FMO) complex. However, others, including Ref. 4 pointed out that the oscillations observed in Ref. 3 could also be interpreted as arising from vibrational coherences. Ultimately it was concluded that the original oscillatory signals reported in Ref. 3 were vibrational in nature, although electronic coherences that dephase rapidly (dephasing time of about 100 fs) can also be observed in FMO [5, 6].

In the limit of no relaxation or dephasing processes, where the linewidths of the excitation and detection peaks are perfect delta functions, the diagonal peaks in 2DPE for the current model system would not oscillate as a function of T. However, Fig. 4.6(c) shows that the diagonal peaks do in fact exhibit small oscillations, which are caused by the shoulders of the homogeneous linewidth of the cross-peaks. This effect is more pronounced in Fig. 4.6(e) due to the higher temperature in the simulation, which increases the linewidths of the peaks. It is also important to remember that when vibrational modes are included in the system Hamiltonian, the diagonal peaks will oscillate due to ground-state bleach effects. Both of these facts mean that diagonal peaks often exhibit some oscillations. We also see in Fig. 4.6(e) that the lower-energy diagonal peak decreases in energy, and that the upper cross-peak shows a corresponding increase as a function of T. When $k_B T_{th} = \hbar\omega_0$, we expect that excitation at the lower energy eigenstate will lead to some population in the higher energy state due to thermalization effects.

One drawback of 2DPE is that the lineshapes of the peaks are a combination of absorptive and dispersive. The purely absorptive and purely dispersive lineshapes can be obtained by adding what is called the non-rephasing spectra, which can be detected using the same setup outlined in Fig. 4.5(a) by swapping the order of pulses a and b, or by moving the detector to the $\mathbf{k}_a - \mathbf{k}_b + \mathbf{k}_c$ direction. The delay between the first two pulses is still denoted τ, and the population time T is defined as the time after the arrival of the first two pulses, and before the third. As shown in the Feynman diagram in Fig. 4.5(c), in the non-rephasing experiment, after the arrival of pulse a, the

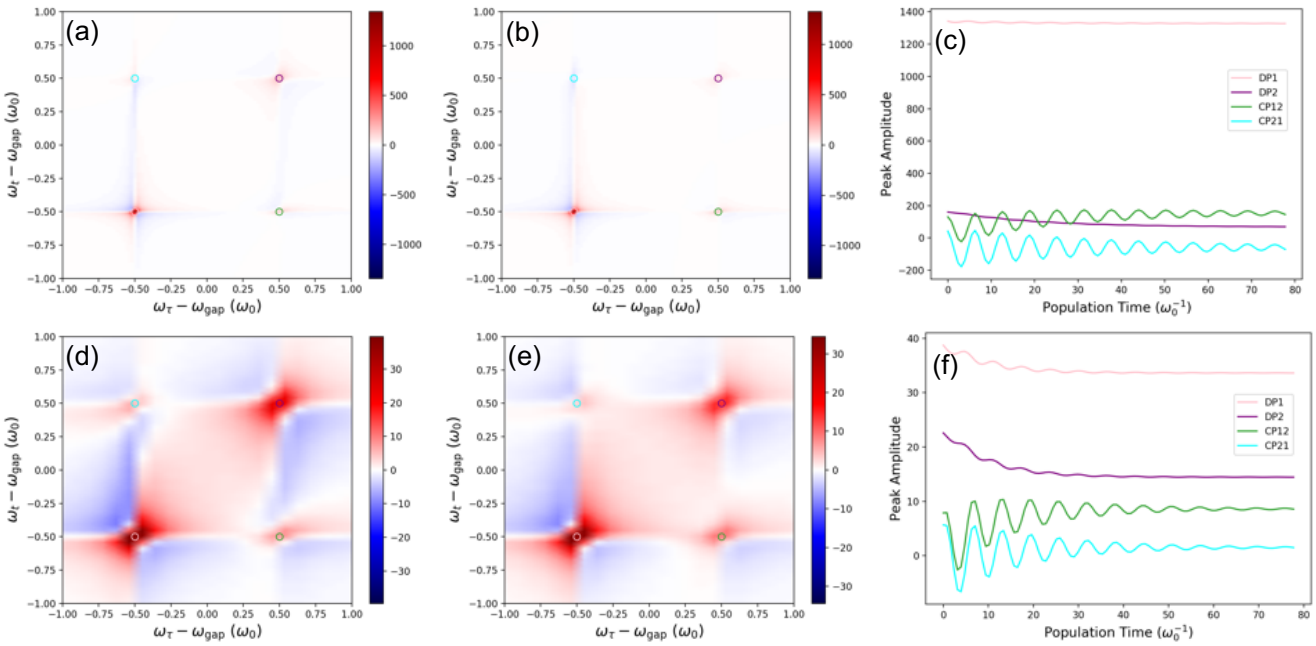

Figure 4.6: Rephasing 2DPE simulations for (a) $k_BT_{th} = 0.25\omega_0$ and 0 population time, (b) $k_BT_{th} = 0.25\omega_0$ and $25\pi/\omega_0$ population time, (d) $k_BT_{th} = \omega_0$ and 0 population time, (e) $k_BT_{th} = \omega_0$ and $25\pi/\omega_0$ population time. Panels (c) and (f) show linecuts at the indicated circles for the diagonal and cross peaks for $k_BT_{th} = 0.25\omega_0$ and $k_BT_{th} = \omega_0$, respectively.

system evolves in the coherence $|e\rangle\langle g|$ and so the bulk polarization evolves as $e^{i(\omega_g-\omega_e)\tau} \approx e^{-i\omega_{\text{gap}}\tau}$ (approximate due to the fact that there are really two states with eigenenergies of $\omega_{\text{gap}} \pm 0.5\omega_0$). As with 2DPE, during the interval T the excited molecules evolve in the excited states. After pulse c, the system is once again evolving in the coherence $|e\rangle\langle g|$, and so the bulk polarization does not flip (as it does for 2PE/2DPE), but again evolves as $e^{-i\omega_{\text{gap}}t}$. For this reason, the bulk polarization decays at roughly the rate $e^{-i\omega_{\text{gap}}(\tau+t)}$, and so the maximum non-rephasing signal is always at $t = 0$, regardless of τ (and therefore the non-rephasing signal does not exhibit a photon echo). Figs. 4.5(e) and (h) show the non-rephasing signal at population time $T = 0$ with and without inhomogeneous broadening, respectively.

By adding the real parts of the rephasing and non-rephasing signal, one obtains the purely absorptive signal (one can similarly obtain the purely dispersive signal by adding the imaginary parts), shown in Figs. 4.5(f) and (i). The sum of the rephasing and non-rephasing signal is often called 2D electronic spectroscopy (2DES). Fig. 4.5(f) shows that with inhomogeneous broadenings, the 2DES retains some of the narrow linewidth along the anti-diagonal. Fig. 4.5(f) shows that the diagonal peaks are much rounder in the absence of inhomgeneous broadening. It is worth noting that in the model system under study the non-rephasing signal has static cross-peaks, while the oscillations occurring at ω_0 as a function of T appear in the diagonal peaks (though other effects can contribute oscillations to the cross-peaks of the non-rephasing signal, such as vibrational coherences from the ground-state bleach signal). Adding the non-rephasing to the rephasing signal thus adds oscillatory signal to the diagonal peaks.

Polarization control of the pulses can be used to enhance cross-peak signals [54, 55], which is an idea borrowed

from vibrational spectroscopy (see pages 93-100 of Ref. 56). However, as pointed out in Ref. 57, this polarization control can also greatly enhance pulse overlap effects, and suggest that long pulse tails can make the interpretation of cross-peak oscillations very difficult for short population times T. We discuss the importance of overlap diagrams in the context of higher-order spectroscopies in Chap. 7.

Chapter 5

Closed UF^2

Reproduced from Journal of Chemical Physics 150 214105 (2019), with the permission of AIP Publishing

Original title: *Numerical Method for Nonlinear Optical Spectroscopies: Ultrafast Ultrafast Spectroscopy*

Peter A. Rose and Jacob J. Krich

Original Abstract: We outline a novel numerical method, called Ultrafast Ultrafast (UF^2) spectroscopy, for calculating the n^{th}-order wavepackets required for calculating n-wave mixing signals. The method is simple to implement, and we demonstrate that it is computationally more efficient than other methods in a wide range of use cases. Resulting spectra are identical to those calculated using the standard response function formalism but with increased efficiency. The computational speed-ups of UF^2 come from (a) non-perturbative and costless propagation of the system time-evolution (b) numerical propagation only at times when perturbative optical pulses are non-zero and (c) use of the fast Fourier transform convolution algorithm for efficient numerical propagation. The simplicity of this formalism allows us to write a simple software package that is as easy to use and understand as the Feynman diagrams that organize the understanding of n-wave mixing processes.

Notes:

- Since the publication of this paper, two errors in Table 5.2 have been corrected. In the original publication, the definition of l_{SEM} was incorrectly given as $N_{c,SEM}/N_{f,SEM}$, and the definition of l_{DEM} was incorrectly given as $N_{c,DEM}/N_{f,DEM}$. These definitions have been corrected in this thesis, so that Table 5.2 now shows that l_{SEM} is defined as $N_{f,DEM}/N_{c,SEM}$ and l_{DEM} is defined as $N_{f,DEM}/N_{c,DEM}$.

- Added a brief discussion that did not appear in the original publication comparing a simulation done using UF^2 to a simulation done using a different piece of software (which uses the split-operator technique). The added discussion is in section 5.3.3.

- Since the publication of this paper, an error in Section 5.3.3 has been identified and corrected. We previously stated that UF^2 calculates spectra at least 10^5 times faster than the code from Ref. 21, for systems of the

type described in Section 5.4. This has been corrected to say that UF^2 calculates spectra at least 10^3 times faster for a particular test case, detailed in the text.

5.1 Introduction

Ultrafast nonlinear optical spectroscopies, in the perturbative light-matter limit, are powerful tools for elucidating details about the electronic structure and ultrafast dynamics of optically active systems. Interpreting such spectra often requires understanding what signals would be produced by a range of parametrized system Hamiltonians; one then seeks the best agreement with experimental results by varying system parameters such as energy levels, couplings, dephasing rates, and more. This kind of fitting, as in Ref. 10, involves repeatedly rederiving the spectroscopic signals as system parameters change, which can be computationally expensive.

Simulations of nonlinear spectroscopies can be particularly challenging when the optical pulse durations are similar to relevant timescales in the system dynamics, especially when one must consider the pulses overlapping in time. The response-function formalism, described further below, provides a powerful way to understand nonlinear spectra and can be computationally efficient when considering the limit of impulsive optical pulses, those with durations shorter than any relevant system dynamics [14]. For spectroscopies that rely on varying pulse durations, simulating those variations can be computationally expensive, limiting the range of systems and pulses one can study [58].

The effects of finite pulse durations can be studied using perturbative [15, 17–19, 21, 30, 58–66] and nonperturbative [67–69] methods [70]. Except in special cases, where the dynamics can be solved analytically [15, 18, 19, 58], generic methods numerically solve the time-dependent Schrodinger equation to determine the system response and spectroscopic observables. In the frequently studied cases of electronic excitations coupled to vibrations or optical cavities, without breaking of chemical bonds, this integration is performed efficiently in a basis of electronic and vibrational/optical excitations, with respect to which the system Hamiltonian is generally highly sparse. The numerical integration is frequently performed using Runge-Kutta (RK) methods [66, 71–76], with other frequently used quantum dynamics packages using Adams-Bashforth [77] or Bulirsch-Stoer [62, 78] methods. An alternative approach, especially valuable in cases with bond breaking, where a continuum of vibrational states become relevant, uses a real- and Fourier-space pseudospectral representation of the wavefunction, frequently using the split-operator method to propagate the system dynamics [64, 69, 79, 80].

There has recently been an increased interest in understanding the role of finite-pulse-duration effects in nonlinear spectroscopies, and there are several analytic methods for the propagation of the states required for calculating 4-wave mixing signals using particular shapes of finite pulses [18, 19, 58]. Each of these methods assumes a particular, idealized pulse shape, such as a Gaussian [19, 58] or Lorentzian [18] profile, and some include pulse overlap effects. These methods require that the eigenvalues and eigenvectors of the unperturbed system are known, allowing use of

analytic results for the time-dependent perturbation theory (TDPT), which can be summed over all required states.

We present a complementary numerical technique we call Ultrafast Ultrafast (UF^2) spectroscopy, which combines features of the analytic and full numerical integration approaches. UF^2 assumes that the eigenvalues and eigenvectors of the unperturbed system are known but solves the integrals arising from TDPT numerically and can therefore treat arbitrary pulse shapes. UF^2 works for arbitrary perturbative order n, including pulse-overlap effects. It is a fast implementation of the standard results in perturbative nonlinear spectroscopy in the electric dipole approximation[14]. The version we present here is for closed systems.[i] It assumes that the system Hamiltonian is time-independent and has a finite relevant eigenbasis. When the pulses are finite, we perform the time propagation using the convolution theorem and the fast Fourier transform (FFT) and therefore benefit from the speed of the FFT. Working in the energy eigenbasis allows non-perturbative and costless time evolution of the system at times when the optical pulses are negligible, and UF^2 also solves the TDPT integrals during the pulses with such a small cost that they may not be computationally limiting even for reasonably large system Hamiltonians. These speed-ups give a dramatic improvement over algorithms that involve numerical integration of the wavefunctions or density matrices for each time-step.

In many cases, the limiting step of UF^2 is evaluating the expectation value of the dipole operator after the pulses have passed. The computational cost of UF^2 is dominated by the cost of matrix-vector multiplication, and so benefits from all the speed of linear algebra optimization on modern computers. As it relies on diagonalizing the Hamiltonian, it is clear that at some system size UF^2 will cease to be a competitive technique.

Conventional wisdom holds that diagonalization of the system Hamiltonian is too expensive to be worthwhile for prediction of nonlinear spectroscopies [70]. Especially for sparse matrices, as occur in e.g., vibronic systems, full or partial diagonalizations to attain the relevant eigenvalues and eigenvectors can be relatively efficient [81]; more importantly, if one desires to study 2D spectra with a large number of different pulse delays, the diagonalization cost must be paid only once. Further, one can consider the effects of different pulse shapes, durations, and polarizations; various dipole coupling matrices; and thermal or rotational averaging, all without needing to rediagonalize. In many such cases, the diagonalization cost can be insignificant compared to the cost of generating the spectra, even for relatively large systems.

In Section 5.2 we derive the mathematical framework of UF^2, show how it is used, and describe some of the techniques that increase its efficiency. As with any perturbative technique, it is most efficient where the rotating wave approximation and phase matching can be assumed (see Section 5.2.4), but it does not require these assumptions. In Section 5.3 we compare the computational cost of UF^2 to two alternatives: a Runge-Kutta-based direct propagation method, and a split-operator pseudospectral method, where we focus on example vibronic Hamiltonians. For sufficiently large systems, the direct propagation methods become more efficient than UF^2, but we show that for Hamiltonians with dimensions up to $O(10^4)$, which includes a large number of widely studied

[i] UF^2 can be extended to work with open systems, but we do not do so here.

cases, UF^2 outperforms the direct propagation methods. In the frequently studied case of systems with dimension less than 100, UF^2 can be two orders of magnitude faster than direct propagation methods. We demonstrate how UF^2 enables computationally efficient studies of the effects of varying system parameters and optical pulse shapes in Section 5.4, where we show transient absorption (TA) spectra for a system with a Hamiltonian of dimension 28 and demonstrate how UF^2 can be used to perform rapid parameter sweeps.

UF^2 is not only efficient but also easy and intuitive to use. The method is built around single-sided Feynman diagrams and their double-sided counterparts, which describe the perturbative pathways contributing to desired spectroscopic signals. Using the method is a simple process of translating a desired Feynman diagram into a set of iterative function calls. A python implementation of UF^2 is available for download at github. This code includes examples of how to translate Feynman diagrams into the language of UF^2 and an example implementation that calculates TA spectra. It also contains a comparison Runge-Kutta integration method, usable with the same convenient interface.

5.2 Algorithm

5.2.1 Overview

We begin with a Hamiltonian of the form

$$\hat{H} = \hat{H}_0 + \hat{H}'(t), \tag{5.1}$$

where the light-matter interaction in the electric-dipole approximation is

$$\hat{H}'(t) = -\hat{\mu} \cdot \mathbf{E}(t), \tag{5.2}$$

and is treated as a perturbation, where $\hat{\mu}$ and $\mathbf{E}(t)$ are the dipole moment of the system and the external electric field, respectively. Bold face symbols indicate Cartesian vectors. $\hat{H}_0$ describes the material system (e.g., a molecule or quantum dot) and has eigenstates $|\phi\rangle$ and eigenvalues $\hbar\omega_\phi$, which are assumed to be known (either analytically or numerically). Therefore the unitary time-evolution operator $\hat{U}_0(t - t_0) = e^{-i\hat{H}_0(t-t_0)/\hbar}$ is known. We assume that the set of eigenstates $\{|\phi\rangle\}$ is finite, and that all relevant wavefunctions in this problem can be expressed using N eigenfunctions as

$$|\psi(t)\rangle = \sum_{\phi=1}^{N} e^{-i\omega_\phi t} c_\phi(t) \, |\phi\rangle \, . \tag{5.3}$$

The electric dipole operator $\hat{\mu}$ must be known in the eigenbasis of $\hat{H}_0$, where we define matrix elements as

$$\mu_{\phi\phi'} = \langle\phi| \, \hat{\mu} \, |\phi'\rangle \, .$$

Note that only ω_ϕ and $\mu_{\phi\phi'}$ are needed. The eigenfunctions themselves are unnecessary if $\mu_{\phi\phi'}$ can be calculated in some other way.

We describe the electric field classically as a sum of pulses, where each pulse is denoted by a lowercase letter starting from a. A typical 4-wave mixing signal would be calculated by using up to 4 pulses. We write the electric field as a sum over L pulses,

$$\mathbf{E}(t) = \sum_{i=a,b,\ldots,L} \mathbf{e}_i \varepsilon_i(t) + \mathbf{e}_i^* \varepsilon_i^*(t),$$

where $\mathbf{e}_i$ is the possibly complex polarization vector, and the amplitude ε_i of each pulse is defined with envelope A_i, central frequency ω_i, wavevector $\mathbf{k}_i$, and phase Θ_i as

$$\varepsilon_i(t) = A_i(t - t_i) e^{-i[\omega_i(t-t_i) - \mathbf{k}_i \cdot \mathbf{r} + \Theta_i]}, \tag{5.4}$$

where t_i is the arrival time of each pulse, and we define the Fourier transform of the pulse as

$$\tilde{\varepsilon}_i(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathrm{d}t \varepsilon_i(t) e^{i\omega t}.$$

Then the light-matter interaction is a sum over rotating (ε_i) and counter-rotating (ε_i^*) terms. We express these terms individually as

$$\hat{H}'_{j^{(*)}}(t) = -\hat{\mu} \cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t) \tag{5.5}$$

so that

$$\hat{H}'(t) = \sum_{i=a,b,\ldots,L} \hat{H}'_i(t) + \hat{H}'_{i^*}(t). \tag{5.6}$$

In general, 4-wave mixing signals are calculated from the third-order perturbed density matrix [14]. For closed systems, as described here, we can decrease the computational complexity of the problem by calculating the necessary perturbed wavefunctions up to third-order. We expand the true time-dependent wavefunction using the usual perturbative expansion

$$|\psi(t)\rangle = \left|\psi^{(0)}(t)\right\rangle + \left|\psi^{(1)}(t)\right\rangle + \left|\psi^{(2)}(t)\right\rangle + \ldots \tag{5.7}$$

where $\left|\psi^{(0)}(t)\right\rangle$ is the initial time-independent (ignoring a trivial phase) state, and $\left|\psi^{(n)}(t)\right\rangle$ is proportional to $(H'(t))^n$. UF^2 is a fast method for calculating $\left|\psi^{(n)}(t)\right\rangle$. The total polarization field is

$$\begin{aligned}\mathbf{P}(t) &= \langle\psi(t)| \hat{\mu} |\psi(t)\rangle \\ &= \mathbf{P}^{(1)}(t) + \mathbf{P}^{(2)}(t) + \mathbf{P}^{(3)}(t) + \ldots\end{aligned}$$

To calculate $\mathbf{P}^{(n)}(t)$, we expand $|\psi(t)\rangle$ using Eq. 5.7 and keep only terms proportional to $\left(\hat{H}'(t)\right)^n$.[56] As an example, the 3$^{\text{rd}}$-order polarization field is

$$\mathbf{P}^{(3)}(t) = \left\langle \psi^{(3)}(t) \right| \hat{\mu} \left| \psi^{(0)} \right\rangle + \left\langle \psi^{(2)}(t) \right| \hat{\mu} \left| \psi^{(1)} \right\rangle + c.c.$$

From $\mathbf{P}^{(3)}(t)$ any 3$^{\text{rd}}$-order spectroscopic signal can be determined. $\mathbf{P}^{(3)}(t)$ is the material response coupled to emitted radiation after three optical interactions. In 2D photon echo (2DPE) spectroscopy, the response is due to a single interaction with each of 3 separate pulses, with arrival times t_a, t_b, t_c. The delay times between pulses are $t_b - t_a = \tau$ and $t_c - t_b = T$. The polarization field is therefore a function of τ and T: $\mathbf{P}^{(3)}(\tau, T; t)$. Typically the quantity of interest is actually the Fourier transform partner

$$\mathbf{P}^{(3)}(\tau, T; \omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathrm{d}t \mathbf{P}^{(3)}(\tau, T; t) e^{i\omega t}.$$

The polarization field produces a radiating optical field, which can be heterodyne detected using a fourth local oscillator (LO) pulse. For the case of the 2DPE rephasing signal, the desired signal is calculated as [14]

$$S^{(3)}_{\text{2D}}(\tau, T, \omega) = \text{Im}\left[\tilde{\varepsilon}^*_{LO}(\omega)\mathbf{e}^*_{LO} \cdot \mathbf{P}^{(3)}_{\text{2D}}(\tau, T; \omega)\right]. \tag{5.8}$$

We derive UF^2 for the classic example given in Eq. 5.8. In Section 5.4, we show example TA spectra, which arise due to the interaction of two pulses, a pump and a probe, separated by the delay time T. In the TA case, the first two interactions happen with the pump pulse, and therefore $\tau = 0$. The third interaction comes from the probe, which also acts as the local oscillator. TA calculations are thus a function of only two variables:

$$S^{(3)}_{\text{TA}}(T, \omega) = \text{Im}\left[\tilde{\varepsilon}^*_{\text{probe}}(\omega)\mathbf{e}^*_{\text{probe}} \cdot \mathbf{P}^{(3)}_{\text{TA}}(T, \omega)\right]. \tag{5.9}$$

Returning to the more general case with τ varying, the canonical formula for calculating the third-order time-dependent polarization field is

$$\begin{aligned} P^{(3)}(\tau, T, t) = &\iiint dt_3 dt_2 dt_1 E(t - t_3) E(t - t_3 - t_2 + T) \\ &\times E(t - t_3 - t_2 - t_1 + T + \tau) R(t_1, t_2, t_3) \end{aligned} \tag{5.10}$$

where we follow Mukamel and suppress the polarization of the fields, and $R(t_1, t_2, t_3)$ is a material response function containing all of the factors arising from $\hat{H}_0$ and $\hat{\mu}$ [14]. This formulation separates the material properties ($\hat{H}_0$ and $\hat{\mu}$) from the shapes of the electric fields. If the electric fields are impulsive (i.e., short compared with any system timescale), then the convolutions in Eq. 5.10 are trivial and spectral signals may be calculated directly from the

response function.

In this paper we focus on the case where the shape of the electric field cannot be ignored, and therefore studying the response function alone is insufficient to predict, interpret or understand spectroscopic observables. The triple-nested convolutions of Eq. 5.10 are so costly that they are rarely carried out. We present a fast method for calculating the n^{th}-order wavefunctions, and thence the desired polarization $\mathbf{P}^{(n)}(t)$, with arbitrary pulse shapes. The wavepackets can be studied in conjunction with the signal, giving intuition for the underlying physics [82].

5.2.2 Derivation

We follow standard time-dependent perturbation theory by considering that at time t_0 the system begins in a time-independent state $\left|\psi^{(0)}(t_0)\right\rangle$, i.e., an eigenstate of $\hat{H}_0$. The perturbation $\hat{H}'(t)$ is zero before t_0 and produces $|\psi(t)\rangle$ as in Eq. 5.7, where $\left|\psi^{(n)}\right\rangle$ has contributions proportional to $(\hat{H}')^n$. Then the standard result has [56]

$$\left|\psi^{(n)}(t)\right\rangle = -\frac{i}{\hbar}\int_{t_0}^{t} ds \hat{U}_0^{-1}(s-t)\hat{H}'(s)\left|\psi^{(n-1)}(s)\right\rangle. \tag{5.11}$$

Since $\left|\psi^{(0)}\right\rangle$ is time-independent (up to a trivial phase), we are free to send $t_0 \to -\infty$. We also substitute $t' = t - s$ to arrive at

$$\left|\psi^{(n)}(t)\right\rangle = -\frac{i}{\hbar}\int_{0}^{\infty} dt' \hat{U}_0(t')\hat{H}'(t-t')\left|\psi^{(n-1)}(t-t')\right\rangle.$$

Using the decomposition of $\hat{H}'(t)$ in Eq. 5.6, we define $\left|\psi^{(n)}(t)\right\rangle$ as a sum over $2L$ terms

$$\left|\psi^{(n)}(t)\right\rangle = \sum_{j=a,b,\ldots,L}\left(\hat{K}_j + \hat{K}_{j^*}\right)\left|\psi^{(n-1)}(t)\right\rangle,$$

where

$$\hat{K}_{j^{(*)}} = -\frac{i}{\hbar}\int_{0}^{\infty} dt' \hat{U}_0(t')\hat{H}'_{j^{(*)}}(t-t').$$

A small rearrangement of terms allows us to perform this integral numerically much more quickly:

$$\hat{K}_{j^{(*)}} = -\frac{i}{\hbar}\hat{U}_0(t)\int_{0}^{\infty} dt' \hat{U}_0^{-1}(t-t')\hat{H}'_{j^{(*)}}(t-t'). \tag{5.12}$$

Most spectroscopic signals are not calculated using the full perturbative wavefunction $\left|\psi^{(n)}(t)\right\rangle$, because only specific pathways give nonzero contributions, which are visualized using Feynman diagrams, as in Fig. 5.1.[14] Each operator $\hat{K}_{j^{(*)}}$ represents a single interaction arrow in a Feynman diagram, as in Fig. 5.2. UF^2 calculates the wavefunctions contributing to each diagram separately. For example, as shown in Fig. 5.1, the excited state absorption (ESA) contribution requires two wavefunctions that we label $|\psi_{bc}(t)\rangle \equiv \hat{K}_c\hat{K}_b\left|\psi^{(0)}\right\rangle$ and $|\psi_a(t)\rangle \equiv \hat{K}_a\left|\psi^{(0)}\right\rangle$.

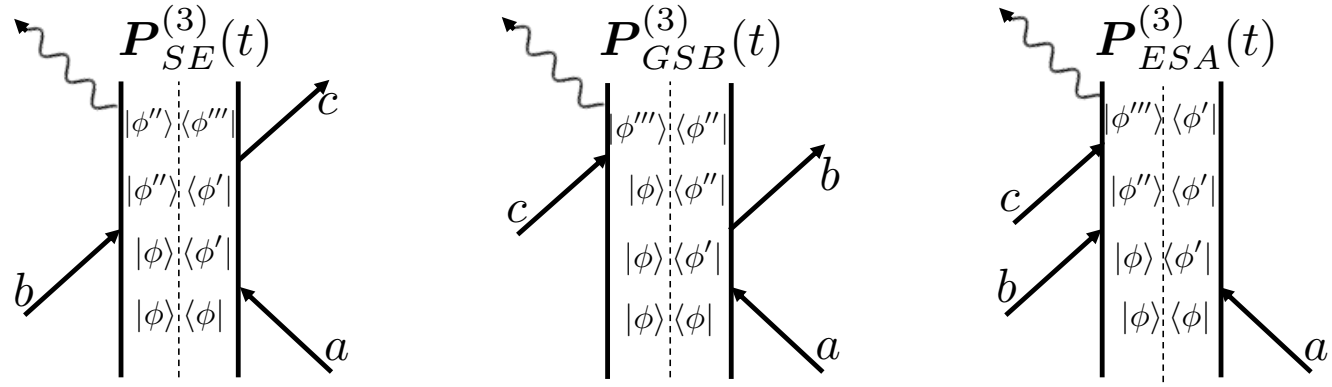

Figure 5.1: From left to right, stimulated emission (SE), ground-state bleach (GSB), and excited-state absorption (ESA) double-sided Feynman diagrams for the rephasing 2D photon echo signal. The dashed lines down the centers of each diagram emphasize that the ket and bra evolve independently for closed systems. We calculate the left and right sides of the diagram separately. These are the only three diagrams that contribute to the rephasing signal when the pulses are well-separated in time, though additional diagrams must be considered when pulses overlap.

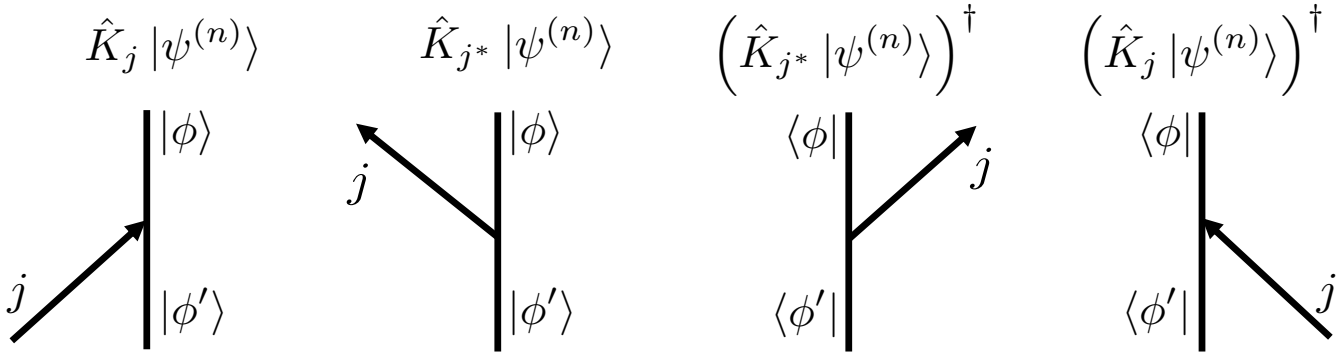

Figure 5.2: Building blocks of Feynman diagrams for the interaction with optical pulse j. The left two diagrams are for kets and the right two are for bras. Time moves from bottom to top, as the system begins in a linear combination of states $|\phi'\rangle$, and evolves to a linear combination of states $|\phi\rangle$. The arrow sloping up to the right is a contribution of the operator K_j, and the arrow sloping up to the left is a contribution of the operator K_{j^*}.

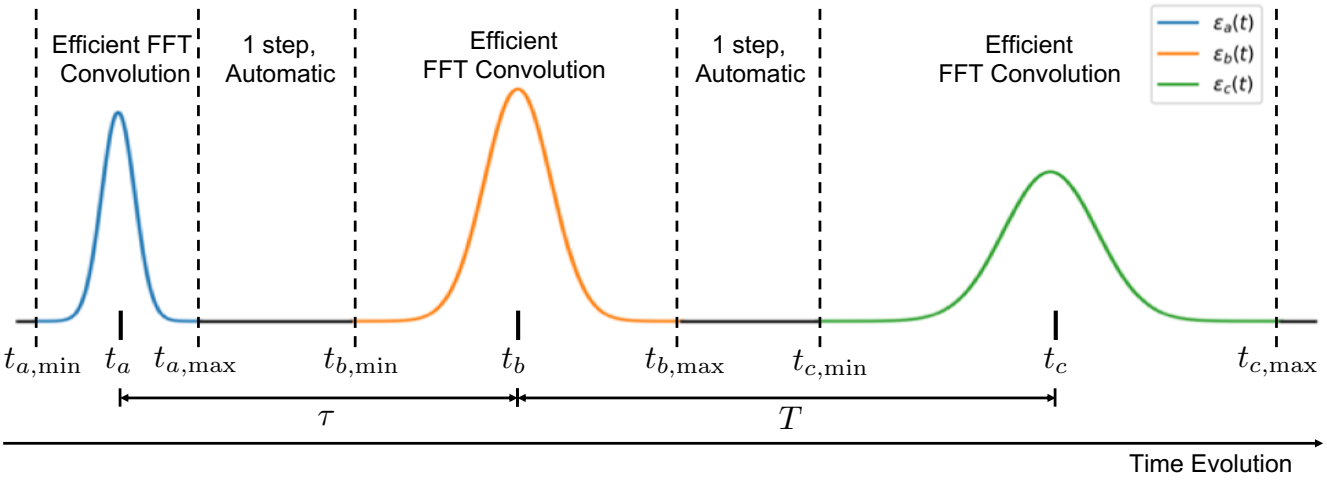

Figure 5.3: Sequence of three pulse envelopes, as would be used in a rephasing 2DPE experiment, showing the time intervals $(t_{j,\min}, t_{j,\max})$ during which the pulses are non-negligible. UF^2 only performs time propagations during these time periods, using efficient FFT convolutions. At other times, the system time evolution operator $\hat{U}_0$ gives exact time evolution.

Following Eq. 5.3, we write

$$|\psi_p(t)\rangle = \sum_{\phi'} e^{-i\omega_{\phi'}t} c_{\phi',p}(t) \, |\phi'\rangle \, ,$$

where p can be a multi-index, such as ac^*. Then we use Eqs. 5.5 and 5.12 to write

$$\begin{aligned}
\left|\psi_{pj^{(*)}}(t)\right\rangle &\equiv \hat{K}_{j^{(*)}} \, |\psi_p(t)\rangle \\
&= \frac{i}{\hbar}\hat{U}_0(t) \int_0^\infty dt' \hat{U}_0^{-1}(t-t') \sum_\phi |\phi\rangle \langle\phi| \left(\hat{\mu}\cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t-t')\right) \sum_{\phi'} e^{-i\omega_{\phi'}(t-t')} c_{\phi',p}(t-t') \, |\phi'\rangle \qquad (5.13) \\
&= \sum_\phi e^{-i\omega_\phi t} \, |\phi\rangle \, \frac{i}{\hbar} \int_{-\infty}^\infty dt' \theta(t') \underbrace{e^{i\omega_\phi (t-t')} \sum_{\phi'} \left(\mu_{\phi\phi'} \cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t-t')\right) e^{-i\omega_{\phi'}(t-t')} c_{\phi',p}(t-t')}_{y_\phi(t-t')}, \qquad (5.14)
\end{aligned}$$

where $\theta(t)$ is the unit step function. We rewrite Eq. 5.14 as

$$\left|\psi_{pj^{(*)}}(t)\right\rangle = \sum_\phi e^{-i\omega_\phi t} \, |\phi\rangle \underbrace{\frac{i}{\hbar} \left[\theta * y_\phi\right](t)}_{c_{\phi,pj^{(*)}}(t)}, \qquad (5.15)$$

where

$$\left[x * y\right](t) = \int_{-\infty}^\infty dt' x(t') y(t-t')$$

is a convolution. Thus we arrive at a compact description of the coefficients $c_{\phi,pj^{(*)}}(t)$. If we make the physical assumption that the incident pulse is localized in time, i.e., $\varepsilon_i(t)$ is negligible for $t < t_{j,\min}$ or $t > t_{j,\max}$, then it is

also true that $y_{\phi,\alpha j^{(*)}}(t)$ is negligible when $t < t_{j,\min}$ or $t > t_{j,\max}$. This assumption implies that

$$c_{\phi,pj^{(*)}}(t) = \begin{cases} 0 & t < t_{j,\min} \\ r_\phi(t) & t_{j,\min} < t < t_{j,\max} \\ C_\phi & t > t_{j,\max} \end{cases}. \tag{5.16}$$

Therefore, we need only calculate this convolution for $t_{j,\min} < t < t_{j,\max}$ (see Fig. 5.3). Note that without the rearrangement of Eq. 5.12, the convolution in Eq. 5.15 would be more difficult to evaluate numerically because the functions $c_{\phi,pj^{(*)}}(t)$ would include the phase factor $e^{-i\omega_\phi t}$ even after $t_{j,\max}$ and thus not be constant. This trivial definitional difference when working symbolically makes a large difference when working numerically, and constitutes one of the primary novel contributions of this work.

While Eq. 5.14 is a rearrangement of the response function formalism, and forms of it appear in References 18, 19, 21, 69, the integral in Eq. 5.14 is usually incorporated into the numerical propagation in direct propagation methods [21, 59, 60, 64, 69], or is handled analytically [17–19, 65]. By expressing this integral as a linear convolution over a small set of points, we are able to include the effects of pulse shapes at a modest cost, as shown in Appendix 5.6, depending on the desired signal accuracy.

Physically, we only need to solve for the time dependence due to the interaction while the pulse is non-zero. The rest of the time-dependence is contained entirely in $\hat{H}_0$, and is therefore known exactly. This realization drastically reduces the computational cost of UF^2 compared to techniques that must use time-stepping for both the system dynamics and the perturbation. We solve for the function $r_\phi(t)$ in Eq. 5.16 numerically using the convolution theorem and FFT. With Gaussian pulse shapes of any duration, for convergence of spectroscopic signals to a tolerance of 1%, we find that 20-40 points is sufficient for calculating each convolution in Eq. 5.15. With such a small number of required points, the evaluation of the dipole expectation value $\vec{P}(t) =< \psi|\hat{\mu}|\psi >$ is as expensive as the propagation, as we detail in Appendix 5.6.

Translating the operators $K_{j^{(*)}}$ into computational algorithms is straightforward. We discretize the time interval $(t_{j,\min}, t_{j,\max})$ to have M equally spaced points with spacing dt, which must be small enough to resolve the pulses. Therefore, the wavefunctions that give rise to the ESA signal are calculated in three steps:

1. Calculate $|\psi_a(t)\rangle = \hat{K}_a \left|\psi^{(0)}\right\rangle$,
2. Calculate $|\psi_b(t)\rangle = \hat{K}_b \left|\psi^{(0)}\right\rangle$ for fixed τ,
3. Calculate $|\psi_{bc}(t)\rangle = \hat{K}_c |\psi_b(t)\rangle$ for fixed T.

The polarization field is then

$$\mathbf{P}^{(3)}_{ESA}(t)|_{\tau,T} = \langle\psi_a(t)|\,\hat{\mu}\,|\psi_{bc}(t)\rangle\,|_{\tau,T}. \tag{5.17}$$

This process is repeated for all τ and T of interest.

We demonstrate how this conceptual procedure is implemented using the syntax of UF^2.

```
psi_a = uf2.up(uf2.psi0,pulse_number = 0)
psi_b = uf2.up(uf2.psi0,pulse_number = 1)
psi_bc = uf2.up(psi_b,pulse_number = 2)
P_ESA = uf2.dipole_expectation(psi_a,psi_bc)
S_ESA = uf2.polarization_to_signal(P_ESA)
```

In the syntax of UF^2, pulse a is pulse 0, pulse b is pulse 1, etc., and the `up` and `down` methods implement K_j and K_{j^*}, respectively. This snippet is adapted from the Jupyter notebook Example.ipynb included with the source code.

One of the strengths of UF^2 is that the user interacts with it by translating Feynman diagrams into nested calls to the $\hat{K}_{i^{(*)}}$ functions. The goal of this algorithm is to make calculating spectra as easy as writing down Feynman diagrams. We have used 3$^{\text{rd}}$-order diagrams for the derivation because they are the most familiar, but calculating arbitrary order diagrams is just as straightforward. UF^2 is therefore easy to use and avoids the necessity of first writing out lengthy symbolic expressions for each diagram and translating each expression into code. UF^2 can also be used in the impulsive limit to calculate response functions, allowing this ease of use to be applied to the study of response functions as well.

5.2.3 Optical dephasing

For closed systems, $\mathbf{P}^{(n)}(t)$ oscillates for all time after the n optical interactions. This undamped oscillation is unphysical. In order to include optical dephasing in general, we would need to solve for the third-order density matrix, $\rho^{(3)}(t)$, instead of the wavepackets. In order to include optical dephasing in our polarization signals, we convolve $\tilde{\mathbf{P}}^{(n)}(\omega)$ with a lineshape function, as in Ref. 69. Whereas that work used a Gaussian shape to model inhomogeneous broadening, in the examples below we use a Lorentzian lineshape to model homogeneous broadening due to an exponential dephasing.

5.2.4 Useful standard approximations

The derivations in Section 5.2.2 do not require the rotating wave approximation (RWA) or phase matching. However, these approximations significantly reduce the cost of the calculations, as in all perturbative spectroscopies.[14, 56, 83] In the RWA, only one of either the rotating or counter-rotating terms of Eq. 5.6 contributes to an interaction.[56] [ii] Since $\left|\psi^{(n)}(t)\right\rangle \propto \left(\hat{H}'(t)\right)^n$, the n^{th}-order wavefunction is composed of $(2L)^n$ terms. The standard RWA and phase-matching conditions reduce the number of terms relevant to spectroscopic signals. The RWA is valid in the limit that the pulse durations are long compared to the optical carrier frequencies, which are roughly degenerate with the optical energy gap of $\hat{H}_0$. If the material system is dispersed over a volume much larger than the wavelength of the

[ii] The rotating term excites a ket and de-excites a bra. The counter-rotating term de-excites a ket and excites a bra.

light then the phase-matching condition ensures that signal fields will be produced only in directions corresponding to sums and differences of wavevectors of the optical pulses [14]. For example Fig. 5.1 shows the signals produced in the rephasing direction $(-\mathbf{k}_a + \mathbf{k}_b + \mathbf{k}_c)$ of 2D photon echo (2DPE) spectroscopy.

In addition to reducing the number of relevant diagrams, the RWA speeds up calculations because we do not need to keep track of the optical carrier frequency. This advantage is significant because UF^2 performs calculations using M time points with spacing dt from $t_{j,\min}$ to $t_{j,\max}$. In the RWA, we set the optical carrier frequency to 0, and therefore dt only needs to resolve the pulse envelope and not the carrier frequency.

5.2.5 Efficiency improvements

We increase the efficiency of UF^2 by decreasing the required dimensionality of the Hilbert space from the full size, N_{total}, to a compressed subspace of dimension N_c, over which the sums in Eq. 5.14 run. We perform the reduction from N_{total} to N_c before running any propagations by determining which elements of the dipole operator $\hat{\mu}$ will not contribute appreciably to the calculation and ignoring them. We can prune the required set of states in two ways. First, many systems have elements $\mu_{\phi'\phi} \approx 0$. Second, we determine which energetic transitions will not be allowed by the electric field shape.

We use Bessel's inequality to define which states $|\phi'\rangle$ are important in order to resolve $\mu_{\phi'\phi}$ accurately. If we begin in the eigenstate $|\phi\rangle$, one interaction with the dipole operator (as occurs in each application of $K_{j^{(*)}}$) yields the state

$$|\Phi\rangle = \hat{\mu}\,|\phi\rangle = \sum_{\phi'=1}^{N_{\text{total}}} \mu_{\phi'\phi}\,|\phi'\rangle\,.$$

We seek to restrict the sum to the smallest number of terms without significantly altering the norm of this state, which implies that we have captured all of the physically relevant states. We write the norm of $|\Phi\rangle$ as

$$\langle\Phi\,|\,\Phi\rangle = \sum_{\phi'=1}^{N_{\text{total}}} |\mu_{\phi'\phi}|^2 = \langle\phi|\,\hat{\mu}^2\,|\phi\rangle = \left(\mu^2\right)_{\phi\phi}.$$

We find the smallest set of states, N_c, such that

$$\frac{\left(\mu^2\right)_{\phi\phi} - \sum{}' |\mu_{\phi\phi'}|^2}{\left(\mu^2\right)_{\phi\phi}} > 1 - \epsilon, \tag{5.18}$$

for small ϵ, where $\sum'$ represents the restricted sum over the required N_c states. We perform this analysis for all required states $|\phi\rangle$.

We make this concept concrete by giving an example using the non-rephasing ground state bleach signal, which is not included in Fig. 5.1. The polarization field produced by that diagram is $\mathbf{P}^{(3)}_{GSB,NR}(t) = \left\langle\psi^{(0)}\right|\mu\,|\psi_{ab^*c}(t)\rangle$. If we

begin, say, in the state $\left|\psi^{(0)}\right\rangle = |1\rangle$, the lowest energy eigenstate of $\hat{H}_0$, then we must determine the smallest number of states, $N_c^{(1)}$, that satisfy Eq. 5.18 for $\phi = 1$. Then we know before calculating Eq. 5.14 that $|\psi_a(t)\rangle = \sum c_{\phi,a}(t)\,|\phi\rangle$ will be composed of $N_c^{(1)}$ terms. To find the states required to describe $|\psi_{ab^*}\rangle$, we use Eq. 5.18 for each of the $N_c^{(1)}$ states needed for $|\psi_a(t)\rangle$, which gives a new set of $N_c^{(2)}$ states, where the superscript indicates the number of times $\hat{\mu}$ has been applied to the initial state.

The next obvious step is to use Eq. 5.18 for each of those $N_c^{(2)}$ states, but that step is actually unnecessary. Such an analysis would allow us to determine the $N_c^{(3)}$ states required to resolve $|\psi_{ab^*c}(t)\rangle$. However, the final signal depends upon $\left\langle\psi^{(0)}\right|\mu\,|\psi_{ab^*c}(t)\rangle$, and therefore the only components of $|\psi_{ab^*c}(t)\rangle$ that matter spectroscopically are those that overlap with $\mu\left|\psi^{(0)}\right\rangle$. These are precisely the $N_c^{(1)}$ states we determined in the first step of this process. Therefore we only need the same $N_c^{(1)}$ states when calculating $|\psi_{ab^*c}(t)\rangle$.

For many systems relevant to optical spectroscopy, there are well-separated manifolds of states with 0, 1, 2, etc. electronic excitations. When there are well-separated manifolds, the RWA allows a further reduction to the relevant size N_c. In the same non-rephasing GSB example, $|\psi_a(t)\rangle$ is in the singly excited manifold (SEM), and $\hat{\mu}\,|\psi_a(t)\rangle$ has components in both the ground state manifold (GSM) and the doubly excited manifold (DEM). In the RWA, however, $|\psi_{ab^*}(t)\rangle$ is only in the GSM, so $N_c^{(2)}$ can be divided into its GSM and DEM portions, with only the GSM portion, $N_{c,GSM}$, required for this diagram. The DEM portion of $N_c^{(2)} = N_{c,DEM}$ is required for the ESA diagram.

In addition, we can use basic knowledge of the shape of the electric field to further restrict N_c. UF^2 performs discrete convolutions for times t with spacing dt (see end of Section 5.2.4). The spacing dt is chosen in order to resolve the shape of the electric field and implies a frequency range in which $\tilde{E}(\omega)$ is non-zero. The maximum frequency resolved is $\frac{\pi}{dt}$. Each element $\mu_{\phi'\phi}$ has an associated frequency difference $\omega_{\phi'} - \omega_\phi$. If $|\omega_{\phi'} - \omega_\phi| \leq \frac{\pi}{dt}$, then the transition is energetically allowed by the pulse. If $|\omega_{\phi'} - \omega_\phi| > \frac{\pi}{dt}$, the transition is not allowed, and we set $\mu_{\phi'\phi} = 0$ for all those energetically inaccessible transitions, further decreasing the relevant size N_c.[iii] For long-duration pulses, this reduction is highly significant and allows treatment of long propagation times without increased cost.

In practice, to obtain the eigenstates of $\hat{H}_0$, it generally must first be truncated from dimension N_{total} to a finite dimension N_f, which must be chosen to be large enough that the N_c required eigenstates are resolved sufficiently accurately. Obtaining the N_c eigenstates requires diagonalizing all or part of the truncated Hamiltonian of size N_f. Full diagonalization is well-known to scale as N_f^3. Since N_f is sparse, iterative methods can be used to obtain the necessary subset N_c, iterative methods can give better scalings in some cases. We will present our method for diagonalizing vibronic systems, including anharmonicities and varying vibrational frequencies, in a separate manuscript. The reduction from dimension N_f to N_c not only makes UF^2 more efficient but also dramatically expands the range of systems for which UF^2 is tractable. In Section 5.4 we calculate TA spectra for a vibronic system, which formally has $N_{\text{total}} = \infty$. In practice we require $N_{f,DEM} = N_{f,GSM} = 15$ and $N_{f,SEM} = 30$. With

[iii] This procedure is not only helpful for speeding up calculations, but actually necessary for accurate calculations. If included, any frequency differences $|\omega_{\phi'} - \omega_\phi| > \frac{\pi}{dt}$ would cause spurious signals to appear in the signal due to aliasing.

$\epsilon = 0.001$ the spectra converge to within 1%. Using the procedure outlined here, we determine that $N_{c,SEM} = 8$ and $N_{c,DEM} = N_{c,GSM} = 10$. $N_{c,SEM}$ would be 35 in order to correctly resolve $|\psi_{ab^*c}(t)\rangle$, but we only need to use $N_{c,SEM} = 8$ of those states to accurately reproduce the TA spectra.

5.3 Computational Cost

In this section we discuss in brief the computational cost of UF^2 and compare it to two alternative methods of obtaining nonlinear spectra: direct propagation (DP) of the Schrodinger equation without diagonalization of $\hat{H}_0$ and a Fourier-space pseudospectral method. These alternative methods are commonly used in predictions of nonlinear spectroscopies and each has a domain where it is the most efficient method.

In contrast to UF^2, DP methods propagate the perturbative wavefunction using the Schrodinger equation in the form

$$\frac{d\left|\psi^{(n)}(t)\right\rangle}{dt} = -\frac{i}{\hbar}\hat{H}_0\left|\psi^{(n)}(t)\right\rangle - \frac{i}{\hbar}\hat{H}'_j(t)\left|\psi^{(n-1)}(t)\right\rangle. \tag{5.19}$$

UF^2 performs the same integration in the eigenbasis, where Eq. 5.11 permits rapid evaluation. There are a number of methods of directly integrating Eq. 5.19, and to give a sense of where UF^2 is most effective, we compare here to an adaptive step-size 4-5$^{\text{th}}$ order Runge-Kutta solver for the $\hat{H}_0$ term and an Euler method for the perturbation, which is described in Section 5.7.1 and which we call the RKE method. We believe the general scaling trends of these results will be similar for other DP methods, whether they be Adams-Bashforth [77], Burlisch-Stoer [62], or others. While the above derivation is general, we perform these comparisons for sets of vibronic systems, which we introduce in Section 5.3.1. In all cases, we use the RWA to increase step sizes as much as possible.

The cost of obtaining all eigenvalues and eigenvectors of a matrix of dimension N_f scales as N_f^3 for large N_f. If N_f is sparse, efficient iterative algorithms exist to extract the subspace N_c, which can scale more favorably. In contrast, direct propagation scales as N_f in the case that $\hat{H}_0$ is sparse and N_f^2 if $\hat{H}_0$ is not sparse. Therefore it is clear that for sufficiently large N_f the DP methods will eventually be the more efficient option. We find, however, that N_f needs to be quite large before this crossover occurs. Defining this break-even size is somewhat difficult, as diagonalization only needs to be performed once, and the timing comparison depends upon how many times the eigensystem is to be reused. In Section 5.4, we compute isotropically averaged TA spectra of a model system for hundreds of different dipole moments and 3 different electric field pulse shapes, which requires thousands of diagrams to be calculated, all with the same $\hat{H}_0$. Even if the diagonalization time were greater than the cost of obtaining a single spectrum, that diagonalization cost may be negligible compared to the cost of producing all of the spectral predictions. It is also therefore important to compare the cost of UF^2 separately from the cost of diagonalization. The cost of propagation with UF^2 is also asymptotically worse than DP methods, but we show in Sec. 5.3.2 that UF^2 remains faster for vibronic systems up to size $\sim 10^4$. For frequently studied systems with dimension less than 100, as in Sec. 5.4, UF^2 can be over 100 times faster than direct propagation methods.

5.3.1 Properties of Vibronic Systems

The relative costs of propagation methods depend on the structure of the Hamiltonian, and we use vibronic systems for our examples. Vibronic systems consist of two or more coupled electronic states, each of which is coupled to one or more vibrational degrees of freedom, which are often harmonic. There is no known general solution to the time-independent Schrodinger equation for such systems. A common choice of basis for DP methods is the number basis of the harmonic oscillator of each vibrational mode, and in this basis, H_0 is highly sparse [70]. This basis is infinite and therefore must be truncated to some finite size, which we call N_f.

For the purposes of cost comparisons, we consider vibronic systems with optically separated manifolds. We focus in particular on s coupled two-level systems (TLS), each with an electronic ground state, $|g\rangle$, and a single optical excitation $|s\rangle$, described by the Hamiltonian

$$\hat{H}_e = |g\rangle\langle g| + \sum_{i=1}^{s} E_i |i\rangle\langle i| + \sum_{i\neq j} J_{ij} |i\rangle\langle j|,$$

where E_i is the site energy and J_{ij} is a Hermitian matrix of couplings. Vibrations are described by the Hamiltonian

$$\hat{H}_{ph} = \frac{1}{2}\left(\sum_{\alpha=1}^{k} \hat{p}_\alpha^2 + \omega_\alpha^2 \hat{q}_\alpha^2\right),$$

where k is the number of independent vibrations, and $\hat{p}_\alpha$ and $\hat{q}_\alpha$ are the generalized momentum and coordinate operators for each vibration, with frequencies ω_α. Standard linear coupling of the position of each oscillator to its site excitation gives

$$\hat{H}_{e-ph} = \sum_{i=1}^{s}\sum_{\alpha=1}^{k} \omega_\alpha^2 d_{\alpha,i} q_\alpha |i\rangle\langle i|,$$

where $d_{\alpha,i}$ are the coupling strengths, corresponding to Huang-Rhys factors

$$S_{\alpha,i} = \frac{1}{2}\omega_\alpha d_{\alpha,i}^2.$$

The total system Hamiltonian is

$$\hat{H}_0 = \hat{H}_e + \hat{H}_{ph} + \hat{H}_{e-ph}. \tag{5.20}$$

This Hamiltonian is block diagonal, and in third-order spectroscopies there are three relevant manifolds: the GSM, SEM, and DEM as described in Section 5.2.5, with higher manifolds becoming relevant in higher-order spectroscopies. Only the perturbation, $\hat{H}'$, mixes these manifolds. Propagation within each manifold can be handled independently, and each manifold can be truncated with dimension $N_{f,X}$, where X can be GSM, SEM, DEM, ... If we consider only the first three manifolds, then $N_f = N_{f,GSM} + N_{f,SEM} + N_{f,DEM}$.

The relevant dimension $N_{f,SEM}$ can be dramatically smaller than $N_{f,GSM}$ and $N_{f,DEM}$. The reason for this

Table 5.1: Summary of symbols used to describe computational cost of a 2DPE rephasing signal. M and M_t are convergence parameters.

Symbol	Definition
$N_{f,X}$	Full size of Hilbert space of manifold X
$N_{c,X}$	Compressed size of Hilbert space of manifold X
M_t	Number of time points to resolve $\mathbf{P}(t)$ in UF^2
M_{RK}	Number of time points to resolve $\mathbf{P}(t)$ in RKE
α	Cost of complex floating point multiplication
m_τ	Number of desired values of $\tau = t_b - t_a$
m_T	Number of desired values of $T = t_c - t_b$

Table 5.2: Summary of symbols used to describe computational cost of a 2DPE rephasing signal. Particular values are stated for the range of k = 2-8 vibrations, corresponding to the systems studied in Figure 5.4.

Symbol	Definition	Range
l_{SEM}	$N_{f,DEM}/N_{c,SEM}$	2-60
l_{DEM}	$N_{f,DEM}/N_{c,DEM}$	≈ 2
M_{RK}/M_t	Ratio of field-free step sizes	2-10
r	Number of nonzero entries per row of $\hat{H}_0$	$2k+1$
q	Sparse matrix overhead	≈ 2

smaller size can be illustrated with the case $s = k = 1$. In this case, the vibrational system can be diagonalized when the electron is in state $|g\rangle$, giving a vibrational basis $|g, n\rangle$, where n is the vibrational quantum number; it can be rediagonalized when the electron is in the excited state $|e\rangle$ to give a vibrational basis $|e, n\rangle$, and the union of the two bases is a basis for the full system. For the first pulse interaction, with small values of d, the ground vibrational state $|g, 0\rangle$ only couples to the first few vibrational levels of the excited state via the dipole interaction (e.g., $|e, 0\rangle, |e, 1\rangle, |e, 2\rangle$), due to the rapidly decaying Franck-Condon overlaps. With the second pulse interaction, these three levels in turn couple to the bottom 5 or 6 ground state levels. Each optical transition couples more vibrational levels. However, as described in Section 5.2.5, 4-wave mixing signals only require correctly resolving amplitudes in the basis states required for the first- and second-order wavepackets. In this case, the required numbers of states in the manifolds are different, so $N_{f,SEM} < N_{f,GSM}$ and $N_{c,SEM} < N_{c,GSM}$. This result remains true for vibronic systems with $s > 1$ and $k > 1$, and similar arguments lead to the conclusion that $N_{c,SEM} < N_{c,DEM}$.

5.3.2 Comparison to Direct Propagation

The most difficult part of calculating any n-wave mixing signal is the calculations that correspond to the last two arrows of a Feynman diagram. In Fig. 5.1, for example, these are the interaction with pulse c and the emission of the polarization field. These steps are the most costly because they must be repeated for all m_τ desired values of delays between pulses a and b and all m_T desired values of delays between pulses b and c.

In Appendix A, we show in Eq. 5.29 that the cost of calculating third-order signals using UF^2 scales asymptot-

ically as

$$C_{\mathrm{UF}^2} = m_\tau m_T \alpha N_{c,SEM} N_{c,DEM} M_t, \tag{5.21}$$

where α is the cost of multiplying two complex floating-point numbers, and symbols used in this section are defined in Table 5.1. In Appendix B, we introduce the RKE method and show in Eq. 5.30 that the same calculation using RKE scales asymptotically as

$$C_{RKE} = 6 m_\tau m_T \alpha q r N_{f,DEM} M_{RK}, \tag{5.22}$$

where q, r, and M_{RK} are defined in Table. 5.2.

We now estimate the break-even system size where RKE becomes less expensive than UF^2 and compare to timings with our sample systems. In each of our comparisons, the nonzero $d_{\alpha,i}$ are identical, and the ω_α are within 10% of one another. We study systems with s ranging from 2 to 8, k ranging from 2 to 8, and k/s ranging from 1 to 4. We also study $S_{\alpha,i}$ ranging from 0.02 to 4.5, with $S_{\alpha,i}$ equal for all modes in a given system. We use a homogeneous linewidth of $0.1\omega_0^{-1}$, where ω_0 is the smallest value of ω_α. We use identical pump and probe pulses with a Gaussian profile centered on the transition from the lowest energy GSM eigenstate to the lowest energy SEM eigenstate. Both pulses have a Gaussian time-domain standard deviation of $\sigma = \omega_0^{-1}$.

To compare Eqs. 5.21-5.22, we must relate $N_{c,X}$ to $N_{f,X}$ and M_t to M_{RK}. In order to correctly determine the $N_{c,X}$ eigenvalues and eigenvectors in manifold X, $N_{f,X}$ must be made large enough. For third-order signals to converge to better than 1%, we find that $N_{c,X} \lesssim N_{f,X}/2$ in the systems we have studied. Stricter convergence requirements tend to leave $N_{c,X}$ unchanged while increasing the required $N_{n,X}$. As a wavefunction propagates with a DP method, it obtains weights in ever higher vibrational states and thus requires $N_{f,X}$ to be sufficiently large to ensure accurate calculations; this requirement is similar to the requirement that $N_{f,X}$ must be large enough to resolve the relevant $N_{c,X}$ eigenvalues of the system, and we assume that, given the accuracy of the predicted spectrum that is desired, the required $N_{f,X}$ are approximately equal for both UF^2 and DP methods. An important factor in determining N_f is the longest propagation time required. The required N_f is smaller for a TA spectrum that includes delay times out to $2T_\alpha$ than it is for a TA spectrum that includes delay times out to $15T_\alpha$, where $T_\alpha = 2\pi/\omega_\alpha$ is the vibrational period.

For the numerical comparisons in Fig. 5.4, we use both UF^2 and RKE to calculate spectra for $m_T = 90$ delay times between 0 to $89\omega_\alpha^{-1}$, which is half the number used for all calculations in Section 5.4. Both UF^2 and RKE scale linearly with the number of delay times, so the number of delay times is important only for the ratio including the cost of diagonalization. For the cost of calculating a 2DPE, we scale the cost per delay time by 1000, corresponding to $m_\tau = 50$ and $m_T = 20$.

We write $N_{c,SEM} = \frac{N_{f,DEM}}{l_{SEM}}$ and $N_{c,DEM} = \frac{N_{f,DEM}}{l_{DEM}}$, where l_{SEM} ranges from 2 for $k = s = 2$ and $S = 0.02$, to 60 for $k = s = 8$ and $S = 0.02$. For all the cases studied here, $l_{DEM} \approx 2$.

Ignoring the small effect from the J_{ij} couplings, given k vibrational modes, $r = 2k + 1$. Generally $M_{RK} > M_t$,

since UF^2 is limited only by the pulse, whereas RKE is limited both by the pulse and the system. In particular, the RKE step size, and therefore M_{RK}, is limited by the largest eigenvalues included in the truncated Hamiltonian of size H_f. We have found in the sparse matrix implementation of scipy running on a MacBook Pro, 2.3 GHz Intel Core i5, that $q \approx 2$. We have also found for systems ranging from $s = k = 2$ to $s = k = 8$ that M_{RK}/M_t ranges from 2 to 10. The ratio of the costs of the two methods is then

$$\frac{C_{RK}}{C_{\mathrm{UF}^2}} = 6qrl_{SEM}l_{DEM}\frac{M_{RK}}{M_t}\frac{1}{N_f}, \tag{5.23}$$

which predicts that UF^2 will outperform the RKE solver until the system Hamiltonian reaches a size in the range of $10^3 - 10^5$. This estimate is rough, and depends upon the details of each system.

In Fig. 5.4 we plot the ratio of observed runtimes for RKE to UF^2 for a range of different vibronic systems. Simulations were run on an Intel Xeon E5-2640 v4 CPU with a 2.40GHz clock speed. We compare UF^2 and RKE by calculating spectra with better than 1% convergence for UF^2, and better than 5% convergence for RKE. We hold RKE to a lower standard because our implementation requires the same time step while the pulse is on and while the pulse is off. We have found that for RKE to reach 1% convergence, dt must be so small that it fails to take advantage of the adaptive RKE step size. We therefore run tests in a regime where RKE can take advantage of its adaptive step size, in order to give a fair comparison. To achieve 1% convergence in its current form, RKE would take about 4-5 times longer to run for all of the cases in Fig. 5.4, and we are not sure that extra time can truly be eliminated, so we are giving the RKE method an advantage in Fig. 5.4.

The results in Figure 5.4 indicate that UF^2 is 150 times faster than RKE for systems with $N_{f,DEM} < 100$ and is competitive until $N_{f,DEM}$ is over 10^4, with the exact break-even point depending on how many different times the eigenvectors can be reused. For all but the largest systems studied, the diagonalization cost is not a significant contribution to runtimes. Including rotational or thermal averaging, or variations in $\hat{\mu}$ would further reduce the importance of the diagonalization costs. Systems with required dimensions over 10^5 would benefit from the RKE method.

We note that the cost of the current implementation of UF^2 is determined by resolving the final polarization with M_t points having the same step size dt as needed to resolve the last pulse interaction. There is no fundamental reason that the polarization must be resolved with the same time spacing as the last optical pulse, and we anticipate that future developments could reduce the UF^2 runtimes by a further factor of 5, until the wavefunction propagation and final polarization evaluation have similar cost.

5.3.3 Comparison to split-operator pseudospectral methods

Another large class of methods useful for numerical prediction of nonlinear spectroscopies is pseudospectral methods, which use real- or Fourier-space representations of nuclear wavepackets, rather than explicit eigenstates or a basis

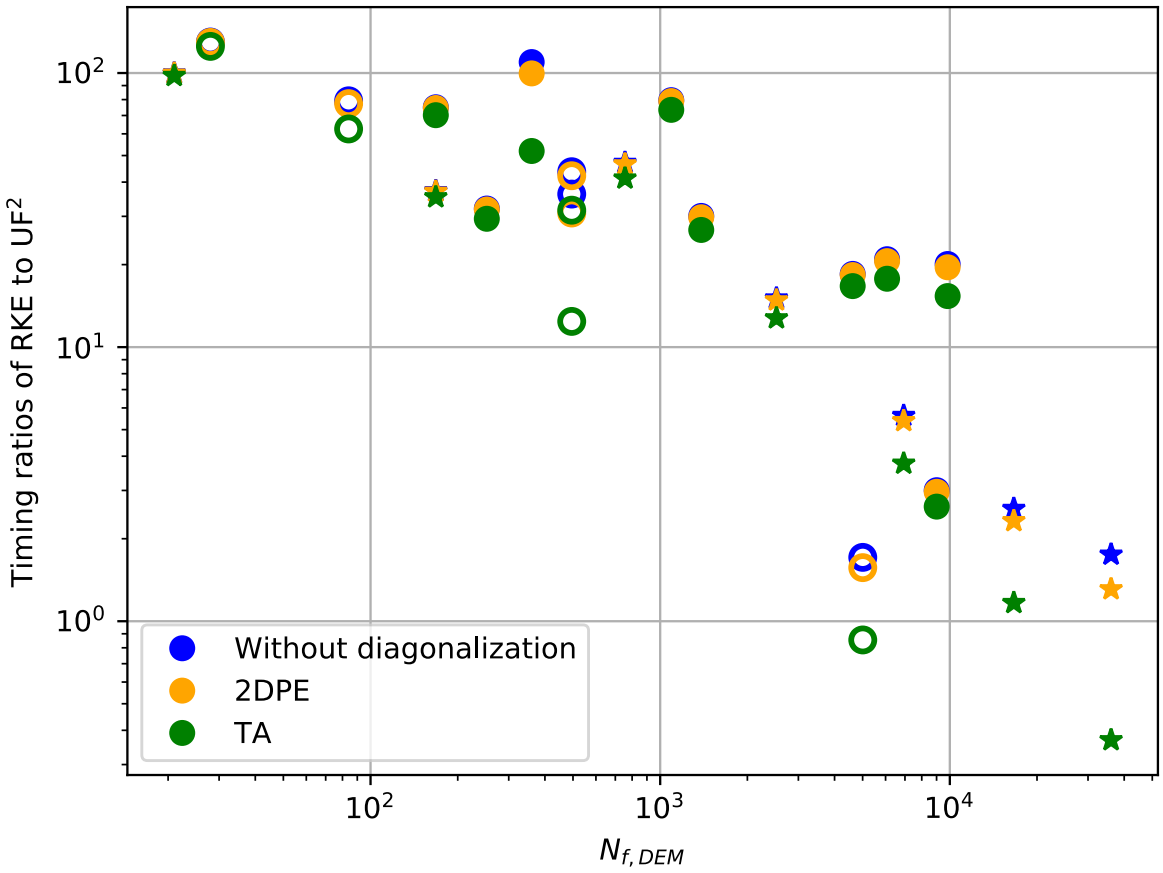

Figure 5.4: Ratio of runtimes of RKE to UF^2 for an array of vibronic systems, as described in the text. We include the diagonalization time when we calculate a single TA spectrum with 90 delay times, and a single 2DPE with 50 values of τ and 20 values of T. In most cases, the cost of diagonalization is negligible. Stars indicate systems with $s = k$ ranging from 2 to 8 with $S = 0.02$. Open circles are systems with $s = 2$ and k ranging from 2 to 8, with S ranging from 0.02 to 0.08. Filled circles are systems with $s = 3$ and $k = 3$ or $k = 6$, with S ranging from 0.02 to 4.5.

of vibrational quanta [70, 79, 80]. The widely used split-operator pseudospectral method (SOP) uses a plane-wave basis and an associated evenly-spaced real-space basis, with the efficient FFT to move between them. Briefly, SOP performs part of the time evolution due to $\hat{H}_0$ in position space and part in momentum space, with FFT operations between. The perturbative interaction with the pulses is also handled numerically, using a Euler method similar to that in Appendix 5.7.1. The time step dt must be short enough to resolve both the system and perturbative dynamics. It is not trivial to compare UF^2 directly to SOP methods, as the convergence parameters (e.g., N_c for UF^2 and the grid-spacing in SOP) are not directly comparable. We compare UF^2 to the SOP implementation detailed in Ref. 21.

For problems where a small number of eigenstates are required to describe the system dynamics, the eigenbasis used by UF^2 is considerably more efficient than a real/Fourier-space representation, which can still need many discretization points to resolve the wavepackets. In such problems, which include most of the class of problems described by Eq. 5.20, we expect UF^2 to greatly outperform SOP methods. Additionally, since UF^2 handles the evolution due to $\hat{H}_0$ exactly, UF^2 can evolve the wavefunctions forward in the absence of the pulses for near zero cost. In addition, the time step dt used by UF^2 to evaluate the convolutions during the pulses is determined solely by the pulse shapes and need not be small enough to accurately resolve the dynamics in $\hat{H}_0$. We have used the code from Ref. 21 to calculate the frequency-integrated TA spectrum for the same system used in Section 5.4 (which corresponds to the smallest value of $N_{f,DEM}$ shown in Fig. 5.4). To achieve 1% convergence on the ℓ_2 norm of the total frequency-integrated spectrum, the code from Ref. 21 took 1000 times longer than it takes UF^2 to compute the frequency-resolved TA spectrum. The code from Ref. 21 does not include functionality to compute frequency-resolved TA spectra. We estimate that the SOP code we are comparing to could take as much as 10 times longer to generate the frequency-resolved TA spectrum, and thus we believe that the SOP code would be approximately 10^4 times slower than UF^2 for frequency-resolved TA simulations. We verified that the simulation done using the SOP method agreed with the equivalent UF^2 simulation to within 1%.

The SOP methods are superior to UF^2 and DP methods in cases where a continuum of eigenstates is required to describe the dynamics, as in bond-breaking, isomerization, or chemical reactions. In such cases, UF^2 easily becomes intractable and SOP is superior. In the large class of energy-transfer problems, where a discrete set of eigenstates captures the dynamics of the system, we believe UF^2 will generally outperform SOP methods.

5.4 Example

In this section we study a system that has the smallest Hilbert space considered in Fig. 5.4, in which wavefunctions may be represented using $N_c = 28$ terms. This small size allows us to generate TA spectra in a matter of seconds, and therefore it is easy to run fast parameter sweeps and map out how spectroscopic observables change.

We explore nonlinear optical spectra for a model system presented by Tiwari and Jonas [84]. Their model system

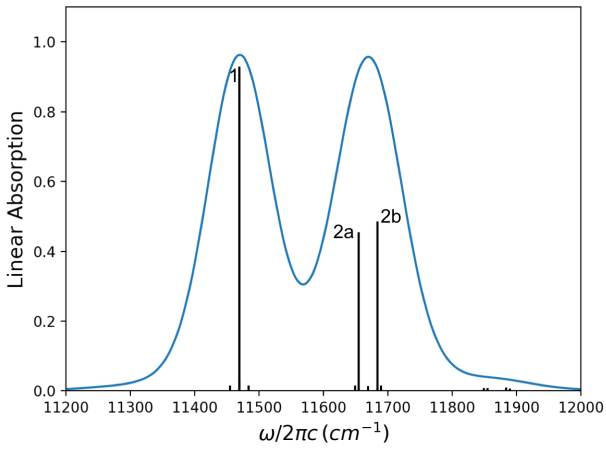

Figure 5.5: Linear absorption stick spectrum calculated using the same eigenstates as in Ref. 84. The frequencies of the stick spectra are the eigenenergies of the Hamiltonian Eq. 5.20 with $k = s = 2$, $\omega_i = 200\,\mathrm{cm}^{-1}$ and $S_i = 0.025$ for each monomer subunit. The two subunits have dipoles $|\mu_{ag}| = |\mu_{bg}|$ and $\mu_{ag} \cdot \mu_{bg} = 0$. The heights of the sticks are $|\hat{\mu}_{\phi'\phi}|^2$. The spectrum is thermally averaged at a temperature of 80 K, using the 3 lowest energy ground vibrational states, $|00\rangle, |01\rangle, |10\rangle$. The blue curve represents the absorption cross-section, which is obtained by broadening the stick spectrum with homogeneous and inhomogenous linewidths of $10\,\mathrm{cm}^{-1}$ and $45\,\mathrm{cm}^{-1}$, respectively.

is inspired by the Fenna-Matthews-Olson complex [85] and consists of a molecular dimer formed from two electronic TLS , each locally coupled to a single harmonic vibrational mode, which is the case of $k = s = 2$ from Eq. 5.20.

The system couples to optical pulses in the electric-dipole approximation, with transition dipole matrix elements μ_{ga} and μ_{gb}. Following Ref. 84, we consider the homodimer to have nearly identical subunits, with identical vibrational frequencies $\omega_a = \omega_b = 200\,\mathrm{cm}^{-1}$, Huang-Rhys factors $S = 0.025$, and with $E_b - E_a = 150$ cm^{-1}. We consider that μ_{ga}, μ_{gb} have the same magnitude but extend the model by varying the angle θ between these transition dipoles, corresponding to varying the angle between the two subunits. Due to the small Huang-Rhys factors considered in Ref. 84, only a small number of vibrational states contribute to third-order spectra, allowing truncation of the vibrational Hilbert space, as described in Sec. 5.2.5. Reference 84 did not calculate TA spectra; first we compare our results for linear absorption and then present our TA predictions.

Reference 84 demonstrates how the Huang-Rhys factors split one exciton peak into two vibronic peaks. We plot the linear absorption spectrum for this system in Fig. 5.5, choosing lineshape parameters to visually reproduce results from Ref. 84. We label the peaks in Fig. 5.5 following Tiwari and Jonas. There are two broad peaks centered around $f_1 = 11469.5\,\mathrm{cm}^{-1}$ and $f_2 = 11669.6\,\mathrm{cm}^{-1}$, with the latter composed of two peaks, $2a$ and $2b$, separated by $\Omega_1 = 29\,\mathrm{cm}^{-1}$. However, this splitting is invisible to linear absorption because peaks $2a$ and $2b$ smear together.

Ω_1 is, however, the dominant beat frequency in a TA experiment, shown in Fig. 5.6, corresponding to oscillations repeating about every 1.15 ps. The TA signal is detected as two broad features centered at f_1 and f_2. Notably, the oscillations centered around f_1 are roughly 90° out of phase with those at f_2, and there is a node halfway between them. This type of nodal feature is widely observed and has been studied in many vibrational systems [15, 86, 87].

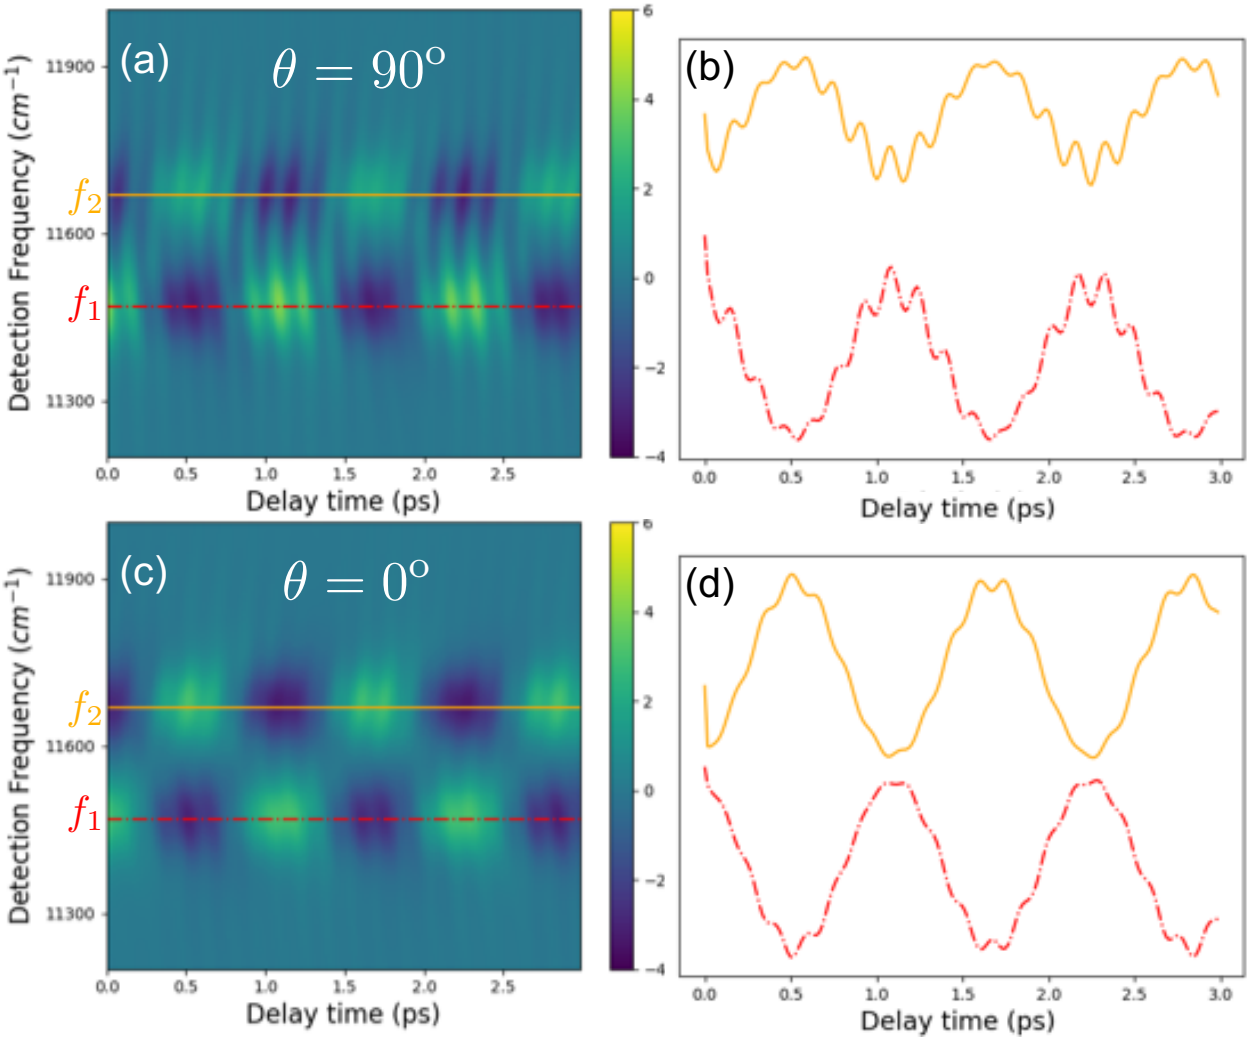

Figure 5.6: Transient absorption with Gaussian pump and probe pulses centered on f_2 and with FWHM of 12 fs. (a) Dipole angle $\theta = 90^o$, (b) linecuts of (a) at f_2 (orange) and f_1 (red, dash-dotted), (c) Dipole angle of $\theta = 0^o$, (d) linecuts of (c) at f_2 (orange) and f_1 (red, dash-dotted). In both cases the low-frequency beats are out of phase between f_1 and f_2. The high-frequency beats are in phase between f_1 and f_2 for $\theta = 90^o$ and are reduced in amplitude and out of phase for $\theta = 0^o$.

As mentioned in Section 5.2.5, to achieve convergence of these TA spectra to within 1%, it is sufficient to have $N_{f,DEM} = N_{f,GSM} = 15$ and $N_{f,SEM} = 30$, while $N_{c,DEM} = N_{c,GSM} = 10$ and $N_{c,SEM} = 8$. Each TA spectrum is isotropically averaged, with the same homogeneous and inhomogeneous linewidths in the linear absorption spectrum. Each TA spectrum, which is of the form $S_{TA}(T,\omega)$, consists of $m_T = 180$ delay times, and $M_t = 286$, 562, 1391 detection frequency points for FWHM pulse durations of 62, 31 and 12 fs, respectively, and takes about 3 seconds to generate on a 2.3GHz Macbook Pro. The equivalent calculation using the SOP code (see Section 5.3.3) would take almost 6 days to converge to the same accuracy if run on the same machine. This system is the leftmost point of Fig. 5.4, and the RKE method would take over 7 min to produce each spectrum.

The efficiency of UF^2 enables rapid study of a wide range of parameters. The frequency-integrated TA signal,

$$S_{FI,TA}(T) = \int \mathrm{d}\omega S_{TA}(T,\omega), \tag{5.24}$$

is particularly sensitive to electric field shape [58, 88]. Figure 5.7(a,b) shows frequency-integrated TA spectra with dipole angles $\theta = 20°$, 58° and three optical pulse durations. The responses with different pulse shapes are quite different for both values of θ. In order to study the differences over a wide range of θ we study the Fourier transform of $S_{FI,TA}(T)$ with respect to the delay time:

$$\tilde{S}_{FI,TA}(\omega_T) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathrm{d}T S_{FI,TA}(T) e^{i\omega_T T},$$

and track the response at $\omega_T = \Omega_1$ and $\omega_T = \Omega_2 = f_{2a} - f_1 = 185.5\,\mathrm{cm}^{-1}$ in Fig. 5.7(c,d). A nearly impulsive pulse (12 fs FWHM - blue curves) gives a dramatically different response as compared to somewhat longer pulses (35 fs - red curves, 70 fs - green curves). Notice that the longest pulse considered here is still short compared to the fastest time scale of the system, Ω_3, corresponding to a period of 156 fs, so the large changes of the signals with pulse duration are a warning sign that accurate study of finite-pulse effects is important to correctly model nonlinear spectra. This study is easily tractable using UF^2, taking only 30 minutes to run, but would have taken 3 days to run using RKE and more than a year using our SOP implementation on the same computer.

5.5 Conclusion

UF^2 is a computationally efficient method for calculating perturbative nonlinear spectroscopies of systems with finite dimensional relevant Hilbert spaces. It includes all finite-duration and pulse-overlap effects, enabling accurate modeling of and fitting to nonlinear spectroscopic data. The current introduction has focused on closed systems, but UF^2 can be readily extended to open systems with time-independent Liouvillians, in which case its cost scales as N_c^4 rather than N_c^2, where N_c is the relevant Hilbert space dimension. We have shown the methods that can be used to aggressively reduce the required dimension N_c, which enables UF^2 to be highly efficient for small systems and

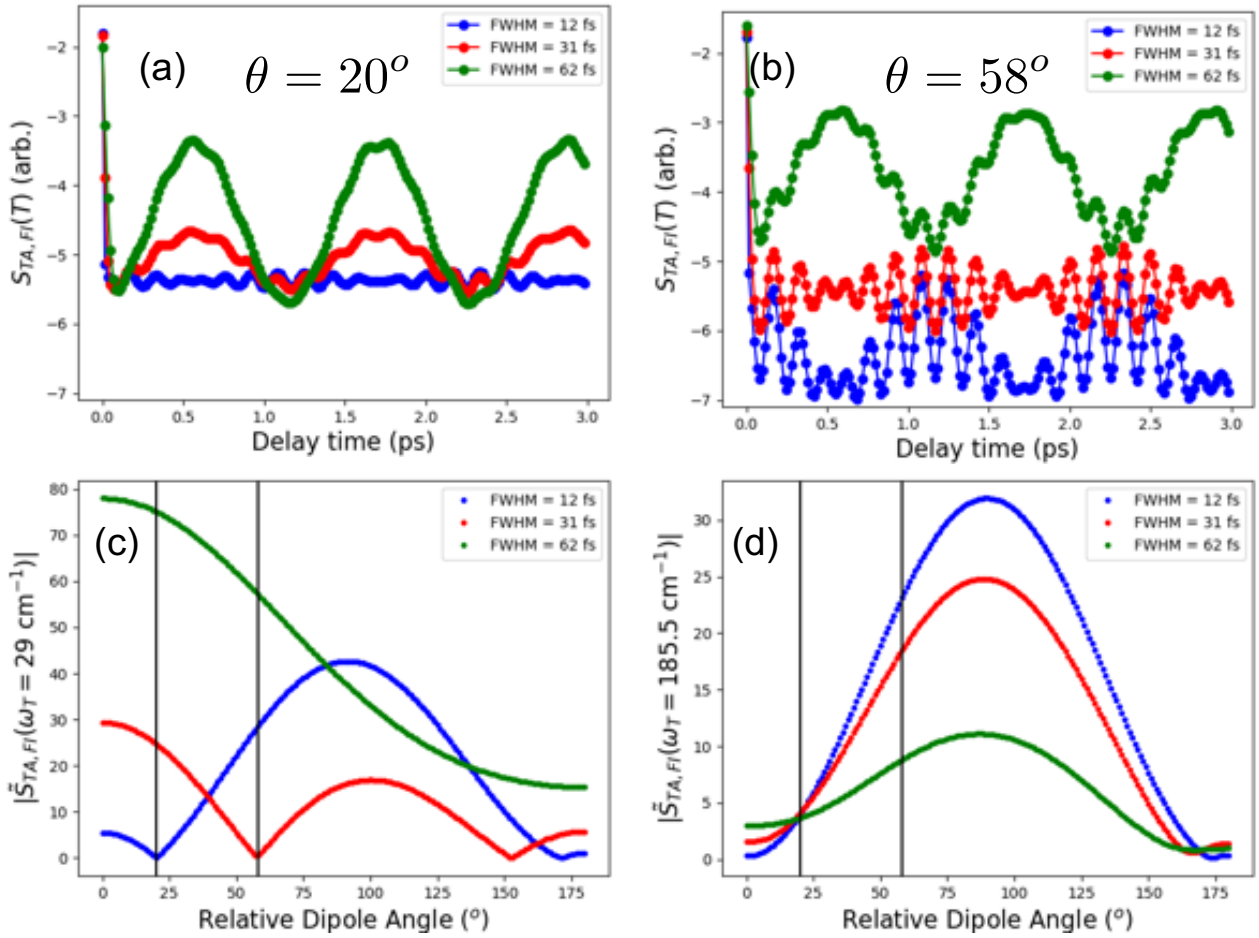

Figure 5.7: Demonstration of parameter variations for the same system as in Fig. 5.6. Panels (a) and (b) show frequency-integrated transient absorption (FITA) spectra for (a) $\theta = 20°$ and (b) $\theta = 58°$ with three different Gaussian pulse durations. There are two prominent time scales present: the Ω_1 beat (corresponding to the 1.15 ps oscillation) and the Ω_2 and Ω_3 beats (corresponding to the 180 fs and 156 fs oscillations). Panels (c) and (d) show the magnitude of these beats of the FITA spectra at (c) $\omega_T = \Omega_1$ and (d) $\omega_T = \Omega_2$ as a function of θ. Black vertical lines indicate θ from panels (a) and (b). Note that at $\omega_T = \Omega_1$, the FITA magnitude shows a dramatic dependence on both dipole angle and electric field shape, demonstrating the importance of taking into account electric field shapes. The signals at $\omega_T = \Omega_2$ are less sensitive.

competitive for surprisingly large problems. For vibronic systems with thousands of relevant states, UF^2 generally outperforms our comparison DP method and is competitive for many systems with tens of thousands of relevant states; for larger systems, DP methods are generally superior.

The intent of the publicly available UF^2 code is to present a package that makes it easy to translate Feynman diagrams into code. Since UF^2 readily enables consideration of arbitrary pulse shapes, the effects of pulse durations, chirps, etc. on nonlinear spectroscopies can now be included as a matter of course in analyses of experiments. Given the large number of problems under active study with tens to hundreds of relevant states, we hope that UF^2 will enable easy and rapid spectroscopic prediction and analysis including pulse shape effects that are often neglected.

Acknowledgments

We acknowledge helpful conversations with Joel Yuen-Zhou, Ivan Kassal, Mark Embree, Luc Robichaud, and Eduard Dumitrescu. We also acknowledge funding from the Natural Sciences and Engineering Research Council of Canada and the Ontario Trillium Scholarship.

5.6 Appendix A: Computational cost of UF^2

We derive the computational cost of UF^2, highlighting which parameters are required for convergence and which are at the user's discretion, focusing on the case of a 2DPE rephasing spectrum, as in Eq. 5.8.

We derive the computational cost of calculating the 2DPE rephasing signal. We briefly outline what is required for this calculation by working backwards from Eq. 5.8. Symbols used in this section are summarized in Table 5.3. Calculation of the signal requires that we determine $\mathbf{P}^{(3)}_{\text{2D}}(\tau, T; \omega)$ at the desired values of τ and T. UF^2 directly calculates $\mathbf{P}^{(3)}_{\text{2D}}(t)$ for one pair of (τ, T) at a time. We calculate $\mathbf{P}^{(3)}_{\text{2D}}(t)$ at sufficient time points, M_t, in order to obtain the desired frequency resolution. The cost of the FFT to obtain $\mathbf{P}^{(3)}_{\text{2D}}(\omega)$ from $\mathbf{P}^{(3)}_{\text{2D}}(t)$ is negligible compared with the other costs of UF^2. Assuming that we require the polarization field at m_τ values of τ and m_T values of T, the cost of the full spectrum is $m_\tau m_T \text{Cost}(\mathbf{P}^{(3)}_{\text{2D}}(t))$. $\mathbf{P}^{(3)}_{\text{2D}}(t)$ is a sum of three Feynman diagrams (see Fig. 5.1), each of which is calculated separately. We focus on $\mathbf{P}^{(3)}_{ESA}(t)$, which is often the dominant cost. The cost of the other diagrams follows directly from this derivation.

Using Eq. 5.17, the cost of $\mathbf{P}^{(3)}_{ESA}(t)$ can be broken into three parts: calculating the two necessary perturbed wavepackets and the cost of the dipole matrix element of those wavepackets. This latter cost turns out to be the dominant cost asymptotically. Since $|\psi_a(t)\rangle = K_a \left|\psi^{(0)}(t)\right\rangle$ and $|\psi_{bc}(t)\rangle = K_c K_b \left|\psi^{(0)}(t)\right\rangle$, the cost of calculating these wavepackets is, at first glance, the cost of three calls to the $K_{j^{(*)}}$ operator, which is the heart of UF^2. However, $|\psi_a(t)\rangle$ corresponds to the interaction with the first pulse, which arrives at a fixed time, and therefore need only be calculated once. Further, $|\psi_b(t)\rangle$ corresponds to the second pulse, which arrives at m_τ different coherence times τ. Therefore $|\psi_b(t)\rangle$ only needs to be calculated m_τ different times. The only wavefunction that must be recalculated

Table 5.3: Summary of symbols used to describe computational cost of a 2DPE rephasing signal. M, M_E, ϵ, and dt are convergence parameters.

Symbol	Definition
M	Number of time points to resolve pulses UF^2
M_E	Number of time points to resolve pulse in RKE
ϵ	Local tolerance of RK45 algorithm
dt	Fixed step size for RKE and UF^2 during pulses
dt_{RK}	Adaptive step size of RK45 algorithm

for every pair τ, T is the one caused by the third interaction, $|\psi_{bc}(t)\rangle$. Therefore the dominant cost of calculating the perturbative wavefunctions for many different values of T is simply one call to K_c.

To determine the cost of $K_{j^{(*)}}$, we first explain how wavefunctions are stored in UF^2. Each wavefunction is represented in the same way. For example, the second-order wavefunction is

$$|\psi_{bc}(t)\rangle = \sum_{\phi=1}^{N_{c,DEM}} e^{-i\omega_\phi t} c_{\phi,ac^*}(t) \, |\phi\rangle \, , \tag{5.25}$$

where each of the $c_{\phi,ac^*}(t)$ are calculated at M evenly spaced time points . The spacing dt is determined by the shape of the pulse amplitude $A_i(t)$ (see Eq. 5.4). Recall from Eq. 5.16 that $c_{\phi,pj^{(*)}}(t)$ only varies in the interval $(t_{j,\min}, t_{j,\max})$. Outside of this interval, $c_{\phi,pj^{(*)}}$ is constant and can thus be trivially extended to any time points needed. This property allows wavefunctions to be stored for re-use without a significant memory cost. For simplicity in this discussion we assume that $t_{j,\max} - t_{j,\min}$ is the same for each pulse. Since we are discretizing a continuous convolution integral, $M = \frac{t_{j,\max} - t_{j,\min}}{dt}$ is a convergence parameter. For any electric field that does not go strictly to zero, the choice of $t_{j,\min}$ and $t_{j,\max}$ must also be checked for convergence.

Inspecting Eq. 5.14, the cost of $\hat{K}_j$ is $N_{c,X}$ times the cost of evaluating $\theta * y_\phi$ to obtain $c_{\phi,pj}$. Calculating y_ϕ is dominated by the cost of taking the dipole-weighted sum over ϕ' (we are neglecting the small cost of multiplying by the pulse amplitude $\varepsilon_j^{(*)}$ and by the time evolution factors $e^{i\omega_\phi(t-t')}$ and $e^{-i\omega_{\phi'}(t-t')}$). The sum has a cost of $\alpha N_{c,X} M$, where α is the cost of multiplying two complex numbers. We calculate $\theta * y_\phi$ using the convolution theorem and the FFT. Since we are interested in the linear convolution, we evaluate θ at $2M-1$ points and zero-pad y_ϕ to be size $2M-1$. We only calculate the FFT of θ once, making that cost negligible. We calculate the FFT of y_ϕ, which has a cost of $\beta'(2M-1)\log_2(2M-1) \approx 2\beta' M \log_2 M + O(M)$, where β' depends upon the implementation of the FFT.[iv] We multiply the two ($\tilde{y}_\phi$ and $\tilde{\theta}$), which has cost $O(M)$, and then take the inverse FFT of the product (cost $2\beta' M \log_2 M + O(M)$). Thus the total cost of the convolution is $4\beta' M \log_2 M + O(M)$. The cost to obtain

[iv] We find that when calculating the convolution between a step-function $\theta(t)$ and another function $y(t)$, we achieve convergence much more quickly when we use an odd number of time points, with half positive, half negative, and the point $t = 0$ with associated value $\theta(0) = 0.5$. The FFT is fastest for arrays of length 2^k, which is impossible in this case. To keep β' small, we use the FFTW library and pick a size $2M-1$ that has only small prime-factors.

one coefficient $c_{\phi,pj^{(*)}}(t)$ to highest order in M and N is then

$$\text{Cost}(c_{\phi,pj^{(*)}}(t)) = \alpha N_{c,X} M + \beta M \log_2 M + O(M), \tag{5.26}$$

where $\beta = 4\beta'$. The cost of calling $K_{j^{(*)}}$ is then $N_{c,Y}\text{Cost}(c_{\phi,pj^{(*)}})$, where $N_{c,X}$ is the dimension of the old manifold and $N_{c,Y}$ is the dimension of the new manifold, giving

$$\text{Cost}(K_{j^{(*)}}) = N_{c,Y} M(\alpha N_{c,X} + \beta \log_2 M) + O(M N_{c,Y}), \tag{5.27}$$

and so $\text{Cost}(K_{j^{(*)}}) \sim M N_{c,X} N_{c,Y}$ for large N_c. In the case of $|\psi_{bc}(t)\rangle$, then, we have a cost $\sim M N_{c,SEM} N_{c,DEM}$. Given that $N_{c,DEM}$ is often the largest value of $N_{c,X}, X \in \{GSM, SEM, DEM\}$, the ESA diagram often dominates the cost of UF^2. The other diagrams involve wavefunctions that move between the GSM and the SEM, and therefore the cost of each K_j scales as $N_{c,GSM} N_{c,SEM}$.

Once we know each of the necessary wavefunctions at its M time points, we evolve the wavefunctions using $\hat{U}_0$ to include M_t times points, in order to obtain the desired frequency resolution of the final spectrum. The M_t points are spaced by the same dt, and span from just before the last pulse arrives, $t_{c,\text{min}}$, until the signal has decayed to an appropriate cut-off due to the optical dephasing described in Section 5.2.3. This evolution is of negligible cost since we know the exact form of $\hat{U}_0(t)$.

We then calculate the expectation value $\boldsymbol{P}^{(3)}_{ESA}(t) = \langle\psi_a(t)|\,\hat{\boldsymbol{\mu}}\,|\psi_{bc}(t)\rangle$, which has a cost of $\alpha N_{c,SEM} N_{c,DEM}(M + M_t)$, since the polarization field must capture both the pulse turn-on described by M points, and the decay of the field described by the additional M_t points, so in total $\text{Cost}(\mathbf{P}^{(3)}_{ESA}(t)) = \alpha N_{c,SEM} N_{c,DEM}(M + M_t) + \text{Cost}(K_j)$. The signal must be calculated at m_τ τ points and m_T T points, so we arrive at the full cost of the ESA signal

$$\begin{aligned}\text{Cost}(S_{ESA}(\tau, T, \omega)) = m_\tau m_T \Big(&\alpha N_{c,SEM} N_{c,DEM}(M + M_t) \\ &+ N_{c,DEM} M\big(\alpha N_{c,SEM} + \beta \log_2 M\big)\Big)\end{aligned} \tag{5.28}$$

to highest order in $N_{c,X}$, M_t and M. In general we find that $M < M_t$, and the cost of evaluating the necessary wavefunctions is sometimes less than half of the cost of obtaining the total signal. In Section 5.4, we calculate the TA signal, which is composed of four Feyman diagrams with costs less than or equal to that in Eq. 5.28, with $m_\tau = 1$, since the TA signal has $\tau = 0$. In Section 5.4 we use $M = 21$ and $M_t \approx 1000$, in order to achieve convergence of better than 1%.[v]

[v] The value of M_t depends upon the pulse shape and the optical dephasing γ. We resolve the polarization field from $t_{c,\text{min}}$ until $t_c + 6.91/\gamma$.

Assuming that $M < M_t$, and that $\alpha N_c \gg \beta \log_2 M$,[vi] the total cost scales as

$$C_{\mathrm{UF}^2}(S_{ESA}) \sim m_\tau m_T N_{c,SEM} N_{c,DEM} M_t. \tag{5.29}$$

Note that UF^2 assumes that the eigenvalues and eigenstates of $\hat{H}_0$ have already been attained. If the dimension of $\hat{H}_0$ is N_f, then the computational cost of that diagonalization can scale as N_f^3, which can exceed the cost of UF^2 itself, especially since UF^2 reduces the relevant Hilbert space dimension N_c so aggressively. For many systems, however, $\hat{H}_0$ is sparse, allowing efficient iterative methods to be used to find its eigenvalues and eigenstates. The scaling of those algorithms is beyond the scope of this manuscript, but we simply note that there are many important cases where solving $\hat{H}_0$ is not computationally limiting. For example, vibronic systems of coupled chromophores, each with local harmonic oscillators, are extremely sparse, and an efficient algorithm to determine their eigenstates will be detailed separately.

5.7 Appendix B: RKE implementation and scaling

5.7.1 Implementation

Direct propagation (DP) methods solve the Schrodinger equation

$$\frac{d\left|\psi^{(n)}(t)\right\rangle}{dt} = -iH_0\left|\psi^{(n)}(t)\right\rangle - iH'_j(t)\left|\psi^{(n-1)}(t)\right\rangle,$$

where we have set $\hbar = 1$, by propagating the ODE forward in time domain.

In order to compare the cost of UF^2 to the cost of DP methods, we have implemented a hybrid DP method that uses the adaptive-step size 4-5$^{\mathrm{th}}$ order Runge-Kutta ODE solver from scipy (called RK45) to propagate the system dynamics

$$\frac{d\left|\psi^{(n)}\right\rangle}{dt} = -iH_0\left|\psi^{(n)}\right\rangle,$$

and uses the first-order Euler method to include the perturbation

$$-iH'_j(t)\left|\psi^{(n-1)}(t)\right\rangle,$$

with a previously calculated lower-order wavepacket $\left|\psi^{(n-1)}(t)\right\rangle$, as described in Refs. 21, 59. As with UF^2, we assume that each pulse $\mathbf{E}_j$ is non-zero only during times $t_{j,\min} < t < t_{j,\max}$. When all of the pulses are zero, propagation is simply handled by RK45. Given a local absolute and relative tolerance ϵ, RK45 can integrate from t_i to t_f using an adaptively determined step size dt_{RK}. The wavefunctions are stored at equally spaced time points with time step dt, which must be short compared to the optical pulses and in our implementation we set to the same

[vi] In our implementation of this algorithm, $\beta \log_2 M/\alpha \approx 100$.

dt used in UF^2. This fixed time step is used for the Euler method and allows the reuse of lower-order wavefunctions for many values of pulse delay times. Note that as mentioned at the end of Section 5.3.2, this choice of dt does not converge the RKE spectra to the same accuracy as UF^2.[vii] Given a state vector $|\psi(t_i)\rangle$, we use the notation $|\psi(t_i + dt)\rangle = U_0(dt)\,|\psi(t_i)\rangle$ to represent the RK45 propagation.

To include the perturbation $H'_j(t)$, we begin by defining the first-order wavefunction $|\psi_a(t < t_{a,\min})\rangle = 0$. We define $|\psi_a(t_{a,\min})\rangle = -iH'_a(t_{a,\min})dt\,\left|\psi^{(0)}(t_{a,\min})\right\rangle$, where $\left|\psi^{(0)}(t_{a,\min})\right\rangle$ is a stationary state (possibly with an evolving phase factor). We propagate until $t_{a,\max}$ in steps of size $dt = (t_{a,\max} - t_{a,\min})/M_E$. Let m run from 0 to M_E. Then

$$|\psi_a(t_{a,\min} + mdt)\rangle = U_0(dt)\,|\psi_a(t_{a,\min} + (m-1)dt)\rangle + i\mu \cdot \mathbf{e}_a \varepsilon_a(t_{a,\min} + mdt)\left|\psi^{(0)}(t_{a,\min} + mdt)\right\rangle.$$

Once we have obtained $|\psi_a(t_{a,\max})\rangle$, we obtain $|\psi_a(t > t_{a,\max})\rangle$, using RK45 alone. This method works in general for the n^{th} order wavefunction generated by an interaction with the j^{th} pulse. Given the $(n-1)^{\text{st}}$ wavefunction $|\psi_p(t)\rangle$, we define $\left|\psi_{pj^{(*)}}(t_{j,\min})\right\rangle = i\mu \cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t_{j,\min})\,|\psi_p(t_{j,\min})\rangle$, and then at all times until $t_{j,\max}$ using

$$\left|\psi_{pj^{(*)}}(t_{a,\min} + mdt)\right\rangle = U_0(dt)\left|\psi_{pj^{(*)}}(t_{a,\min} + (m-1)dt)\right\rangle + i\mu \cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t_{a,\min} + mdt)\,|\psi_p(t_{a,\min} + mdt)\rangle.$$

Once again, for all times after $t_{j,\max}$, we only use RK45 to propagate the wavefunction in its manifold.

5.7.2 Computational cost analysis

We derive the computational cost of calculating the 2DPE rephasing signal, now using the RK45-Euler hybrid (RKE) method outlined in Section 5.7.1, following a similar path as our derivation of the cost of UF^2, working backwards from Eq. 5.8. Symbols used in this section are summarized in Tables 5.1-5.3. As for UF^2, we begin by considering the ESA diagram, and derive the cost of $\mathbf{P}^{(3)}_{ESA}(t)$.

The cost of RKE is the cost of obtaining the necessary wavefunctions, $|\psi_a(t)\rangle$ and $|\psi_{bc}(t)\rangle$, and computing $\mathbf{P}_{ESA}(t) = \langle\psi_a(t)|\,\mu\,|\psi_{bc}(t)\rangle$. In the Condon approximation, μ is a sparse rectangular matrix consisting of diagonal blocks, so this expectation value has a cost of $\alpha q r_\mu N_{f,DEM}$. Even outside of the Condon approximation, μ is still sparse, and therefore the cost is still $O(N_f)$; the number of entries per row, r_μ, is simply a little larger.

The remaining cost is the cost of obtaining the necessary wavefunctions. The cost of obtaining a perturbative wavefunction breaks down into two parts: the cost of the RK45 algorithm to propagate H_0 and the cost of the Euler method to include the perturbation H'_j. We first derive the cost of RK45 for times after $t > t_{j,\max}$, where RK45 can take full advantage of its adaptive step size. In this regime, the user sets a local tolerance ϵ, and RK45 determines

[vii] For problems with a small number of delay times, it may be more efficient to use the RK45 solver for both system and perturbation propagation, but this choice would make wavefunction reuse more complicated. As this manuscript is interested in the case where many different experimental configurations are considered, we have chosen an implementation that we believe gives the best performance, though we do not claim to have optimized all aspects.

what step size dt_{RK} satisfies that tolerance, and we approximate it as a constant. To take each step, RK45 must make 6 function evaluations of the form $H_0 |\psi\rangle$ and sum the results with the required weights. Each evaluation has the cost of a sparse matrix-vector multiplication, $\alpha q r N_{n,X}$, where r is the average number of non-zero entries of H_0 per row and q is an additional overhead factor for sparse matrix operations. Ignoring the extra costs when steps are rejected, we thus find that the cost to take a step dt_{RK} is $6\alpha(qr+1)N_f \approx 6\alpha q r N_{f,X}$. Assuming we need to know the wavefunction at some final time $t_f > t_{j,\max}$, the cost of propagating the wavefunction when the pulse is off,

$$C_{RK} = 6\alpha q r N_{f,X} M_{RK},$$

where $M_{RK} = (t_f - t_{j,\max})/dt_{RK}$ is the required number of steps after the pulse ends. This wavefunction can be relatively inexpensively evaluated on a mesh with spacing dt using the RK45 interpolator, but the computational cost is set by the adaptively optimized dt_{RK}.

During the time when the pulse is non-zero, the cost of the RK45 portion of the evolution is

$$6\alpha q r N_{n,X} \max\left(M_E, (t_{j,\max} - t_{j,\min})/dt_{RK}\right),$$

where $M_E = (t_{j,max} - t_{j,min})/dt$. If dt is smaller than dt_{RK}, we must pay an additional cost of taking smaller steps than field-free evolution using RK45 would require. We must also pay the cost of adding in the perturbation. Again, since μ is a sparse matrix with r_μ entries per row, the cost of $\mu |\psi\rangle$ is $\alpha q r_\mu N_f$. The cost of evaluating ε_j at a single time point is negligible, and so the cost of adding in the perturbation is $\alpha q r_\mu N_{n,X} M_E$. Assuming that $dt < dt_{RK}$, we can easily add these costs to obtain the cost of propagating when the pulse is on,

$$C_E = \alpha q (6r + r_\mu) N_{n,X} M_E.$$

Thus the cost of calculating a single wavefunction is

$$C_{RKE}(\left|\psi_{pj^{(*)}}\right\rangle) = C_E + C_{RK}$$

Despite the fact that dt is generally smaller than dt_{RK}, it is still often the case that $M_{RK} > M_E$. In general we find that $M_E = 100$ is usually sufficient to resolve the pulse interaction and converge the spectroscopic signals to within 1%. In contrast, we often find that $M_{RK} > 100$, and therefore the cost of obtaining a perturbative wavefunction is usually dominated by the propagation cost in the absence of the electric field. This observation depends upon the imposed decay of the polarization field, γ, as described in Section 5.2.3. As for UF^2, once the polarization field is obtained, we multiply by $e^{-\gamma|t-t_{probe}|}$, and therefore we generally need to resolve the wavefunction to $t_f = t_{\text{probe}} + \frac{5}{\gamma}$.

We developed this hybrid method for the purposes of making a fair comparison with UF^2. We assert that if

$C_E > C_{RK}$, there is likely a more efficient algorithm, for example using RK45 both while the pulse is on and when the pulse is off. For the purposes of comparison, we make sure to operate in a regime where $C_E < C_{RK}$, and therefore we approximate $C_{RKE} \approx C_{RK}$.

This process must be repeated 3 times in order to calculate the ESA at a single pair of pulse delays τ, T. However, $|\psi_a\rangle$ only needs to be calculated once, and can be stored and re-used because we are working with a regularly spaced time grid. Similarly, $|\psi_b\rangle$ only needs to be calculated once for each value of τ but can be reused for all values of T. The only wavefunction that must be recalculated each time is $|\psi_{bc}\rangle$, and therefore this cost dominates the cost of obtaining $\mathbf{P}_{ESA}$. Therefore

$$C_{RKE}(S_{ESA}) \approx m_\tau m_T C_{RK} \tag{5.30}$$

Note that the costs of the other diagrams contributing to $\mathbf{P}^{(3)}(t)$ have a similar form, but scale as $N_{f,SEM}$ or $N_{f,GSM}$. Since $N_{f,SEM}, N_{f,GSM} < N_{f,DEM}$ in general, the cost of obtaining the rephasing signal is dominated by the cost of $\mathbf{P}_{ESA}$.

Chapter 6

Open UF^2

Reproduced from Journal of Chemical Physics 154 034108 (2021), with the permission of AIP Publishing

Original title: *Efficient numerical method for predicting nonlinear optical spectroscopies of open systems*

Peter A. Rose and Jacob J. Krich

Original abstract: Nonlinear optical spectroscopies are powerful tools for probing quantum dynamics in molecular and nanoscale systems. While intuition about ultrafast spectroscopies is often built by considering impulsive optical pulses, actual experiments have finite-duration pulses, which can be important for interpreting and predicting experimental results. We present a new freely available open source method for spectroscopic modeling, called Ultrafast Ultrafast (UF^2) Spectroscopy, which enables computationally efficient and convenient prediction of nonlinear spectra, including treatment of arbitrary finite duration pulse shapes. UF^2 is a Fourier-based method that requires diagonalization of the Liouvillian propagator of the system density matrix. We also present a Runge-Kutta Euler (RKE) direct propagation method. We include open-systems dynamics in the secular Redfield, full Redfield, and Lindblad formalisms with Markovian baths. For non-Markovian systems, the degrees of freedom corresponding to memory effects are brought into the system and treated nonperturbatively. We analyze the computational complexity of the algorithms and demonstrate numerically that, including the cost of diagonalizing the propagator, UF^2 is 20-200 times faster than the direct propagation method for secular Redfield models with arbitrary Hilbert space dimension; that it is similarly faster for full Redfield models at least up to system dimensions where the propagator requires more than 20 GB to store; and that for Lindblad models it is faster up to Hilbert space dimension near 100, with speedups for small systems by factors of over 500. UF^2 and RKE are part of a larger open source Ultrafast Software Suite, which includes tools for automatic generation and calculation of Feynman diagrams.

Notes:

- Added a brief discussion that did not appear in the original publication about comparing the results of simulations done using RKE and UF^2. Discussion is added to captions of Figures 6.3 and 6.4, and also added to the main text in Section 6.4.

6.1 Introduction

Nonlinear optical spectroscopies (NLOS) are widely used tools for probing the excited state dynamics of a wide range of systems [70, 89]. The signals that can be measured using NLOS contain a wealth of information, but correctly interpreting that information generally requires making a model of the system and predicting the spectra that result. Such analysis can require repeated lengthy computations in order to fit multiple parameters to the collected data [7–10]. Fast methods for simulating spectra of model system enable better interpretation of experimental results.

NLOS are often calculated in the impulsive limit of infinitely short optical pulses. Recent work has shown that finite pulse effects can have dramatic effects on measured NLOS, and that fitting experimental data using intuition developed in the impulsive limit can lead to incorrect conclusions [57], adding to the existing body of work exploring the effects of finite pulse shapes [13, 15, 17–19, 88, 90–95]. The effects of Gaussian and exponential pulse shapes have been treated analytically for various types of NLOS, providing valuable insights into the effects of pulse shapes and durations [18, 19]. However, real experimental pulses are often not well represented by Gaussian or other analytical shapes. Ideally, modeling of NLOS should include actual experimental pulse shapes rather than approximate forms, and a number of numerical methods have this capability [21, 59, 64, 71, 74, 90, 91].

In Ref. [96] we introduced a novel fast algorithm based on Fourier convolution, called Ultrafast Ultrafast (UF^2) spectroscopy, capable of simulating any order NLOS using arbitrary pulse shapes. We compared it to our own implementation of a standard direct propagation method that we called RKE (Runge-Kutta-Euler) and demonstrated that UF^2 shows a significant speed advantage over RKE for systems with a Hilbert space dimension smaller than 10^4. However, that work is based upon wavefunctions and is only valid for closed systems. Condensed-phase systems consist of too many degrees of freedom to treat them all explicitly, leading to essential dephasing and dissipation, and making wavefunction methods of limited use in interpretation of experiments.

In this work we present the extension of both UF^2 and RKE to open quantum systems with Markovian baths. Degrees of freedom corresponding to memory in the bath can be included explicitly in the system Hamiltonian, while the rest of the bath is assumed to be weakly coupled and treated perturbatively using Redfield or Lindblad formalisms. We show that UF^2 is over 200 times faster than RKE for small system sizes, and we believe this result is representative of the advantage that UF^2 provides over direct propagation methods. With a secular Redfield model, UF^2 outperforms RKE for all system sizes. Hereafter, the terms UF^2 and RKE refer to the new open extensions of the old algorithms of the same name, with the understanding that the closed system algorithms are now contained as special cases.

UF^2 works in the eigenbasis of the Liouvillian that propagates system density matrices and thus requires diagonalization of this Liouvillian. We show that, surprisingly, the cost of this diagonalization is negligible for the system sizes where UF^2 outperforms RKE, despite the Liouvillian having dimension N^2. Diagonalization yields fast, exact propagation of the unperturbed system and allows the optical pulses to be included using the computational efficiency of the fast Fourier transform (FFT) and the convolution theorem. UF^2 requires only that the

pulse envelope be known at a discrete set of time points, and thus is able to study any pulse shape of interest, including experimentally measured pulse shapes. As few as 25 points are required with Gaussian pulses to obtain 1% convergence of spectra.

UF^2 and RKE are part of a software package we call the Ultrafast Spectroscopy Suite (UFSS), outlined in Fig. 6.1, which is designed to simplify the process of predicting spectra or fitting spectra to models. UFSS is designed in particular to facilitate inclusion of finite pulse effects with low computational cost. There are two distinct effects of finite pulses. First is the inclusion of additional Feynman diagrams that must be calculated when pulses overlap in time. UFSS includes an automated Feynman diagram generator (DG), described in Ref. 97, which automates the construction of these diagrams and determination of which ones give non-negligible contributions. Second is the calculation of the contribution from each diagram. Both UF^2 and RKE take in diagrams and calculate their contributions including the effects of pulse shapes. UFSS also contains a Hamiltonian and Liouvillian generator (HLG), described in this manuscript, which parametrically constructs models for vibronic systems. Each of the packages in UFSS can be used independently. In this work we demonstrate how UF^2 and RKE can be used separately, as well as with the HLG and DG. UFSS is free and open-source software written in Python, available for download from github.

UF^2 and RKE are numerical methods for including effects of optical pulse shapes in the perturbative limit, given that the equations of motion for an open quantum system in the absence of the pulses are known. The efficient inclusion of finite pulse durations in UF^2 relies on having a time-independent propagation superoperator for the density matrix in the absence of optical fields. There are many methods for describing the field-free dynamics of open quantum systems, including Lindblad theory [98], Redfield theory [99], multi configurational time-dependent Hartree (MCTDH) [100, 101], the hierarchical equations of motion (HEOM) [28, 102], and more [103]. Both Lindblad and Redfield theory allow treatment of dephasing and relaxation due to a Markovian bath, resulting in time-independent system propagators, and we implement both in UF^2. While both HEOM and MCTDH include non-Markovian bath dynamics, they do not yield time-independent Liouvillians and are not amenable to the techniques in UF^2; for those methods, slower direct propagation methods are still required.

In Sec. 6.2 we briefly review the formalism of NLOS calculated using time-dependent perturbation theory and then derive the UF^2 and RKE algorithms. The computational complexity of these methods is shown in Appendix 6.7. While UF^2 can propagate many types of systems, in Sec. 6.3 we describe the HLG built-in to UFSS. In Sec. 6.4 compare the computational cost of UF^2 and RKE for a range of system sizes generated by the HLG. In Sec. 6.5 we show the accuracy of UF^2 by comparing to analytical expressions for the 2D photon echo signal of the optical Bloch equations perturbed by Gaussian pulses from Ref. 19. We demonstrate that UF^2 quantitatively agrees with the analytical results, including effects of finite pulses, using just 25 evenly spaced points to represent the Gaussian pulse shape.

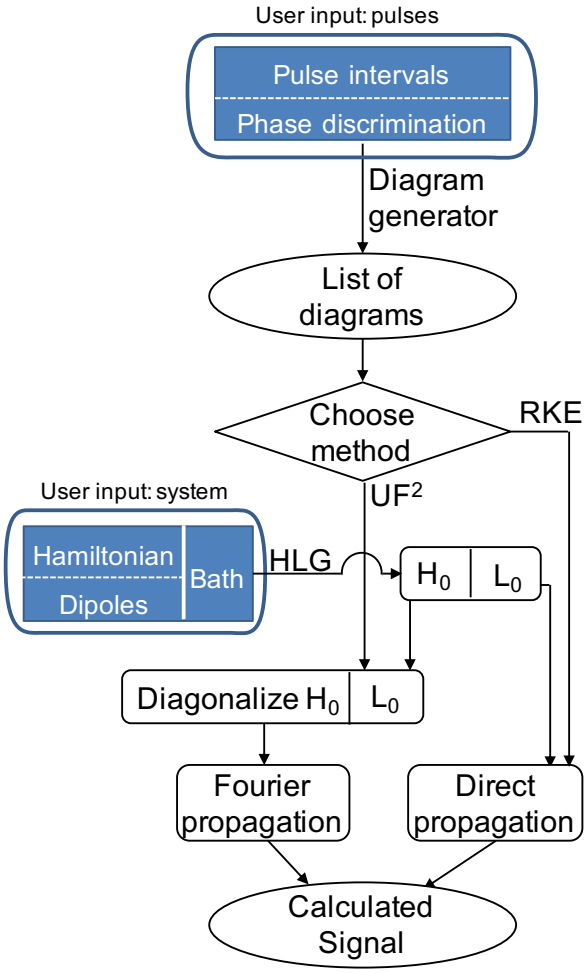

Figure 6.1: Logical flow of the UFSS package. Users input information about the experiment (pulse intervals and pulse-discrimination condition) and the system (states, optical dipoles, and system-bath interaction). UFSS consists of the diagram generator (DG), the Hamiltonian/Liouvillian generator (HLG), and two choices of propagators: UF^2 and RKE. The DG produces a list of Feynman diagrams, as described in Ref. [97]. This list and the Hamiltonian or Liouvillian of the system are inputs to either UF^2 or RKE, which calculate the contribution of each diagram to the resulting signal. The DG updates the list of diagrams as the pulse delay times change, so that UF^2 and RKE only calculate causal diagrams for each set of pulse delays.

6.2 Algorithm

We begin this section by outlining the standard results of time-dependent perturbation theory, and how it is applied to nonlinear optical spectroscopies [14], in order to introduce our notation and derive the formal operators that we use to describe signals. In Sec. 6.2.1 we build on this foundation to derive a novel open-systems algorithm called UF^2 for calculating perturbative spectroscopies. In Sec. 6.2.2 we briefly present a direct propagation method called RKE that is included in UFSS, which is used as a benchmark for timing comparisons with UF^2.

We begin with a Hamiltonian of the form

$$H = H_0(t) + H'(t), \tag{6.1}$$

where the light-matter interaction with a classical field $\mathbf{E}(t)$ is treated perturbatively in the electric-dipole approximation as

$$H'(t) = -\boldsymbol{\mu} \cdot \mathbf{E}(t), \tag{6.2}$$

where $\boldsymbol{\mu}$ is the electric dipole operator. Cartesian vectors are indicated in bold. We include a time-independent system-bath interaction in the equations of motion for the system density matrix ρ, so

$$\frac{\mathrm{d}\rho}{\mathrm{d}t}(t) = -\frac{i}{\hbar}[H(t), \rho(t)] + D\rho(t), \tag{6.3}$$

where D is a superoperator that describes dephasing and dissipation. The UF^2 algorithm can be applied with any time-independent operator D. Separating the perturbation $H'(t)$ yields two superoperators, $\mathscr{L}_0$ and $\mathscr{L}'(t)$, which are defined as

$$\frac{\mathrm{d}\rho}{\mathrm{d}t}(t) = \underbrace{\frac{-i}{\hbar}\left([H_0, \rho(t)] + i\hbar D\rho(t)\right)}_{\mathscr{L}_0|\rho(t)\rangle\!\rangle} + \underbrace{\frac{-i}{\hbar}[H'(t), \rho(t)]}_{\mathscr{L}'(t)|\rho(t)\rangle\!\rangle}. \tag{6.4}$$

ρ can be considered as an operator in the Hilbert space of the material system $\mathbb{H}$ and as a vector in the Liouville space $\mathbb{L}$, which is the vector space of linear operators on $\mathbb{H}$. We denote vectors in $\mathbb{L}$ by $|\cdot\rangle\!\rangle$. For linear operators A and B acting on $\mathbb{H}$, we write the operator $A \otimes B^T$ in $\mathbb{L}$ such that $A \otimes B^T|\rho\rangle\!\rangle$ is equivalent to $A\rho B$.[i] Using this transformation, we rewrite Eq. 6.4 as

$$\frac{\mathrm{d}|\rho(t)\rangle\!\rangle}{\mathrm{d}t} = \mathscr{L}_0|\rho(t)\rangle\!\rangle + \mathscr{L}'(t)|\rho(t)\rangle\!\rangle, \tag{6.5}$$

where, in terms of operators on $\mathbb{H}$,

$$\mathscr{L}_0 = -\frac{i}{\hbar}H_0 \otimes \mathbb{I} + \frac{i}{\hbar}\mathbb{I} \otimes H_0^T + D \tag{6.6}$$

[i] Note that the Liouville space is also a Hilbert space, with an inner product (v_1, v_2), which can be expressed in terms of the inner product on $\mathbb{H}$. If $|v_1\rangle\!\rangle = |a\rangle\langle b|$ and $|v_2\rangle\!\rangle = |c\rangle\langle d|$, then the inner product of v_1 and v_2 is $\mathrm{Tr}[v_1^\dagger v_2] = \langle d\,|\,b\rangle\langle a\,|\,c\rangle$, where the trace is taken with respect to a Hilbert-space basis. All other cases following by linearity.

and

$$\mathscr{L}'(t) = -\frac{i}{\hbar}\boldsymbol{\mu}^K \cdot \mathbf{E}(t) + \frac{i}{\hbar}\boldsymbol{\mu}^B \cdot \mathbf{E}(t), \tag{6.7}$$

with

$$\boldsymbol{\mu}^K = \boldsymbol{\mu} \otimes \mathbb{I} \quad \text{and} \quad \boldsymbol{\mu}^B = \mathbb{I} \otimes \boldsymbol{\mu}^T.$$

In a closed system, $D = 0$, and this formulation becomes equivalent to the closed case, which can be expressed with wavefunctions rather than density matrices [96].

We describe the electric field as a sum over L pulses, where each pulse is denoted by a lowercase letter starting from a. A typical 3^{rd}-order signal is produced by up to 4 pulses. We write the electric field as

$$\mathbf{E}(t) = \sum_{j=a,b,\ldots,L} \mathbf{e}_j \varepsilon_j(t) + \mathbf{e}_j^* \varepsilon_j^*(t) \tag{6.8}$$

where $\mathbf{e}_j$ is the possibly complex polarization vector, and the amplitude ε_j of each pulse is defined with envelope A_j, central frequency ω_j, wavevector $\mathbf{k}_j$, and phase ϕ_j as

$$\varepsilon_j(t) = A_j(t - t_j)e^{-i(\omega_j(t-t_j) - \mathbf{k}_j \cdot \mathbf{r} - \phi_j)},$$

where t_j is the arrival time of pulse j. We make the physical assumption that each pulse is localized in time so $\varepsilon_j(t)$ is zero outside $t \in [t_{j,\min}, t_{j,\max}]$. For the purposes of UFSS, $A_j(t)$ does not need to be a closed-form expression; it only needs to be known on a regularly spaced time grid in $[t_{j,\min}, t_{j,\max}]$. We define the Fourier transform of the pulse as

$$\tilde{\varepsilon}_i(\omega) = \int_{-\infty}^{\infty} \varepsilon_i(t) e^{i\omega t}.$$

The light-matter interaction, Eq. 6.7, is a sum over the rotating (ε_i) and counter-rotating (ε_i^*) terms. We express these terms individually as

$$\mathscr{L}'_{K\,j^{(*)}}(t) = \frac{i}{\hbar}\boldsymbol{\mu}^K \cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t) \tag{6.9}$$

$$\mathscr{L}'_{B\,j^{(*)}}(t) = -\frac{i}{\hbar}\boldsymbol{\mu}^B \cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t) \tag{6.10}$$

so that

$$\mathscr{L}'(t) = \sum_{i=a,b,\ldots} \mathscr{L}'_{K\,i}(t) + \mathscr{L}'_{K\,i^*}(t) + \mathscr{L}'_{B\,i}(t) + \mathscr{L}'_{B\,i^*}(t). \tag{6.11}$$

In the rotating wave approximation (RWA), the rotating terms, $\mathscr{L}'_{K\,i}$ and $\mathscr{L}'_{B\,i}$, excite the ket-side and de-excite the bra-side of the density matrix, respectively. The counter-rotating terms, $\mathscr{L}'_{K\,i^*}$ and $\mathscr{L}'_{B\,i^*}$, de-excite the ket side and excite the bra-side, respectively [14, 21].

We treat the effect of $\mathscr{L}'(t)$ using standard time-dependent perturbation theory and assume that at time t_0 the

system is in a stationary state of $\mathscr{L}_0$, which is $|\rho^{(0)}\rangle\!\rangle$. Equation 6.5 is easily integrated in the absence of perturbation to give the time-evolution due to $\mathscr{L}_0$,

$$\mathcal{T}_0(t) = \exp\left[\mathscr{L}_0 t\right]. \tag{6.12}$$

The perturbation $\mathscr{L}'(t)$ is zero before t_0 and produces a time-dependent density matrix $|\rho(t)\rangle\!\rangle$, which is expanded perturbatively as

$$|\rho(t)\rangle\!\rangle = |\rho^{(0)}\rangle\!\rangle + |\rho^{(1)}(t)\rangle\!\rangle + |\rho^{(2)}\rangle\!\rangle + ... \tag{6.13}$$

where the n^{th} term can be expressed as [14]

$$|\rho^{(n)}(t)\rangle\!\rangle = \mathcal{T}_0(t)\int_0^\infty dt'\mathcal{T}_0^{-1}(t-t')\mathscr{L}'(t-t')|\rho^{(n-1)}(t-t')\rangle\!\rangle.$$

Using the decomposition of $\mathscr{L}'(t)$ in Eq. 6.11, we write $|\rho^{(n+1)}(t)\rangle\!\rangle$ as a sum over four types of terms

$$|\rho^{(n+1)}(t)\rangle\!\rangle = \sum_j \left(K_j + K_{j^*} + B_j + B_{j^*}\right)|\rho^{(n)}(t)\rangle\!\rangle,$$

where all four terms are compactly defined as

$$O_{j^{(*)}} = \eta_O \frac{i}{\hbar}\mathcal{T}_0(t)\int_0^\infty dt'\mathcal{T}_0^{-1}(t-t')\left(\boldsymbol{\mu}^O\cdot\mathbf{e}_j^{(*)}\varepsilon_j^{(*)}(t-t')\right), \tag{6.14}$$

with $O = K, B$, $\eta_K = 1$ and $\eta_B = -1$, and the asterisk denotes the counter-rotating term.

From $\rho^{(n)}(t)$, perturbative signals can be determined. The full perturbative density matrix is given by

$$|\rho^{(n)}(t)\rangle\!\rangle = \left[\sum_j \left(K_j + K_{j^*} + B_j + B_{j^*}\right)\right]^n |\rho^{(0)}\rangle\!\rangle, \tag{6.15}$$

which gives $(4L)^n$ different terms, each of which is represented as a double-sided Feynman diagram. The number of diagrams that must be calculated can be dramatically reduced when considering the phase matching or phase cycling conditions in a particular spectrum, which are sensitive only to some of these contributions to $\rho^{(n)}(t)$. Further, many calculations are zero in the RWA. Time ordering also greatly reduces the number of required diagrams when the pulses do not overlap. There are well established methods to minimize the number of diagrams required to predict a spectrum, and Ref. [97] demonstrates how to automate that process.

Once the desired diagrams have been determined, the sum in Eq. 6.15 can be evaluated with only the relevant diagrams to produce the contributions to $\left|\rho^{(n)}\right\rangle$ that produce the desired signal. For example, in the case of a phase-matching experiment with detector in the direction $\mathbf{k}_d = \sum_j m_j\mathbf{k}_j$, where m_j are integers, we call the portion of

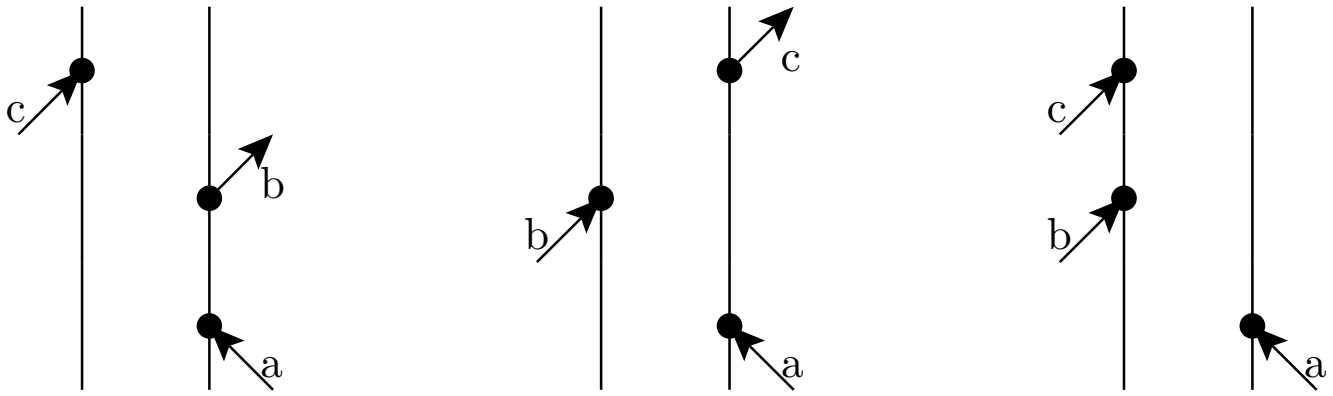

Figure 6.2: Time-ordered Feynman diagrams that contribute to the rephasing 2DPE, measured in the $\mathbf{k}_d = -\mathbf{k}_a + \mathbf{k}_b + \mathbf{k}_c$ direction. Up to 13 other diagrams contribute to the signal when one or more of the pulses overlap [97].

the density matrix that contributes to the signal $\rho^{(n)}_{k_d}$. Then the signal $S^{(n)}_{k_d}$ is calculated using

$$\begin{aligned}
\mathbf{P}^{(n)}_{\mathrm{k}_d}(t) &= \langle \boldsymbol{\mu} \rho^{(n)}_{\mathrm{k}_d}(t) \rangle \\
\tilde{\mathbf{P}}^{(n)}_{\mathrm{k}_d}(\omega) &= \int_{-\infty}^{\infty} dt e^{i\omega t} \mathbf{P}^{(n)}_{\mathrm{k}_d}(t) \\
S^{(n)}_{\mathrm{k}_d}(\omega) &= \mathrm{Im}\left[\tilde{\varepsilon}^*_d(\omega) \mathbf{e}_d \cdot \tilde{\mathbf{P}}^{(n)}_{\mathrm{k}_d}(\omega)\right]
\end{aligned} \tag{6.16}$$

where $\mathbf{P}^{(n)}_{k_d}$ is the n^{th}-order polarization contributing to the desired signal and the final pulse, with electric field $\mathbf{E}_d$, is the local oscillator used to detect the radiated field. Figure 7.1 shows the diagrams contributing to the calculation of the rephasing two-dimensional photon echo (2DPE) signal when none of the pulses overlap.

The UF^2 and RKE methods each implement an operation of $O_{j^{(*)}}$ on a density matrix. When they are given a diagram to evaluate, they compute the required successive $O_{j^{(*)}}$ operations, for example $B_c K_b B_{a^*} |\rho^{(0)}\rangle\!\rangle$, which is the second diagram in Fig. 7.1.

6.2.1 Novel open systems algorithm: UF^2

We now describe the open systems algorithm we call UF^2 for the operators $\{O_{j^{(*)}}\}$, which is an extension of the closed systems algorithm of the same name presented in Ref. 96. UF^2 requires that $\mathscr{L}_0$ be time-independent and therefore that the bath be Markovian. All degrees of freedom corresponding to non-Markovian effects must be brought into the system, where they are treated non-perturbatively. With modest computational resource, we can include several explicit vibrational modes in the system, effectively giving highly accurate non-Markovian effects to a system that is formally treated as having a Markovian bath.

We diagonalize $\mathscr{L}_0$ by finding the right and left eigenvectors. The right eigenvectors $|\alpha\rangle\!\rangle$ form a basis and have eigenvalues z_α, as

$$\mathscr{L}_0 |\alpha\rangle\!\rangle = z_\alpha |\alpha\rangle\!\rangle.$$

The left eigenvectors are defined using overbars as

$$\langle\!\langle\bar{\alpha}|\mathscr{L}_0 = \langle\!\langle\bar{\alpha}|z_\alpha.$$

Since $\mathscr{L}_0$ need not be Hermitian, $|\alpha\rangle\!\rangle^\dagger \neq \langle\!\langle\bar{\alpha}|$. We normalize the left and right eigenvectors to satisfy

$$\langle\!\langle\bar{\alpha}|\beta\rangle\!\rangle = \delta_{\alpha\beta}. \tag{6.17}$$

In the absence of $\mathscr{L}'(t)$, Eq. 6.12 gives $|\rho(t)\rangle\!\rangle = \mathcal{T}_0(t)|\rho(0)\rangle\!\rangle$. $\mathcal{T}_0$ is diagonal in the basis $\{|\alpha\rangle\!\rangle\}$, so we have

$$|\rho(t)\rangle\!\rangle = \sum_\alpha^{N^2} e^{z_\alpha t} c_\alpha |\alpha\rangle\!\rangle, \tag{6.18}$$

where N is the dimension of $\mathbb{H}$, which we take to be finite. If the physical system has an infinite dimensional $\mathbb{H}$, as in the case of a harmonic oscillator, we truncate $\mathbb{H}$ to dimension N, and therefore truncate $\mathscr{L}_0$ to dimension N^2.

The electric dipole operator acting from the left, $\boldsymbol{\mu}^K$, and from the right, $\boldsymbol{\mu}^B$, must be known in the eigenbasis of $\mathscr{L}_0$, where we define matrix elements

$$\boldsymbol{\mu}^K_{\alpha\beta} = \langle\!\langle\bar{\alpha}|\boldsymbol{\mu}^K|\beta\rangle\!\rangle.$$
$$\boldsymbol{\mu}^B_{\alpha\beta} = \langle\!\langle\bar{\alpha}|\boldsymbol{\mu}^B|\beta\rangle\!\rangle.$$

The derivation of UF^2 for open systems is formally similar to that for closed systems in Ref. 96, with replacements of U by $\mathcal{T}$, the wavefunction $|\psi\rangle$ by the density vector $|\rho\rangle\!\rangle$, and the dipole operator $\boldsymbol{\mu}$ by $\boldsymbol{\mu}^K$ and $\boldsymbol{\mu}^B$. Because the action of the dipole operator on the ket and bra must be considered separately, the operator $K_{j^{(*)}}$ is joined in the open systems case by its counterpart $B_{j^{(*)}}$.

We represent $|\rho^{(n)}(t)\rangle\!\rangle$ with coefficients $c^{(n)}_\alpha(t)$ that contain only the time dependence induced by the perturbation, while keeping the evolution due to $\mathscr{L}_0$ separate as

$$|\rho^{(n)}(t)\rangle\!\rangle = \sum_\alpha e^{z_\alpha t} c^{(n)}_\alpha(t)|\alpha\rangle\!\rangle. \tag{6.19}$$

With this notation, Eq. 7.7 gives

$$\begin{aligned} O_{j^{(*)}}|\rho^{(n)}(t)\rangle\!\rangle &= \eta_O \frac{i}{\hbar}\mathcal{T}_0(t)\int_0^\infty dt'\mathcal{T}_0^{-1}(t-t')\sum_\beta |\beta\rangle\!\rangle\langle\!\langle\bar{\beta}|\left(\boldsymbol{\mu}^O\cdot\mathbf{e}_j^{(*)}\varepsilon_j^{(*)}(t-t')\right)\sum_\alpha e^{z_\alpha(t-t')}c^{(n)}_\alpha(t-t')|\alpha\rangle\!\rangle \\ &= \eta_O\frac{i}{\hbar}\sum_\beta e^{z_\beta t}|\beta\rangle\!\rangle\int_{-\infty}^\infty dt'\theta(t')\underbrace{e^{-z_\beta(t-t')}\sum_\alpha\left(\boldsymbol{\mu}^O_{\beta\alpha}\cdot\mathbf{e}_j^{(*)}\varepsilon_j^{(*)}(t-t')\right)e^{z_\alpha(t-t')}c^{(n)}_\alpha(t-t')}_{y_\beta(t-t')}. \end{aligned} \tag{6.20}$$

The integral in Eq. 6.20 is a convolution, and we express it in the compact form

$$O_{j^{(*)}}|\rho^{(n)}(t)\rangle\!\rangle = \eta_O \frac{i}{\hbar} \sum_\beta e^{z_\beta t}|\beta\rangle\!\rangle \left[\theta * y_\beta\right](t), \tag{6.21}$$

where

$$\left[x * y\right](t) = \int_{-\infty}^{\infty} dt' x(t') y(t-t').$$

Assuming that $\varepsilon_j(t)$ is zero outside the interval $[t_{j,\min}, t_{j,\max}]$,

$$\left[\theta * y_\beta\right](t) = \begin{cases} 0 & t < t_{j,\min} \\ r_\beta(t) & t_{j,\min} < t < t_{j,\max} \\ C_\beta & t > t_{j,\max} \end{cases}.$$

for constant C_β. Therefore, we need only calculate this convolution for $t_{j,\min} < t < t_{j,\max}$.

Physically, we only need to solve for the time dependence due to the interaction with the pulse while the pulse is nonzero. The rest of the time dependence is contained in $\mathscr{L}_0$ and is therefore known exactly. This realization drastically reduces the computational cost of UF^2 compared to techniques that must use time stepping for both the system dynamics and the perturbation.

We evaluate the convolution $\left[\theta * y_\beta\right](t)$ numerically to solve for the function $r_\beta(t)$ using the FFT and the convolution theorem. Each electric field envelope $A_j(t-t_j)$ is represented using M_j equally spaced time points, where $M_j = (t_{\mathrm{j,max}} - t_{\mathrm{j,min}})/dt_j$ and dt_j is the spacing between points. Before convolving y_β is zero-padded up to $2M_j - 1$ points, and after the convolution is performed we retrieve only the M_j points corresponding to a linear convolution. Appendix 6.7.1 describes the computational cost of UF^2 and shows how it scales with N and M.

6.2.2 RKE

The RKE method is an alternative algorithm for evaluating the operators $\{O_{j^{(*)}}\}$ and is also included in UFSS. It was introduced in Ref. 96 for closed systems. RKE uses the Runga-Kutta 45 (RK45) adaptive time step algorithm to propagate the evolution due to $\mathscr{L}_0$ and a fixed-step Euler method to include the perturbation $\mathscr{L}'(t)$. It is a direct propagation method, meaning that it propagates $|\rho\rangle\!\rangle$ forward one step at a time using the differential form of the equations of motion, Eq. 6.5. RKE is a simple example of a direct-propagation method, and we intend it to be representative of the computational scaling differences between UF^2 and direct-propagation methods; more efficient and higher-order methods than RKE are possible [59, 62, 64, 66, 70, 76–78].

In the absence of pulses, the RK45 method advances the density matrix $|\rho\rangle\!\rangle$ forward in time according to

$$|\dot{\rho}\rangle\!\rangle = \mathscr{L}_0|\rho\rangle\!\rangle \tag{6.22}$$

where we represent the time evolution due to $\mathscr{L}_0$ as an $N^2 \times N^2$ matrix acting on $\mathbb{L}$, rather than using $N \times N$ operators on $\mathbb{H}$ as in Eq. 6.3. We represent a step using the RK45 algorithm alone as $|\rho(t_i + dt)\rangle\!\rangle = \mathcal{T}_0(dt)|\rho(t_i)\rangle\!\rangle$.

Starting from $|\rho^{(0)}\rangle\!\rangle$, RKE evaluates diagrams by successive $O_{j^{(*)}}$ operations. RKE calculates $|\rho_\beta\rangle\!\rangle \equiv O_{j^{(*)}}|\rho_\alpha\rangle\!\rangle$ for some state $|\rho_\alpha\rangle\!\rangle$ as

$$\begin{aligned}|\rho_\beta(t_{j,\min} + mdt_E)\rangle\!\rangle &= \mathcal{T}_0(dt_E)|\rho_\beta(t_{j,\min} + (m-1)dt_E)\rangle\!\rangle \\ &+ \mathscr{L}'_{Oj^{(*)}}(t_{j,\min} + mdt_E)|\rho_\alpha(t_{j,\min} + mdt_E)\rangle\!\rangle, \end{aligned} \tag{6.23}$$

where we propagate using fixed step size dt_E from $t_{j,\min}$ to $t_{j,\max}$. This method accumulates error proportional to dt_E. It is possible to construct analogous methods that accumulate error proportional to dt_E^2 [64]. Defining $M_E = (t_{j,\max} - t_{j,\min})/dt$, m runs from 0 to M_E. Once we obtain $|\rho_\beta(t_{j,\max})\rangle\!\rangle$, the remainder of the time evolution for $t > t_{j,\max}$ is obtained using the standard RK45 method alone, with a variable time step. Section 2 of the Appendix describes the computational cost of RKE and shows how it scales with N and M_E.

6.3 Hamiltonian/Liouvillian generator

Here we outline the Hamiltonian and Liouvillian generator (HLG) included as part of UFSS. Note that both UF^2 and RKE are compatible with any time-independent Hamiltonian or Liouvillian that can be expressed as a finite matrix. One need not use the HLG in order to take advantage of the other modules in UFSS.

HLG is a vibronic model generator, designed to create a Hamiltonian for a network of s two-level systems (2LS) coupled linearly to k harmonic vibrational modes. The HLG constructs a Liouvillian by including coupling of each degree of freedom to a Markovian bath using either Redfield (full or secular) or diabatic Lindblad formalisms. Models of this type have been used to describe many systems including conical intersections in pyrazine and energy transfer in photosynthetic complexes [34, 101, 104–110].

6.3.1 Hamiltonian Structure

We begin with an electronic system described by s 2LS,

$$H_e = E_0 + \sum_{n=1}^{s} E_n a_n^\dagger a_n + \sum_{m \neq n} J_{mn} a_m^\dagger a_n,$$

where a_n is the annihilation operator for the excited state in the n^{th} 2LS, E_m is the site energy, E_0 is the ground state energy, and J_{mn} is a Hermitian matrix of electronic couplings. The system includes k explicit harmonic

vibrational modes of frequency ω_α, generalized momentum p_α and coordinate q_α with Hamiltonian

$$H_{ph} = \frac{1}{2}\left(\sum_{\alpha=1}^{k} p_\alpha^2 + \omega_\alpha^2 q_\alpha^2\right).$$

We treat standard linear coupling of these modes to the electronic system as

$$H_{e-ph} = \sum_{\alpha=1}^{k}\sum_{n=1}^{s} \omega_\alpha^2 d_{\alpha,n} q_\alpha a_n^\dagger a_n,$$

where $d_{\alpha,n}$ indicates the coupling of each vibrational mode to each 2LS. It is related to the Huang-Rhys factor by

$$S_{\alpha,n} = \frac{1}{2}\omega_\alpha d_{\alpha,n}^2.$$

The total system Hamiltonian is

$$H_0 = H_e + H_{ph} + H_{e-ph}. \tag{6.24}$$

If we work in the number basis of the vibrational modes, using the ladder operators b_α, with $q_\alpha = \frac{1}{\sqrt{2}}(b_\alpha + b_\alpha^\dagger), p_\alpha = \frac{i}{\sqrt{2}}(b_\alpha - b_\alpha^\dagger)$, then H_0 is highly sparse. H_{e-ph} has $2k+1$ entries per row. H_{ph} is formally infinite in size, so we truncate H_0 to size N by fixing the total vibrational occupation number. Note that Eq. 6.24 is block diagonal with $s+1$ blocks. Each of these blocks is an optically separated manifold, and we index manifolds using X and Y, where X, Y can refer to the ground-state manifold (GSM), the singly excited manifold (SEM), the doubly excited manifold (DEM), and so on. Each block has a size N_X, and $N = N_{GSM} + N_{SEM} + N_{DEM} + ...$ The block diagonal form of H_0 is not required by UF^2, but it allows useful simplifications in certain cases, which are discussed briefly at the end of this section and in Appendix 6.7 .

6.3.2 Liouvillian Structure

Using the Hamiltonian from Eq. 6.24, we construct the unitary part of the Liouvillian, $\mathscr{L}_U = \mathscr{L}_0 - D$, where $\mathscr{L}_0$ is defined in Eq. 6.6. UFSS allows coupling of a Markovian bath to all degrees of freedom of H_0 using either the Redfield [99] or Lindblad formalisms [98] or a user-specified combination of them, should that be desirable.

6.3.2.1 Redfield

Redfield theory arises from microscopic derivations of the properties of the bath and system-bath coupling, in contrast to diabatic Lindblad theory, which requires phenomenological relaxation and dephasing parameters. Given a system-bath coupling Hamiltonian of the form

$$H_{SB} = \sum_r O_{r,S} \otimes O_{r,B},$$

where $O_{r,S}$ is an operator defined in the Hilbert space of the system, and $O_{r,B}$ is an operator defined in the Hilbert space of the bath, the dissipation tensor is defined as

$$Y_{ijkl}(O_{r,S}) = \sum_{r,r'} \langle i| O_{r,S} |k\rangle \langle l| O_{r',S} |j\rangle C_{rr'}(\omega_{ki}),$$

where $H_0 |i\rangle = \hbar\omega_i |i\rangle$ defines the eigenvectors of H_0, given by Eq. 6.24, and

$$C_{rr'}(t) = \langle O_{r,B}(t) O_{r',B}(0)\rangle$$

are the two-point correlation functions of the phonon modes of the bath. The index r specifies either the site index n or the vibrational mode index α. We assume that the cross-correlation terms are zero, and so $C_{rr'}(t) = \delta_{rr'} C_{rr'}(t)$. The Fourier transform of $C(t)$ is specified using a spectral density $J(\omega)$ as

$$\begin{aligned} \Re\left[C(\omega)\right] &= \frac{1}{2}\hbar J(\omega) \coth\left(\frac{\hbar\omega}{2k_B T}\right) \\ \mathrm{Im}\left[C(\omega)\right] &= \frac{1}{\pi} P \int_{-\infty}^{\infty} d\omega' \frac{\Re\left[C(\omega')\right]}{\omega - \omega'}, \end{aligned}$$

where P indicates the Cauchy principal value. In the examples that follow, we use an Ohmic spectral density with the Drude-Lorentz cut-off function

$$J(\omega) = 2\omega\lambda \frac{\gamma}{\omega^2 + \gamma^2}$$

where λ is the strength of the system-bath coupling and γ is the cut-off frequency of the bath.Users are also free to specify any spectral density function that is appropriate to their system.

The Redfield tensor is

$$\begin{aligned} R_{ijkl}(O) = &- \left(Y_{ijkl}(O) + Y^*_{jilk}(O)\right) + \delta_{jl} \sum_{\sigma} Y_{\sigma ik\sigma}(O) \\ &+ \delta_{ik} \sum_{\sigma} Y^*_{\sigma jl\sigma}(O) \end{aligned}$$

We consider system-bath coupling to each site n via the Redfield tensor $R(a_n^\dagger a_n)$, and coupling to each vibrational mode via the Redfield tensor $R(q_\alpha)$, where $q_\alpha = \frac{1}{\sqrt{2}}\left(b_\alpha^\dagger + b_\alpha\right)$, so that the dissipative part of $\mathscr{L}_0$ (see Eq. 6.6) is

$$D = -R(a_n^\dagger a_n) - R(q_\alpha).$$

We also optionally include in D the relaxation of electronic excitations to the ground state via the Redfield tensor $R(a_n^\dagger + a_n)$. The microscopic derivation of the two-time correlation function $C(t)$ that mediates such relaxation processes is more difficult, as it involves non-adiabatic coupling terms [111] and is rarely used. As such, by default

we have included a flat spectral density for inter-manifold relaxation processes such that the associated $C(\omega)$ is

$$C(\omega) = \begin{cases} \gamma_r & \omega > 0 \\ \gamma_r e^{\hbar\omega/k_B T} & \omega < 0 \end{cases},$$

where γ_r provides a phenomenological inter-manifold relaxation rate. In the secular approximation (see below), this phenomenological form reduces to a commonly used approach [108, 112, 113]. Users are free to provide more complicated spectral densities to describe this type of process.

The commonly used secular approximation sets $R_{ijkl} = 0$, except when $|\omega_{ij} - \omega_{kl}| = 0$, which guarantees positivity of the density matrix. The remaining simple form of $\mathscr{L}_0$ only couples the populations of the density matrix couple to one another (see Appendix 6.7 for exceptions)[34, 99, 107]. In this case, $\mathscr{L}_0$ consists of a single $N \times N$ block for the populations and is otherwise diagonal. This simplification has important implications on the computational complexity of both the UF^2 and RKE algorithms. In particular, the cost of diagonalizing $\mathscr{L}_0$ becomes $\sim N^3$, which is the same scaling as the cost of diagonalizing H_0, and is in stark contrast to the the cost of diagonalizing a generic $\mathscr{L}_0$, which is $\sim N^6$.

6.3.2.2 Diabatic Lindblad

The Lindblad formalism is widely used to describe open quantum systems using a small number of phenomenological dephasing and relaxation constants. It guarantees the complete positivity of $\mathscr{L}_0$. Secular Redfield theory can be mapped to a Lindblad structure using operators in the eigenbasis of the Hamiltonian. Recently it has been shown that full Redfield theory can be mapped to a Lindblad structure in some cases (again this mapping is done in the eigenbasis of the Hamiltonian) [35]. Here, we present a Lindblad formalism in the diabatic basis, which can be fairly accurate for short time-scales, though it does not produce a thermal distribution at infinite time, so must break down for long time-scales [104, 105].

If O is an operator on $\mathbb{H}$, the Lindblad superoperator is

$$L[O]\rho = O\rho O^\dagger - \frac{1}{2}\left(O^\dagger O\rho + \rho O^\dagger O\right).$$

We consider dissipation in vibrational modes and both inter- and intra-manifold dephasing and relaxation of the electronic modes.

For vibrational mode α, we describe coupling to the rest of the bath with the dissipation operator

$$D_v = \sum_\alpha \gamma_{v,\alpha}\left(N_{\text{th}} L[b_\alpha^\dagger] + (N_{\text{th}} + 1) L[b_\alpha]\right), \tag{6.25}$$

where $\gamma_{v,\alpha}$ is the thermalization rate of the α^{th} mode and $N_{\text{th}} = \langle b_\alpha^\dagger b_\alpha \rangle = 1/(\exp(\beta\hbar\omega_\alpha) - 1)$ is the average number

of excitations at equilibrium of a mode with energy $\hbar\omega_\alpha$ coupled to a Markovian bath of temperature T, with $\beta = 1/k_B T$, with k_B the Boltzmann constant [98].

We describe inter-manifold electronic relaxation and the complementary incoherent thermal excitation processes with

$$D_{r1} = \sum_{n=1}^{s} \gamma_{r1,n} \left(C_{gn} L[a_n] + C_{ng} L[a_n^\dagger] \right),$$

where $C_{nm} = \frac{e^{-\beta E_n}}{e^{-\beta E_n} + e^{-\beta E_m}}$. Intra-manifold relaxation processes are described by

$$D_{r2} = \sum_{m \neq n} \gamma_{r2,nm} C_{nm} L[a_n^\dagger a_m].$$

Since $C_{nm}/C_{mn} = e^{-\beta(E_n - E_m)}$, we obtain a thermal distribution of eigenstates as $t \to \infty$ if $J_{mn} = 0$.

Relaxation processes necessarily give rise to dephasing. We include additional intra-manifold dephasing using

$$D_{d2} = \sum_{n \neq m} \gamma_{d2,nm} L \left[a_n^\dagger a_n - a_m^\dagger a_m \right].$$

All of the above processes also give rise to inter-manifold (optical) dephasing. We include additional, pure inter-manifold dephasing using

$$D_{d1} = \gamma_{d1} L \left[\sum_{n=1}^{s} a_n^\dagger a_n \right]. \tag{6.26}$$

In the common case where γ_{d1} is larger than all the other bath coupling rates, the homogeneous linewidth(s) are dominated by the D_{d1} term. Putting all of these operators together we arrive at the total dissipation operator

$$D = D_{r1} + D_{r2} + D_{d1} + D_{d2} + D_\nu. \tag{6.27}$$

When modeling optical spectroscopies, H_0 is often block diagonal, and therefore the system is composed of distinct manifolds, as constructed in Sec. 6.3.1. In the case where we can also neglect inter-manifold relaxation processes ($\gamma_{r1,n} = 0$), the total Liouvillian $\mathscr{L}_0$ is also block diagonal. Under these assumptions $\mathscr{L}_0$ can be arranged into blocks of size $N_X^2 \times N_Y^2$, allowing us to save both computational cost and memory, both for diagonalization (if applicable) as well as for use with UF^2 or RKE. The cost of diagonalizing $\mathscr{L}_0$ is then $\sim N_X^3 N_Y^3$.

6.4 Computational advantage

We compare the computational costs of the two UFSS propagation methods, UF^2 and RKE, for three methods of including the bath: secular Redfield, full Redfield, and diabatic Lindblad. We show that the convolution-based UF^2 is over 200 times faster than the direct-propagation RKE method for small systems. Asymptotically the relative performance of UF^2 depends strongly on the method of treating the bath. Appendix 6.7 derives the asymptotic

Table 6.1: Scaling of computational cost with the Hamiltonian dimension N. "$\mathscr{L}_0$ Diag" is the diagonalization of $\mathscr{L}_0$. Derivations are in Appendix 6.7.

	UF^2	RKE	$\mathscr{L}_0$ Diag.	UF^2 advantage
Full Redfield	N^4	N^4	N^6	up to large N
Secular Redfield	N^3	N^3	N^3	all N
Diabatic Lindblad	N^4	N^2	N^6	$N \lesssim 100$

computational complexity of these methods, and the results are summarized in Table 6.1. These results predict that the UF^2 method is always more efficient than the direct propagation method for secular Redfield. While these computational complexity results do not include the memory requirements of the algorithms, we demonstrate below that for full Redfield, UF^2 is more efficient than RKE up until system sizes where $\mathscr{L}_0$ requires at least 20 GB to store, as summarized in the last column of the table.

UF^2 has a one-time cost of diagonalizing $\mathscr{L}_0$. While $\mathscr{L}_0$ is generally a matrix of size $N^2 \times N^2$ with diagonalization cost $\sim N^6$, in the secular approximation this cost is only $\sim N^3$, as explained in Sec. 6.7. Despite diagonalization being an expensive calculation for full Redfield and the diabatic Lindblad models, it does not necessarily contribute significantly to the total cost of calculating spectra, because the diagonal form is reused for each set of pulse delays, pulse shapes, pulse polarizations, etc. Consider a sample 2DPE signal $S^{(3)}(\tau, T, \omega_t)$ with 100 coherence times τ and 20 population times T, as well as a sample TA signal $S^{(3)}(T, \omega_t)$ with 100 delay times. Since the 2DPE spectrum requires 3-16 diagrams evaluated with 2000 different pulse delays, the cost of the diagonalization is effectively amortized over > 6000 calculations. The TA calculation requires 6-16 diagrams evaluated at only 100 delay times, so amortizes the diagonalization cost over ≈ 600 calculations. The RKE method does not require diagonalization, so the system size at which it becomes cost effective to use the RKE method in principle depends on what type of spectrum is being considered. In diabatic Lindblad, diagonalization cost begins to be limiting for TA spectra around the same size that RKE becomes more efficient, regardless.

The computational costs of predicting spectra depend upon both the size and structure of $\mathscr{L}_0$. To make concrete comparisons between UF^2 and RKE, we use a vibronic Hamiltonian H_0 coupled to a Markovian bath, as outlined in Section 6.3. Figures 6.3 and 6.4 show the ratio of the computation time of RKE to the computation time of UF^2 for TA spectra with Redfield and Lindblad models, respectively. We consider systems with the number of sites and number of vibrational modes equal ($s = k$) and varying from 2 to 4. The energy scale of the problem is defined by the k nearly identical vibrational frequencies ω_α, which are all within about 1% of ω_0. The modes are detuned for convenience, to avoid degeneracies in the ground state manifold, which makes the structure for the Redfield tensor simpler in the secular approximation. We use the RWA and, after rotating away the optical gap, the site energies E_i vary from $0 - 1.5\omega_0^{-1}$, and the coupling terms J_{mn} vary from $0 - 0.5\omega_0^{-1}$. We use the same value $d_{\alpha,n} = d\delta_{\alpha,n}$ for each α, n pair, where $\delta_{i,j}$ is the Kronecker delta, and choose values of d from 0 to 1.5. Larger values of d require a larger truncation size N for the spectra to converge. For Redfield theory we use the same bath parameters for both

the sites and the vibrational modes, $\lambda = 0.05\omega_0$ and $\gamma = \omega_0$, and have taken $\mathrm{Im}[C(\omega)] = 0$. In the diabatic Lindblad model, we include a Markovian bath using $\gamma_{r2,i} = 0.05\omega_0$, $\gamma_{\nu,\alpha} = 0.05\omega_0$ and $\gamma_{d1} = 0.2\omega_0$ (with $\gamma_{r1} = \gamma_{d2} = 0$). The optical pulses have Gaussian envelopes with standard deviation $\sigma = \omega_0^{-1}$, centered on the transition $E_1 - E_0$. All sites i have parallel dipole moments, with magnitudes varying from 0.7 to 1, and we use the Condon approximation that μ is independent of vibrational coordinate. We choose the number of vibrational states in the simulations to be sufficiently large by using UF^2 to generate a TA signal and seeking the truncation size N that converges the resulting spectra within 1% using an ℓ_2 norm over the full spectrum. We perform this convergence separately for each case of s, d. For each choice of s,d, we use the same N for the RKE calculations. The optical field parameters (M and dt for UF^2 and M_E and dt_E for RKE) were determined by testing on some of the smaller systems and were held constant for all s, d. The spectra produced by UF^2 and RKE were compared for each s, d by computing the ℓ_2 norm difference between the two. For each system for which we calculated spectra using both UF^2 and RKE (see Figures 6.3 and 6.4), we verified that spectra agree to within 2% relative error. This agreement justifies the decision to hold M, dt for UF^2 and M_E, dt_E for RKE constant for all s, d. This agreement also gives confidence in the spectra produced by both algorithms. Values of t were selected in order to resolve all optical oscillation frequencies in the RWA and to resolve the homogeneous linewidth. We took the inter-manifold relaxation rate $\gamma_{r1} = 0$, so that the optical manifolds are separable and $\mathscr{L}_0$ is block diagonal in all cases. All calculations were performed on an Intel Xeon E5-2640 v4 CPU with a 2.40 GHz clock speed and 96 GB of RAM.

UF^2 always outperforms RKE for secular Redfield theory, saturating at 20 times faster for large N_{SEM}, as predicted in Appendix 6.7. In all three formalisms, UF^2 is 200-500 times faster than RKE for small N_{SEM}. For full Redfield, with large N_{SEM}, the cost of diagonalization should eventually cause RKE to outperform UF^2, but we see that UF^2 is approximately 100 times faster even at the larger system sizes studied. The memory requirements of constructing $\mathscr{L}_0$ using full Redfield theory become limiting, as $\mathscr{L}_0$ is a full $N^2 \times N^2$ matrix of complex floats. The largest $\mathscr{L}_0$ studied was 20 GB. For diabatic Lindblad the memory requirements of diagonalizing $\mathscr{L}_0$ are similar.

The insets of both figures show the wall-clock runtimes for each value of N_{SEM}, not including the diagonalization cost for UF^2. The dashed lines in the insets show that as N_{SEM} increases, the expected asymptotic scalings from Table 6.1 are obeyed. For the diabatic Lindblad model, UF^2 is more efficient than RKE with sufficiently small systems, the superior cost scaling of RKE leads to a crossover in runtimes near $N_{SEM} = 100$, corresponding to $\mathscr{L}_0$ having blocks of dimension $N_{SEM}^2 = 10^4$. This crossover is consistent with our result with closed systems in Ref. 96, where the UF^2 method was more efficient than RKE for N_{SEM} smaller than $10^4 - 10^5$. For small N_{SEM}, the cost of diagonalization is negligible, as shown by the blue and orange dots in Figure 6.4 all overlapping for $N_{SEM} < 10$. As N_{SEM} increases, the cost of diagonalization becomes apparent as the colors separate. With or without diagonalization, Fig. 6.4 shows that the crossover occurs with $N_{SEM} \approx 100$.

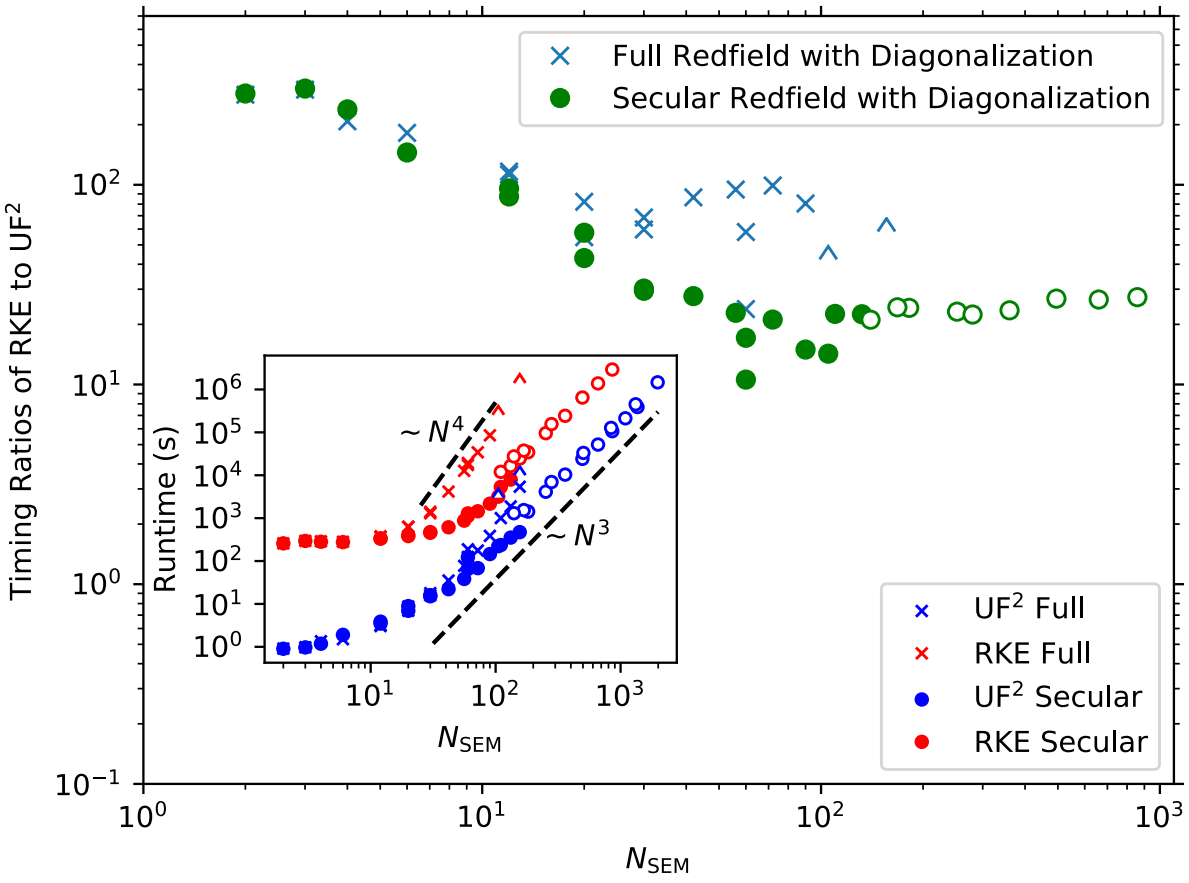

Figure 6.3: Timing ratios of RKE to UF^2 for TA spectra with 100 time delays using the Redfield formalism for a range of systems and parameters (closed circles and x's). For large systems, we simulate a single time delay (open circles and carats) with only one ESA diagram and multiply the runtime by 600 to effectively treat 100 delay times and 6 diagrams, due to the long runtimes. N_{SEM} is the dimension of the Hamiltonian describing the truncated singly excited manifold, large enough to converge the spectra. Ratios include the cost of diagonalizing $\mathscr{L}_0$ in the UF^2 costs. The cost of diagonalization is never important in the secular approximation, and the UF^2 advantage plateaus at a factor of 20 for large N_{SEM}, matching the predictions from Appendix 6.7. With full Redfield, ignoring diagonalization, asymptotically UF^2 has a theoretical relative advantage of 40, in good agreement with the results (including diagonalization costs) shown here. Inset shows the associated runtimes without diagonalization costs. Dashed lines show the predicted asymptotic scalings from Appendix 6.7. Note that for each timing ratio displayed, we compared the spectra produced by UF^2 and RKE by taking the ℓ_2 norm difference between the two. We verified that UF^2 and RKE agree to within less than 2% difference. Since UF^2 and RKE use different numerical algorithms both for the propagation of the system dynamics and for the inclusion of finite pulse effects, these comparisons give confidence that both methods produce accurate spectra.

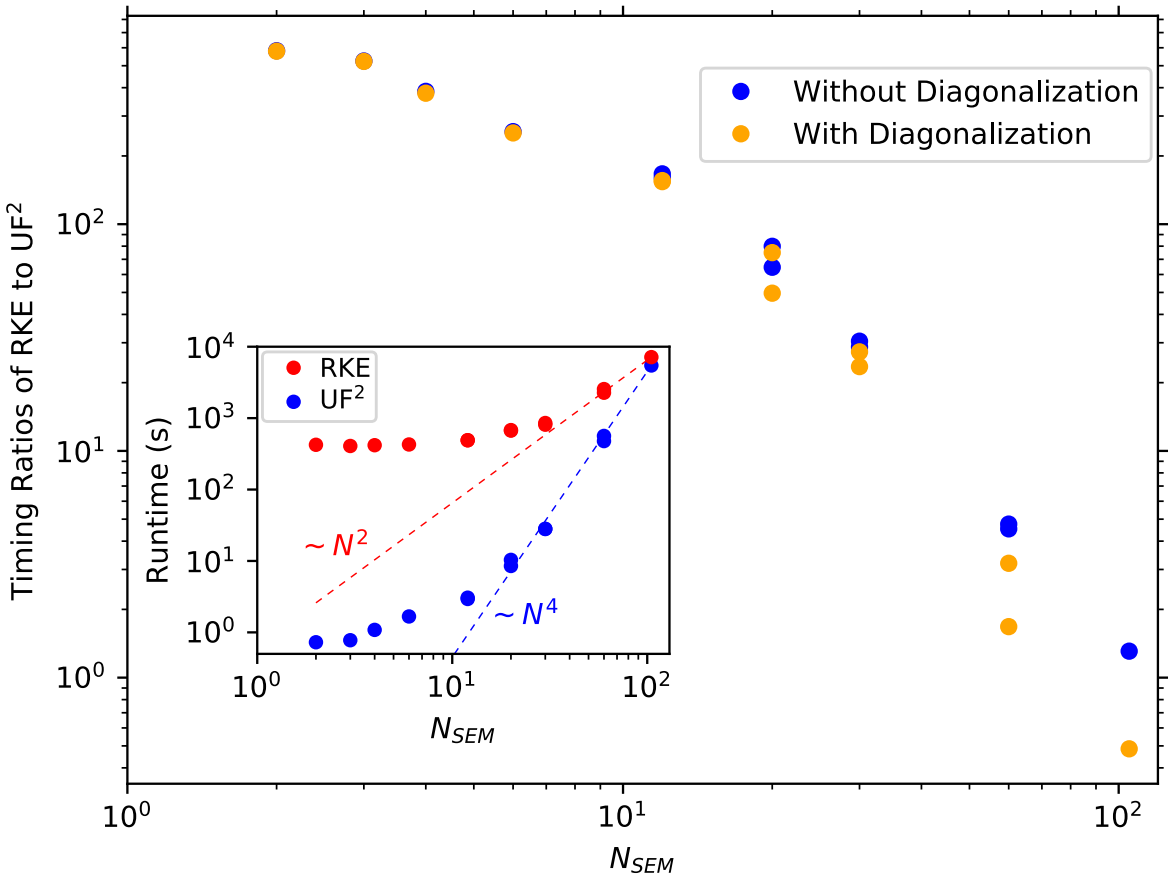

Figure 6.4: Timing ratios of RKE to UF^2 methods for TA spectra with 100 time delays for a range of systems and parameters, described in the text. N_{SEM} is the dimension of the Hamiltonian describing the truncated singly excited manifold, large enough to converge the spectra. Blue points show the timing ratios for the evaluation of the spectra but not the diagonalization of $\mathscr{L}_0$. Orange points include the cost of the diagonalization, which is insignificant with small N_{SEM}. Near $N_{SEM} = 100$, the direct-propagation RKE method becomes more efficient than UF^2. Inset shows the time required to calculate the TA spectrum using RKE and UF^2, without the diagonalization cost included. The dashed lines show the slopes of the predicted asymptotic scaling of each algorithm. Note that for each timing ratio displayed, we compared the spectra produced by UF^2 and RKE by taking the ℓ_2 norm difference between the two. We verified that UF^2 and RKE agree to within less than 2% difference. Since UF^2 and RKE use different numerical algorithms both for the propagation of the system dynamics and for the inclusion of finite pulse effects, these comparisons give confidence that both methods produce accurate spectra.

6.5 Comparison of UF^2 to analytic results

We now demonstrate that the signals produced by UF^2 reproduce an analytical solution for the rephasing 2D photon echo (2DPE) signal for the optical Bloch equations using Gaussian pulses [19]. Reference 19 considered a small ($N = 3$) system that can be mapped to the Hamiltonian described by $s = 2$ and $k = 0$ from Sec. 6.3, with the doubly excited state removed. Although this is a small system, these comparisons are some of the only available analytical solutions including finite pulse durations, and thus provides a benchmark to show that UF^2 calculates spectra with a high degree of accuracy. UF^2 converges to within 1% of the analytical result using just $M = 25$ points to discretize $\varepsilon_j(t)$.

In this model, the energy difference between the two excited states is $\hbar\omega_0 = E_2 - E_1$. All pulses are taken to have identical Gaussian envelopes

$$A(t) = \frac{1}{\sqrt{2\pi}\sigma} e^{-t^2/2\sigma^2},$$

where $\sigma = \omega_0^{-1}$ and have central frequency $\omega_j = (E_2 + E_1)/(2\hbar)$. In the RWA, we are free to set $\omega_j = 0$ for all pulses, which we do. The model includes phenomenological dephasing rates and population decay rates of $0.2\omega_0$ and $0.1\omega_0$, respectively. This bath coupling is similar to a Lindblad formalism like the one outlined in Sec. 6.3.2, except that it does not conserve the total probability of the density matrix. Rather than using the HLG included in UFSS, for this comparison we separately created the model described in Ref. 19. The construction of $\mathscr{L}_0$ and the evaluation of the resulting spectra is demonstrated in the Jupyter notebook `Smallwood2017Comparison.ipynb`, available in the UFSS repository.

The 2DPE signal is the result of three pulses, which gives rise to an emitted field $P(\tau, T, t)$. UF^2 is designed to calculate $P(\tau, T, t)$, while the result from Ref. 19 is for the quantity

$$P(\omega_\tau, T, \omega_t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} dt e^{i\omega_t t} \int_{-\infty}^{\infty} d\tau e^{-i\omega_\tau \tau} P(\tau, T, t).$$

To compare UF^2 to the analytical solutions, we calculate $P(\tau, T, t)$ for a discrete set of τ, T, t, and take a 2D discrete Fourier transform with respect to τ and t. Since $A(t)$ is symmetric, we evaluate $A(t - t_j)$ on the interval $t \in [t_{j,\min}, t_{j,\max}]$ with spacing dt, and choose $t_{j,\max} - t_j = t_j - t_{j,\min}$, where t_j is the arrival time of the j^{th} pulse. For symmetric pulses, UF^2 converges most rapidly when the value $t - t_j = 0$ is included at the center of the discretization interval and the endpoints of the interval are also included. We define the duration of the pulse, $\Delta_j \equiv t_{j,\max} - t_{j,\min}$ and the number of points $M_j = \Delta_j/dt + 1$. All the pulses are identical, and so we use $M = M_j$, and therefore $\Delta = \Delta_j$ for each pulse.

Figure 6.5(a) shows the analytical result for $P(\omega_\tau, 0, \omega_t)$, which provides a benchmark for quantifying the convergence behavior of UF^2. Fig. 6.5(b) shows the ℓ_2 norm of the difference between the analytical solution shown in (a) and the result from UF^2. The white contours in Fig. 6.5(b) show that the errors due to Δ and dt

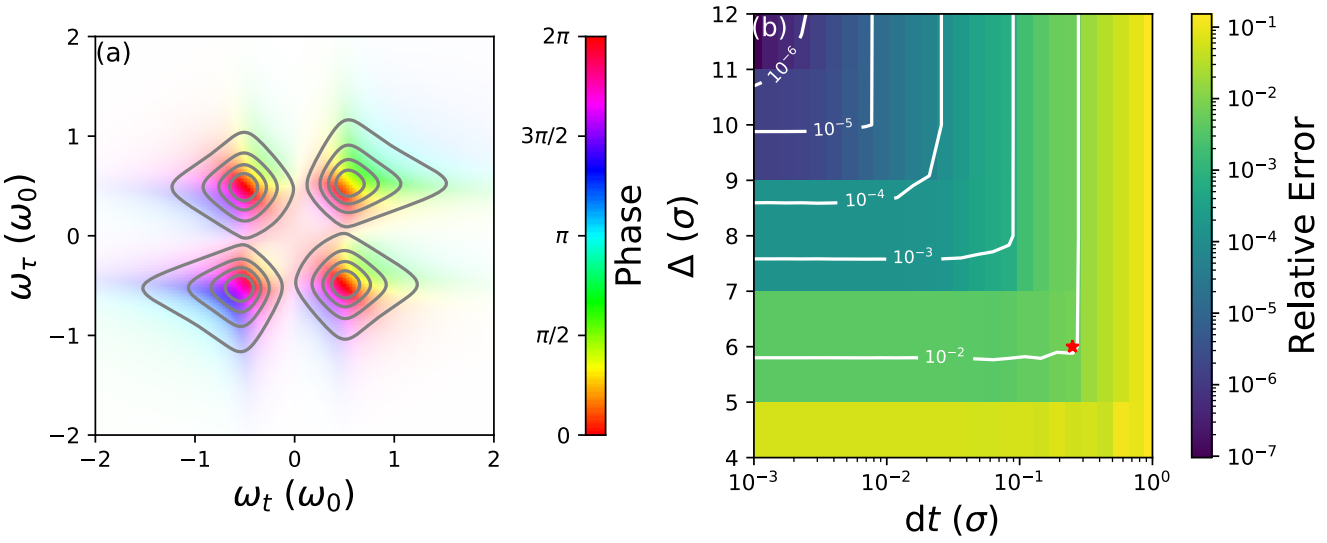

Figure 6.5: (a) Reproduction of Smallwood's Figure 3e) using their analytical forms. The color shows the phase of the complex signal $P^{(3)}(\omega_\tau, T, \omega_t)$, while the intensity of the color shows the magnitude of the signal. The numerical result using UF^2 appears visually identical so is not shown. (b) ℓ_2 norm of the difference between the analytical solution shown in (a) and the result from UF^2, both sampled on a mesh of 801×801 ω_t, ω_τ points, as a function of dt and the pulse duration Δ used in the numerical convolutions of Eq. 6.21. We evaluate $P^{(3)}(\tau, T, t)$ for t and τ ranging from -100σ to 100σ in steps of 0.25σ. The red star in (b) indicates the smallest value of $M = \Delta/dt + 1$ needed to reach 1% agreement with the analytical solution shown in (a), and corresponds to $M = 25$ points. Note that for $\Delta = 12\sigma$, the figure shows that UF^2 converges to the analytical signal as dt^2.

are nearly independent, since the contours are approximately composed of horizontal and vertical lines, leading to the appearance of terraces in the color plot. Inspection of the top of that plot shows that UF^2 converges to the analytical signal as dt^2, when Δ is sufficiently large that only the error from dt is significant. UF^2 reproduces the analytical result to within 1% by using $\Delta = 6\sigma$ and $dt = 0.25\sigma$, corresponding to $M = 25$. The spectra attained using these parameters is not shown, as it is visually identical to Fig. 6.5(a). We conclude that UF^2 accurately predicts nonlinear optical spectra including finite-pulse duration effects in systems with small M.

6.6 Conclusion

We have presented three separate components of the Ultrafast Spectroscopy Suite (UFSS), which is a modular suite of tools designed for predicting nonlinear optical spectra. We have presented a novel algorithm called UF^2 that uses the convolution theorem to efficiently propagate the time evolution of eigenstates of the system Liouvillian. UF^2 is designed for evaluating the contributions to spectroscopic signals from the Feynman diagrams that organize perturbative calculations of nonlinear optical spectra. UF^2 is the open-systems extension of the closed-system algorithm of the same name presented in Ref. 96. We have also presented a direct propagation technique called RKE and a Hamiltonian/Liouvillian generator (HLG), which creates Hamiltonians and Liouvillians for vibronic systems coupled to a Markovian bath. Using the HLG, we have demonstrated that UF^2 can be over 500 times faster than RKE for systems with small Hilbert space dimension N. Using a secular Redfield model, UF^2 is always

faster than RKE, and with full Redfield UF^2 is faster up to system sizes where $\mathscr{L}_0$ rqeuires more than 20 GB of memory. In the diabatic Lindblad model, UF^2 outperforms RKE for $N \lesssim 100$. Both UF^2 and RKE methods are available with UFSS and can be used where appropriate.

A fourth module of UFSS, called the diagram generator (DG), is presented in Ref. [97]. The DG automatically generates all of the necessary Feynman diagrams that are needed to calculate a spectroscopic signal given the phase-matching (or phase-cycling) condition, the pulse shapes and pulse arrival times. Taken all together, the UFSS allows fast and automated calculations of nonlinear optical spectra of any perturbative order, for arbitrary pulse shapes, since UF^2 and RKE can both automatically calculate spectra given a list of Feynman diagrams. If desired, a user of UFSS need not concern themselves with the details of the perturbative calculations carried out by UF^2 and RKE. They simply must input the phase-matching conditions and pulse shapes of interest.

Each module of UFSS presented here can also be used separately. UF^2 and RKE can calculate the signal due to only a single Feynman diagram, or only the time-ordered diagrams, as is done when comparing to the analytical results of Ref. 19. UF^2 and RKE are compatible with any Hamiltonian or Liouvillian that is time-independent and can be expressed as or well-approximated by a finite matrix. Thus users are free to input their own model systems, and we include helper functions for saving other Hamiltonians and Liouvillians into a format compatible with UF^2 and RKE.

UFSS is available under the MIT license at github. The repository includes Jupyter notebooks that generate Figures 7.1 and 6.5 from this manuscript and scripts that generate Figures 6.3 and 6.4.

Acknowledgments

We thank an anonymous reviewer for the suggestion to consider the Redfield methods and acknowledge support from the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Ontario Trillium Scholarship.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request. In addition, the code to generate all of the figures in this manuscript are available at github.

6.7 Appendix: Computational cost

Here we derive the asymptotic computational costs of the UF^2 and RKE methods for the secular Redfield, full Redfield, and diabatic Lindblad models. For an arbitrary n^{th}-order spectroscopy, both UF^2 and RKE must calculate the same number of diagrams. Since the total number of diagrams affects the total runtime of each algorithm and not the ratio of the runtimes, we derive the computational cost of calculating the signal $S_d^{(n)}$ due to a single Feynman

diagram for each algorithm, which we call C_{UF^2} and C_{RKE}. This cost is the cost of calling $O_{j^{(*)}}$ n times to arrive at $|\rho_d^{(n)}(t)\rangle\!\rangle$, plus the additional cost of calculating the signal $S_d^{(n)}$ from $|\rho_d^{(n)}(t)\rangle\!\rangle$, as in, for example, Eq. 6.16. If the RWA holds, and there are well-defined manifolds (ground-state, singly excited, doubly excited, etc.) such that only the optical perturbations couple between them on the timescales of interest, then the density matrix and Liouvillian can be broken into smaller pieces. We derive the cost of the general case where the manifolds are not separable and then extend the result to separable manifolds.

6.7.1 UF^2

For UF^2, each evaluation of the $O_{j^{(*)}}$ operators is dominated by two operations: (1) multiplying the old state by the dipole operator to obtain $y_\beta(t)$ (see Eq. 6.20), and (2) performing the convolution $\theta(t) * y_\beta(t)$ using the FFT (see Eq. 6.21). The cost of both of these operations depends on the model system being studied. In the general case with inter-manifold relaxation processes, all density matrices are expressed as vectors of length N^2:

$$|\rho\rangle\!\rangle = \sum_\alpha c_\alpha(t) e^{z_\alpha t} |\alpha\rangle\!\rangle,$$

and all the operators on this space, $\mathscr{L}_0$, $\boldsymbol{\mu}^K \cdot \mathbf{e}_i$, $\boldsymbol{\mu}^B \cdot \mathbf{e}_i$, are $N^2 \times N^2$ matrices. Given $|\rho^{(n-1)}(t)\rangle\!\rangle$, the first step in determining $|\rho_d^{(n)}(t)\rangle\!\rangle = O_{j^{(*)}}|\rho^{(n-1)}(t)\rangle\!\rangle$ is to determine the coefficients $y_\beta(t)$ at M time points, where M must be large enough to well-represent the pulse envelope shape $A(t)$, see Fig. 6.5. This cost is the cost of the matrix-vector multiplication $\mu^O|\rho\rangle\!\rangle$ performed M times,

$$C(\mu^O|\rho(t)\rangle\!\rangle) \sim M C_\mu^{\mathrm{UF}^2},$$

where we show that for all of the cases we study here, $C_\mu^{\mathrm{UF}^2}$ scales as either N^3 or N^4. The next step, calculating the convolutions $\theta(t) * y_\beta(t)$ using the FFT, is

$$C\left(\theta(t) * y_\beta(t)\right) \sim N^2 M \log_2 M,$$

where the factor of N^2 arises from the N^2 values of β. Since $C\left(\theta(t) * y_\beta(t)\right) \sim N^2$, this cost is lower order than C_μ, and we disregard it. Asymptotically we thus find that

$$C(O_{j^{(*)}}) \sim M C_\mu^{\mathrm{UF}^2}, \tag{6.28}$$

where we retain the scaling with M for the purpose of comparison to RKE later.

For an n^{th}-order spectroscopy, each Feynman diagram describes an n^{th}-order density matrix. Starting from the unperturbed density matrix $|\rho^{(0)}\rangle\!\rangle$, we require n calls to the $O_{j^{(*)}}$ operator, and so the cost of obtaining $|\rho_d^{(n)}(t)\rangle\!\rangle$

from $|\rho^{(0)}\rangle\!\rangle$ is $nC(O_{j^{(*)}})$. This cost accrues for a single set of pulse delays. Since the calculations of $|\rho^{(n-1)}(t)\rangle\!\rangle$ can be reused, the most expensive part of calculating a multidimensional spectrum is varying the last pulse delay. We discussed this scaling in Appendix A of Ref. 96, and the same arguments apply here.

In order to calculate the desired signal S_d from $|\rho_d^{(n)}(t)\rangle\!\rangle$, the density matrix must be evaluated at a single time point (in the case of integrated measurements, as in phase-cycling experiments) or at the $M + m_t$ time points that determine $P_d^{(n)}(t)$ (as in Eq. 6.16 for phase-matching experiments). M time points are needed to resolve the turn-on of the signal, governed by the pulse envelope shape $A(t)$, and m_t is determined by the optical dephasing rate(s) and the desired frequency resolution of the final signal. Since $P_d^{(n)}(t) = \mathrm{Tr}[\boldsymbol{\mu}\rho_d^{(n)}(t)]$, we define the cost $C\left(\mathrm{Tr}[\boldsymbol{\mu}\rho]\right) = C_{\langle\mu\rangle}^{\mathrm{UF}^2}$ at a single time. Taking the trace at $M + m_t$ points has a cost of $C_{\langle\mu\rangle}^{\mathrm{UF}^2}(M + m_t)$. For polarization-based signals,

$$C(P_d^{(n)}(t)) \sim \underbrace{MC_{\mu}^{\mathrm{UF}^2}}_{C(O_{j^{(*)}})} + \underbrace{(M+m_t)C_{\langle\mu\rangle}^{\mathrm{UF}^2}}_{C\left(\mathrm{Tr}[\boldsymbol{\mu}\rho_d^{(n)}(t)]\right)}. \tag{6.29}$$

The cost of taking the FFT of $P_d^{(n)}(t)$ to obtain $S_d^{(n)}(\omega)$ does not depend on N and is negligible. Thus

$$C_{\mathrm{UF}^2} \sim MC_{\mu}^{\mathrm{UF}^2} + (M+m_t)C_{\langle\mu\rangle}^{\mathrm{UF}^2} \tag{6.30}$$

per set of pulse delay times. Depending upon the structure of $\mathscr{L}_0$, it is possible that $C_{\mu}^{\mathrm{UF}^2} = C_{\langle\mu\rangle}^{\mathrm{UF}^2}$, in which case calculating the polarization from $|\rho_d^{(n)}(t)\rangle\!\rangle$ can be more expensive than constructing $|\rho_d^{(n)}(t)\rangle\!\rangle$, while in other cases, $C_{\langle\mu\rangle}^{\mathrm{UF}^2}$ may be negligible. For phase-cycling cases, C_{UF^2} is only $MC_{\mu}^{\mathrm{UF}^2}$.

In order to use the UF^2 algorithm, we must also diagonalize $\mathscr{L}_0$. We call the cost of this operation C_D. The scaling of this cost with N is important for understanding for which Hamiltonian sizes UF^2 has an advantage over RKE. However, the precise size N at which the diagonalization cost becomes important depends upon how many calculations are done using the diagonalization, since C_D is amortized over each diagram, each time delay, each electric field shape studied, and each molecular-frame electric field polarization considered, as described in Sec. 6.4.

6.7.2 RKE

For RKE, each call to $O_{j^{(*)}}$ (1) uses the Euler method to connect $\rho^{(n-1)}$ to $\rho^{(n)}$ via Eq. 6.23 while the pulse is non-zero, and (2) extends the density matrix beyond $t_{j,\max}$ using the RK45 method to solve the ODE given by Eq. 6.22.

Given any state $|\rho\rangle\!\rangle$, the cost of the evolution according to Eq. 6.5 when $\mathscr{L}'(t) = 0$ is the cost of multiplying the vector $|\rho\rangle\!\rangle$ by the matrix $\mathscr{L}_0$. Given a local tolerance ϵ, the RK45 algorithm takes adaptive steps of size dt_{RK}. We approximate dt_{RK} as a constant and neglect the additional cost incurred when a step is rejected. For each step dt_{RK}, the RK45 algorithm must evaluate $\mathscr{L}_0|\rho\rangle\!\rangle$ 6 times. We take the cost per time step to be $C_{RK} = 6C(\mathscr{L}_0|\rho\rangle\!\rangle)$.

Given $|\rho^{(n-1)}(t)\rangle\!\rangle$, the cost of determining $|\rho^{(n)}(t)\rangle\!\rangle$ from $t_{j,\min}$ to $t_{j,\max}$ is the cost of the two main ingredients: including the pulse via the operation $\mu^O|\rho^{(n-1)}\rangle\!\rangle$, which has cost C_μ^{RKE}, and a call to the RK45 algorithm with cost C_{RK}. We divide the interval $[t_{j,\min}, t_{j,\max}]$ into M_E time points with equal spacing dt_E. Thus the cost of determining $|\rho^{(n)}(t)\rangle\!\rangle$ from $t_{j,\min}$ to $t_{j,\max}$ is

$$C_{\text{pulse}} = \left(C_\mu^{RKE} + C_{RK}\right) M_E,$$

where we have assumed that $dt_E < dt_{RK}$.

From $t_{j,\max}$ to some final time t_f, the RK45 algorithm advances the density matrix forward in time with cost $C_{RK}M_{RK}$, where $M_{RK} = (t_f - t_{j,\max})/dt_{RK}$. Thus

$$\begin{aligned} C_{RKE}(O_{j^{(*)}}) &= C_{\text{pulse}} + C_{RK}M_{RK} \\ C_{RKE}(O_{j^{(*)}}) &= M_E C_\mu^{RKE} + (M_E + M_{RK})\, C_{RK}. \end{aligned}$$

As with UF^2, RKE must also resolve the polarization field, which involves the cost $C_{\langle\mu\rangle}$. However, for RKE, $C_{\langle\mu\rangle}$ is always a lower-order cost. In diabatic Lindblad, the cost of $\mathrm{Tr}[\mu\rho]$ scales linearly with N, because μ is sparse. For both full and secular Redfield, the cost of $\mathrm{Tr}[\mu\rho]$ scales as N^2. In all cases these are lower order than other scaling costs (as summarized in Table 6.1 and derived below). The cost of a signal for RKE is thus

$$C_{RKE} \sim M_E C_\mu^{RKE} + (M_E + M_{RK})\, C_{RK}$$

6.7.3 Open systems models

6.7.3.1 Diabatic Lindblad

In the diabatic basis, $\mathscr{L}_0$ and μ^O are represented in the site and vibration number basis. In this basis, $\mathscr{L}_0$ and μ^O are sparse matrices, so that for RKE, C_μ^{RKE} and C_{RK} both scale as N^2. For UF^2, the cost of diagonalization is $C_D \sim N^6$. Transforming μ^O into the eigenbasis of $\mathscr{L}_0$ causes μ^O to become dense, so that for UF^2, both $C_\mu^{\text{UF}^2}$ and $C_{\langle\mu\rangle}^{\text{UF}^2}$ scale as N^4. Therefore, even without C_D, RKE outperforms UF^2 for large enough N. We find that RKE starts to outperform UF^2 around $N \approx 100$ in our test cases shown in Fig. 6.4.

6.7.3.2 Full Redfield

In the full Redfield formalism, $\mathscr{L}_0$ is expressed in the eigenbasis of H_0, and is a dense $N^2 \times N^2$ matrix, which requires both RKE and UF^2 to work in this eigenbasis; RKE then no longer has the advantage of a sparse μ operator. We define the eigenstates of H_0 to be $H_0\,|i\rangle = \epsilon_i\,|i\rangle$. In $\mathbb{H}$, when μ is transformed into the eigenbasis of H_0 it becomes a dense $N \times N$ matrix, and thus μ^O in $\mathbb{L}$ is a sparse matrix with N entries per row ($\mu^O|\rho\rangle\!\rangle$ can be reexpressed as

$\mu\rho$ or $\rho\mu$, which shows more transparently that this operation is the cost of multiplying two dense matrices, with cost scaling as N^3).

Therefore all of the dipole-multiplications have the same scaling, $C_{RK}, C_\mu^{\mathrm{UF}^2}, C_{\langle\mu\rangle}^{\mathrm{UF}^2} \sim N^4$ and $C_\mu^{RKE} \sim N^3$. RKE is then dominated by the cost of propagating the density matrix using the RK45 method, and both C_{RKE} and C_{UF^2} scale as N^4. For UF^2, $C_D \sim N^6$, as before.

At large N, we find

$$\frac{C_{RKE}}{C_{\mathrm{UF}^2}} \approx \frac{(M_E + M_{RK})\, C_{RK}}{(2M + m_t) C_\mu^{\mathrm{UF}^2}}.$$

C_{RK} is 6 times the cost of matrix-vector multiplication, while $C_\mu^{\mathrm{UF}^2}$ is the cost of a single matrix-vector multiplication, and so in this case $C_{RK} \approx 6C_\mu^{\mathrm{UF}^2}$. In the cases that we have tested, with parameters chosen to achieve 1% agreement in the resulting spectra, we typically find that $M_E \approx 20M$ and that $M_{RK} \approx m_t \approx M$. With these substitutions, for large N,

$$\frac{C_{RKE}}{C_{\mathrm{UF}^2}} \approx 40.$$

Since UF^2 has better prefactors than RKE, RKE does not outperform UF^2 until C_D becomes dominant, though memory constraints (not included in this calculation) likely constrain N before this crossover is reached. Figure 6.3 shows that for $N \approx 100$, $C_{RKE}/C_{\mathrm{UF}^2} \approx 80$, exceeding the estimate here, even when including the cost of diagonalization in the UF^2 cost.

6.7.3.3 Secular Redfield

In secular Redfield formalism, $\mathscr{L}_0$ is expressed in the eigenbasis of H_0, but nearly all of the entries of this $N^2 \times N^2$ matrix are zero. The unitary part of $\mathscr{L}_0$ is diagonal in this basis, and so the only off-diagonal terms come from the Redfield tensor R. The secular approximation sets all terms of R_{ijkl} to zero, except those that satisfy the condition that $|\omega_{ij} - \omega_{kl}| = 0$, where $\omega_{ij} = \omega_i - \omega_j$ specifies the time evolution frequency of the density matrix element ρ_{ij} due to the unitary part of $\mathscr{L}_0$. All populations ρ_{ii} evolve at $\omega_{ii} = 0$, and thus the secular approximation preserves all of the terms of R_{iikk}. All coherence-coherence and coherence-population terms are zero unless there are degeneracies in H_0 or harmonic ladders of eigenstates [108]. Even in those cases, the subsets of coupled coherences form additional blocks in $\mathscr{L}_0$ that are of a negligible size compared to the $N \times N$ populations block (see note below). Thus, in the secular approximation, we have $C_D \sim N^3$, regardless of the structure of H_0.

In the general case without degenerate eigenstates and harmonic ladders, $\mathscr{L}_0$ has a single $N \times N$ block coupling populations and is otherwise already diagonal. Since $\mathscr{L}_0$ is block diagonal, it is also sparse, and therefore $C_{RK} \sim N^2$. As for full Redfield theory, $C_\mu^{RKE} \sim N^3$ because the dipole operator must be represented in the eigenbasis of H_0.

For UF^2, we diagonalize $\mathscr{L}_0$ by finding the right and left eigenvectors, $|\alpha\rangle\!\rangle$ and $\langle\!\langle\bar{\alpha}|$, respectively. Let V_R be the matrix whose columns are $|\alpha\rangle\!\rangle$, and let V_L be the matrix whose rows are the $\langle\!\langle\bar{\alpha}|$. Just as with $\mathscr{L}_0$, V_R and V_L each have a dense $N \times N$ block, with the rest of each matrix being an identity (since $\mathscr{L}_0$ was otherwise already

diagonal). For UF^2, we represent ρ in the eigenbasis of $\mathscr{L}_0$ as in Eq. 6.19. We can also represent ρ in the basis that arises naturally from the eigenbasis of H_0 as

$$\rho = \sum_{i,j} c_{ij} |ij\rangle\!\rangle,$$

where $|ij\rangle\!\rangle = |i\rangle \langle j|$. V_R and V_L allow us to move between these two bases. For general $\mathscr{L}_0$, as in the full Redfield or diabatic Lindblad cases, V_R and V_L are dense, and so the cost of evaluating $V_L|\rho\rangle\!\rangle$ or $V_R|\rho\rangle\!\rangle$ scales as N^4; in those cases, we do not move between bases in order to compute $\mu^O|\rho\rangle\!\rangle$, but instead transform the dipole operator into the $|\alpha\rangle\!\rangle$ basis once. However, in the secular approximation, the cost of $V_L|\rho\rangle\!\rangle$ and $V_R|\rho\rangle\!\rangle$ scales as N^2 and is thus a negligible asymptotic cost. UF^2 can then propagate in the $|\alpha\rangle\!\rangle$ basis and apply μ in the $|ij\rangle\!\rangle$ basis, giving a large performance improvement. Then $C_\mu^{\mathrm{UF}^2}$ is identical to C_μ^{RKE}, scaling as N^3. Note that in the $|ij\rangle\!\rangle$ basis, $C_{\langle\mu\rangle}^{\mathrm{UF}^2} \sim N^2$, so is negligible. We then find

$$\frac{C_{RKE}}{C_{\mathrm{UF}^2}} \approx \frac{C_\mu^{RKE} M_E}{C_\mu^{\mathrm{UF}^2} M} \approx 20,$$

where we have once again used $M_E \approx 20M$ from the studies in Fig. 6.3. This result shows that UF^2 always outperforms RKE, regardless of N. Thanks to the block-diagonal structure of $\mathscr{L}_0$, C_D is unimportant, regardless of N.

6.7.3.4 Secular Redfield with harmonic modes

We now briefly justify the claim that $C_D \sim N^3$ even for the case of harmonic ladders. Let us take a system of k harmonic modes with unique frequencies ω_α. Representing each mode in the number basis, we truncate each mode at an occupation number of n. In this case $N = n^k$. In the secular approximation, all of the populations are coupled, and so, as stated above, $\mathscr{L}_0$ has a dense $N \times N$ block describing population dynamics. The secular approximation only couples the coherences of a single harmonic mode to other coherences of the same mode. Furthermore, it only couples coherences that oscillate at the same frequency. In the stated truncation scheme, each mode has $n-1$ coherences that oscillate at frequency $\omega_{ij} = \omega_\alpha$. Each mode has $n-2$ coherences that oscillate at frequency $\omega_{ij} = 2\omega_\alpha$. In general, each mode has $n-\nu$ coherences that oscillate at frequency $\omega_{ij} = \nu\omega_\alpha$. Therefore, for each harmonic mode, there are ν blocks of size $(n-\nu) \times (n-\nu)$ for $\nu = 1, 2, ...n$. The cost of diagonalizing each block scales as $\sim (n-\nu)^3$, and the total cost of diagonalizing all of the blocks is therefore

$$k \sum_{\nu=1}^{n} (n-\nu)^3 \approx kn^4.$$

Now recall that the cost of diagonalizing the population block is $O(N^3)$, which is $O(n^{3k})$. Thus for $k > 1$, the cost of diagonalizing all of the smaller coherence-coupling blocks is negligible for computational complexity analyses. For the case of $k = 1$, an analytical solution exists for diagonalizing $\mathscr{L}_0$ and thus $C_D = 0$ [114]. The derivation of

Table 6.2: Summary of asymptotic computational cost for each algorithm and bath formalism. This table has the same form as Table 6.1, with N replaced with N_X and N_Y, where $X \neq Y$. The cost of a general n^{th}-order spectroscopy scales with the size of the two largest accessible manifolds. For 3^{rd}-order spectroscopies of polymers with $s > 2$, the two largest accessible manifolds are $X = SEM, DEM$.

	UF^2 Scaling	RKE Scaling		$\mathscr{L}_0$ Diagonalization Cost	UF^2 advantage
		C_{RK45}	C_μ^{RKE}		
Full Redfield	$N_X^2 N_Y N_{Y'}$	$N_X^2 N_Y^2$	$N_X^2 N_Y$	$N_X^3 N_Y^3$	Up to large N_X
Secular Redfield	$N_X^2 N_Y$	$N_X N_Y$	$N_X^2 N_Y$	N_X^3	All N_X
Diabatic Lindblad	$N_X^2 N_Y N_{Y'}$	$N_X N_Y$	$N_X N_Y$	$N_X^3 N_Y^3$	$N_X \lesssim 100$

the analytical solution is in the Lindblad formalism; however, secular Redfield can be mapped onto the Lindblad formalism. Therefore, even for the case of harmonic ladders, we have that $C_D \sim N^3$.

6.7.4 Separable manifolds

We briefly describe how both UF^2 and RKE scale when there is no inter-manifold relaxation process, and therefore $\mathscr{L}_0$ breaks down into blocks of size $N_X N_Y \times N_X N_Y$, and thus for dense $\mathscr{L}_0$, $C_{RK45} \sim N_X^2 N_Y^2$. When $X = Y$, the block describes population and coherence dynamics within a manifold. When $X \neq Y$, the block describes the evolution of coherences between manifolds. In general, the dipole operator μ^O connects blocks of $\mathscr{L}_0$ by changing either X or Y, and thus has a shape $N_X N_{Y'} \times N_X N_Y$ or $N_{X'} N_Y \times N_X N_Y$. For dense μ^O, $C_\mu \sim N_X^2 N_Y N_{Y'}$. The actual scalings depend upon the sparsity structure (or lack thereof) of $\mathscr{L}_0$ and μ^O. The asymptotic costs in terms of manifold sizes N_X are summarized in Table 6.2. We plot the scaling of C_{RK} and C_{UF^2} and their ratios as a function of N_{SEM} because the ratio C_{RK}/C_{UF^2} (1) depends upon N_{SEM} alone for diabatic Lindblad, (2) depends upon N_{DEM}/N_{SEM} for full Redfield, or (3) is a constant for secular Redfield. The runtime costs of C_{UF^2} and C_{RKE} individually depend upon both N_{SEM} and N_{DEM}, as shown below.

6.7.4.1 Full Redfield

In full Redfield theory, RKE is dominated by the RK45 algorithm, with $C_{RK45} \sim N_X^2 N_Y^2$. For 3^{rd}-order spectroscopies, the most expensive diagram is the excited-state absorption (ESA), which evolves $\rho^{(3)}$ in the coherence between the $X = SEM$ and the $Y = DEM$, so $C_{RK45} \sim N_{SEM}^2 N_{DEM}^2$.

UF^2 is dominated by $C_\mu^{\mathrm{UF}^2} \sim N_X^2 N_Y N_{Y'}$. For the ESA, UF^2 is dominated by the cost of $\mu^K |\rho^{(2)}\rangle\!\rangle$, where $|\rho^{(2)}\rangle\!\rangle$ is in the SEM with length N_{SEM}^2. μ^K connects the SEM block to the SEM/DEM coherence block, and so has shape $N_{SEM} N_{DEM} \times N_{SEM}^2$. Thus $C_\mu^{\mathrm{UF}^2} \sim N_{SEM}^3 N_{DEM}$.

6.7.4.2 Secular Redfield

Both UF^2 and RKE are dominated by the same operation in the asymptotic limit, evaluating $\mu\rho^{(n-1)}$. In secular Redfield both UF^2 and RKE perform this multiplication with μ and ρ represented as matrices in the eigenbasis of

H_0. In general μ has size $N_{X'} \times N_Y$ and $\rho^{(n-1)}$ has size $N_X \times N_{Y'}$. For the ESA, μ has size $N_{DEM} \times N_{SEM}$, and $\rho^{(2)}$ has size $N_{SEM} \times N_{SEM}$. The cost of this operation then scales in general as $C_\mu^{\mathrm{UF}^2} \sim C_\mu^{RKE} \sim N_X^2 N_Y$ and for the ESA as $\sim N_{SEM}^2 N_{DEM}$.

6.7.4.3 Diabatic Lindblad

In diabatic Lindblad, $\mathscr{L}_0$ is sparse. RKE represents μ^O in the diabatic site and vibration-number basis, so that it is also sparse. Thus $C_{RK45} \sim C_\mu^{RKE} \sim N_X N_Y$. The cost of the ESA then goes as $\sim N_{SEM} N_{DEM}$.

In the eigenbasis of $\mathscr{L}_0$, μ^O is dense, and so, just as for full Redfield theory, $C_\mu^{\mathrm{UF}^2} \sim N_X^2 N_Y N_{Y'}$, and for the ESA, $C_\mu^{\mathrm{UF}^2} \sim N_{SEM}^3 N_{DEM}$.

Chapter 7

Diagram Generator

Reproduced from Journal of Chemical Physics 154 034109 (2021), with the permission of AIP Publishing

Original title: *Automatic Feynman diagram generation for nonlinear optical spectroscopies and application to fifth-order spectroscopy with pulse overlaps*

Peter A. Rose and Jacob J. Krich

Original Abstract: Perturbative nonlinear optical spectroscopies are powerful methods to understand the dynamics of excitonic and other condensed phase systems. Feynman diagrams have long provided the essential tool to understand and interpret experimental spectra and to organize the calculation of spectra for model systems. When optical pulses are strictly time ordered, only a small number of diagrams contribute, but in many experiments pulse-overlap effects are important for interpreting results. When pulses overlap, the number of contributing diagrams can increase rapidly, especially with higher order spectroscopies, and human error is especially likely when attempting to write down all of the diagrams. We present an automated Diagram Generator (DG) that generates all of the Feynman diagrams needed to calculate any n^{th}-order spectroscopic signal. We characterize all perturbative nonlinear spectroscopies by their associated phase-discrimination condition as well as the time intervals where pulse amplitudes are nonzero. Although the DG can be used to automate impulsive calculations, its greatest strength lies in automating finite pulse calculations where pulse overlaps are important. We consider third-order transient absorption spectroscopy and fifth-order exciton-exciton interaction 2D (EEI2D) spectroscopy, which are respectively described by 6 or 7 diagrams in the impulsive limit but 16 or 240 diagrams, respectively, when pulses overlap. The DG allows users to automatically include all relevant diagrams at relatively low computational cost, since the extra diagrams are only generated for the inter-pulse delays where they are relevant. For EEI2D spectroscopy, we show the important effects of including the overlap diagrams.

7.1 Introduction

Nonlinear optical spectroscopies (NLOS) are powerful tools for elucidating excited state dynamics of a wide variety of condensed phase systems and have been particularly important for determining the evolution of excitonic systems [4, 7, 66, 89, 95, 113, 115–120]. Interpretation of data-rich NLOS signals is centered on Feynman diagrams, which conveniently graphically summarize time-dependent perturbation theory contributions to the signals [14]. These diagrams give a visual understanding of what types of excited- and ground-state dynamics and/or coherences are probed and can be straightforwardly turned into calculations of contributions to signals. In many cases, diagrams can be determined to give zero contribution without a complicated calculation, making theoretical interpretation easier and calculation less expensive.

When pulse durations are shorter than system dynamics – the impulsive limit – the number of contributing diagrams is often small, especially for third-order spectroscopies. It has been easiest to build intuition about the experimental signatures of particular excitonic processes by considering the impulsive limit [4, 7, 10, 66, 116, 117, 121–127]. When pulse durations are similar to the timescale of system evolution, however, the effects of pulse overlaps – where the tail of a nominally earlier pulse arrives after the beginning of a nominally later pulse – become important to accurately model and understand experimental results [16–19, 57, 88, 90–94, 128, 129]. Considering such processes requires many more contributing diagrams. Higher-order spectroscopies, both in the impulsive and pulse-overlap limits, also involve rapidly increasing numbers of contributing diagrams. Human error and fatigue in determining these diagrams accurately become increasingly likely as their numbers proliferate.

We present here an automated Feynman Diagram Generator (DG), which allows convenient, fast, and accurate determination of the diagrams contributing to a particular spectroscopy. We describe all perturbative nonlinear spectroscopies by their associated phase-discrimination conditions. These conditions allow description of non-colinear phase-matching, colinear phase-cycling experiments, and combinations of the two [10, 14, 39, 41, 56, 83, 84, 127, 130]. Users input those conditions and the time intervals in which their pulses are nonzero, and the DG produces a list of contributing diagrams. That list can be passed directly to calculation engines to predict the associated spectroscopic signals or can be drawn in the standard diagrammatic format for review. The DG is part of a set of NLOS tools called the Ultrafast Spectroscopy Suite (UFSS), which also contains methods to generate Hamiltonians or Liouvillians for vibronic systems and to calculate the contributions from each diagram [96, 131], but its output diagrams can be used with other computational tools or analytic methods to calculate spectroscopic signals [18, 19, 21, 30, 59, 64, 66, 76–78, 91, 125, 132–134]. The DG is free and open-source software written in Python, available for download from github.

The DG (and the larger UFSS) is designed to make simple the inclusion of the effects of finite-duration pulses in NLOS. While modeling often considers the impulsive limit, pulse-overlap effects can dominate third-order signals such as two-dimensional photon echo (2DPE) spectroscopy, even outside of what is commonly thought of as the pulse-overlap window [57]. 2DPE is usually calculated using 6 time-ordered diagrams (3 for the rephasing and 3 for

the non-rephasing pathway). When pulses overlap, up to 16 diagrams (shown in Fig. 7.1) contribute to the signal [13]. These 16 diagrams have been written down by hand and used for calculations in the past.

In higher-order methods, such as fifth-order exciton-exciton interaction 2D spectroscopy (EEI2D), there are 7 Feynman diagrams in the impulsive limit [126]. We show that outside of the impulsive limit, EEI2D requires up to 240 Feynman diagrams. This number of diagrams is too large for generation by hand, and as a result we do not believe a calculation with all of these diagrams has previously been attempted in the perturbative limit. References 20, 95 consider finite pulse effects from time-ordered diagrams for several 5^{th}-order spectroscopies. We use EEI2D as our key example of the application of automated diagram generation, showing both 1) that these extra overlap diagrams can be generated and their contributions calculated with surprisingly small extra computational cost and 2) that finite pulse effects arising from non-time-ordered diagrams can provide significant modifications to a sample EEI2D spectra, which can be important for interpretation of experimental results.

The DG automatically creates all of the diagrams that satisfy a given phase-discrimination condition. Since the DG is computationally inexpensive compared to evaluating the contribution of each diagram, in UFSS the DG is run for each set of desired pulse delays. It determines whether pulses overlap, and thus whether overlap diagrams contribute, allowing computation of only causal diagrams for each set of pulse delays. This determination provides a significant computational time savings, since the large number of overlap diagrams only need to be calculated at the (usually small) proportion of pulse delays where they contribute. We further show that heavy-tailed pulse envelopes, such as Lorentzians, require calculation of overlap diagrams out to longer delay times than Gaussian pulse envelopes.

We begin with an overview of perturbative spectroscopy calculations in Sec. 7.2 in order to establish the construction of Feynman diagrams, with a unified perspective on both phase matching and phase cycling spectroscopies. We describe the algorithm of the DG in Sec. 7.3 including optional methods to reduce the number of diagrams, depending on the system and spectroscopy considered. We demonstrate the utility of the DG by exploring EEI2D spectroscopy in Sec. 7.4. We consider the same model system as in Ref. 112, and show that when optical pulses have slightly longer durations than considered in that work, neither the impulsive limit nor the time-ordered diagrams with pulse-shape effects included accurately predict spectra. The calculation with all 240 diagrams requires less than twice the time as using only the 7 time-ordered diagrams despite considering 34 times as many diagrams. These results underscore the importance of including all of the additional 233 overlap diagrams, as well as the utility of the DG in not only generating these diagrams, but also automatically determining in which conditions they contribute.

7.2 Nonlinear spectroscopy and Feynman diagrams

We begin by establishing the standard perturbative framework of nonlinear optical spectroscopy, from which Feynman diagrams are defined [14]. Consider a system with density matrix ρ, which evolves in the absence of perturbation according to

$$\rho(t) = \mathcal{T}_0(t, t')\rho(t'),$$

where $\mathcal{T}_0$ is a time evolution operator. We restrict the following discussion to the case of Hamiltonian systems or Markovian open systems, in which $\mathcal{T}_0(t, t') = \mathcal{T}_0(t - t')$, which must be known or approximated in order to complete calculations, but the resulting diagrams are broadly applicable to non-Markovian situations, as well. For the purposes of diagram generation, however, we simply need to assume $\mathcal{T}_0$ exists. We also define the differential time evolution operator $\mathscr{L}_0$ so

$$\frac{d\rho(t)}{dt} = \mathscr{L}_0\rho(t).$$

The perturbative optical fields are described as classical electric fields $\mathbf{E}(t)$, which interact with the system in the electric-dipole approximation through the perturbation Hamiltonian

$$H'(t) = -\boldsymbol{\mu} \cdot \mathbf{E}(t), \tag{7.1}$$

where $\boldsymbol{\mu}$ is the electric dipole operator. Then the time evolution of the system is

$$\frac{d\rho(t)}{dt} = \mathscr{L}_0\rho(t) - \frac{i}{\hbar}\left[H'(t), \rho(t)\right]. \tag{7.2}$$

This form is the basis for diagramatic perturbation theory in $\mathbf{E}(t)$.

We describe $\mathbf{E}(t)$ as a sum over L pulses, where each pulse is denoted by a lowercase letter starting from a. A typical 3^{rd}-order signal is calculated using up to 4 pulses. We write the electric field as

$$\mathbf{E}(t) = \sum_{j=a,b,\ldots,L} \mathbf{e}_j\varepsilon_j(t) + \mathbf{e}_j^*\varepsilon_j^*(t) \tag{7.3}$$

where $\mathbf{e}_j$ is the possibly complex polarization vector, and the amplitude ε_j of each pulse is defined with envelope A_j, central frequency ω_j, wavevector $\mathbf{k}_j$, and phase ϕ_j as

$$\varepsilon_j(t) = A_j(t - t_j)e^{-i(\omega_j(t-t_j) - \mathbf{k}_j\cdot\mathbf{r} - \phi_j)},$$

where t_j is the arrival time of each pulse. We make the physical assumption that each pulse is zero outside the finite interval $[t_{j,\min}, t_{j,\max}]$. The DG uses this range to determine when pulses overlap; the form of $A_j(t)$ is unimportant for diagram generation. The light-matter interaction is a sum over the rotating (ε_i) and counter-rotating (ε_i^*)

terms. In the rotating wave approximation (RWA), the rotating terms excite the ket-side and de-excite the bra-side of the density matrix, respectively, and the counter-rotating terms excite the bra-side and de-excite the ket side, respectively [14].

We treat the effect of the optical fields using standard time-dependent perturbation theory and assume that at time t_0 the system is in a stationary state of $\mathscr{L}_0$, which is $\rho^{(0)}$. Then

$$\rho(t) = \rho^{(0)} + \rho^{(1)}(t) + \rho^{(2)}(t) + ... \tag{7.4}$$

where [14]

$$\rho^{(n+1)}(t) = \int_0^\infty dt' \mathcal{T}_0(t') \left[\frac{-i}{\hbar} \boldsymbol{\mu} \cdot \mathbf{E}(t-t'), \rho^{(n)}(t-t') \right]. \tag{7.5}$$

Using Eq. 7.3, we define $\rho^{(n+1)}(t)$ as a sum over four types of terms,

$$\rho^{(n+1)}(t) = \sum_j \left(K_j + K_{j^*} + B_j + B_{j^*} \right) \rho^{(n)}(t), \tag{7.6}$$

where K_j and B_j are superoperators representing ket- and bra-side actions, respectively, of the rotating terms of pulse j on ρ, while K_{j*} and B_{j*} give the equivalent counter-rotating terms. By inspection of Eq. 7.5, all four types of terms in Eq. 7.6 can be compactly defined as

$$O_{j^{(*)}} = \eta_O \frac{i}{\hbar} \int_0^\infty dt' \mathcal{T}_0(t') \left(\boldsymbol{\mu}^O \cdot \mathbf{e}_j^{(*)} \varepsilon_j^{(*)}(t-t') \right), \tag{7.7}$$

where $O = K, B$, $\eta_K = 1$ and $\eta_B = -1$, and we define dipole superoperators $\mu^K \rho \equiv \mu\rho$ and $\mu^B \rho \equiv \rho\mu$. The operators $\{O_{j^{(*)}}\}$ are the building blocks for all perturbative spectroscopies. The full $\rho^{(n)}$ is constructed from the unperturbed state $\rho^{(0)}$ as

$$\rho^{(n)}(t) = \left[\sum_{j=a,b,...,L} \left(K_j + K_{j^*} + B_j + B_{j^*} \right) \right]^n \rho^{(0)}, \tag{7.8}$$

which involves $(4L)^n$ terms when the exponent and sum are fully expanded. Each of these terms is a sequence of n applications of $O_{j^{(*)}}$ and can be represented as a Feynman diagram. Most of these diagrams are unimportant for any given spectroscopy, and the subset of diagrams that contributes to a particular experiment is determined by the set of optical pulses and a phase-discrimination condition.

Regardless of the computational details used to calculate the $\{O_{j^{(*)}}\}$, all perturbative calculations can be represented and organized using the same Feynman diagrams, and thus the DG is useful for any perturbative spectroscopy algorithm. For example, UFSS contains two methods for calculating the action of the $O_{j^{(*)}}$ operators, which are derived for closed systems in Ref. 96 and for open systems in Ref. 131. The integral form of the $\{O_{j^{(*)}}\}$ in Eq. 7.7 is convenient for UF^2 and is the open-systems analogue of similar expressions derived for wavefunctions [21, 59].

The two most widely used phase-discrimination conditions are phase matching and phase cycling. Phase-matching conditions are achieved by using pulses that travel along different directions denoted by wavevectors $\mathbf{k}_j$, converging to interact with a sample that is assumed to be uniform over a volume large compared to the wavelength of the pulses [14]. The pulses induce a polarization field in the sample $\mathbf{P}(t) = \langle \boldsymbol{\mu}\rho(t)\rangle$, which produces radiation in all directions. A detector placed in a direction that satisfies the phase-matching condition $\mathbf{k}_d = \sum_j m_j \mathbf{k}_j$, where m_j are integers, is sensitive to a polarization field that is described by only a subset of diagrams. Often one is interested in the lowest-order signal in perturbation theory that contributes to the given phase-discrimination condition. Heterodyne detection with a local oscillator (LO) allows full determination of amplitude and phase of the emitted radiation.

For example 2D photon echo (2DPE) signals involve three pulses, a, b, c, that interact with the sample and a fourth LO pulse d. The 2DPE rephasing and non-rephasing signals are measured with detectors placed in the $\mathbf{k}_d = -\mathbf{k}_a + \mathbf{k}_b + \mathbf{k}_c$ (see Fig. 7.1) and $\mathbf{k}_d = \mathbf{k}_a - \mathbf{k}_b + \mathbf{k}_c$ directions, respectively. These signals are calculated using

$$\begin{aligned}
\mathbf{P}^{(3)}_{\mathrm{k}_d}(t) &= \langle \boldsymbol{\mu}\rho^{(3)}_{\mathrm{k}_d}(t)\rangle \\
\tilde{\mathbf{P}}^{(3)}_{\mathrm{k}_d}(\omega) &= \int_{-\infty}^{\infty} dt e^{i\omega t}\mathbf{P}^{(3)}_{\mathrm{k}_d}(t) \\
S^{(3)}_{\mathrm{k}_d}(\omega) &= \mathrm{Im}\left[\tilde{\varepsilon}^*_d(\omega)\mathbf{e}_d \cdot \tilde{\mathbf{P}}^{(3)}_{\mathrm{k}_d}(\omega)\right]
\end{aligned}$$

where $S^{(3)}_{\mathrm{k}_d}(\omega)$ are the signals and $\rho^{(3)}_{\mathrm{k}_d}(t)$ is the portion of $\rho^{(3)}(t)$ that produces radiation in the $\mathbf{k}_d$ direction; the primary purpose of diagrammatic perturbation theory is to organize the efficient calculation of $\rho^{(3)}_{\mathrm{k}_d}(t)$ without needing to calculate all contributions to $\rho^{(3)}(t)$.

An alternative phase discrimination method uses phase cycling over the relative phases of collinear pulses. This method generally detects a signal proportional to an excited state population, such as fluorescence or photocurrent [10, 39, 41, 130]. A fourth-order signal $S^{(4)}_d(t)$ in such a setup is calculated as

$$S^{(4)}_d(t) = \left\langle Q\rho^{(4)}_d(t)\right\rangle,$$

where Q is a projection operator onto the relevant excited electronic states and $\rho^{(4)}_d(t)$ includes only those contributions to $\rho^{(4)}(t)$ that contribute to the chosen phase-cycling condition. Combinations of phase-matching and phase-cycling are also possible and just as easily described in the DG [135].

For example, the 2DPE rephasing signal is composed of the 16 diagrams in Fig. 7.1, which were automatically generated and drawn using the DG. If, however, $t_{b,\mathrm{min}} > t_{a,\mathrm{max}}$ and $t_{c,\mathrm{min}} > t_{b,\mathrm{max}}$ (that is, the pulses are time ordered), then only three diagrams contribute, shown in the black box.

The operators $\{O_{j^{(*)}}\}$ as an abstract concept have previously been introduced in various forms [21, 59]. To our knowledge, they have not been used for the purposes of automated calculations before. In Sec. 7.3, we describe how

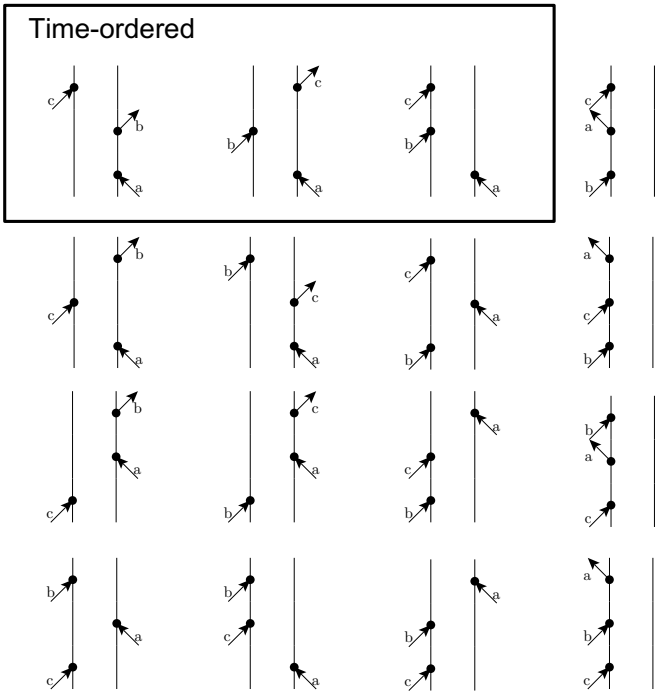

Figure 7.1: All diagrams needed to calculate the rephasing 2DPE signal. The box encloses the 3 time-ordered diagrams that are typically calculated. The other diagrams contribute to the signal when two or more of the pulses overlap. These diagrams were automatically generated and drawn using the diagram generator (DG). The DG can produce both double-sided diagrams, as shown, and pairs of single-sided diagrams for equivalent wavefunction calculations.

users inputs the pulse intervals and phase discrimination conditions and present the algorithm that produces the relevant diagrams. In Sec. 7.4 we demonstrate the DG by using up to 240 diagrams and show the importance of including the full set of diagrams when performing simulations with finite pulse durations.

7.3 Diagram Generator

The user of the DG inputs the desired phase discrimination condition and pulse intervals, and the DG automatically generates all Feynman diagrams. The DG is much less computationally expensive than the evaluation of the contribution from each diagram, so it is intended to be rerun for each set of desired pulse delays. The DG then returns pulse-overlap diagrams only when the pulse overlaps allow them to contribute. This package generates diagrams as a list of the operators $\{O_{j^{(*)}}\}$ defined in Eq. 7.7. The Ultrafast Ultrafast (UF^2) and Runge-Kutta-Euler (RKE) tools outlined in Ref. 131 and included in UFSS are designed to interpret this list and produce spectra. Alternatively, the generated diagrams can be exported to another calculation engine. The DG can also draw them for inspection.

Feynman diagrams are determined by the number of pulses and the phase-discrimination condition. Phase-discrimination determines the number of interactions with the rotating or counter-rotating part of each pulse that contribute to the measured signal. The user inputs the number of each type of interaction as a list of tuples

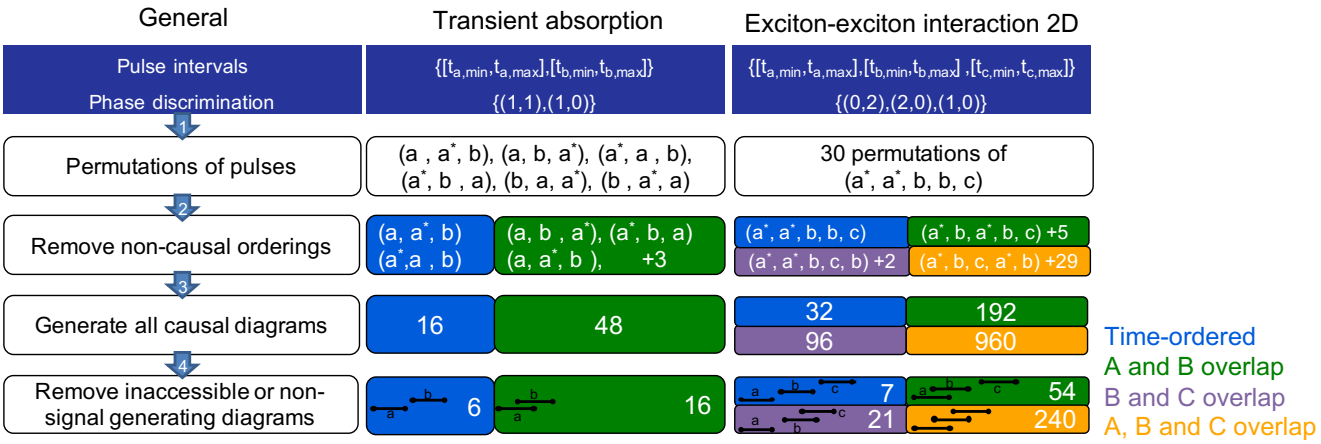

Figure 7.2: Graphical outline of automated diagram generation. Left column is the general case, while the center and right columns show specific examples. User inputs of pulse intervals and phase discrimination of the desired experiment are in the top box. The steps of the algorithm, numbered in the arrows, are described in the main text. After step 3, the numbers in the boxes indicate the number of diagrams at that stage. Blue boxes show results with non-overlapping pulses. Green shows results with pulses a and b overlapping, purple shows results with pulses b and c overlapping, and gold shows results with pulses a, b and c all overlapping. Inset black bars in the bottom row graphically indicate these pulse overlaps.

$[(n_r^a, n_c^a), (n_r^b, n_c^b), ..., (n_r^L, n_c^L)]$. This list corresponds to detection with $\mathbf{k}_d = \sum_{j=a,b,...,L}(n_r^j - n_c^j)\mathbf{k}_j$ in the case of phase matching and to detection of signals in phase with $\sum_{j=a,b,...,L}(n_r^j - n_c^j)\phi_j$, where ϕ_j are the modulated phases, in the case of phase cycling [39]. The order of spectroscopy is given by $n_{\text{total}} = \sum_{j=a,b,...,L} n_r^j + n_c^j$. In order to determine which diagrams are causally allowed, the user must also input the time intervals $[t_{j,\min}, t_{j,\max}]$ when each pulse is nonzero as a list. The user updates this list for each set of pulse delay times desired. The number of diagrams that must be considered when dealing with finite pulses expands dramatically as one considers higher-order signals, such as exciton-exciton interaction 2D spectra (the 3-pulse 5^{th}-order signal measured in the $\mathbf{k}_d = -2\mathbf{k}_1 + 2\mathbf{k}_2 + \mathbf{k}_3$ direction, see Sec 7.4), where the number of time-ordered diagrams is 7 [112]. However, when all pulses overlap, there are 240 diagrams.

We outline the steps for determining all diagrams that contribute to the signal for a given pulse configuration in Fig. 7.2, which includes the general scheme and two examples.

(1) Starting from $[(n_r^i, n_c^i)]_{i=a,b,...,L}$, we begin with a canonical list of interactions, where we list the pulses in order of pulse number starting from a, with rotating terms coming before counter-rotating terms (see Fig. 7.2 for examples). The length of the resulting list is n_{total}. For each pulse there are n_r^j repetitions of "j" with the rotating term and n_c^j repetitions of "j^*" with the counter-rotating term, as shown in Fig. 7.2. We then generate all unique permutations of this canonical list of interactions. The number of unique permutations is

$$\frac{n_{\text{total}}!}{\Pi_i n_r^i! n_c^i!}.$$

(2) Using the list of pulse intervals $\{[t_{j,\min}, t_{j,\max}]\}_{j=a,b,...L}$, non-causal orderings are removed. For each permu-

tation of the time-ordered list, we check whether $t_{\min}$ of each pulse occurs before $t_{\max}$ of each following pulse in the list and remove the permutation if not.

(3) For each permutation from step 2, we generate all of the allowed diagrams that satisfy the phase-discrimination conditions. Each interaction can occur either on the ket-side ($K_{j^{(*)}}$) or the bra-side ($B_{j^{(*)}}$), giving $2^{n_{\text{total}}}$ diagrams associated with each permutation from step 2. For example, the interaction a^* can act as K_{a^*} or as B_{a^*}, while the interaction b can act as either K_b or B_b.

At this point we have the maximum number of diagrams that could contribute to the calculation, given the phase-discrimination condition and pulse intervals. However, many of these diagrams do not contribute under common assumptions. Using some minimal information about the material system that is being modeled, many of these diagrams can be removed in the optional step 4.

(4) (Optional)

(a) Keep only diagrams that remain in accessible states. For instance, consider optical spectroscopy of a system that has three optically separated manifolds, each separated by an energy gap E_g. If the temperature is much less than E_g, we can approximate that the initial thermal state is entirely in the ground-state manifold. Any diagram that includes excitation above the highest manifold or de-excitation from the lowest manifold is removed. We track the number of optical excitations by assigning manifold indices for both the ket and bra sides of a density matrix. We say that the initial density matrix $\rho^{(0)}$ is entirely composed of ground-state populations, and therefore it has manifold indices $[0, 0]$. We then assign the following rules describing the action of the $O_{j^{(*)}}$ operators:

$$K_j : [+1, 0]$$
$$K_{j^*} : [-1, 0]$$
$$B_j : [0, -1]$$
$$B_{j^*} : [0, +1]$$

We apply these rules in succession for a diagram and track the indices $[i, l]$ after each interaction. If either i or l drop below 0 or rise above the maximum allowed manifold, the diagram is removed. When the manifolds are coupled by relaxation processes, we no longer remove diagrams that rise above the allowed maximum manifold, since population can decay to a lower manifold and then be excited up again by a subsequent interaction. We do still remove diagrams where i or l drop below zero.

(b) The integer logic of part 4(a) is also helpful in determining which diagrams contribute to the final signal. For spectroscopies that measure the emitted polarization field, we are interested in the object $\text{Tr}[\mu\rho]$. Typically in optical spectroscopy, μ connects only adjacent manifolds (either because this is an accurate model for the dipole

operator or because the measurement bandwidth only supports 1-manifold transitions). The components of ρ that contribute are then coherences between adjacent optical manifolds. Therefore, we filter out all diagrams except those that end in a state $[i+1, i]$ (Note that the diagrams that end in $[i, i+1]$ are physically valid; however, we do not calculate them, as they are the Hermitian conjugate pairs of the calculated diagrams [14]). Note that we again cannot apply this filter when $\mathcal{T}$ includes inter-manifold relaxation.

(c) If the final observable is instead linked to excited-state populations, as in the case of fluorescence or photocurrent detection, we look only for diagrams that end in a population $[i, i]$ where $i > 0$. Again, we cannot apply this filter when $\mathcal{T}$ includes inter-manifold relaxation.

Each of the diagram reductions from step 4 can be turned off with a flag or modified to suit an accurate model for the system in question. For example, in vibrational spectroscopy, rule 4a) does not apply, but rules 4b) or 4c) (or modifications of them) may still apply. We include a variety of examples of step (4) in the Jupyter notebook `DiagramGeneratorExample.ipynb`, including an example of infrared vibrational transient absorption that cannot remove any of the causal diagrams and must use all 16 (48 for overlapping pulses) diagrams that come from step (3).

Thus far we have described how the DG creates the double-sided Feynman diagrams associated with density-matrix-based calculations. The DG can also create the one-sided Feynman diagrams used for wavefunction-based calculations [14, 21]. This procedure begins by creating all relevant double-sided diagrams. Each double-sided diagram is converted to a pair of one-sided diagrams associated with the bra and ket wavefunctions. Since wavefunction-based calculations do not impose time-ordering between the bra and ket sides [14], several of the double-sided diagrams map onto the same pair of single-sided diagrams. After creating all of the pairs of single-sided diagrams, we eliminate duplicates to avoid over-counting some pathways.

The DG both produces a full set of contributing diagrams to be calculated and updates that list as the pulse timings change. In the example 3-pulse 5^{th}-order spectroscopy considered in Sec. 7.4, the pulse delays vary to produce frequency-domain signals. Only a small number of pulse configurations have all three pulses overlapping. In those cases, 240 diagrams must be calculated, but as the pulse delays change, calculations can be reduced to 54, 21, or 7 diagrams, which all occurs automatically. Including overlap diagrams increases the number of required diagrams by a factor of over 30, but since diagrams are only included when pulse delays require, calculations of frequency-domain spectra are only 1.5-5 times longer than those including only the 7 time-ordered diagrams, depending upon which pulse delays are calculated.

7.4 Importance of overlap diagrams

Spectra are often calculated assuming impulsive pulses. Even when finite pulse shapes are considered, the overlap diagrams are often neglected, as in the only two studies of finite pulses in 5^{th}-order spectroscopies of which we

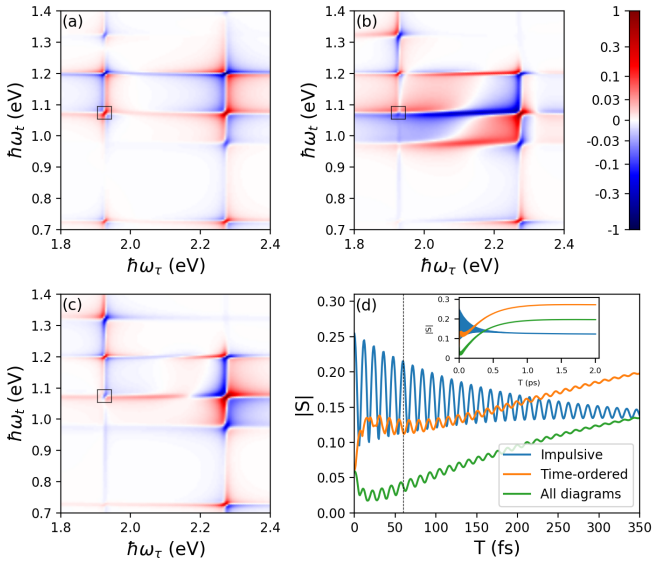

Figure 7.3: (a)-(c) Real part of the EEI2D spectrum at a delay time of $T = 60\,\text{fs}$ for a dimer of three-level systems with a relaxation rate of $0.015\ \text{fs}^{-1}$, as described in the text, with different electric field shapes. Horizontal axes show the Fourier transform of the delay time τ between the first and second pulses. Vertical axes show the detection frequency. EEI2D spectra were calculated for 275×275 values of τ, t and then Fourier transformed. Note the signed-logarithmic color scale. (a) Impulsive calculation, which requires only time-ordered diagrams, (b) 15 fs FWHM Gaussian pulses using only time-ordered diagrams, (c) 15 fs FWHM Gaussian pulses including all 54 diagrams that contribute at this T. (a)-(c) are normalized independently so that the largest peak has magnitude 1. (d) Maximum of the absolute value of the signal contained within the black boxes in (a)-(c) as a function of delay time T. Dashed line shows T used in (a-c). The inset shows a longer time window and the steady-state that is reached by 1 ps. The impulsive and time-ordered calculations both quantitatively and qualitatively deviate from the full calculation.

are aware [20, 95]. In some cases these approximations are warranted, but sometimes the overlap diagrams are important to understand signals both quantitatively and qualitatively. Here, we demonstrate an example where neglecting overlap diagrams leads to significant artifacts in predicted spectra.

We consider a model system used in Ref. 112 to study exciton-exciton annihilation, using three-pulse exciton-exciton interaction 2D (EEI2D) spectroscopy. The model consists of a dimer of two coupled three-level systems (3LS). The Hamiltonian takes the form

$$\begin{aligned} H_0 = & E_{gg}\left(|gg\rangle\langle gg|\right) \\ & + E_{eg}\left(|eg\rangle\langle eg| + |ge\rangle\langle ge|\right) + J\left(|eg\rangle\langle ge| + h.c.\right) \\ & + E_{ee}\left(|ee\rangle\langle ee|\right) + E_{fg}\left(|fg\rangle\langle fg| + |gf\rangle\langle gf|\right) \\ & + K\left(|fg\rangle\langle ee| + |gf\rangle\langle ee| + h.c.\right) \\ & + E_{ef}\left(|fe\rangle\langle fe| + |ef\rangle\langle ef|\right) + L\left(|fe\rangle\langle ef| + h.c.\right) \\ & + E_{ff}\left(|ff\rangle\langle ff|\right) \end{aligned}$$

where $E_{ij} = E_i + E_j$, all of the constants are defined in Table 7.1, $|g\rangle, |e\rangle, |f\rangle$ are the ground, singly excited, and doubly excited states, respectively, of each 3LS, and $|uv\rangle = |u\rangle_1 \otimes |v\rangle_2$. Following Ref. 112, we consider relaxation processes at zero temperature for each isolated monomer unit from $|f\rangle_i$ to $|e\rangle_i$ at rate k_M and neglect relaxation from $|e\rangle_i$ to $|g\rangle_i$, since it is not important for exciton-exciton interactions. Reference 112 included relaxation using the stochastic Schrodinger equation, while we include relaxation using the Lindblad formalism with

$$\dot{\rho} = -\frac{i}{\hbar}[H_0, \rho] - \sum_{n \neq m} k_{nm} L\left[|m\rangle\langle n|\right]\rho$$

where $|m\rangle$ and $|n\rangle$ are eigenstates of H_0 and the Lindblad superoperator $L[O]$ is defined by

$$L[O]\rho = O\rho O^\dagger - \frac{1}{2}\left(O^\dagger O\rho + \rho O^\dagger O\right).$$

We follow Ref. 112 by projecting the monomer relaxation rates into the eigenstates $|n\rangle$ with energy E_n. New relaxation rates coupling the eigenstates of H_0 are defined as

$$\begin{aligned} k_{nm} = & \Big(\left|\langle n\,|\,fg\rangle\right|^2 \left|\langle eg\,|\,m\rangle\right|^2 + \left|\langle n\,|\,gf\rangle\right|^2 \left|\langle ge\,|\,m\rangle\right|^2 \\ & + \left|\langle n\,|\,fe\rangle\right|^2 \left|\langle ee\,|\,m\rangle\right|^2 + \left|\langle n\,|\,ef\rangle\right|^2 \left|\langle ee\,|\,m\rangle\right|^2 \Big) k_M. \end{aligned}$$

The resulting rates range from 0.0075 fs^{-1} to 0.0016 fs^{-1} where the smallest rate corresponds to decay from the lowest-energy double-exciton state to the lowest-energy single-exciton state.

Table 7.1: Values used in the Hamiltonian H_0 and bath coupling rates for the model dimer system studied using EEI2D, adapted from Ref. 112.

	E_g	E_e	E_f
Energy (eV)	0.0	1.0	2.2
	J	K	L
Coupling (eV)	0.2	0.1	0.05
	k_M		
Rate (fs^{-1})	0.015		

EEI2D is designed to probe the dynamics of the doubly excited states. Diagonalizing H_0 shows that the two optically bright doubly excited states are separated by 0.35 eV, corresponding to an oscillation period of $T_o = 12$ fs, which is the fastest important oscillation in this system. One generally assumes that pulses significantly shorter than this period will be well-approximated by impulsive pulses but that pulses that interact with the system on a comparable timescale require more careful treatment. In this section we study the effects of using Gaussian pulses with a full-width half-maximum (FWHM) of 15 fs, with a comparison to Lorentzian pulses at the end.

The perturbative calculations of optical signals were performed using the UF^2 method detailed in Ref. 131. UF^2 calculates the fifth-order polarization signal as a function of the delay time τ between pulses a and b and the delay time T between pulses b and c as

$$\mathbf{P}^{(5)}_{\mathrm{k}_d}(\tau, T, t) = i\langle \boldsymbol{\mu} \rho^{(5)}_{\mathrm{k}_d}(t) \rangle,$$

where t is the time measured after the arrival of pulse c. The 2D frequency-frequency correlation spectrum is calculated as

$$\tilde{\mathbf{P}}^{(5)}_{\mathrm{k}_d}(\omega_\tau, T, \omega_t) = \frac{1}{2\pi} \int_{-\infty}^{\infty} d\tau e^{-i\omega_\tau \tau} \int_{-\infty}^{\infty} dt e^{i\omega_t t} \mathbf{P}^{(5)}_{\mathrm{k}_d}(\tau, T, t),$$

which is approximated using the discrete Fourier transform. We use $\tilde{\mathbf{P}}^{(5)}_{\mathrm{k}_d}(\omega_\tau, T, \omega_t)$ as a proxy for the signal field, which would in practice be detected using heterodyne detection with a local oscillator.

We compare calculations of EEI2D spectra using three approximations: impulsive pulses, finite pulses with only time-ordered diagrams, and finite pulses including all diagrams. Figure 7.3(a) shows the impulsive limit with a delay time of $T = 60$ fs. In keeping with Ref. 112, we perform all calculations with τ and t each ranging from 0 to 823 fs, and multiply the τ and t axes with a Gaussian window function of $\sigma = 200$ fs to avoid ringing effects from the discrete Fourier transform. We calculate spectra for 275 values of τ and t, which are more than sufficient to produce well-resolved frequency-frequency spectra. Reference 112 performs calculations in the impulsive limit and states that calculations using time-ordered diagrams with 5 fs FWHM pulses are visually nearly identical to the impulsive limit, which we also find. Finite pulse effects are, unsurprisingly, unimportant when the pulse durations are shorter than T_o, the fastest timescale in the system. We use this model system to illustrate the differences that

occur with only modestly longer pulses.

We consider Gaussian pulses with FWHM of 15 fs, similar to T_o. Figure 7.3(b) shows calculations using finite pulses but only the 7 time-ordered diagrams, showing clear differences from the impulsive limit in 7.3(a). Figure 7.3(c) shows results using the same finite pulses but including all of the 233 additional pulse overlap diagrams at delay times when they are required. The visual difference between Fig. 7.3(b) and (c) demonstrates the importance of including pulse-overlap diagrams in addition to finite pulse effects in time-ordered diagrams. Simply adding finite pulse effects to the time-ordered diagrams is not sufficient for making good spectroscopic predictions.

Figure 7.3(d) shows the magnitude of the cross peak contained in the solid box in panels (a)-(c), and the oscillation period T_o is clearly visible. The oscillations at that peak correspond to absorption into one and emission from another doubly excited eigenstate of H_0. This peak is clearly visible in the impulsive limit in Fig. 7.3(a), but is dominated by significant horizontal streaks in Fig. 7.3(b), which bleed over from the stronger peak at $\hbar\omega_\tau = 2.3$ eV. This significant extension of the peaks in the ω_τ direction is an artifact of neglecting the overlap diagrams, demonstrated by its removal in Fig. 7.3(c). Figure 7.3(d) shows that both the visibility of the oscillation and the overall evolution of the envelope are qualitatively different in the three studied approximations. We vary T from 0 to 2 ps, for a total of 2667 different values of T. At long T (shown in the inset), the doubly excited states decay and there is a static excited state absorption signal from the singly excited population, with the timescale of saturation and oscillation decay matching the slowest k_{nm}. Neither impulsive nor time-ordered calculations with finite pulses accurately predict this EEI2D spectrum despite using an ultrafast optical pulse.

Note that the differences between the full calculation and the one using only the time-ordered diagrams persist even with delay times T much greater than the pulse durations. In constructing $\tilde{\mathbf{P}}^{(5)}_{\mathrm{k}_d}(\omega_\tau, T, \omega_t)$, contributions with τ, t smaller than the pulse duration are always included, making pulse-overlap effects apparent even at long T and requiring calculation of the overlap diagrams. Since the DG only produces the extra diagrams for time delays that merit their evaluation, the full calculation is not much more expensive than the case with 7 time-ordered diagrams. For example, the calculation of the green and orange curves in Fig. 7.3(d), out to 350 fs, took 33 and 17 minutes on a 2017 MacBook Pro, respectively, and included over 120,000 combinations of τ and T. The equivalent curves out to 2 ps, as in the inset, required 155 and 93 minutes with the same density of T points; those involved over 700,000 pulse delay combinations. Even though the full calculation includes contributions from 34 times more diagrams, it required less than twice the time to run than the time-ordered calculation.

The DG determines which diagrams contribute based solely upon the user-chosen intervals where the pulses are nonzero. For well-behaved pulses such as Gaussians with standard deviation σ, convergence of spectra within 1% is obtained for the results in Fig. 7.3 when the pulses are declared to be zero after 4σ. However, for pulses with heavy tails, extra care must be taken to determine the correct pulse interval. We compare Gaussian pulses, as used in Fig. 7.3, to Lorentzian pulses, which have heavy tails. Figure 7.4 shows the weight $D \int d\omega_t \left|\tilde{\mathbf{P}}^{(5)}_d(\tau, T, \omega_t)\right|$ for each of the 240 diagrams d as a function of τ at $T = 100\,\mathrm{fs}$ for the system in Fig. 7.3, where $\tilde{\mathbf{P}}^{(5)}_d$ is the signal due to a

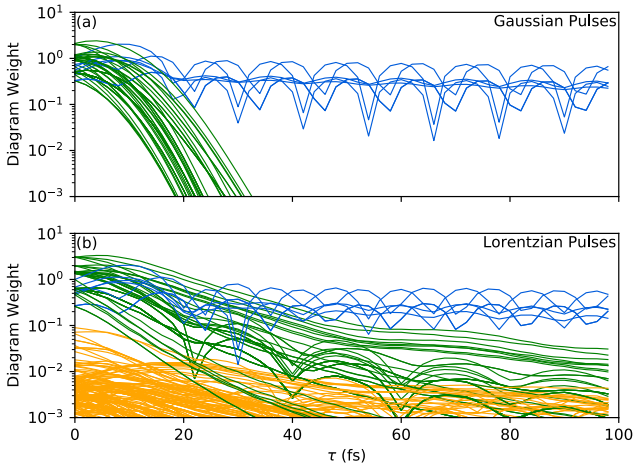

Figure 7.4: Weight of all 240 EEI2D diagrams for the same system as in Fig. 7.3 at $T = 100$ fs as a function of τ for 15 fs FWHM pulses with (a) Gaussian and (b) Lorentzian time-domain envelopes. Blue lines show the weight of the 7 time-ordered diagrams. Green lines show the weight of the 47 additional diagrams that contribute when pulses a and b overlap. Panel (b) also shows the weights of 176 diagrams that contribute only when all three pulses overlap (gold), whose weights are too small to appear in (a); 10 more orange lines have weights less than 10^{-3}. Diagram weights are normalized so that the strongest time-ordered diagram has a weight of 1 at $\tau = 0$. Line colors for each set of diagrams match those in Figure 7.2.

single diagram and D is chosen so that at $\tau = 0$, the largest weight of a time-ordered diagram is 1. We consider Gaussian and Lorentzian pulse envelopes with the same FWHM of 15 fs.

We observe that for both sets of pulses, at $\tau = 0$ all 47 a, b overlap diagrams (green) have similar weight and therefore must be calculated, and the same is true of all 233 overlap diagrams at $\tau = T = 0$ (not shown). With Gaussian pulses, all of the overlap diagrams decay rapidly when τ exceeds the nominal pulse duration and are negligible when $\tau > 30$ fs. In contrast, with Lorentzian pulses of the same nominal duration, even some diagrams where pulse c arrives before pulse a or b (shown in gold) are surprisingly close in weight to the time-ordered diagrams at $\tau = 0$, even though the center of pulse c arrives 100 fs after the first two. Further, some of the a,b overlap diagrams (green) continue to have weight equal to time-ordered diagrams (blue) for τ approaching 60 fs, four times the pulse FWHM; given the large number of overlap diagrams, their contributions must be calculated even for considerably larger τ. It is clear that in the case of Lorentzian or other heavy-tailed pulse envelopes, care must be taken in choosing the pulse intervals for use with the DG, and users are encouraged to ensure that their results are converged. Choosing a large pulse interval ensures accuracy but increases computational cost as more overlap diagram contributions must be calculated.

Many ultrashort optical pulses have heavy tails in the time domain, and Ref. 57 showed experimental evidence that pulse overlap effects can extend for up to 100 fs when using 17 fs FWHM pulses in 2DPE experiments. They attributed this effect to the significant wings on the pulse envelopes, which is consistent with our analysis of Lorentzian pulses.

7.5 Conclusion

We have presented the diagram generator (DG), a tool for automatically generating the Feynman diagrams that contribute to perturbative nonlinear optical spectroscopies. The DG automatically determines when pulses overlap and only generates extra overlap diagrams when required. This automated process allows users to get the full advantage of including all causally allowed diagrams with a low computational cost.

We have shown that including these overlap diagrams can be important to correctly predicting or interpreting spectra when pulses are not in the impulsive limit. Using EEI2D as an example, we have shown significant errors when using only the time-ordered diagrams with pulses whose duration is similar to the dynamics of the system. We have also shown that overlap diagrams can make significant contributions to the signal for delay times that are large relative to the pulse durations, when those pulses have heavy temporal tails, as in the case of Lorentzian pulses.

The DG is one module of the larger Ultrafast Spectroscopy Suite. The other components in UFSS are described in Refs. [96, 131]. Taken together, UFSS is a tool for automatically calculating arbitrary-order spectroscopic signals while accounting for effects of finite pulse shapes, which can be of arbitrary form. The diagrams produced by the DG can also be used in other analytical or numerical tools. The DG can be used for determining all of the diagrams that contribute to any order spectra, whether in the impulsive limit or with finite pulses. The DG may open the door to more easily calculating higher-order corrections to commonly used 3$^{\text{rd}}$-order spectroscopies and may lend itself to developing intuition for higher-order spectroscopic techniques.

Acknowledgments

We acknowledge support from the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Ontario Trillium Scholarship.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon reasonable request. In addition, the code to generate all of the figures in this manuscript is available at github.

Chapter 8

Composite Diagrams

We now demonstrate how we can increase the computational efficiency of calculations done using UF^2 and RKE, by automatically combining diagrams into what we call composite diagrams. This composite diagram method can be used for any phase-discrimination condition and any pulse sequence. The composite diagram method is a general tool that replaces the diagram generator. The composite diagram method can speed up most perturbative calculations, and makes it easy for a user to push beyond the perturbative limit, and quickly simulate non-perturbative effects. We begin by taking equations of motion of the form

$$\dot{\rho} = \mathscr{L}_0\rho + \mathscr{L}'(t)\rho,$$

where $\mathscr{L}'(t) = \sum_j \sum_O \mu^O(\varepsilon_j + \varepsilon_j^*)$, and $\mu^K|\rho\rangle\!\rangle \equiv \mu\rho$ and $\mu^B|\rho\rangle\!\rangle \equiv \rho\mu$, as first introduced in Eq. 6.7. We will begin with the assumption that the system Liouvillian, $\mathscr{L}_0$, cannot be broken into optically separable manifolds, and consider the separable case later. Manifolds can be separated by optical excitation if we ignore all physical processes that cause the system to relax from the singly excited states to the ground state, from the doubly excited state to the singly excited states, and so on. The inseparable case occurs when we do include these relaxation processes. Given an initial condition $\rho^{(0)}$, we have shown in chapter 6 how we can efficiently compute a single Feynman diagram using the operators $O_{j^{(*)}}$ (defined in eq. 6.21). We have further demonstrated in chapter 7 that we can automatically generate all relevant diagrams that contribute to a given phase-discrimination condition. With these two ingredients we can automate the process of calculating any desired spectroscopic signal. We show in chapter 7 that for higher-order spectroscopies such as EEI2D, that the number of Feynman diagrams required to properly account for pulse overlap effects is much larger than for a lower-order case such as 2DES. We further show that including the pulse-overlap diagrams does not greatly increase the calculation time, as these diagrams are only calculated when pulses are overlapping, which is generally not a large portion of the pulse delays considered in a spectroscopy. However, as one considers even higher-order calculations, the number of time-ordered and pulse-

overlap diagrams grows geometrically. In high-order cases, simply enumerating the required number of diagrams can become prohibitively expensive. In this chapter, we demonstrate how we can geometrically reduce the cost of higher-order spectroscopic calculations by combining many diagrams into composite diagrams. We then give a practical demonstration of the power of this technique by simulating Rabi oscillations, generally a non-perturbative effect, by going to sufficient order in perturbation theory.

8.1 Inseparable manifolds

Each Feynman diagram can be expressed symbolically using the $O_{j^{(*)}}$ super-operators. For example, the three time-ordered diagrams that contribute to the rephasing 2DPE signal (see Fig. 7.1) are

$$\begin{aligned}
|\rho^{(3)}_{ESA}\rangle\!\rangle &= K_c K_b B_{a^*} |\rho^{(0)}\rangle\!\rangle \\
|\rho^{(3)}_{SE}\rangle\!\rangle &= B_c K_b B_{a^*} |\rho^{(0)}\rangle\!\rangle \\
|\rho^{(3)}_{GSB}\rangle\!\rangle &= K_c B_b B_{a^*} |\rho^{(0)}\rangle\!\rangle.
\end{aligned}$$

The total time-ordered third-order density matrix that is needed for calculating the rephasing signal is

$$|\rho^{(3)}\rangle\!\rangle = K_c K_b B_{a^*} |\rho^{(0)}\rangle\!\rangle + B_c K_b B_{a^*} |\rho^{(0)}\rangle\!\rangle + K_c B_b B_{a^*} |\rho^{(0)}\rangle\!\rangle. \tag{8.1}$$

Each of the diagrams is typically interpreted as a separate calculation. However, we can factor out K_c from the ESA and GSB contributions:

$$|\rho^{(3)}\rangle\!\rangle = K_c \left(K_b B_{a^*} |\rho^{(0)}\rangle\!\rangle + B_b B_{a^*} |\rho^{(0)}\rangle\!\rangle \right) + B_c K_b B_{a^*} |\rho^{(0)}\rangle\!\rangle, \tag{8.2}$$

Eq. 8.2 is composed of one standard diagrammatic calculation, and one composite diagram calculation. Eq. 8.2 represents a more efficient computational strategy as compared to eq. 8.1. To understand the true savings we need to consider the number of unique calculations that must be done to calculate the signal for a large number of coherence times τ (time delay between pulses a and b) and population times T (time delays between pulses b and c). The arrival time of pulse a is never changed (it is considered to arrive at $t = 0$), and thus the object $B_{a^*}|\rho^{(0)}\rangle\!\rangle$ is calculated once and stored for re-use. The objects $K_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ and $B_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ are recalculated for each unique τ. Finally the three objects $K_c K_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$, $B_c K_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$, $K_c B_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ are recalculated for each value of T (and for each value of τ). Amortized over all values of τ and T, the cost of calculating $B_{a^*}|\rho^{(0)}\rangle\!\rangle$, as well as recalculating $K_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ and $B_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$, is negligible. The most expensive step for any spectroscopic calculation is usually varying the delay between the last two pulses. Therefore Eq. 8.2 actually represents a time savings of about 33% compared with Eq. 8.1. At this point it is worth pointing out the factorization in Eq. 8.2 cannot be performed if

the optical manifolds are separable. However, composite diagrams can still be useful in many cases (see Section 8.2 for more about this).

The savings is more dramatic when one considers the case where all three pulses are overlapping, and 16 diagrams must be calculated. In this case, the composite calculation looks like

$$\begin{aligned}|\rho^{(3)}\rangle\!\rangle = K_c\left(K_bB_{a^*} + B_bB_{a^*} + B_{a^*}K_b + K_{a^*}K_b\right)|\rho^{(0)}\rangle\!\rangle + B_c\left(K_bB_{a^*} + B_{a^*}K_b\right)|\rho^{(0)}\rangle\!\rangle \\ + K_b\left(K_cB_{a^*} + B_cB_{a^*} + B_{a^*}K_c + K_{a^*}K_c\right)|\rho^{(0)}\rangle\!\rangle + B_b\left(K_cB_{a^*} + B_{a^*}K_c\right)|\rho^{(0)}\rangle\!\rangle \\ + B_{a^*}\left(K_cK_b + K_bK_c\right)|\rho^{(0)}\rangle\!\rangle + K_{a^*}\left(K_cK_b + K_bK_c\right)|\rho^{(0)}\rangle\!\rangle,\end{aligned}$$

which only consists of 6 composite diagrams rather than 16 standard diagrams. In this case, fewer calculations can be re-used as τ and T vary, but 6/16 is still 37.5%.

In order to generalize this approach, we define

$$\begin{aligned}\rho_{a^*} &= B_{a^*}|\rho^{(0)}\rangle\!\rangle + K_{a^*}|\rho^{(0)}\rangle\!\rangle, \\ \rho_b &= B_b|\rho^{(0)}\rangle\!\rangle + K_b|\rho^{(0)}\rangle\!\rangle, \\ \rho_c &= B_c|\rho^{(0)}\rangle\!\rangle + K_c|\rho^{(0)}\rangle\!\rangle,\end{aligned}$$

where $K_{i^*}|\rho^{(0)}\rangle\!\rangle$ and $B_i|\rho^{(0)}\rangle\!\rangle$ are both zero in the RWA. We then define

$$\begin{aligned}\rho_{a^*b} &= B_{a^*}\rho_b + K_{a^*}\rho_b + B_b\rho_{a^*} + K_b\rho_{a^*} \\ \rho_{a^*c} &= B_{a^*}\rho_c + K_{a^*}\rho_c + B_c\rho_{a^*} + K_c\rho_{a^*} \\ \rho_{bc} &= B_b\rho_c + K_b\rho_c + B_c\rho_b + K_c\rho_b\end{aligned}$$

and finally

$$|\rho^{(3)}\rangle\!\rangle = \rho_{a^*bc} = B_c\rho_{a^*b} + K_c\rho_{a^*b} + B_b\rho_{a^*c} + K_b\rho_{a^*c} + B_{a^*}\rho_{bc} + K_{a^*}\rho_{bc}, \tag{8.3}$$

where ρ_{a^*bc} is the entire third-order density matrix that obeys the phase-discrimination condition associated with the rephasing 2DPE signal, without assuming the RWA or time-ordering (note that if calculations outside the RWA are desired, calculations must be done using short enough time-resolution to resolve the carrier frequencies of the pulses). The bulk polarization is then calculated as $P_s(t) = \langle\!\langle\mu|\rho^{(3)}(t)\rangle\!\rangle$, from which the signal can be calculated using Eq. 3.7. In section 3.3, we derived that without these two approximations, there are 48 diagrams, with each diagram consisting of 1-3 calls to $O_{j^{(*)}}$ (3 if nothing can be reused; <3 if some calculations have been stored for reuse). In contrast, the cost of Eq. 8.3 consists of 6 calls to $O_{j^{(*)}}$ to construct $|\rho^{(1)}\rangle\!\rangle$, 12 calls to $O_{j^{(*)}}$ to construct $|\rho^{(2)}\rangle\!\rangle$, and 6 calls to $O_{j^{(*)}}$ to construct $|\rho^{(3)}\rangle\!\rangle$, with a total cost of 24 calls to $O_{j^{(*)}}$ if nothing can be reused. This is only 17% of the $48 \times 3 = 144$ calls to $O_{j^{(*)}}$ required by the diagrammatic approach. Thus the composite technique

is 6 times faster than the diagrammatic approach.

The fully general expression for $|\rho^{(n)}\rangle\!\rangle$ can be expressed using the recursive formula

$$\rho_{i^{(*)}} = B_{i^{(*)}}|\rho^{(0)}\rangle\!\rangle + K_{i^{(*)}}|\rho^{(0)}\rangle\!\rangle$$

$$|\rho^{(n)}\rangle\!\rangle = \rho_{i^{(*)}j^{(*)}\ldots m^{(*)}} = B_{i^{(*)}}\rho_{j^{(*)}\ldots m^{(*)}} + K_{i^{(*)}}\rho_{j^{(*)}\ldots m^{(*)}} + B_{j^{(*)}}\rho_{i^{(*)}\ldots m^{(*)}} + K_{j^{(*)}}\rho_{i^{(*)}\ldots m^{(*)}} + \ldots + B_{m^{(*)}}\rho_{i^{(*)}j^{(*)}\ldots} + K_{m^{(*)}}\rho_{i^{(*)}j^{(}$$

with the order n of $\rho_{i^{(*)}j^{(*)}\ldots n^{(*)}}$ being given by counting the number of subscripts. This recursive formula will give the entire $\rho^{(n)}$ that contributes to the given phase-discrimination condition. Thus, from this procedure, we may calculate the desired quantity (for example $P^{(n)}(t) = \mathrm{Tr}\left[\mu\rho^{(n)}\right]$). This general expression does not automatically take advantage of the RWA and time-ordering, when they apply. In the case of inseparable manifolds, we do not take advantage of the RWA. When manifolds are inseparable, all of the terms $K_i|\rho^{(n)}\rangle\!\rangle, K_{i^*}|\rho^{(n)}\rangle\!\rangle, B_i|\rho^{(n)}\rangle\!\rangle, B_{i^*}|\rho^{(n)}\rangle\!\rangle$ are non-zero in the composite approach (typically some of those terms would be zero for particular diagrams, and could be skipped). Thus we can, surprisingly, do calculations outside of the RWA without doing any additional perturbative calculations (though the calculations themselves will be more costly, since a smaller time mesh will be required in order to resolve the optical carrier frequency). We do skip improperly-ordered calculations when they do not apply, by checking to see whether pulses overlap, and in what order they arrive. We thus only calculate overlap diagrams when they are necessary.

Although I have described this process as a recursive procedure, the actual implementation is constructive. We build up from $\rho^{(0)}$ by using what we call partial phase-discrimination conditions (pPDCs). The full phase-discrimination condition (PDC), $[(n_r^a, n_c^a), (n_r^b, n_c^b), ..., (n_r^L, n_c^L)]$, as described in Section 7.3 can be expressed as a $L \times 2$ matrix, where L is the number of pulses. For example EEI2D spectroscopy has a PDC of $[(0,2),(2,0),(1,0)]$, which is recast in matrix form as

$$\begin{bmatrix} 0 & 2 \\ 2 & 0 \\ 1 & 0 \end{bmatrix}.$$

The first entry in each row is the number of interactions with the rotating term, n_r, and the second entry in the row is the number of interactions with the counter-rotating term, n_c. The procedure is as follows:

1. Given a phase-discrimination condition (PDC), produce a list of all $O_{j^{(*)}}$ operators that will contribute to the calculation. For each row j in the matrix-form PDC, add K_j and B_j to the list if there is a non-zero entry in the first column, and add K_{j^*} and B_{j^*} to the list if there is a non-zero entry in the second column. Using EEI2D as an example, the list of relevant operators is $[K_{a^*}, B_{a^*}, K_b, B_b, K_c, B_c]$.

2. Begin with $\rho^{(0)}$, which has a partial phase-discrimination condition (pPDC) of the zero matrix (of size $L \times 2$):

$$\begin{bmatrix} 0 & 0 \\ \vdots & \vdots \\ 0 & 0 \end{bmatrix}.$$

Call it the current density matrix.

3. Until the algorithm is complete, repeat the remaining steps. For each ρ in the set of current density matrices (in the first round, there is only $\rho^{(0)}$), check whether the calculation $O_{j^{(*)}}|\rho\rangle\!\rangle$ for each $O_{j^{(*)}}$ from the list constructed in step 1 can produce a valid calculation by constructing its pPDC. If it can, perform the calculation.

   (a) Produce a new pPDC for each potential calculation by taking the pPDC of ρ, and adding the effect of either O_j, which adds 1 to the first entry of the j^{th} row, or O_{j^*}, which adds 1 to the second entry of the j^{th} row. Then subtract the new pPDC from the PDC, to produce a remainder PDC (rPDC).

   (b) If the rPDC contains any negative entries, do not do the calculation (this case corresponds to a calculation that involves too many interactions with either the rotating or counter-rotating term of a given pulse). Using EEI2D as an example, imagine that the incoming pPDC was $[(0,2),(0,0),(0,0)]$, and that the trial calculation is K_{a^*}. This would lead to a new pPDC of $[(0,3),(0,0),(0,0)]$. The rPDC is $[(0,-1),(2,0),(1,0)]$. Since there is a negative number in the rPDC, we do not proceed with the calculation.

   (c) If the pulse sequence forbids a contribution, do not do the calculation. The trial calculation involves an interaction with the j^{th} pulse, which is non-zero on the domain $[t_{j,min}, t_{j,max}]$. The trial calculation would therefore be non-zero on the domain $t \geq t_{j,min}$. We next make a list of all relevant pulse endpoints. For each row j in the rPDC that contains a non-zero entry, we add to the list the value $t_{k,\max}$. From this list, we take the minimum value of $t_{k,\max}$, and compare it to $t_{j,\min}$. If $\min_k t_{k,\max} < t_{j,\min}$, we do not perform the calculation, since it would have zero contribution to the final signal.. Using EEI2D as an example, imagine that the incoming pPDC was $[(0,1),(0,0),(0,0)]$, and that the trial calculation is K_b, which results in a new pPDC of $[(0,1),(1,0),(0,0)]$. The rPDC is then $[(0,1),(1,0),(1,0)]$. The trial calculation involves an interaction with pulse b, so we know that the trial calculation would be nonzero for $t > t_{b,\min}$, thus in this case $t_{j,\min} = t_{b,\min}$. The rPDC involves interactions with all pulses. Let us assume that none of the pulses overlap, and that all of the pulses are ordered (so in the case of EEI2D $t_{a,\max} < t_{b,\min}$ and $t_{b,\max} < t_{c,\min}$). In this case, we have $\min_k t_{k,\max} = t_{a,\max}$. But since $t_{a,\max} < t_{b,\min}$, we do not proceed with the calculation.

4. Given all of the new density matrices $\left\{O_{j^{(*)}}|\rho^{(0)}\rangle\!\rangle\right\}$ with their associated pPDC's that pass the checks in step 3, add together all density matrices with the same pPDC, i.e., $\rho_{a^*} = K_{a^*}|\rho^{(0)}\rangle\!\rangle + B_{a^*}|\rho^{(0)}\rangle\!\rangle, \rho_b = K_b|\rho^{(0)}\rangle\!\rangle +$

$B_b|\rho^{(0)}\rangle\!\rangle, \ldots$. Store the resulting density matrices as the list of current density matrices. Stop when there is a single current density matrix with pPDC equal to the PDC.

See Sec. 8.3.3 below for a comparison of the cost of calculating each diagram individually vs. doing composite calculations. The composite diagram method gives an exponential speed-up for higher-order calculations.

8.2 Separable Manifolds

In the case that manifolds are separable, the above procedure can be used as already introduced, if one is willing to pay the cost of representing the density matrix in all manifolds at once. (Note: in such a case, one would still take advantage of the block diagonal nature of the system Liouvillian when preforming the diagonalization.) In Sec. 6.7 we present a computational complexity analysis of UF^2, showing how the cost of calculations scale with the dimension of the Hamiltonian, N (results summarized in Tables 6.1 and 6.2 for inseparable and separable manifolds, respectively). The cost scales as either N^3 for secular Redfield or N^4 for full Redfield and diabatic Lindblad. By separating manifolds, one gains the advantage of expressing the density matrix in manifolds of size N_X^2 and optical coherences of size $N_X N_Y$, where X, Y index the optical manifolds. The benefits of this procedure vary based upon the model system. For example, for a trimer of 2LSs without vibrations ($s = 3$ and $k = 0$ for the model in Eq. 2.4), the total size including all manifolds is $N = 8$, while the individual manifold sizes are $N_{GSM} = 1$, $N_{SEM} = 3$, $N_{DEM} = 3$ and $N_{TEM} = 1$ (where GSM is the ground-state manifold, SEM is the singly-excited manifold, DEM is the doubly-excited manifold, and TEM is the triply-excited manifold). Thus the relative cost of doing calculations with UF^2 using separated manifolds, relative to using unseparated manifolds, is $(3/8)^3$ or $(3/8)^4$ (note that the ratios are preserved for any number of vibrational modes k).

Given the reduced cost of calculations, in many cases it is desirable to take advantage of the separation of manifolds by representing the density matrix only in a single manifold at a time. This can also save memory (in addition to computation time). Implementing composite diagrams in this case requires some care, and results in more individual calculations than in the general case, though still reduced compared to the individual-diagram method. Thus for small systems (small N), it may actually be advantageous to ignore the separability and proceed using the above scheme. However, for sufficiently large N, it is either necessary, or faster, to separate the manifolds. At the time of writing, composite diagrams for separable manifolds has not been implemented in UFSS. But it is well within reach, and thus deserves a brief proposal.

In the case of separable manifolds, we can no longer take advantage of the factorization shown in Eq. 8.2, because $K_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ exists in the singly excited manifold, while $B_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ exists in the ground-state manifold, and thus the calculations $K_c K_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ and $K_c B_b B_{a^*}|\rho^{(0)}\rangle\!\rangle$ must be done separately. In the case of the time-ordered rephasing calculations, no speed-up is possible using composite diagrams. However, when including overlap diagrams, composite diagrams may still be used to speed up the calculations. For example, we may combine

$$P_1\rho_{a^*b}P_1 = K_bB_{a^*}\rho^{(0)} + B_{a^*}K_b\rho^{(0)}$$

and

$$P_0\rho_{a^*b}P_0 = B_bB_{a^*}\rho^{(0)} + K_{a^*}K_b\rho^{(0)},$$

where $P_n = |n\rangle\langle n|$ are projection operators on the electronic excitation space.

Looking to higher-order examples, we return to the case of EEI2D (phase-matching condition $-2k_a + 2k_b + k_c$). As an incomplete example that illustrates the case where the density matrix is in an optical coherence between manifolds, we can combine

$$P_0\rho_{a^*a^*b}P_1 = K_{a^*}K_bB_{a^*}\rho^{(0)} + K_{a^*}B_{a^*}K_b\rho^{(0)} + B_{a^*}K_{a^*}K_b\rho^{(0)}$$
$$P_1\rho_{a^*a^*b}P_2 = B_{a^*}K_bB_{a^*}\rho^{(0)} + B_{a^*}B_{a^*}K_b\rho^{(0)} + K_bB_{a^*}B_{a^*}\rho^{(0)}.$$

To implement the composite approach in this case, we propose following the same procedure outlined for inseparable manifolds, but only combining calculations that result in density matrices in the same manifold. Each density matrix must then be stored using both the partial phase-matching condition and the number of electronic excitations accrued by both the bra and ket sides. As outlined in Section 7.3, we track the number of optical excitations to the bra and ket sides separately, and so it should be straightforward to collect calculations by number of excitations. In the case of separable manifolds, we can once again take advantage of the RWA and knowledge of the maximum number of optical excitations/de-excitations. This can be easily done by rejecting calculations that reduce the ket or bra excitation index below 0 or above the maximum relevant to the system under study.

8.3 Rabi Oscillations

To give a concrete example of the power of the composite diagrams, we demonstrate calculations of Rabi oscillations in pump-probe spectra in a two-level system and a three-level system, using TDPT calculations up to 27^{th} order. We demonstrate that simply enumerating all of the diagrams is prohibitively expensive (much less actually calculating them). However, using composite diagrams the whole calculation can be done in a matter of seconds.

8.3.1 Two-level system

We begin with an electronicHamiltonian of the form

$$H = \hbar\omega_0 \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix},$$

with no explicit vibrations and include a phenomenological dephasing time of 30 fs, along with a phenomenological relaxation rate of 30 ps, using the Lindblad formalism (see Section 6.3.1 for details of Lindblad implementation). We consider pump-probe spectroscopy with the pump pulse indexed as a and the probe pulse indexed as b:

$$\boldsymbol{E}(t) = A_a(t)e^{-i(\omega_a t - \boldsymbol{k}_a\cdot\boldsymbol{r}+\phi_a)} + A_b(t)e^{-i(\omega_b t - \boldsymbol{k}_b\cdot\boldsymbol{r}+\phi_b)} + c.c.$$

with $\omega_a = \omega_b = \omega_0$. We use Gaussian pulses with an envelope function

$$A_j(t) = \frac{\eta_j}{\sqrt{2\pi}\sigma}e^{-t^2/2\sigma^2}$$

centered at frequency $\omega = \omega_0$, with σ the pulse duration and η_j its amplitude. We calculate the frequency-integrated pump-probe spectra at a delay time of 500 fs up to 27^{th} order in perturbation theory, by calculating each order separately. This is achieved by calculating the signal that results from the phase-discrimination condition $[(m,m),(1,0)]$, which is the $2m+1^{\text{th}}$ order contribution to the pump-probe signal. We write the pump-probe signal as a partial power series in the intensity of the pump pulse

$$S_{PP}(I,T) \propto I_b \sum_{m=1}^{m_{\max}} s^{(2m+1)}(T)\left(\frac{I_a}{I_0}\right)^m,$$

where $I_j \propto \eta_j^2$. Note that we are doing a partial summation up to 26^{th}-order in the pump amplitude η_a, and 1^{st}-order in the probe intensity. Thus we are modeling a situation where a non-perturbative pump pulse is used, but we assume that a much weaker probe pulse is used, so that we are only sensitive to the lowest order contributions from the probe pulse. The great advantage of this technique, as compared to non-perturbative calculations, is that once all of the individual signals $s^{(2m+1)}(T)$ are calculated, we can vary I_a with essentially no computatioanl cost, though of course there is some I_{max} determined by the maximum m we have included above which the calculation is not accurate. At any given I_a, we must check for convergence with respect to $m_{\max}$. We show the partial series with $m_{\max}$ varying from 10 to 13 in Figure 8.1b. Visually, we see that the results appear well-converged up to around $I = 16I_0$.

We study three different full-width half maxima (FWHM $= 2\sqrt{2\log 2}\sigma$): 0.1 fs, 10 fs, and 30 fs. In the impulsive limit of FWHM equal to 0.1 fs (in this case the only relevant system time-scale is the dephasing rate of 30 fs), we find the expected behavior where the pump-probe signal strength varies as $\sin^2\sqrt{I_a/I_0}$. As the pulse duration increases, we see that the Rabi oscillations damp (see Fig. 8.1).

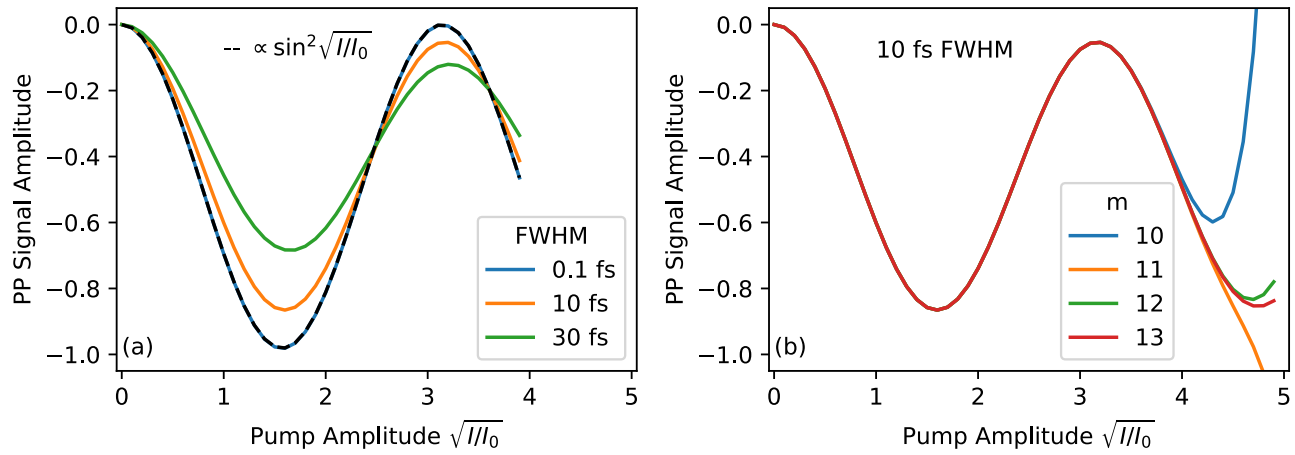

Figure 8.1: (a) Rabi oscillations for varying pump peak electric field amplitude up to $m = 13$. (b) Partial summation for the pump-probe signal up to order m as a function of pump pulse amplitude with a pump FWHM of 10 fs. The probe pulse is assumed to be in the perturbative regime, and thus we include only the lowest order effects of the probe pulse. Panel b demonstrate visually that we have converged the calculation with respect to TDPT, up to a pump intensity of about $I = 16I_0$, using Gaussian pump and probe pulses each with FWHM of 10 fs.

8.3.2 3LS

Next we consider a 3LS with Hamiltonian of the form

$$H = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1.904 \end{pmatrix},$$

a phenomenological dephasing time of 30 fs, and relaxation rates of 30 fs from $|f\rangle \rightarrow |e\rangle$ and 30 ps from $|e\rangle \rightarrow |g\rangle$. We consider a dipole operator

$$\mu = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0.33 \\ 0 & 0.33 & 0 \end{pmatrix},$$

and the same pulses as used for the 2LS. Figure 8.2 shows our results for a 3LS.

8.3.3 Comparison to calculations with individual Feynman diagrams

These calculations would not have been possible without using the composite diagrams outlined in this chapter. The reason is that the number of Feynman diagrams grows at least exponentially with the perturbative order. To understand this, recall from Section 3.3 that the number of response functions that contribute to a given order spectroscopy grows as 2^n, where in this case $n = 2m + 1$. Further, when doing calculations with finite pulses, you must include all unique permutations of the pulse envelopes. For pump-probe spectroscopy, each interaction with

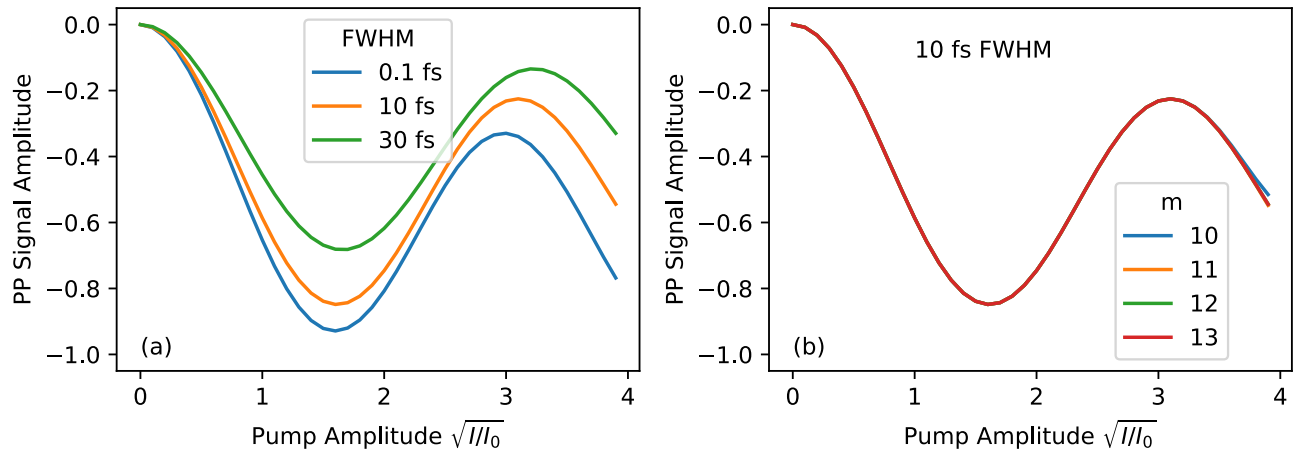

Figure 8.2: (a) Rabi oscillations for varying pump peak electric field amplitude for a 3LS, showing a slightly more complicated dependence on pulse amplitude. (b) Partial summation for the pump-probe signal as a function of pump pulse amplitude. Again the probe pulse is assumed to be in the perturbative regime, and thus we include only the lowest order effects of the probe pulse. Panel b demonstrates visually that we have converged the calculation with respect to TDPT, up to a pump intensity of about $I = 16I_0$, using Gaussian pump and probe pulses each with FWHM of 10 fs.

the rotating term (ε_a) and the counter-rotating term (ε_a^*) are distinguishable. For the $2m+1^{\text{th}}$ order contribution, we have the phase discrimination condition $[(m,m),(1,0)]$, which involves m interactions with ε_a and m interactions with ε_a^*. There is also one interaction with ε_b. Thus when the pump and probe pulses overlap, we have an upper bound on the number of diagrams, r, of

$$r_o \leq 2^{2m+1} \frac{(2m+1)!}{(m!)^2} \tag{8.4}$$

different diagrams, and when the pump and probe pulses do not overlap, we have

$$r_{to} \leq 2^{2m+1} \frac{(2m)!}{(m!)^2}. \tag{8.5}$$

Using WolframAlpha, we find that $\frac{(2m+1)!}{(m!)^2} = 2^{2m+1} \frac{\Gamma(m+\frac{3}{2})}{\sqrt{\pi}\Gamma(m+1)}$. Since $\frac{\Gamma(m+\frac{3}{2})}{\sqrt{\pi}\Gamma(m+1)} > 1$ for large m, we can take 2^{4m+2} as an underestimate for the upper bound on r_o, which gives some sense of how quickly the number of diagrams increases with m.

We have equality in Eqs. 8.4 and 8.5 when the RWA approximation does not hold. When the RWA holds, and the entire electronic population of $\rho^{(0)}$ is in the ground-state, we can eliminate diagrams that involve de-excitation below the ground-state (generally at least half of the diagrams may be eliminated, since half of the diagrams will begin with a de-excitation of either the ket- or bra-sides of the density matrix). We do not have a formula to predict how many diagrams may be removed, so we use the Diagram Generator to enumerate all of the time-ordered diagrams up to $m = 6$, and extrapolate how many diagrams would be needed to reach $m = 13$. We also demonstrate how long it takes to run the Diagram Generator just to enumerate the diagrams and find that it would likely take well

over a year to enumerate all of the diagrams on a single computer. We find numerically that in the case of a 3LS, if manifolds were separable (and therefore diagrams involving more than 2 consecutive excitations may be ignored), that the number of diagrams seems to scale as $\propto 2^{3m+1}$ (see Fig 8.3d). Assuming the Diagram Generator's runtime scales proportional to the number of diagrams, we find that the Diagram Generator would take about 2.4×10^7 seconds to run. However, it appears that the Diagram Generator's runtime scaling is worse than linear with the number of diagrams, and thus the actual runtime would likely be much greater than 1 year (see Fig. 8.3c).

In contrast, we plot the effective number of Feynman diagrams required when using composite diagrams for a time-ordered calculation in Fig. 8.3b, which demonstrates that the effective number of diagrams grows linearly in the composite diagram approach. To calculate the effective number of diagrams, we keep track of the number of times that the $O_{j^{(*)}}$ operator is called, and divide by $2m+1$, the perturbative order of the calculation, since a single Feynman diagram involves $2m+1$ interactions. Thus, the total cost of the calculation should scale quadratically in m (number of diagrams multiplied by the order). This agrees well with Fig. 8.3a, which shows runtime as a function of m for calculating the $2m+1^{\text{th}}$ contribution to the pump-probe signal at a single delay time.

Figure 8.3 demonstrates that the composite diagram approach is an example of a fast algorithm, where we have reduced the computational complexity of the problem from an exponential scaling to a quadratic scaling in m. At this point it is worth pointing out that we have not investigated whether some diagrams at each order m could be avoided, because they are in some sense redundant with other diagrams. It is possible that some of the 2^{3m+1} diagrams would not have to be calculated, with clever identification of identical/related diagrams. Looking at $\tau = 0$ in Fig. 7.4, we can see that for the EEI2D signal at $\tau = 0$ (related to the $m = 2$ contribution to pump-probe spectra), that many diagrams have the same absolute magnitude, suggesting that they may (or may not) be related at $\tau = 0$. However, the figure shows that as one moves away from $\tau = 0$, each diagram appears to become unique. The composite diagram approach outlined here is applicable to any phase-discrimination condition and any pulse sequence, and is thus a useful and general tool for speeding up any calculation. The composite diagram method can speed up most perturbative calculations, and makes it easy for a user to push beyond the perturbative limit, and quickly simulate non-perturbative effects.

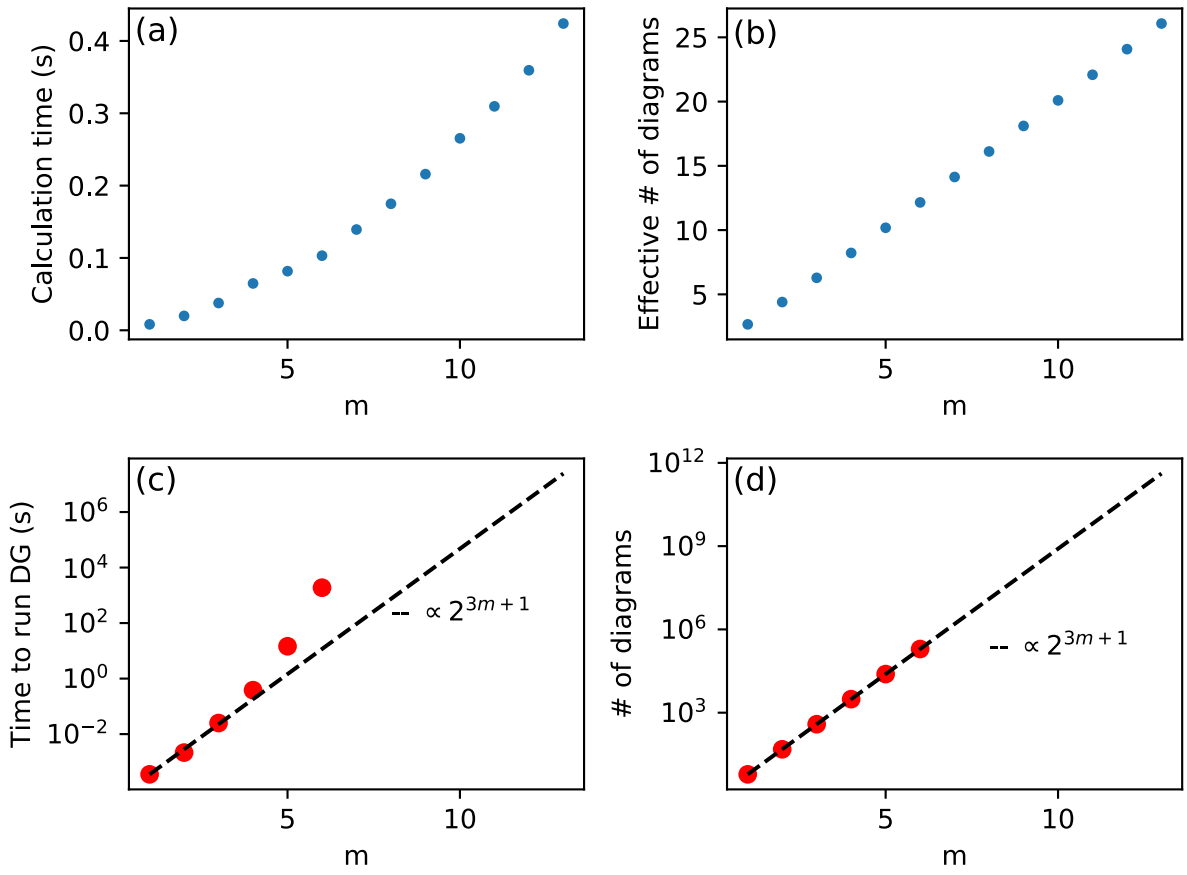

Figure 8.3: (a) Calculation runtime for calculating a single pump-probe signal at a delay time of 500 fs using the composite diagram method on a 2017 MacBook Pro. (b) Effective number of diagrams that must be calculated in the composite diagrams approach. (c) Runtime for using the DG to enumerate all Feynman diagrams that contribute to a $2m+1$ order calculation of a 3LS where manifolds are separable. (d) Number of individual diagrams needed to calculate the $2m+1$ order signal for a 3LS where manifolds are separable. Larger systems and/or systems with inseparable manifolds would require more diagrams.

Chapter 9

Efficient Methods for Diagonalizing Polymer Hamiltonians

In this chapter we describe in more detail the basis that we use to represent vibronic Hamiltonians. In section 2.2 we introduce a class of vibronic Hamiltonians, Eq. 2.4, which consists of s two-level systems (2LS), collectively coupled to k vibrational modes. As mentioned in section 2.2 we represent the Hamiltonian in the number basis of the simple harmonic oscillator using ladder operators, and truncate the infinite dimensional Hilbert space to a size that is sufficient to give good convergence of the spectra we are interested in simulating. Attaining good convergence of spectroscopic signals is fundamentally contingent upon converging the relevant subset of eigenvectors and eigenvalues of H_0. The number of eigenvalues needed, and the required accuracy of those eigenvalues, will depend upon the pulse shapes used, and the spectroscopic signals of interest. However, the convergence behavior of simulated spectra as a function of basis truncation will mirror the convergence behavior of the eigenvalues of H_0. We therefore find it useful to study the convergence of eigenvalues of H_0 in order to determine the best basis and truncation scheme to use when simulating spectra.

In this chapter, we describe the basis and truncation scheme that we use in more detail. We first compare two different choices of basis, which we refer to as the ground-state basis and the halfway-shifted basis, and demonstrate that the halfway-shifted basis requires fewer basis terms to converge eigenvalues of H_0. We then compare two different methods of truncating the basis, which we refer to as square truncation and triangular truncation, and demonstrate that the triangular truncation scheme requires fewer terms to converge the eigenvalues of H_0. Finding the most efficient method of representing H_0 is of crucial importance to UF^2 because diagonalizing H_0 scales as N^3 (where N is the dimension of the Hamiltonian), and diagonalizing $\mathscr{L}_0$ scales as either N^3 or N^6.

To begin, we start with a simple case of the vibronic Hamiltonian, Eq. 2.4, with a single 2LS ($s = 1$) linearly

coupled to k harmonic modes. In first-quantized notation for the electronic states, the Hamiltonian takes the form

$$H = \frac{1}{2}\sum_{\alpha=1}^{k} p_\alpha^2 + \frac{1}{2}\sum_{\alpha=1}^{k} \omega_{\alpha,g}^2 q_\alpha^2 \left|g\right\rangle \left\langle g\right| + \frac{1}{2}\sum_{\alpha=1}^{k} \left(\omega_{\alpha,e}^2 (q_\alpha - d_\alpha)^2 + E_e\right) \left|e\right\rangle \left\langle e\right|, \tag{9.1}$$

where $|e\rangle$ is the electronic excited state, $\omega_{\alpha,g}$ are the normal mode frequencies in the ground state $|g\rangle$, $\omega_{\alpha,e}$ are the normal mode frequencies in the excited state $|e\rangle$, E_e is the energy gap between the minima of the ground-state and excited-state potential energy surfaces, and the Huang-Rhys factors are given by $S_\alpha = \frac{\omega_{\alpha,g}\omega_{\alpha,e}d^2}{1+\omega_{\alpha,e}/\omega_{\alpha,g}}$ [136]. The ground-state Hamiltonian $H_g = \langle g| H |g\rangle$ and excited-state Hamiltonian $H_e = \langle e| H |e\rangle$ are both simply representations of a set of k independent harmonic oscillators, and are thus trivially diagonalizable when considered separately. By making the substitution $q_\alpha \to \sqrt{\frac{\hbar}{2\omega_{\alpha,g}}}\left(b_\alpha + b_\alpha^\dagger\right)$ and $p_\alpha \to i\sqrt{\frac{\hbar\omega_{\alpha,g}}{2}}\left(b_\alpha - b_\alpha^\dagger\right)$, we find that $H_g = \hbar\omega_{\alpha,g}\sum_{\alpha=1}^{k}\left(b_\alpha^\dagger b_\alpha + \frac{1}{2}\right)$, with eigenvectors $\left|\nu^{(g)}\right\rangle$. By making a similar substitution that accounts for the different frequency and displacement of the excited state, we can similarly obtain $H_e = \hbar\omega_{\alpha,e}\sum_{\alpha=1}^{k}\left(\tilde{b}_\alpha^\dagger \tilde{b}_\alpha + \frac{1}{2}\right)$, with eigenvectors $\left|\nu^{(e)}\right\rangle$, where $\tilde{b}_\alpha = \sqrt{\frac{\omega_{\alpha,e}}{\omega_{\alpha,g}}}\left(b_\alpha - d_\alpha\sqrt{\frac{\omega_{\alpha,g}}{2\hbar}}\right)$. However, diagonalizing the full Hamiltonian using a single basis is more difficult. Ref. [137] provides an analytical formula for the overlaps $\left\langle \nu^{(e)} \middle| \nu^{(g)} \right\rangle$, which allows us to diagonalize H_e using the eigenbasis of H_g.[i] In practice, we find that the formulas from Ref. [137] can only be used with arbitrary precision arithmetic for large values of S_α (we have found this to be the case for $S_\alpha \gtrapprox 3$), and that numerically diagonalizing H_e is actually faster in some cases. In either case, there is in general no solution when $s > 1$ and/or when anharmonicities are included. The most straight-forward choice of basis is to express the entire Hamiltonian using the eigenbasis for the ground-state. We refer to this basis choice as the ground-state basis. In section 9.1 we introduce an alternative basis that we call the halfway-shifted basis, and show that this alternative basis is a more efficient method for representing H_0, because it requires a smaller Hilbert space dimension to achieve convergence of the eigenvalues.

9.1 Vibronic Basis Choice

We use the Hamiltonian in Eq. 9.1 to motivate our choice of basis, and then show how that basis can be used to diagonalize the full vibronic Hamiltonian from Eq. 2.4. One of the simplest choices is to represent the vibrations in the eigenbasis of the ground state Hamiltonian H_g. In this basis H_g is already diagonal, but H_e is not. We have found that using a different basis can make the diagonalization of the overall Hamiltonian H much more efficient. We choose ladder operators corresponding to a fictitious potential energy surface that has a minimum that is halfway between the minima of V_g and V_e, and that has a curvature (set of vibrational frequencies) corresponding to the geometric mean of each of the individual vibrational frequencies $\omega_{\alpha,\text{eff}} = \sqrt{\omega_{\alpha,g}\omega_{\alpha,e}}$.

We first define dimensionless p_α and q_α via substitution $q_\alpha \to \sqrt{\frac{\hbar}{\omega_{\alpha,g}}}q_\alpha$ and $d_\alpha \to \sqrt{\frac{\hbar}{\omega_{\alpha,g}}}d_\alpha$, which implies that

[i] Ref.[**?**] expands upon the work of ref. [137] to include the Duschinsky effect, which is the case where the normal modes in the excited state $|e\rangle$ are linear combinations of $\{q_\alpha\}_{\alpha=1}^{k}$, the normal modes in the ground state.

$p_\alpha \to \sqrt{\hbar\omega_{\alpha,g}}p_\alpha$, which allows us to write Eq. 9.1 as

$$H = \frac{1}{2}\sum_{\alpha=1}^{k} \hbar\omega_{\alpha,g}\left(p_\alpha^2 + q_\alpha^2 \left|g\right\rangle\left\langle g\right| + \left(y_\alpha^2(q_\alpha - d_\alpha)^2 + E_e\right)\left|e\right\rangle\left\langle e\right|\right),$$

where $y_\alpha = \frac{\omega_{\alpha,e}}{\omega_{\alpha,g}}$.

We next rewrite H using the transformations $q_\alpha \to \frac{q_\alpha}{y_\alpha^{1/4}} + \frac{d_\alpha}{2}, p_\alpha \to y_\alpha^{1/4}p_\alpha$, where $y_\alpha = \frac{\omega_{\alpha,e}}{\omega_{\alpha,g}}$, which yields

$$H = \frac{1}{2}\sum_{\alpha=1}^{k} \hbar\omega_{\alpha,\text{eff}}\left(p_\alpha^2 + y_\alpha^{-1}\left(q_\alpha + \frac{y_\alpha^{1/4}d_\alpha}{2}\right)\left|g\right\rangle\left\langle g\right| + \left(y_\alpha\left(q_\alpha - \frac{y_\alpha^{1/4}d_\alpha}{2}\right)^2 + E_e\right)\left|e\right\rangle\left\langle e\right|\right). \tag{9.2}$$

This transformation effectively splits the computational difficulty equally between diagonalizing H_g and H_e. Although this is not of fundamental importance, since analytical solutions are known for $s = 1$, it is very useful for the case where $s > 1$, where the singly-excited manifold (SEM) mixes excitations where a single site i is excited, while the other sites are in their electronic ground states. We find that diagonalizing the SEM is most efficiently done in the "halfway basis" described here (which is similarly true of the doubly-excited manifold for $s > 2$, and so on). This same transformation is also helpful for $s \geq 1$ when the modes are anharmonic. In general we can apply the same transformation for the anharmonic case where

$$\begin{aligned}
V_{g,\alpha}(q_\alpha) &= \frac{1}{2}(\omega_{g,\alpha}^2 q_\alpha^2 + \beta_{g,\alpha}q_\alpha^3 + \gamma_{g,\alpha}q_\alpha^4 + ...) \\
V_{e,\alpha}(q_\alpha) &= \frac{1}{2}(\omega_{e,\alpha}^2 (q_\alpha - d_\alpha)^2 + \beta_{e,\alpha}(q_\alpha - d_\alpha)^3 + \gamma_{e,\alpha}(q_\alpha - d_\alpha)^4 + ...).
\end{aligned}$$

Note that convergence of the algorithm requires that the highest-order term in $V(q)$ must have even order, and that the coefficient for that term must be positive, so as to guarantee that all states are bound. If this condition is not met, the eigenstates are no longer bound, but rather are resonances. These resonances may be solved for numerically, but care must be taken not to include too many basis terms.

Thus far we have discussed the most efficient method for representing the Hamiltonian in terms of the quantities y_α and d_α. To diagonalize the Hamiltonian, we express it using the raising and lowering operators of the quantum harmonic oscillator, defined for the unitless displacements and momenta as

$$\begin{aligned}
q_\alpha &= \frac{1}{\sqrt{2}}\left(b_\alpha + b_\alpha^\dagger\right) \\
p_\alpha &= \frac{i}{\sqrt{2}}\left(b_\alpha - b_\alpha^\dagger\right),
\end{aligned}$$

where $b_\alpha^\dagger$ are the raising operators for vibrational modes of frequencies $\omega_{\alpha,\text{eff}}$ that are centered at positions $y_\alpha^{1/4}d_\alpha/2$ relative to the ground-state equilibrium displacement. In the basis of the raising and lowering operators, H is an infinite sparse matrix, which must be truncated at some maximum occupation number. In order to resolve the

lowest m eigenvectors of H to some tolerance, we must include $N > m$ basis states. For large d_α, y_α and/or anharmonic terms, it may be that $N \gg m$, and it can be much more efficient to use sparse eigensolvers (such as Lanczos) to determine the necessary m eigenstates. Arbitrary precision may be obtained by increasing N.

9.2 Truncation schemes

For $k = 1$, there is only one possible truncation scheme, which is to include all basis elements $|\nu\rangle$ for $\nu < N$. However, for $k > 1$, there are a variety of possibilities. The simplest truncation scheme is to truncate each oscillator at the same maximum occupation number n, where we require that $\left\langle b_\alpha^\dagger b_\alpha \right\rangle < n$ for each α. This can be thought of as a square truncation scheme, because it truncates the infinite basis space into a square for $k = 2$, a cube for $k = 3$, and hypercubes for $k > 3$. The hypercube has edges of size n, and so encloses a total hypervolume of n^k basis states. The square scheme therefore gives a total dimension of $N = n^k$. We have found that this truncation scheme is not a very efficient one.

We motivate our alternative truncation scheme by considering a special case: a system of k identical oscillators where $\omega = \omega_{\alpha,g}$, $y = y_\alpha$ and $d = d_\alpha$. In such a case, we argue that a truncation scheme where the maximum total occupation number is held constant is more physically intuitive. We define the total occupation number of a state as

$$n_{\text{total}} = \sum_\alpha \left\langle b_\alpha^\dagger b_\alpha \right\rangle, \tag{9.3}$$

and truncate our basis such that $n_{\text{total}} < n$. The number of basis states with constant $n_{\text{total}} = i$ is given by $\begin{pmatrix} k+i-1 \\ i \end{pmatrix}$, and therefore the total number of basis states with $n_{\text{total}} < n$ is

$$N = \sum_{i=0}^{n-1} \begin{pmatrix} k+i-1 \\ i \end{pmatrix},$$

$$N = \begin{pmatrix} k+n-1 \\ n-1 \end{pmatrix}.$$

As n becomes large, $N \sim \frac{n^k}{k}$, and so we have saved by roughly a factor of k (in reality a little less, because we are neglecting lower-order terms that are all additive). The scaling is still exponential with the number of modes, but it represents an important savings, since the cost of diagonalizing the Hamiltonian goes as N^3.

As a concrete example, if $k = 3$ and $n_{\text{total}} < 2$, then only the states $|000\rangle, |100\rangle, |010\rangle, |001\rangle$ are included in the triangular truncation scheme. This is in contrast with the square truncation scheme where $|110\rangle, |101\rangle, |011\rangle, |111\rangle$ would also be included. Since the coupling of each individual oscillator is equal in this example case, we expect the states $|100\rangle, |010\rangle, |001\rangle$ to all contribute equally to the lowest-energy eigenstate. Likewise the states

Table 9.1: Coefficients of the lowest energy eigenstate in the SEM using the halfway-shifted basis for a monomer ($s = 1$) with two vibrational modes ($k = 2$) with $d = 0.4$ and $y = 1$. Anti-diagonals of table correspond to constant n_{total}. Each anti-diagonal consists of entries that are all within a factor of 2. In contrast, moving along any diagonal gives entries that diminish by an order of magnitude at each step. In the constant n_{total} truncation scheme, only entries out to a given anti-diagonal are kept.

$\langle \nu_1 \nu_2 \mid \psi_0 \rangle$	ν_1			
ν_2	0	1	2	3
0	0.98020	0.13862	0.01386	0.00113
1	0.13862	0.01960	0.00196	0.00016
2	0.01386	0.00196	0.00020	0.00002
3	0.00113	0.00016	0.00002	$< 10^{-5}$

$|110\rangle, |101\rangle, |011\rangle$ also contribute equivalently. However, the square truncation scheme misses the states $|200\rangle, |020\rangle, |002\rangle$, which should contribute at roughly the same level as the other $n_{\text{total}} = 2$ states, and should contribute more strongly than the $|111\rangle$ state (which is only one of many states with $n_{\text{total}} = 3$). As an example of this, see Table 9.1 for the $k = 2$ case, which shows the amplitudes of various basis states for the lowest energy eigenvector. This truncation scheme still produces Hilbert space dimensions that scale exponentially with the number of vibrational modes, but the dimension grows more slowly than with naive truncation, and allows us to treat more vibrational modes than we would otherwise.

Figure 9.1 shows the convergence of the 11^{th} lowest energy eigenvalue in the SEM as a function of the number of basis elements included for a trimer with three harmonic vibrational modes ($s = k = 3$). This figure demonstrates that fewer basis elements are required when using the halfway-shifted basis, as compared with the naive basis. The figure further demonstrates that even fewer basis elements are required when using the truncation scheme described by Eq. 9.3. Figure 9.2 shows that the same convergence behavior holds for a dimer with two anharmonic vibrational modes ($s = k = 2$).

If the couplings between sites are coordinate dependent ($dV_{ij}/dq_\alpha \neq 0$), then this coupling is expressed using the ladder operators b_α as well. Such couplings can be used to describe, for example, a conical intersection, as in molecules like Pyrazine.

Next consider the case where the vibrational frequencies are identical ($\omega = \omega_{\alpha,g}$), but the values of y_α and/or d_α are not identical. In such a case, the most efficient truncation scheme would likely be to keep basis states such that the quantity

$$\tilde{n}_{\text{total}} = \max \sum_\alpha r_\alpha \langle b_\alpha^\dagger b_\alpha \rangle \tag{9.4}$$

is held constant, where $r_\alpha = f(d_\alpha, y_\alpha)$. We have not implemented a truncation scheme of this style, but it is something that should be fairly easy to include if a useful form of f could be determined. We propose a simple algorithm that could be implemented to automatically determine r_α. The algorithms is as follows:

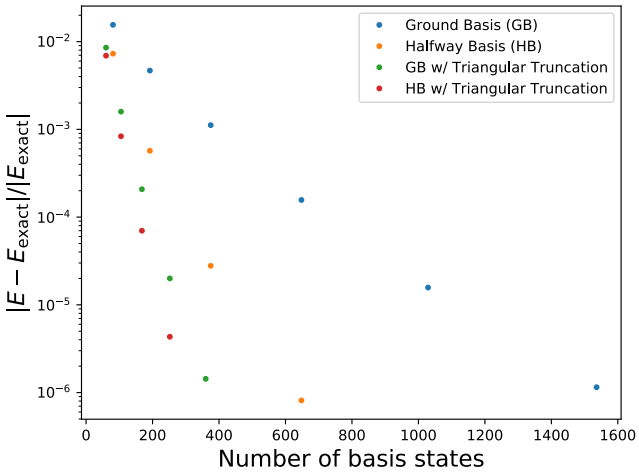

Figure 9.1: (a) Relative difference from a converged reference value of the 11$^{\text{th}}$ lowest energy eigenvalue in the SEM for a trimer ($s = k = 3$) as a function of the number of basis states included using two different bases (blue and green dots) without the truncation scheme of Eq. 9.3, and for the Halfway-shifted basis using the truncation scheme of Eq. 9.3 (green dots). The reference value of 2.99029 for all three curves is obtained using the halfway-shifted basis with truncation by total occupation number. The trimer has site energies of 0, 0.5, and $0.8\omega_0$ and site couplings of 0.1, 0.15, and $0.2\omega_0$. Each site is linearly coupled to harmonic vibrational modes with coupling $d = 0.8$. The ground and excited state vibrational frequencies are identical, with values $\omega = 1, 1.01, 1.0101\omega_0$. We have set the values of the vibrational frequencies to be different from one another in order to avoid degeneracies in the GSM.

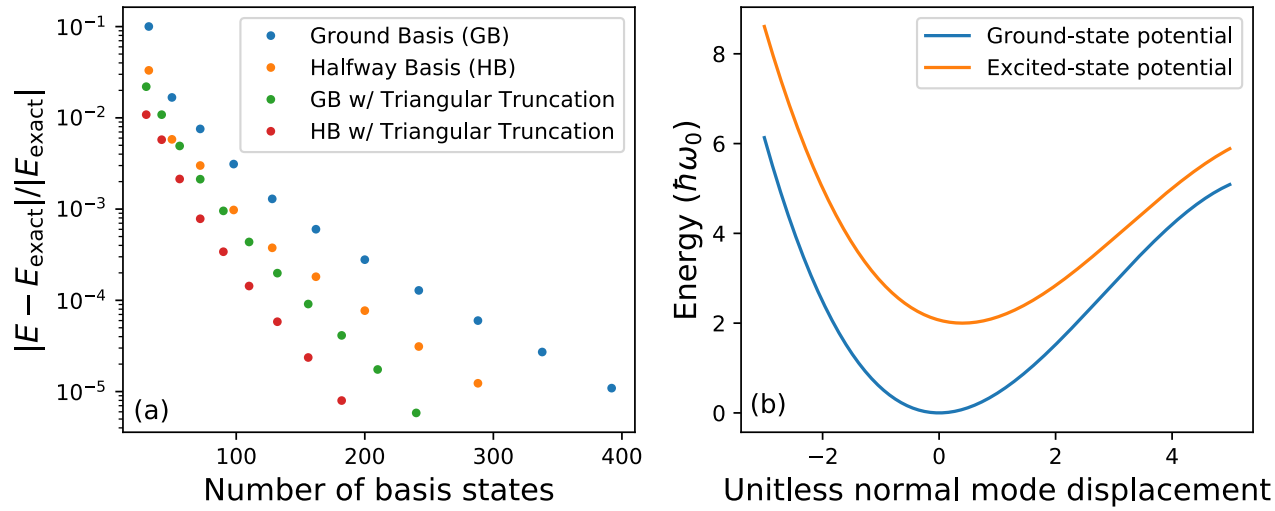

Figure 9.2: (a) Relative difference from a converged reference value of the 16$^{\text{th}}$ lowest energy eigenvalue in the SEM for an anharmonic dimer ($s = k = 2$) as a function of the number of basis states included using two different bases (blue and orange dots) using the square truncation scheme of Eq. 9.3, and for the triangular truncation scheme (green and red dots). The reference value for all four curves is obtained using the halfway-shifted basis with triangular truncation. The dimer is composed of two monomers with site energies 0, $0.75\omega_0$ and coupling $J = 0.3307\omega_0$. Each site is coupled to nearly identical anharmonic vibrational modes with $V_g(q) = \frac{1}{2}\hbar\omega_{g,i}\left(q^2 - 0.12q^3 + 0.0003q^4\right)$ and $V_e(q) = 0.81V_g(q - 0.4) + E_g$, where E_g is assumed to be larger than any of the other relevant energy scales, and $\omega_{g,1} = \omega_0$ and $\omega_{g,2} = 1.01\omega_0$ (we have set the vibrational frequencies to be slightly different in order to avoid degeneracies in the GSM). The factor of 0.81 corresponds to shifting the frequency $\omega_e = 0.9\omega_g$. The excited-state potential curve has a minimum shifted by $d = 0.4$. Panel (b) shows V_g (blue) and V_e (orange) for a single monomer. Note that we dramatically reduce the energy shift E_g between the orange and blue curves in order to plot both curves together.

1. For each α, construct the Hamiltonian for only vibrational mode α in the halfway-shifted basis:

$$H_\alpha = \frac{1}{2}\hbar\omega_{\alpha,\text{eff}}\left(p_\alpha^2 + y_\alpha^{-1}\left(q_\alpha + \frac{y_\alpha^{1/4}d_\alpha}{2}\right)\right),$$

   using the ladder operators to represent the matrix.

2. Converge the lowest energy eigenvalue (or alternatively all eigenvalues below some threshold) of H_α as a function of the number of basis states to some desired tolerance. The number of basis states required to meet this tolerance is n_α.

3. Define $n_{\text{ave}} = \left(\prod_{\alpha=1}^{k} n_\alpha\right)^{1/k}$, and $r_\alpha = \frac{n_{\text{ave}}}{n_\alpha}$. (The arithmetic mean could also be used in lieu of the geometric mean.)

The next step would then be to use Eq. 9.4 to truncate the full vibronic Hamiltonian, and converge over the single parameter $\tilde{n}_{\text{total}}$. Note that Eq. 9.4 reduces to Eq. 9.3 when $r_\alpha = 1$.

It is worth noting that Eqs. 9.3 and 9.4 represent flat hyper-surfaces in the k-dimensional space spanned by $\left\{\langle b_\alpha^\dagger b_\alpha\rangle\right\}$, with intercepts at $\tilde{n}_{\text{total}}/r_\alpha$. Using all basis states that are inside the hyper-volume bounded by this hyper-surface (with conditions that $n_\alpha \geq 0$) seems like a natural choice when $r_\alpha = 1$, since it is related to the idea of a maximum total occupation number. However, generalizing to Eq. 9.4, where $\tilde{n}_{\text{total}}$ need no longer be an integer, opens up the the question as to whether curved hyper-surfaces might provide faster convergence. It is possible that using all of the basis states to the interior of a hyper-sphere (or hyper-ellipsoid) could be more favorable. Or, even more generally, one could consider truncation schemes

$$\tilde{n}'_{\text{total}} = \max\left(\sum_\alpha r_\alpha \langle b_\alpha^\dagger b_\alpha\rangle^q\right)^{1/q}.$$

As $q \to \infty$, one would retrieve the naive truncation method, and as $q \to 0$ one would arrive at a truncation scheme where only basis states that exist along the axes $\left\{\langle b_\alpha^\dagger b_\alpha\rangle\right\}$ would be included, which is neither physical, nor mathematically useful, since a basis containing only elements of the variety $\{|n_1, 0, 0...\rangle, |0, n_2, 0, 0, ...\rangle, ...\}$ does not span the Hilbert space of the system. We have not tested any truncation schemes beyond that of Eq. 9.3, but expect that the ideal truncation scheme is likely to be somewhere between $q = \frac{1}{2}$ and $q = 2$. Looking at the anti-diagonals of Table 9.1 shows that not all of the terms at fixed n_{total} have equal magnitude , and that coefficients of terms where all $n_\alpha \neq 0$ are slightly larger. This points towards the possibility that $q > 1$ could yield slightly faster convergence, particularly for larger k.

This chapter has introduced the most efficient basis that we have found for representing and diagonalizing vibronic Hamiltonians with vibrational modes that are very similar. We also point out that modifications of this truncation scheme may be developed for other cases, and show where future work could be done to attempt to improve the truncation scheme, if desired. Both choices of basis and both truncation schemes are included in

the Hamiltonian/Liouvillian generator (HLG) that is included in UFSS and that is described in Section 6.3. The triangular truncation and halfway basis are set as the default methods for representing the vibrational degrees of freedom in the HLG.

Chapter 10

Conclusion

This thesis describes the development of a fully functional software suite, called ultrafast spectroscopy suite (UFSS), that can automatically predict nonlinear optical spectroscopies (NLOSs) for any phase-matching/phase-cycling (collectively phase-discrimination) condition, including the finite pulse effects. The diagram generator (DG), outlined in Chapter 7, automatically generates all of the Feynman diagrams that give a non-zero contribution to a given phase-discrimination condition, allowing us to automate perturbative calculations, and visually inspect any diagram or set of diagrams. Chapter 8 demonstrates how to automatically combine subsets of Feynman diagrams into composite diagrams, which gives us an exponential speedup as we go to higher orders in perturbation theory. This allows us to go well beyond low-order perturbation theory, including studying multiple Rabi oscillations as a function of pulse amplitude, by summing many orders of perturbation theory (where we must check that the summation converges as a function of the order of perturbation theory included). These two contributions represent the most general contributions that we have made to the field of NLOSs, as they can be applied to any method of performing perturbative calculations in the light-matter interaction.

We have also introduced a new and efficient algorithm for performing time-dependent perturbation theory calculations in the context of NLOSs, called Ultrafast Ultrafast spectroscopy (UF^2), described for closed systems in Chapter 5 and for open systems in Chapter 6. Given a model for the system (e.g. a molecular Hamiltonian or Liouvillian), the electronic dipole operator, and pulse parameters, UF^2 is designed to automatically calculate the contribution of any diagram or composite diagram, making the prediction of NLOSs as easy as defining the pulse shapes and phase-discrimination condition. In order to perform calculations, UF^2 requires as input a model for the system dynamics, which can be any time-independent Hamiltonian or Liouvillian that can be well approximated by a discrete and finite number of basis states. Users can define their own model systems, but we have also included a model system generator, called the Hamiltonian/Liouvillian generator (HLG), which is introduced in Chapter 2 and explained in further detail in Chapter 6. We describe the details of how vibronic Hamiltonians are represented, including the basis states and truncation method used, in Chapter 9.

Taken altogether UFSS, comprising the HLG, UF^2 and DG (along with the composite diagram method), represents a fully functional tool for predicting NLOSs, including arbitrary pulse shapes. Once a user has defined the material system under study, UFSS makes it straightforward to predict any NLOS they like. Once a user understands how to generate a transient absorption spectra, generating 2D photon echo spectra, or 5th-order exciton-exciton interaction 2D spectra, is as simple as changing less than 10 lines of code. This user-friendliness, combined with the speed of the calculations, make UFSS a useful tool that can be used to explore how experimental pulse shapes effect spectra, to run many repeated simulations in order to fit experimental data to models, to generate large data sets with varying system and pulse parameters in order to train neural nets, and to explore the connection between perturbative and non-perturbative limits of the optical interaction with model systems, and hopefully others.

Bibliography

[1] U. Morgner, F. X. Kärtner, S. H. Cho, Y. Chen, H. A. Haus, J. G. Fujimoto, E. P. Ippen, V. Scheuer, G. Angelow, and T. Tschudi, Opt. Lett. **24**, 411 (1999).

[2] D. H. Sutter, G. Steinmeyer, L. Gallmann, N. Matuschek, F. Morier-Genoud, U. Keller, V. Scheuer, G. Angelow, and T. Tschudi, Opt. Lett. **24**, 631 (1999).

[3] G. S. Engel, T. R. Calhoun, E. L. Read, T.-K. Ahn, T. Mančal, Y.-C. Cheng, R. E. Blankenship, and G. R. Fleming, Nature **446**, 782 (2007).

[4] N. Christensson, H. F. Kauffmann, T. Pullerits, and T. Mančal, J. Phys. Chem. B **116**, 7449 (2012).

[5] E. Thyrhaug, R. Tempelaar, M. J. P. Alcocer, K. Žídek, D. Bína, J. Knoester, T. L. C. Jansen, and D. Zigmantas, Nature Chemistry **10**, 780 (2018).

[6] J. Cao, R. J. Cogdell, D. F. Coker, H.-G. Duan, J. Hauer, U. Kleinekathöfer, T. L. C. Jansen, T. Mančal, R. J. D. Miller, J. P. Ogilvie, V. I. Prokhorenko, T. Renger, H.-S. Tan, R. Tempelaar, M. Thorwart, E. Thyrhaug, S. Westenhoff, and D. Zigmantas, Science Advances **6** (2020), 10.1126/sciadv.aaz4888.

[7] M. Cho, H. M. Vaswani, T. Brixner, J. Stenger, and G. R. Fleming, The Journal of Physical Chemistry B **109**, 10542 (2005).

[8] J. Adolphs, F. Müh, M. E.-A. Madjet, and T. Renger, Photosynthesis Research **95**, 197 (2007).

[9] F. Müh, M. E.-A. Madjet, J. Adolphs, A. Abdurahman, B. Rabenstein, H. Ishikita, E.-W. Knapp, and T. Renger, Proceedings of the National Academy of Sciences **104**, 16862 (2007).

[10] A. Perdomo-Ortiz, J. R. Widom, G. A. Lott, A. Aspuru-Guzik, and A. H. Marcus, J. Phys. Chem. B **116**, 10757 (2012).

[11] J. Dobler, W. Zinth, W. Kaiser, and D. Oesterhelt, Chemical Physics Letters **144**, 215 (1988).

[12] R. Schoenlein, L. Peteanu, R. Mathies, and C. Shank, Science **254**, 412 (1991).

[13] D. M. Jonas, Annual Review of Physical Chemistry **54**, 425 (2003).

[14] S. Mukamel, *Principles of Nonlinear Optical Spectroscopy* (Oxford University Press, 1999).

[15] J. A. Cina, P. A. Kovac, C. C. Jumper, J. C. Dean, and G. D. Scholes, The Journal of Chemical Physics **144**, 175102 (2016).

[16] X. Leng, S. Yue, Y.-X. Weng, K. Song, and Q. Shi, *Chemical Physics Letters*, **667** (2016).

[17] T. N. Do, M. F. Gelin, and H.-S. Tan, The Journal of Chemical Physics **147**, 144103 (2017).

[18] V. Perlík, J. Hauer, and F. Šanda, J. Opt. Soc. Am. B **34**, 430 (2017).

[19] C. L. Smallwood, T. M. Autry, and S. T. Cundiff, J. Opt. Soc. Am. B **34**, 419 (2017).

[20] T. N. Do, L. Chen, A. K. Belyaev, H.-S. Tan, and M. F. Gelin, Chemical Physics **515**, 119 (2018).

[21] J. Yuen-Zhou, J. J. Krich, I. Kassal, A. S. Johnson, and A. Aspuru-Guzik, *Ultrafast Spectroscopy* (IOP Publishing, 2014).

[22] W. Kohn, Rev. Mod. Phys. **71**, 1253 (1999).

[23] J. Erhard, P. Bleiziffer, and A. Görling, Phys. Rev. Lett. **117**, 143002 (2016).

[24] P. R. Bunker and P. Jensen, *Molecular symmetry and spectroscopy*, 2nd ed. (NRC Research Press, 2006).

[25] F. Di Maiolo, G. A. Worth, and I. Burghardt, J. Chem. Phys. **154**, 144106 (2021).

[26] C. Schnedermann, A. M. Alvertis, T. Wende, S. Lukman, J. Feng, F. A. Y. N. Schröder, D. H. P. Turban, J. Wu, N. D. M. Hine, N. C. Greenham, A. W. Chin, A. Rao, P. Kukura, and A. J. Musser, Nature Communications **10**, 4207 (2019).

[27] G. A. Worth, H.-D. Meyer, and L. S. Cederbaum, J. Chem. Phys. **109**, 3518 (1998).

[28] Y. Tanimura and R. Kubo, J. Phys. Soc. Jpn. **58**, 101 (1989).

[29] C. Kreisbeck, T. Kramer, M. Rodríguez, and B. Hein, J. Chem. Theory Comput. **7**, 2166 (2011).

[30] Y. Tanimura, The Journal of Chemical Physics **137**, 22A550 (2012).

[31] H. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2007).

[32] V. Gorini, A. Kossakowski, and E. C. G. Sudarshan, Journal of Mathematical Physics **17**, 821 (1976).

[33] G. Lindblad, Communications in Mathematical Physics **48**, 119 (1976).

[34] P. Malý, O. J. G. Somsen, V. I. Novoderezhkin, T. Mančal, and R. van Grondelle, ChemPhysChem **17**, 1356 (2016).

[35] G. McCauley, B. Cruikshank, D. I. Bondar, and K. Jacobs, npj Quantum Information **6**, 74 (2020).

[36] I. I. Rabi, N. F. Ramsey, and J. Schwinger, Rev. Mod. Phys. **26**, 167 (1954).

[37] A. G. Redfield, Phys. Rev. **98**, 1787 (1955).

[38] J. I. Sudyka, S. Pustelny, and W. Gawlik, New Journal of Physics **21**, 023024 (2019).

[39] P. F. Tekavec, G. A. Lott, and A. H. Marcus, J. Chem. Phys. **127**, 214307 (2007).

[40] F. D. Fuller and J. P. Ogilvie, Annu. Rev. Phys. Chem. **66**, 667 (2015).

[41] A. A. Bakulin, C. Silva, and E. Vella, The Journal of Physical Chemistry Letters **7**, 250 (2016).

[42] E. U. Condon, Phys. Rev. **32**, 858 (1928).

[43] M. Toutounji, Chemical Physics **525**, 110402 (2019).

[44] J. Seibt and T. Mančal, Chemical Physics **515**, 129 (2018).

[45] R. P. Feynman, *Quantum Electrodynamics*, Frontiers in Physics (Westview Press, Philadelphia, PA, 1998).

[46] J. F. Ward, Rev. Mod. Phys. **37**, 1 (1965).

[47] A. M. Bonch-Bruevich and V. A. Khodovoĭ, Soviet Physics Uspekhi **8**, 1 (1965).

[48] A. Omont, E. W. Smith, and J. Cooper, Astrophys. J. **175**, 185 (1972).

[49] T. K. Yee and T. K. Gustafson, Phys. Rev. A **18**, 1597 (1978).

[50] A. Salam, *Molecular Quantum Electrodynamics: Long-Range Intermolecular Interactions* (John Wiley and Sons, Ltd, 2009).

[51] J. Hahn and B. S. Ham (Optical Society of America, Kauai, Hawaii, 2011) p. OWC3.

[52] S. V. Poltavtsev, M. Salewski, Y. V. Kapitonov, I. A. Yugova, I. A. Akimov, C. Schneider, M. Kamp, S. Höfling, D. R. Yakovlev, A. V. Kavokin, and M. Bayer, Phys. Rev. B **93**, 121304 (2016).

[53] S. V. Poltavtsev, M. Reichelt, I. A. Akimov, G. Karczewski, M. Wiater, T. Wojtowicz, D. R. Yakovlev, T. Meier, and M. Bayer, PRB **96**, 075306 (2017).

[54] G. S. Schlau-Cohen, A. Ishizaki, T. R. Calhoun, N. S. Ginsberg, M. Ballottari, R. Bassi, and G. R. Fleming, Nature Chemistry **4**, 389 (2012).

[55] S. Westenhoff, D. Paleček, P. Edlund, P. Smith, and D. Zigmantas, J. Am. Chem. Soc. **134**, 16484 (2012).

[56] M. Z. Peter Hamm, *Concepts and Methods of 2D Infrared Spectroscopy* (Cambridge University Press, 2011).

[57] D. Paleček, P. Edlund, E. Gustavsson, S. Westenhoff, and D. Zigmantas, The Journal of Chemical Physics **151**, 024201 (2019).

[58] A. S. Johnson, J. Yuen-Zhou, A. Aspuru-Guzik, and J. J. Krich, The Journal of Chemical Physics **141**, 244109 (2014).

[59] V. Engel, Computer Physics Communications **63**, 228 (1991).

[60] S. Meyer, M. Schmitt, A. Materny, W. Kiefer, and V. Engel, Chemical Physics Letters **301**, 248 (1999).

[61] T. Kato and Y. Tanimura, Chemical Physics Letters **341**, 329 (2001).

[62] D. V. Tsivlin, H.-D. Meyer, and V. May, The Journal of Chemical Physics **124**, 134907 (2006).

[63] Y.-C. Cheng, H. Lee, and G. R. Fleming, The Journal of Physical Chemistry A **111**, 9499 (2007).

[64] K. Renziehausen, P. Marquetand, and V. Engel, Journal of Physics B: Atomic, Molecular and Optical Physics **42**, 195402 (2009).

[65] J. D. Bell, R. Conrad, and M. E. Siemens, Opt. Lett. **40**, 1157 (2015).

[66] J. H. Fetherolf and T. C. Berkelbach, The Journal of Chemical Physics **147**, 244109 (2017).

[67] L. Seidner, G. Stock, and W. Domcke, The Journal of Chemical Physics **103**, 3998 (1995).

[68] T. Mančal, A. V. Pisliakov, and G. R. Fleming, The Journal of Chemical Physics **124**, 234504 (2006).

[69] J. Seibt, K. Renziehausen, D. V. Voronine, and V. Engel, The Journal of Chemical Physics **130**, 134318 (2009).

[70] W. Domcke and G. Stock, "Theory of ultrafast nonadiabatic excited-state processes and their spectroscopic detection in real time," in *Advances in Chemical Physics* (John Wiley & Sons, Ltd, 2007) pp. 1–169.

[71] M. F. Gelin, D. Egorova, and W. Domcke, The Journal of Chemical Physics **123**, 164112 (2005).

[72] M. Gelin, D. Egorova, and W. Domcke, Chemical Physics **312**, 135 (2005).

[73] M. F. Gelin, D. Egorova, and W. Domcke, Accounts of Chemical Research **42**, 1290 (2009).

[74] M. F. Gelin, D. Egorova, and W. Domcke, The Journal of Chemical Physics **131**, 194103 (2009).

[75] Y. Tanimura and Y. Maruyama, The Journal of Chemical Physics **107**, 1779 (1997).

[76] Y.-a. Yan, Chinese Journal of Chemical Physics **30**, 277 (2017).

[77] J. R. Johansson, P. D. Nation, and F. Nori, Computer Physics Communications **183**, 1760 (2012).

[78] M. Beck, A. Jackle, G. Worth, and H.-D. Meyer, Physics Reports **324**, 1 (2000).

[79] M. Feit, J. Fleck, and A. Steiger, Journal of Computational Physics **47**, 412 (1982).

[80] D. Kosloff and R. Kosloff, Journal of Computational Physics **52**, 35 (1983).

[81] D. Calvetti, L. Reichel, and D. C. Sorensen, Electron. Trans. Numer. Anal. **2**, 1 (1994).

[82] E. J. Heller, Accounts of Chemical Research **14**, 368 (1981).

[83] M. Cho, *Two-Dimensional Optical Spectroscopy* (CRC Press, 2009).

[84] V. Tiwari and D. M. Jonas, The Journal of Chemical Physics **148**, 084308 (2018).

[85] R. E. Fenna and B. W. Matthews, Nature **258**, 573 (1975).

[86] M. Liebel, C. Schnedermann, T. Wende, and P. Kukura, The Journal of Physical Chemistry A **119**, 9506 (2015).

[87] S. D. McClure, D. B. Turner, P. C. Arpin, T. Mirkovic, and G. D. Scholes, The Journal of Physical Chemistry B **118**, 1296 (2014).

[88] J. Yuen-Zhou, J. J. Krich, and A. Aspuru-Guzik, The Journal of Chemical Physics **136**, 234501 (2012).

[89] D. Abramavicius, B. Palmieri, D. V. Voronine, F. Šanda, and S. Mukamel, Chemical Reviews **109**, 2350 (2009).

[90] S. M. Gallagher Faeder and D. M. Jonas, The Journal of Physical Chemistry A **103**, 10489 (1999).

[91] N. Belabas and D. M. Jonas, Opt. Lett. **29**, 1811 (2004).

[92] P. F. Tekavec, J. A. Myers, K. L. M. Lewis, F. D. Fuller, and J. P. Ogilvie, Opt. Express **18**, 11015 (2010).

[93] H. Li, A. P. Spencer, A. Kortyna, G. Moody, D. M. Jonas, and S. T. Cundiff, J. Phys. Chem. A **117**, 6279 (2013).

[94] A. Anda and J. H. Cole, "Two-dimensional spectroscopy beyond the perturbative limit: the influence of finite pulses and detection modes," (2020), arXiv:2011.04343 .

[95] J. Süß and V. Engel, The Journal of Chemical Physics **153**, 164310 (2020).

[96] P. A. Rose and J. J. Krich, The Journal of Chemical Physics **150**, 214105 (2019).

[97] P. A. Rose and J. J. Krich, J. Chem. Phys. **154**, 034109 (2021).

[98] C. Gardiner and P. Zoller, *Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics (Springer Series in Synergetics)* (Springer, 2004).

[99] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems* (Oxford University Press, Oxford New York, 2002).

[100] A. Raab, I. Burghardt, and H.-D. Meyer, J. Chem. Phys. **111**, 8759 (1999).

[101] A. Raab and H.-D. Meyer, The Journal of Chemical Physics **112**, 10718 (2000).

[102] Y. Tanimura, Phys. Rev. A **41**, 6676 (1990).

[103] C. Liang and T. L. C. Jansen, J. Chem. Theory Comput. **8**, 1706 (2012).

[104] D. Egorova, A. Kühl, and W. Domcke, Chemical Physics **268**, 105 (2001).

[105] U. Kleinekathöfer, I. Kondov, and M. Schreiber, Chemical Physics **268**, 121 (2001).

[106] G. Katz, R. Kosloff, and M. A. Ratner, Israel Journal of Chemistry **44**, 53 (2004).

[107] A. Ishizaki and G. R. Fleming, The Journal of Chemical Physics **130**, 234110 (2009).

[108] V. May and O. Kühn, *Charge and energy transfer dynamics in molecular systems* (Wiley-VCH, Weinheim, 2011).

[109] F. Caycedo-Soler, A. W. Chin, J. Almeida, S. F. Huelga, and M. B. Plenio, The Journal of Chemical Physics **136**, 155102 (2012).

[110] N. Killoran, S. F. Huelga, and M. B. Plenio, The Journal of Chemical Physics **143**, 155102 (2015).

[111] B. Brügemann and V. May, The Journal of Chemical Physics **118**, 746 (2003).

[112] J. Süß, J. Wehner, J. Dostál, T. Brixner, and V. Engel, The Journal of Chemical Physics **150**, 104304 (2019).

[113] P. Malý, S. Mueller, J. Lüttig, C. Lambert, and T. Brixner, The Journal of Chemical Physics **153**, 144204 (2020).

[114] P. A. Rose, A. C. McClung, T. E. Keating, A. T. C. Steege, E. S. Egge, and A. K. Pattanayak, "Dynamics of non-classicality measures in the decohering harmonic oscillator," (2012), arXiv:1206.3356 [quant-ph] .

[115] W. M. Zhang, T. Meier, V. Chernyak, and S. Mukamel, The Journal of Chemical Physics **108**, 7763 (1998).

[116] T. Brixner, J. Stenger, H. M. Vaswani, M. Cho, R. E. Blankenship, and G. R. Fleming, Nature **434**, 625 (2005).

[117] G. Panitchayangkoon, D. V. Voronine, D. Abramavicius, J. R. Caram, N. H. C. Lewis, S. Mukamel, and G. S. Engel, Proceedings of the National Academy of Sciences **108**, 20908 (2011).

[118] S. J. Jang and B. Mennucci, Rev. Mod. Phys. **90**, 035003 (2018).

[119] C. D. Cruz, J. Yuan, C. Climent, N. T. Tierce, P. R. Christensen, E. L. Chronister, D. Casanova, M. O. Wolf, and C. J. Bardeen, Chem. Sci. **10**, 7561 (2019).

[120] A. J. Kiessling and J. A. Cina, The Journal of Chemical Physics **152**, 244311 (2020).

[121] L. Yang, I. V. Schweigert, S. T. Cundiff, and S. Mukamel, Phys. Rev. B **75**, 125302 (2007).

[122] P. Huo and D. F. Coker, J. Phys. Chem. Lett. **2**, 825 (2011).

[123] J. Dostál, T. Mančal, R.-n. Augulis, F. Vácha, J. Pšenčík, and D. Zigmantas, J. Am. Chem. Soc. **134**, 11611 (2012).

[124] K. J. Karki, J. R. Widom, J. Seibt, I. Moody, M. C. Lonergan, T. Pullerits, and A. H. Marcus, Nature Communications **5**, 5869 (2014).

[125] J. Provazza, F. Segatta, M. Garavelli, and D. F. Coker, Journal of Chemical Theory and Computation **14**, 856 (2018).

[126] J. Dostál, F. Fennel, F. Koch, S. Herbst, F. Würthner, and T. Brixner, Nature Communications **9**, 2466 (2018).

[127] P. Malý and T. Mančal, J. Phys. Chem. Lett. **9**, 5654 (2018).

[128] M. K. Yetzbacher, N. Belabas, K. A. Kitney, and D. M. Jonas, The Journal of Chemical Physics **126**, 044511 (2007).

[129] V. Tiwari, W. K. Peters, and D. M. Jonas, Proceedings of the National Academy of Sciences **110**, 1203 (2013).

[130] G. Nardin, T. M. Autry, K. L. Silverman, and S. T. Cundiff, Opt. Express **21**, 28617 (2013).

[131] P. A. Rose and J. J. Krich, J. Chem. Phys. **154**, 034108 (2021).

[132] J. Albert, M. Falge, M. Keß, J. G. Wehner, P.-P. Zhang, A. Eisfeld, and V. Engel, The Journal of Chemical Physics **142**, 212440 (2015).

[133] Y. Ke and Y. Zhao, The Journal of Chemical Physics **149**, 014104 (2018).

[134] T. Mančal, "Quantarhei: Molecular open quantum systems package," `http://github.com/tmancal74/quantarhei` (2020).

[135] Z. Zhang, K. L. Wells, E. W. J. Hyland, and H.-S. Tan, Chemical Physics Letters **550**, 156 (2012).

[136] J.-L. Chang, Journal of Molecular Spectroscopy **232**, 102 (2005).

[137] J.-L. Chang, The Journal of Chemical Physics **128**, 174111 (2008).