

Communicating Research to the General Public

At the March 5, 2010 UW-Madison Chemistry Department Colloquium, the director of the Wisconsin Initiative for Science Literacy (WISL) encouraged all Ph.D. chemistry candidates to include a chapter in their Ph.D. thesis communicating their research to non-specialists. The goal is to explain the candidate's scholarly research and its significance to a wider audience that includes family members, friends, civic groups, newspaper reporters, state legislators, and members of the U.S. Congress.

Ten Ph.D. degree recipients have successfully completed their theses and included such a chapter, less than a year after the program was first announced; each was awarded \$500.

WISL will continue to encourage Ph.D. chemistry students to share the joy of their discoveries with non-specialists and also will assist in the public dissemination of these scholarly contributions. WISL is now seeking funding for additional awards.



The dual mission of the Wisconsin Initiative for Science Literacy is to promote literacy in science, mathematics and technology among the general public and to attract future generations to careers in research, teaching and public service.

UW-Madison Department of Chemistry
1101 University Avenue
Madison, WI 53706-1396
Contact: Prof. Bassam Z. Shakhshiri
bassam@chem.wisc.edu
www.scifun.org

January 2011

Picosecond and Femtosecond Coherent
Multidimensional Spectroscopy of Colloidal PbSe
Quantum Dot Structure and Dynamics

By

Stephen Benjamin Block

A Dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

(Chemistry)

at the

University of Wisconsin-Madison

2012

The dissertation is approved by the following members of the Final Oral Committee:

Date of final oral examination: June 27, 2012

Month and year degree to be awarded: August 2012

The dissertation is approved by the following members of the Final Oral Committee:

John Wright, professor, Chemistry

Mark Ediger, professor, Chemistry

Robert Hamers, professor, Chemistry

Randall Goldsmith, assistant professor, Chemistry

Deniz Yavuz, associate professor, Physics

Appendix F: Ultrafast Coherent Multidimensional Spectroscopy of Quantum

Dots for a Non-Technical Audience

The Wisconsin Initiative for Science Literacy (WISL) has established a program that provides incentives for graduating PhD students to include a non-technical description of their research in their dissertations. As much of the work done in doctoral programs interacts only indirectly with participants outside the individual sub-fields, the WISL objective is admirable. I welcome their invitation to include such content in this document and am grateful for the award WISL offers for participation in this part of their programming. The research I have conducted over the course of the past few years under the direction of John Wright is fundamental research—a step or two (if I'm generous) removed from commercial or industrial applications. Such foundational research, however, feeds into the growing body of knowledge possessed by the scientific community and is essential for the development of new technologies. Over the course of this overview, I intend to provide you with a general description of both my research and some of the connections these insights have and may come to have within the broader technological and scientific world. In a nutshell, I (and the members of the Wright group with whom I have worked) have developed a technique that uses pulsed lasers to interrogate the energetic structure and behavior of molecules and have adapted that technique so that it can be used to study semiconductor materials that are of potential use in solar energy collection devices.

Introduction to Spectroscopy

Spectroscopy is the study of the interaction between light and matter. Any discussion of why a substance reflects, absorbs, or emits a certain color of light stems from an understanding of spectroscopic phenomena. For example, the ozone layer is known to absorb harmful

ultraviolet radiation, which has a shorter wavelength, and thus more energy per photon, than visible light or infrared light. Ozone accomplishes this task because light in that range of wavelengths matches the separation between energy levels in the ozone molecules (we would say that the light is resonant with the transition) and the probability of the molecule absorbing the right-colored light that happens to hit it is high (we would say that the transition dipole moment is large). A blue shirt contains molecules that absorb light in the red and green ranges of the spectrum, but still reflect and transmit blue. A microwave oven works because the long wavelength microwaves emitted by the device are absorbed by molecules in the food. That transfer of energy causes the food to heat up. The wavelengths of light absorbed by a substance are characteristic of the components of the substance—its atoms and their bonding arrangement. With the right tools, it is possible to learn about an unidentified substance by measuring the colors of light it absorbs.

Spectroscopy studies not only the energetic structure of a molecule, but also how each of its high energy states behaves. Consider a car left in the sun on a summer day. Sunlight is absorbed by the upholstery in the car because it is less than perfectly reflective; in fact it is often dark, which means that it absorbs light very effectively. The light that is absorbed doesn't just leave the molecules in their excited states forever. The molecules bump into their neighbors and have a chance of transferring some of their energy. Over the course of these collisions, energy is often dispersed from high energy states to many low energy states, such as molecular vibrations and bends. Given enough time, any of these states can relax back down to its initial state (the ground state) by emitting light again. Excited vibrational states tend to live longer than the higher excited electronic states. If any excited state has enough interactions with its

surroundings, or even other parts of the molecule, then it doesn't emit light of the same color it absorbed. Rather it emits at the lower energy wavelengths of the vibrational modes (infrared instead of visible, for example). In the case of the car, the infrared light, which our skin detects as heat, rather than our eyes detecting it as light, is not transmitted by the windows of the car. Hence, visible light passes through the windows and is absorbed by molecules that relax to molecular vibrations before they can reemit at the same color they absorbed. The molecules then emit infrared radiation that is reflected by the windows and stays trapped in the car. Over time, the infrared radiation builds up and the car's interior temperature rises to an uncomfortably high level. Spectroscopy studies not only the color the upholstery absorbs (and hence the high energy states present in the molecules), but also the color it emits (the low energy states present), as well as how long it takes for a molecule to relax down from high to low energy and any states that the molecule may have passed through on its way down.

Why Use Lasers for Spectroscopy?

In order to learn about the structure and dynamics of molecules with certainty, a great deal of precision is required, and there is frequently no better source of well-defined light than lasers. Laser beams can be made to travel with very little divergence so you can direct all of the energy of the beams to the specific destinations desired. Lasers can be tuned to access wavelengths across the infrared, the visible, and the ultraviolet regions of the spectrum. Further, laser systems can be constructed so that they are not continuous beams, but rather pulses that are short enough in time to be used to measure phenomena as brief as the movement of an electron from one part of a molecule to another (this is the attosecond time scale—one billionth of a billionth of a second). Lasers also have well-defined phase, which is important for any

experiments that need to observe phenomena that can only be described with the wave functions of quantum mechanics.

Spectroscopy in the Wright Group

The particular brand of spectroscopy used in my research is a type of non-linear spectroscopy. Linear spectroscopies are techniques that study interactions with a single light source. Our experiments use three laser pulses to gain a more complete picture of the structure and dynamics of the system. A schematic of the systems we use is shown in Figure 1, along with a photograph of one of the optical tables, at the end of this document. The schematic starts at a part of our laser system after the main laser pulse is created. That main pulse is split into two paths and directed into the boxes labeled “Optical Parametric Amplifier” (OPA). OPAs are devices that allow the wavelength of light to be changed. Once we’ve chosen the color of light that we want to use, we split one of them again and focus all three beams into the sample.

In order to observe and measure any behavior that is time-dependent, it is important to control the sequence and separation in time between the pulses. The schematic also shows a couple of delay stages. These stages have retroreflectors (a combination of mirrors that send incoming beams out exactly parallel to their angle of entrance) and precise motors that control their positions. By changing the position of a stage, we control the length of the path that each pulse takes to the sample. As we shorten the path, the pulse arrives earlier. I mentioned above that precision is required. The phenomena we wish to observe take place on time scales between ten femtoseconds (ten millionths of a billionth of a second) and ten picoseconds (ten millionths of a millionth of a second). To change the relative timing of something slow-moving by such a small amount of time, you would have to change its path by an absurdly small amount.

Conveniently, light is fast. Still, the 93 million mile journey from the sun to Earth takes light only eight minutes. If we want the light pulses to arrive at the sample one picosecond earlier, we have to make the path just 300 microns shorter (less than a third of a millimeter). To measure some of the faster dynamics, we need reliable control over delay stage position down to a few microns—about the spacing between the grooves on a CD.

The Basics of Non-Linear Spectroscopy Investigated in the Wright Group

I mentioned above that our type of spectroscopy is a non-linear spectroscopy, which means that we measure sample behavior as it interacts with multiple light sources. In particular, the three beams we focus into our samples each interact with the sample, and the effects produced by each beam interfere with each other to produce the complicated effects we measure. To get a feel for this kind of interference, consider dropping small rocks into a perfectly still pond. A single rock will create a ripple pattern that spreads out from the place it landed. If two rocks were dropped into the pond next to one another, their ripple patterns would interfere with each other—in some locations the waves they create would combine to make larger waves (constructive interference); in other places they would cancel each other out (destructive interference). If the second rock was dropped into the pond a short while after the first, the overall pattern would be the same, but the locations of each type of interference would be shifted.

The pond is similar to our sample and each rock dropped is like the interaction between a laser pulse and a molecule. As the pulses travel through the sample, it's as if they're dropping rocks across stretches of the pond. Connecting to an earlier idea, working with either a more intense laser or a larger transition dipole moment is like dropping a heavier rock in the pond.

Working with a resonant transition is like dropping rocks in the same location with timing set to make the ripple larger with every drop. So choosing resonant transitions with large dipole moments makes large waves. If the beam angles and wavelengths are chosen well, then the interference from multiple beam lines will cooperate to create large waves that travel out in certain directions. This process of cooperation is called phase matching. The red beam coming out of the sample in Figure 1 is one such “large wave” direction. It only appears when each beam’s interaction with the sample is strong and they all cooperate with each other. We block the portions of each other beam that are transmitted through the sample after they’ve interacted with it so that we only measure this one phase-matched output.

The non-linear output beam we choose is directed into a monochromator before reaching the detector. The monochromator spreads the incoming light into its component colors in an effect that appears like a prism separating sunlight into a rainbow. The mechanics of this spectral dispersion are more like the interference effects described in the pond analogy (or letting a shaft of sunlight reflect off a CD) than they are like a prism, but the final effects are similar. By adjusting the angles of a grating inside the monochromator (like changing the angle of the CD relative to the beam of light), we choose which color of light reaches the detector.

Experimental Overview and Technique Development on RDC

I have now summarized the primary variables used in each of our experiments—the two colors of light generated by the OPAs, the relative pulse timing (controlled by the pair of delay stages), and the color observed by our detector as selected by the monochromator. By changing one or two of these variables and fixing the rest we generate the spectra that give us insight into the structure and dynamics of our system. Additionally, we can change the intensity of the

beams or rotate their polarization (loosely like swishing water back and forth instead of dropping a rock into the water from above) to study additional effects. The technique was developed over the course of a couple graduate student generations of investigation on molecules with interesting vibrations. I'll tell you about some of the primary capabilities of the technique using some of those experiments as examples.

Rhodium Dicarbonyl Acetylacetonate (RDC) was one of the molecules we studied when developing our technique. A picture of this molecule is shown in Figure 2a. We studied the behavior of this molecule's vibrational modes, specifically those involving the stretching and compressing of bonds in the carbonyl groups (carbon/oxygen pairs that aren't in the ring structure) connected to the rhodium atom. The two stretching modes of interest were the symmetric stretch and the antisymmetric stretch, as cartooned in Figure 2b. For RDC to start vibrating in these modes, the molecule must absorb infrared light that is resonant with the transitions. The simple linear absorption scan for RDC in this range of the infrared is shown in Figure 2c. These scans provide little more than the basic energetic structure of the molecule and the relative strengths of the transitions, but they also tell us the frequency ranges to use for our more complex experiments. The frequency of light is simply the multiplicative inverse of the wavelength, so the higher frequency corresponds to higher energy and shorter wavelengths.

The result of one of these more complex experiments is shown in Figure 3. Here the axes are the frequencies of the two lasers, as chosen by the OPAs. The relative delay between the pulses is fixed such that what are called ω_2 and ω_2' (the two beams that come from the same OPA output) arrive at the sample first. (Note that ω_2 is shown as ω_2 in Figure 1; the added negative sign reminds us how this pulse interacts in the particular alignment scheme used in

these experiments.) They create an interference pattern (like the two different sequences of rocks dropped in the pond). Where the electric field intensity is high (where the waves are the strongest), these two beams can together put the molecule in an excited vibrational state, but only if their frequency/wavelength is resonant with a transition. That frequency—the frequency of the pulses that can be used to put the molecule into an excited state—is shown along the vertical axis. The horizontal axis is the ω_1 frequency (the beam that is not split again after being produced by the OPA). This laser interacts with the population states created by the other two beams, either triggering relaxation (stimulated emission) or adding to the molecular vibrations. To the left of the plot and above it, I've juxtaposed the linear absorption spectrum from Figure 2c for comparison.

It is immediately apparent that the spectrum in Figure 3 has more information than the spectrum in Figure 2c. When both lasers are tuned to the same frequency (along the diagonal), peaks appear at the same frequencies as appear in the linear absorption plot. On these peaks, all pulses interact with the same vibrational mode, either putting the molecule in its fundamental excited state or stimulating emission of light from that mode. The peaks that are *not* along the diagonal start to provide new information. When one laser is resonant with one transition and the other laser is resonant with the other transition (when $\omega_2 = 2015\text{cm}^{-1}$ and $\omega_1 = 2084\text{cm}^{-1}$, for example), a peak may appear, but only if the two modes are coupled. In this case, a peak does appear, indicating that the vibrational modes are coupled (which isn't surprising, given that they involve the same atoms). To say that the two modes are coupled simply means that one mode behaves differently when the other mode is populated. When the symmetric stretch of the molecule is excited, the amount of energy required to also get the molecule in the excited

antisymmetric stretching mode decreases. Consider an analogy of two people in a room. Each person can be put into a good mood if they hear a certain amount of good news. Suppose one of the two hears enough good news for him/her to become rather giddy. If the other person in the room can be made happy more (or less) easily because the other person is in a good mood, then their moods are coupled like these two vibrational modes are coupled. If the second person is perfectly indifferent to the mood of the first, then their moods are not coupled. In this spectroscopy, no peak would appear on the cross peak of their moods.

For each of the four fundamental peaks visible in Figure 3 (peaks for which both lasers are resonant with one of the transitions visible in the linear absorption plot), there is also a second peak shifted to the left. These red-shifted peaks appear at the frequency required to create an overtone (the same type of vibration, but roughly two-times faster) or a combination band state (both types of vibration active at the same time). The peaks appear at lower energies than the fundamental transitions because it takes less energy to make the molecule vibrate more than it did to get it vibrating in the first place. An imperfect spring can be imagined to work the same way. When you first try to stretch a really tight spring, it takes a lot of force. If that spring is already loosened up, it's easier to stretch it more because the metal starts to get a little bent out of its original shape. Here as the bonded atoms vibrate more, the increase of energy means that the bond is getting closer and closer to breaking, so it isn't holding the atoms together as tightly. So each of the additional peaks in this spectrum appear when ω_1 can jump from the excited state created by ω_2 and ω_2' to an even higher excited state. From scans like this one, then, we can learn which modes interact with each other and we can start to learn about higher energy modes than just the fundamental transitions. In some cases, a linear absorption scan can have many

peaks next to each other so that it is hard to identify which modes are present and what their energies are. When cross peaks only appear for coupled modes, it is much easier to resolve complicated portions of an absorption spectrum.

Dynamics and Coherences

Figure 3 shows one way that we access more energetic structure information. To study the behavior of any of these combinations of modes, we keep the frequencies of the lasers fixed, but scan the delays between them, producing plots like the one shown in Figure 4. Here the origin of the plot corresponds to all three pulses being overlapped in time. As ω_2 and ω_2' arrive at the sample before or after ω_1 , the axis value becomes negative or positive, respectively. An example pulse ordering is shown above the plot; the timing corresponds to the delay values indicated in the plot by the black circle. Any type of molecular behavior in time can be shown in plots like this, but this one shows plainly a couple of important basic types of information. The signal stays bright (the red color on the plot) the furthest away from (0,0) when the two delay values are equal and negative. These values correspond to pulse sequences where ω_2 and ω_2' arrive at the sample simultaneously and create a population state. As we travel away from the origin in other directions, we instead measure the coherence lifetimes.

Coherences are quantum-mechanical superpositions between two different states. Being able to selectively create these superpositions and watch their behavior gives this technique a lot of its specificity. When populations relax (as in the case of a high-energy state dispersing its energy to low-energy states as a car heats up in the sun), the third pulse we send in to explore the state can still generate signal, even if the state we created with the first two pulses isn't the same as the state present when the third arrives. When we work with coherences only, any disruption

of the energy level of the participating states either makes the signal disappear entirely or in certain circumstances makes it appear at a different wavelength, which we choose for or against with the monochromator setting.

As with many concepts in quantum mechanics, the nature of coherences isn't immediately intuitive. One way to think about a coherence is by considering a light with a dimmer switch. For the purposes of this analogy, suppose lights are normally thought of as either off or fully illuminated—those settings would correspond to population states. Coherences could be any of the partially illuminated settings from very dim (mostly “off” character) to rather bright (mostly “on” character). Signal from coherences is short lived because it is easy to disrupt. Unlike populations, coherences oscillate. When one thinks about molecular vibrational modes on a quantum mechanical level, they are regarded as stationary states—the probability of finding electrons in certain locations doesn't change for as long as the molecule is in that state. The probability of finding electrons in certain locations when the system is in a coherent superposition state does change in time. It's as if the light on the dimmer switch also flickered at a well-known rate whenever it was set to something other than “off” or “on”. The signal from coherences disappears so quickly because while not every collision may be able to rearrange the energy of the molecule (like the collisions in the upholstery of the car on a summer day), collisions *are* able to easily disrupt the phase of the oscillation, making the molecules get out of sync with each other. So if many lights on dimmer switches are blinking in time with each other, it is easy to see the flicker. If they are out of sync with each other, then they all average out and you don't really see the flickering at all.

The added complexity of coherences may seem needless for research studying the behavior of excited states, but there are some very important phenomena that can only be described at the quantum mechanical level of coherences. Though this topic will arise in later descriptions, an example of behavior that can only be described on a quantum mechanical level is the energy transfer from light-absorbing molecules in plants to the parts of the cells that use the energy to convert carbon dioxide to sugars and oxygen. In order to have a detailed understanding of the molecules we study, it is important for us to be able to watch both population- and coherence-based phenomena. Coherent and incoherent behaviors create different effects in our data. Figure 4 shows a basic picture of lifetime measurements. Figure 5 shows the same type of scan, but then we changed the monochromator settings to watch output from a different state—a type of result that can only be observed when coherences retain their phase, but change to a different state. It's as if those blinking lights all changed from flickering at one rate to another rate, but did so in sync. The take-home message from scans like the one shown in Figure 5 is that we can detect coherent energy transfer, and the pattern that appears in scans detecting it tells us the energy separation between the initial and final states of that transfer.

Semiconductors and Quantum-Confinement

All of the research described above relates to vibrational transitions and modes. Over time, it became clear that the techniques developed provided a powerful and unique way to learn about different molecules. Our attention turned, then, to systems that have become of interest to the field of materials science—quantum-confined semiconductors. Before explaining the specifics of our results, I'll explain what quantum-confinement is and the aspects of

semiconductors to which quantum-confinement is relevant. I will also provide something of an explanation of one (of many) important possible applications of these materials—dye-sensitized solar cells.

Semiconductors are given their designation because they can be electrical insulators or conductors, depending on their energetic condition and elemental composition. Instead of discussing individual states, such as the symmetric and anti-symmetric stretching modes of RDC, the identifying features of semiconductors revolve around two collections of electronic states—each of which blur into an effective continuum of states called the valence and conduction bands. The valence band is comprised of electronic states that are localized on the atoms in the material. Any electrons in the conduction band are in energetic states that are delocalized across the lattice of the material, like the electrons in metals that are free to flow around a wire. The energetic separation between the valence and conduction bands is called the bandgap, and can be changed by changing the composition of the semiconductor or, in some interesting cases, by changing the physical size and shape of the material.

If all electrons in a semiconductor structure are in the valence band, then the structure does not conduct electricity. If by injection of energy (be it in the form of light or voltage or another) equal to or greater than the bandgap, an electron is promoted to the conduction band, then the material will conduct... at least a little. When that electron is promoted, however, it is leaving its ground state localized orbital behind. The absence of the negative charge of the electron leaves a positive charge behind. This vacancy carries the positive charge and is called a hole. (I say that it carries the charge because the actual positive charge is in the nucleus of the atom, but if another electron fills in the hole, then the vacancy is associated with a different

orbital, possibly on a different atom.) Even though the electron is now free to roam about the lattice, it still feels attraction to the positive charge it left behind, so it never wanders far. The electron/hole pair is called an “exciton”. How far the electron will wander is a function of the atoms and bonds that make up the semiconductor, so each different material has a native exciton size, called the “Bohr radius”.

Semiconductors can be synthesized in such a way that the actual lattice of atoms is smaller than the exciton Bohr radius of the material in at least one dimension. In this situation, the semiconductor is said to be quantum-confined and starts to take on new properties. First, an exciton has a native size in the first place because that size is the lowest energy configuration. By making the material smaller than the lowest energy exciton Bohr radius, the system can no longer access the lowest energy configuration and is forced to be in a higher energy state. This manifests as an increase in the bandgap energy. As the exciton container size decreases further, the bandgap energy increases further. With enough synthetic customizability, materials can be made with tunable bandgaps. Another effect of reducing the number of atoms (and hence also reducing the container size) in the semiconductor is that the number of electron states decreases. Quantum-confined materials tend to have a few clumps of states at low energies before the continuous band structure of the bulk material reappears. Because these are the quantum-confined excitonic states, the exciton fills the container and hence takes on the shape of the material.

Figure 6 shows a number of the characteristics of quantum-confined semiconductors described above. In Figure 6a, the band structure of a generic bulk semiconductor is shown next to the energetic structure of a hypothetical quantum-confined structure of the same material. The

quantum-confined structure has a larger bandgap and discrete states. A diagram of a quantum dot—a spherical semiconductor that is smaller than the Bohr radius in all three dimensions—is shown in Figure 6b with a partially-transparent sphere representing the exciton Bohr radius. Data showing the effect of different amounts of confinement are shown in Figure 6c. Here a series of linear absorption scans shows the increasing energy of the first absorption peak of lead selenide (PbSe) quantum dots of decreasing radius.

Solar Cells and Quantum-Confined Semiconductors

I have described how changing the size of a semiconductor can change its energetic properties and that the ability to change those properties is important, but I have not yet explained why this flexibility is so important. Before I do, it is important to note that these structures can be made from very abundant and inexpensive materials. Some semiconductors have rare elements in the composition like indium or gallium; others require a great deal of energy to purify—silicon is ubiquitous in the form of silicon dioxide (glass, sand, quartz), but requires processing to reduce to pure silicon. Alternatively, iron oxide (rust, hematite) and titanium dioxide (TiO_2 , present in such luxurious consumer products as white paint and sunscreen) can gain additional high tech uses when their energetic structure is tuned to meet a particular need. One such need is the development of devices that can be used to efficiently harvest solar energy.

Dye-sensitized solar cells make use of a molecule that can efficiently absorb some portion of the solar spectrum and transfer an electron into a conducting material. Figure 7 shows a basic schematic of such an assembly at a molecular level. TiO_2 is shown as the electron acceptor because it is well-suited to the task—it has a large bandgap so it is statistically unlikely

for enough collisions to cooperate to get it to relax down to distant vibrational modes and it has many vacant electronic states in its conduction band, so there's always room for donation of electrons. With the right dye molecule attached to the surface of the electron acceptor, solar cells with low efficiencies can be constructed. The dyes, though, require some means of binding to the TiO_2 and have to be able to move an electron across that additional assembly before relaxing back down to their own ground state. Additionally these molecules must absorb as much of the solar spectrum as possible in order to avoid wasting the energy in the transmitted/reflected wavelength ranges.

It is in the creation of a donor/acceptor assembly that the use of quantum-confined materials is anticipated to shine. Quantum dots can be grown to sizes that absorb light all across the solar spectrum; in a full solar cell, if one size of dot only absorbs a portion of the spectrum, another size can absorb much of the rest. These dots can be grown directly onto the surface of an electron accepting material (and have been—see Figure 8a), meaning that there is no need for additional bridging assemblies. Additionally, as the exciton already wants to be larger than the quantum dot, the electron will already be stretching into the acceptor (diagrammed in Figure 8b), making transfer all the faster and more efficient. If both the donor and acceptor can be adapted to match energetically, then the transfer itself shouldn't require significant loss in energy. The alignment shown in Figure 8c, for example, would waste a lot of energy in the transfer.

The synthetic flexibility and energetic properties of quantum-confined materials already make their use for prototype dye-sensitized solar cells appealing. The question remains how to optimize their creation, and the answer lies in understanding the details of the charge transfer. What states are involved in the transfer? Are there additional electronic states formed at the

interface between the two material types? What kind of synthesis produces the least disruptive interfaces? These are the types of questions that our type of spectroscopy can address.

Coherent and Incoherent Electron Transfer

Before providing an overview of our progress toward answering the questions above, there is one additional line of distinction I would like to draw regarding the nature of the charge transfer. Earlier I discussed coherent energy transfer and said that the distinction between incoherent and coherent dynamics would be important; it is in this discussion that the distinction becomes critical. One way to think about this distinction is to consider a swimming pool with a shallow end and a deep end. If a bunch of water is added to the empty pool, it rather quickly ends up in the deep end. The water spreads out under the influence of gravity until it finds that it can continue to drop down until none of the water is in the shallow end of the pool. This scenario is like coherent electron transfer. The electron orbit is larger than the quantum dot (remember Figure 8b), and the electron acceptor has lower energy states for the electron to fall into, so the electron cloud just flows right into the proverbial deep end. Incoherent energy transfer is more like putting a beach ball in the same pool—it will roll into the deep end and stay there, but it may bounce around for a while first; how quickly it gets to the deep end depends on the slope of the shallow end of the pool (or how much other pull there may be to the deep end), and one could definitively say “now it’s in the shallow end; now it’s in the deep end”. Because coherent energy transfer happens by spreading out and simultaneously exploring the options, the fastest charge transfer phenomena end up being describable in only this way.

Most non-linear spectroscopy techniques (pump-probe and transient absorption are the names of the two most common) establish an excited state population and observe where that

energy is at a later time. These techniques can be very good at determining initial and final states of a system, but they usually cannot provide any information about short-lived intermediate states and they cannot discriminate between different paths the system had taken to reach its final state. Experiments like the one that produced the plot in Figure 5 show the states involved in the transfer, and any intermediate state would contribute to the interference pattern created. So if we are able to measure coherence lifetimes and selectively excite different states, then we can explore every possible situation that could lead to charge transfer. But there's a lot of tailoring and development that must be done first.

Initial Wright Group Spectroscopy of Quantum Dots

Our first forays into the spectroscopy of quantum-confined semiconductor systems focused on PbSe quantum dots. PbSe dots have exciton states that are in an easily accessible wavelength range for our lasers, their synthesis is well-known, and other spectroscopy techniques have already been used to study them, so there is a body of work in existence for comparison. After making some changes to our OPAs to allow them to access the correct ranges, we collected frequency scans (see Figure 9) from a few different batches of dots. We found that we could resolve the peaks and gain some decent structural information about the system. We could confirm coupling between two different quantum-confined exciton states by noting the presence of a cross-peak. When we tried to study dynamics, however, we found that our one-picosecond pulses were too long to measure any coherent phenomena. Figure 10 shows a delay scan like the one in Figure 4. The population lifetimes are certainly long enough for us to measure (along the negative portion of the $x = y$ diagonal and along the positive x -axis of Figure 10), but as soon as the pulses that create the population state aren't overlapped, we have

no non-linear signal light at all. To confirm, we tried simulating our data from quantum mechanical calculations and we found that we could recreate the data with a wide range of coherence lifetimes—only the laser pulse duration mattered. Frequency scans can be matched in a similar way and yield lifetime information with some analysis and processing, but without the ability to select fully coherent processes with our pulses, much of the technique that was available to studies of molecular vibrations was no longer usable.

Quantum Dot Spectroscopy With Femtosecond Pulses

At the end of the line of picosecond system experiments, it became clear that we needed to use much shorter pulses if we wanted to observe any coherent phenomena. The next step, then, was to assemble a laser system that could give us tunable 40-femtosecond pulses—25 times shorter than one picosecond. (That duration was selected to be a match to the coherence lifetimes calculated from the frequency scans on the older system, which are also in accord with measurements from around the scientific community.) We used an experimental setup functionally the same as the one shown in Figure 1. Pulses this short are difficult to create and keep stable, so the optical table had to be designed with all high-end reflective optics on mounts that would dampen any vibrations that might run through the heavy steel table. As of late Spring 2012, we confirmed that we can collect the frequency scans that were available from the picosecond system (Figure 11 provides an example) and measure coherence lifetimes. Figure 12 shows a delay scan that seems similar to the one shown in Figure 10 (with picosecond pulses), but here the width of the peaks is larger than the pulses, which means that we are measuring coherence lifetimes directly. To confirm that we could generate signal from pulse sequences that did not create an intermediate population, we did another frequency scan at a different time

ordering. Figure 13a shows the data next to two simulations (13b and 13c) that include and exclude signal from the fully coherent processes. The fact that the main peak area has a shape that appears more like the simulation with fully coherent contributions than the one without means that we are able to suppress the processes that have intermediate populations in favor of those that do not.

The signs that we have restored our technique's functionality are subtle. For example, Figure 12 (femtosecond system, quantum dots) still looks more like Figure 10 (picosecond system, quantum dots) than Figure 4 (picosecond system, RDC), even though it contains the same information as Figure 4 and more than Figure 10. Even so, the results are conclusive—we can once again selectively measure both energetic structure and excited state dynamics, be they coherent or incoherent. This restoration means that we are now positioned to start characterizing the types of charge transfer events that are essential for any efficient quantum-dot-sensitized solar cell cores. Future Wright group research directions include investigation of charge transfer systems and improvements to the technique, including cryogenic cooling (which increases the coherence lifetimes, making resolution of fully coherent dynamics easier) and OPA tuning range extension (which allows the investigation of a wider range of energy levels and materials).

Acknowledgments

Thank you for taking the time to read this addendum to my dissertation. I am pleased to have been able to see this project evolve over the course of the past few years. Though my formal dissertation acknowledgments are more thorough, I really must provide some additional recognition of the folks who were involved with the work discussed here. Doctors Andrei Pakoulev, Mark Rickard, Kathryn Kornau, and Nathan A. Mathew led the project through the

entirety of the RDC work and associated vibrational studies. Dr. Lena Yurs joined the research group when I did and worked on vibrational studies and on the picosecond quantum dot experiments. Dr. Rachel Selinsky (from the research group of Song Jin) either made or trained the person who made all of the quantum dot samples shown in this research. The quantum dot samples studied on the femtosecond system were synthesized by Dan Kohler, who also helped devise and run many of the corresponding spectroscopy experiments. Doctors Igor Stiopkin and Andrei Pakoulev worked with Schuyler Kain to build the femtosecond laser system. Schuyler also developed the software that integrated and controlled the hardware that allowed us to run our experiments. Doctors Song Jin and Robert Hamers collaborated with our group, providing us with the opportunity to work with nanomaterials with much less blindness than might otherwise have been the case. Dr. John Wright provided direction, connections, recommendations, encouragement, analysis, and perspective throughout each of these projects. The US populace, by way of the National Science Foundation and Department of Energy, provided our group with the funding required to acquire and maintain the elaborate experimental setups used. I am grateful to each of these individuals and groups for their contributions and assistance, just as I am, again, grateful to you for the time you have taken to read this document.





Figure 1—Diagram and photograph of the optical table used in non-linear spectroscopy experiments. OPAs change laser pulse color. Delay stages change the relative timing of the pulses at the sample cell position. The red line traces the path of the non-linear signal whose intensity we measure.




Figure 2—(a) Rhodium dicarbonyl acetylacetonate (RDC) molecular structure. Bonds are indicated by straight lines. (b) The symmetric (left) and anti-symmetric (right) stretching modes of the carbonyl groups. (c) The linear absorption profile of RDC carbonyl stretches. The absorption increases at the frequencies of light that are resonant with the two stretching modes.

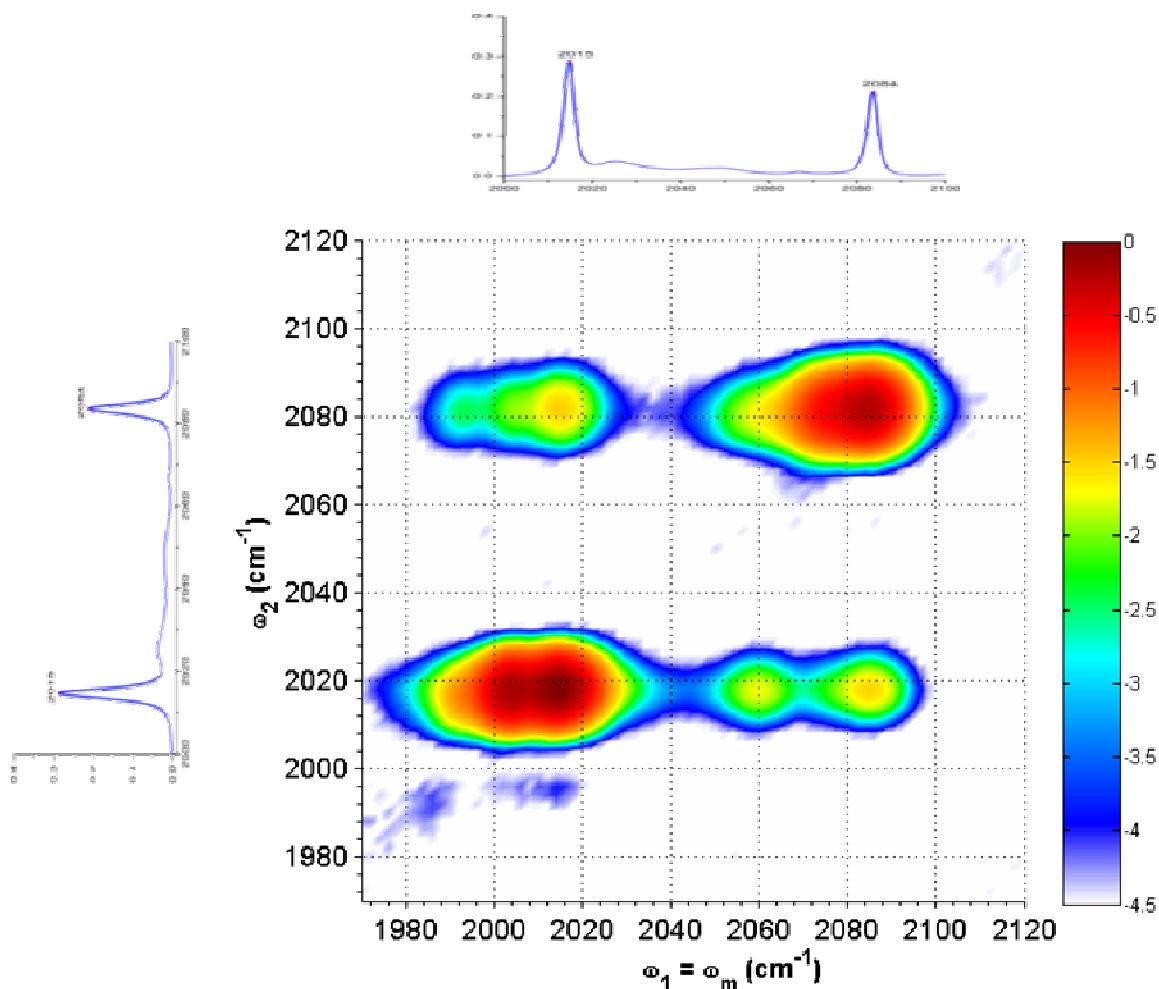


Figure 3—A multi-dimensional frequency scan of RDC with linear absorption scans juxtaposed for comparison. ω_2 and ω_1 create an excited state at the frequency shown on the y-axis. ω_1 (x-axis) interacts with that population. Signal appears when each laser is resonant with an available transition. The z-axis (see color bar to the right) shows the intensity of signal light; this plot shows a logarithmic scale. The presence of off-diagonal peaks indicates coupling between the modes. Peaks that are shifted to the left of primary peaks appear at frequencies matching the transition to overtones (same mode, vibrating faster) and combination bands (both modes vibrating).

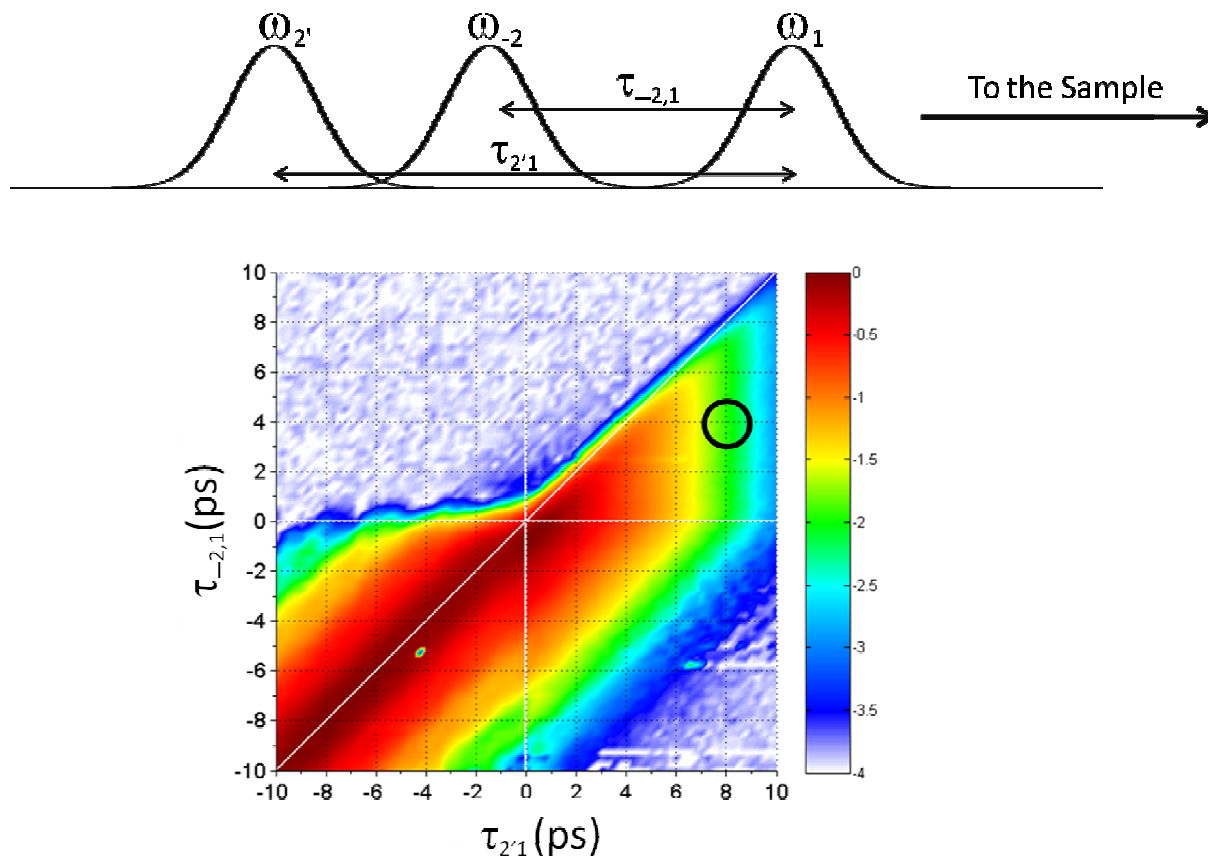


Figure 4—A multi-dimensional delay scan of RDC on a logarithmic scale. Here the laser frequencies correspond to the lower-rightmost peak in Figure 3. At the origin all the pulses are overlapped in time. The black ring shows the time-ordering shown in the diagram above the scan. Interpretation of x-axis and y-axis values is also available pictorially in the diagram. The decrease in signal in such directions as the positive x-axis traces out coherence relaxation (explained later in the main text). The steady signal along the negative $x = y$ diagonal shows the long-lived population lifetime.




Figure 5—A multidimensional delay scan with all lasers tuned to the same frequencies as those used in Figure 4, but with the monochromator set to pass a different wavelength of light. Here signal only appears when there is coherent energy transfer from one mode to another. The beating pattern has a frequency corresponding to the different between the initial and final states of the transfer.

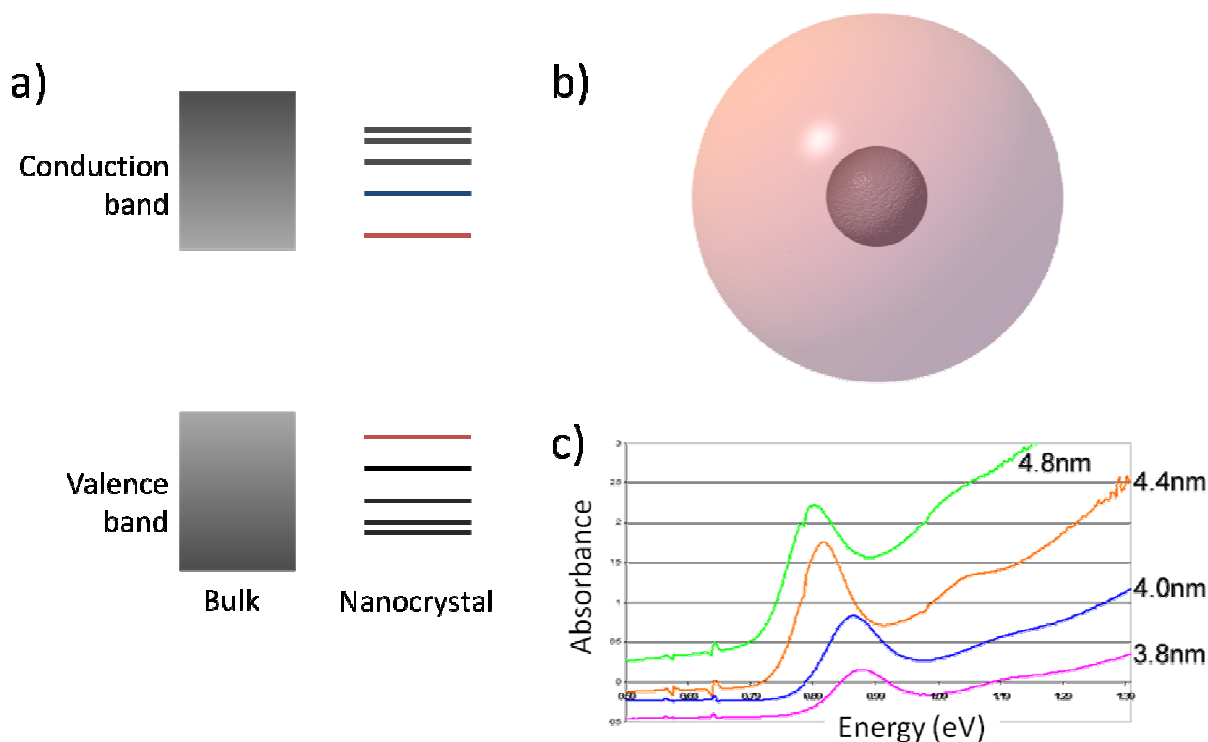


Figure 6—Depicting the effects of quantum-confinement. (a) The energetic structure of a non-confined semiconductor (bulk) and a quantum-confined structure (nanocrystal) of the same material. Notice in particular (1) the development of discrete states and (2) the increased separation between highest energy valence and lowest energy conducting states in the confined structure. (b) A graphical depiction of a quantum dot (inner sphere) compared to the native exciton orbital size (partially-transparent larger sphere). (c) Absorption scans of PbSe quantum dots of different diameters, collected by Rachel Selinsky. Notice that as the dots get smaller, the absorption profile shifts to higher energy. Greater confinement leads to higher energy states.




Figure 7—A simplified schematic of a dye-sensitized solar cell core. Here the chromophore absorbs incident sun light. The excited state electron transfers into the conduction band of the electron-accepting titanium dioxide (TiO_2). To make the transfer, the electron must pass the anchor that holds the dye molecule on the surface of the TiO_2 . The transferred electron then moves through the TiO_2 to a device and completes the circuit by eventually replenishing the chromophore.




Figure 8—Using quantum dots as the “dyes” for solar cells. (a) A microscope (transmission electron microscope) image of PbSe quantum dots grown onto the surface of iron (III) oxide, Fe₂O₃ (image taken from the work of Rachel Selinsky and Mark Lukowski). (b) Conceptual image of a quantum dot on the surface of an electron-accepting material. Recall the exciton size relative to the quantum dot and consider the ease with which an electron might relocate into the TiO₂ (the blue solid). (c) An energy level diagram corresponding to the electron transfer event. An incident photon moves an electron from the highest occupied molecular orbital (HOMO) into an excited state. The excited electron then relaxes into the conduction band of the electron-acceptor.

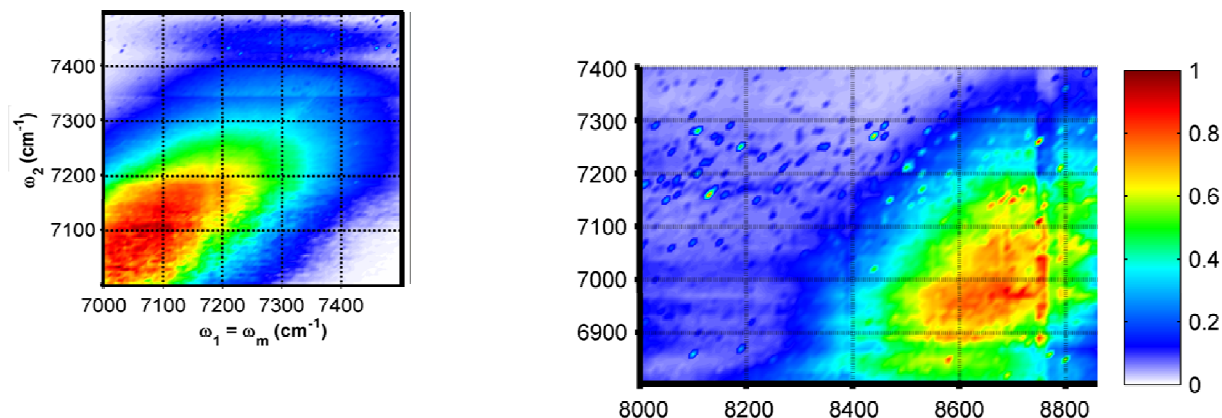


Figure 9—Multidimensional frequency scans taken of PbSe quantum dots with the picoseconds laser system. These two scans have been aligned to show their relative locations. Together they are analogous to the lower half of the scan of RDC shown in Figure 3. The peak on the left appears when all lasers are tuned to the lowest energy feature in Figure 6c (called the 1S exciton). The peak on the right is a cross peak between the lowest energy exciton feature and the higher energy hump (called the 1P exciton) on the rising hill in those absorption scans. The rising hill itself does not appear in this scan because those states are not coupled to the low energy feature, but the 1P exciton is.




Figure 10—Multidimensional delay scan of PbSe quantum dots collected with picosecond pulses. All lasers are tuned to the wavelength of the 1S exciton (see Figure 9 caption for explanation). The population lifetime traces (here shown along the negative $x = y$ diagonal and the positive x -axis) show long-lived signal, but coherence dephasing processes cannot be measured because the pulses are too long. The width of the peak indicated by the two black arrows should depend on the coherence dephasing rate, but instead it corresponds only to the width of the pulses.




Figure 11—A multidimensional frequency scan of PbSe quantum dots collected on the newer femtosecond laser system. The two peaks shown would match the top half of the scan shown in Figure 3. The feature near the diagonal appears when all lasers are tuned to the 1P exciton (see Figure 9 caption for explanation). The upper left peak appears when a population is created in the 1P exciton and the 1S exciton is coupled to the state that exists when it reaches the sample (which was 600fs later in this scan).




Figure 12—A multidimensional delay scan of PbSe quantum dots collected with 45fs pulses. As in Figure 10, lasers are tuned to the 1S exciton transition. The long population lifetimes are still apparent, but now the widths of these features can measure the coherence lifetimes—they are sufficiently wider than the excitation pulses.




Figure 13—Multidimensional frequency scans verifying the measurable contribution of fully-coherent processes. (a) The data collected. (b) A simulation of (a) with signal contributions from fully-coherent processes. (c) A simulation of (a) without signal contributions from fully-coherent processes. The purpose of comparing the data to these two simulations is that the lobe in the spectral data located at $(x,y) = (7050,6800)$ does not appear when no fully-coherent processes are allowed to contribute (c), but does appear when they are allowed (b). This subtle appearance confirms that we are able to once again access the full range of capabilities demonstrated in vibrational mode experiments like those of RDC.