

## PhD TUTORIAL

# Using feedback for coherent control of quantum systems

T C Weinacht<sup>1</sup> and P H Bucksbaum

University of Michigan, Ann Arbor, MI, USA

E-mail: weinacht@jilau1.colorado.edu

Received 19 December 2001

Published 4 April 2002

Online at [stacks.iop.org/JOptB/4/R35](http://stacks.iop.org/JOptB/4/R35)**Abstract**

A longstanding goal in chemical physics has been the control of atoms and molecules using coherent light fields. This paper provides a brief overview of the field and discusses experiments that use a programmable pulse shaper to control the quantum state of electronic wavepackets in Rydberg atoms and electronic and nuclear dynamics in molecular liquids. The shape of Rydberg wavepackets was controlled by using tailored ultrafast pulses to excite a beam of caesium atoms. The quantum state of these atoms was measured using holographic techniques borrowed from optics. The experiments with molecular liquids involved the construction of an automated learning machine. A genetic algorithm directed the choice of shaped pulses which interacted with the molecular system inside a learning control loop. Analysis of successful pulse shapes that were found by using the genetic algorithm yield insight into the systems being controlled.

**Keywords:** Coherent control, feedback, ultrafast, quantum state preparation, learning algorithm

(Some figures in this article are in colour only in the electronic version)

**1. Introduction**

A major goal in chemical physics since the invention of the laser has been to move from observing to controlling the dynamics of atoms and molecules using laser light. One of the original motivations for control is laser-selective chemistry [1, 2], in which tailored optical fields are used to selectively break bonds in molecules and initiate desired chemical reactions. This goal has helped to generate interest and activity in the field, which is now seeing many diverse applications in addition to controlling chemical reactions.

Early attempts at laser-selective bond breaking, which relied on tuning monochromatic laser light to a local mode frequency, met with little success because of rapid intramolecular energy redistribution [3]. Recent efforts have focused on using the coherent nature of laser radiation to effect

‘coherent control’ over atoms and molecules. Two major routes to coherent control were proposed in the mid-1980s. One scheme, as suggested by Paul Brumer and Moshe Shapiro, is a frequency domain approach, which describes the route to control via interfering pathways and relies on explicitly controlling the phases of the contributing pathways [4]. A second approach, proposed by David Tannor and Stuart Rice, takes a time domain view of the process, and aims to control atoms and molecules with tailored laser fields that have field strengths comparable to atomic and molecular forces, and temporal variations which are fast compared with atomic and molecular timescales [5].

At the heart of coherent control experiments is the idea of transferring the coherence inherent in laser light to quantum systems. Control over the preparation of a final state is the result of constructive and destructive interference between several mutually coherent quantum waves or paths in each atom or molecule. Frequency domain experiments

<sup>1</sup> Present address: State University of New York, Stony Brook, Stony Brook NY 11794-3800, USA. Telephone: ++ (631) 632-8100.

require two or more narrowband lasers which have a well defined and adjustable phase relative to one another. Most frequency domain experiments are performed with two lasers, where the phase between the two lasers controls the phase of two interfering pathways leading to the population of a final state, or the yield of a product. Generalization of this approach to multidimensional control requires several phase-locked lasers, which can become experimentally cumbersome. This naturally leads to a time domain approach, in which a single short pulse contains a broad continuum of frequency components whose relative phases are necessarily locked to one another. With programmable pulse shaping techniques [6], one can individually address hundreds of frequency components, and rapidly adjust their relative phases and amplitudes. This makes time domain control attractive for multidimensional control spaces, where several parameters need to be adjusted in order to effect control. Furthermore, with amplified laser systems, the intensity of a focused laser pulse can be used to increase the effective laser bandwidth far beyond its weak field value, extending the range of control even further. However, the huge parameter space afforded by intense, broadband pulses leads to a design problem of how to determine the optimal control field for a given system.

The design of appropriate fields for controlling atoms and molecules on their natural timescales is a difficult task at best. The problem is complicated by many things including decoherence, intramolecular energy redistribution and imperfect knowledge of the atomic or molecular Hamiltonian. Recently, a proposal by Rabitz and Judson suggested the use of feedback to let the atom or molecule design its own optimal field [7]. In their paper they suggest using a genetic algorithm (GA) [8] and a device capable of shaping optical fields inside a feedback or learning control loop. The shaped optical pulses interact with the system being studied and a feedback signal evaluates the performance of each pulse shape in reaching some predetermined goal, such as populating a given rotational state, or breaking a desired bond in a large molecule. Pulses that do well in reaching this goal are used to generate even better pulses as directed by the GA. At some point, the new pulses do not outperform the previous generation and the GA is said to have converged on a solution. The converged solution can be useful not only because it is more effective than an unshaped pulse at generating the desired outcome during its interaction with the system, but also because it can contain new information on the system and its evolution under the influence of the pulse. Thus, coherent control is perhaps not only a useful tool for producing new quantum states and products that are otherwise inaccessible, but it can also be used as a new form of spectroscopy, allowing one to explore the behaviour of atomic and molecular systems driven by laser fields in a new way. The experiments discussed below highlight both aspects.

### 1.1. Experimental implementations

Starting in the late 1980s, there have been a wealth of frequency domain control experiments where the phase of two laser beams has been adjusted in order to adjust the phase of two quantum paths to a single final state. In early experiments by Robert Gordon and coworkers [9], the phase of two laser

beams controlled the one-photon and three-photon paths to ionization of HI in order to control the fraction of dissociative versus nondissociative ionization. This experiment illustrated the power of coherent control in influencing the outcome of a simple chemical reaction. Later implementations of this experiment provide a demonstration of the spectroscopic information available in coherent control experiments. The phase of the oscillations in product yield versus the phase difference for the two lasers has been used to characterize details of the molecular potential, and has been dubbed phase lag spectroscopy [10].

Generalization of the two-path control concept to many varied laser-driven processes has been carried out with control of currents in bulk semiconductors [11], control of the nonlinear optical properties of atoms, liquids and solids and control of photoionization, in both weak and strong fields [12, 13].

Following the advent of frequency domain experiments, and with the development of ultrafast lasers [14], many time domain experiments have been pursued. Control over vibrations in gas phase molecules [15–20] as well as bulk crystals [21–25] and molecules in solution [26] has been demonstrated. There are many time domain approaches to control. One representative approach is the so called ‘pump–dump’ scheme, where the evolution of a molecular wavepacket is controlled with pulse sequences. In this scheme, a pump pulse generates a vibrational wavepacket on an excited state potential energy surface, and a second time-delayed ‘dump’ pulse interacts with this wavepacket to stimulate it down to the ground state surface. The transition back to the ground state will proceed with greater efficiency if the dump pulse arrives at a time when there is a favourable Franck–Condon overlap between the excited state and the ground state.

In addition to controlling molecular vibrations, there have been time domain experiments demonstrating control over molecular rotations [27, 28], excited states in quantum dots [29], exciton orientation and density in quantum wells [30, 31] and many other systems.

Time and frequency domain experiments are closely related. Work by Silberberg on multiphoton transitions in atoms using shaped femtosecond pulses illustrates the value of both time and frequency interpretations [32]. In his experiments, the frequency components of a femtosecond pulse were adjusted to control the multiphoton population transfer from one electronic state to another. The experiments may be interpreted in the time domain since they involve nonlinear processes driven by intense short pulses. On the other hand, the control mechanism is easiest to understand in the frequency domain, where the control depends on the nonlinear power spectrum [33]. The importance of interpreting experiments in both domains is one of the reasons that time–frequency spectrograms are now commonly used to represent coherent optical control fields [34, 35].

A few years after the proposal of Rabitz and Judson to close the loop on coherent control and use feedback and a learning algorithm to discover optimal pulse shapes for control, the first learning control experiment was performed by the Wilson group [36]. The goal of the feedback was to maximize the integrated fluorescence from a dye molecule in solution, pumped by an intense, shaped femtosecond pulse.

The fluorescence yield was optimized inside a learning control loop with a GA dictating the choice of pulse shapes to test. After several generations, the algorithm found pulse shapes that were more efficient for exciting the dye molecules than an unshaped, transform-limited pulse.

There have been several subsequent demonstrations of feedback for coherent control. These include, but are not limited to, selective bond cleaving in large molecules using tailored femtosecond pulses [37]; dissociation of metal clusters [38]; dissociation and rearrangement of organic molecules [39]; and high-harmonic radiation [40]. These experiments have shown that it is possible to coherently control processes and prepare desired quantum states in a large and diverse array of complex systems, even when it is difficult or impossible to calculate the field that would be required to effect control *a priori*.

The success of these control experiments has led to a growing recognition that controlling quantum systems does not necessarily produce a deeper understanding of the mechanisms for control. A more challenging goal is to gain insight into the control process by searching for and understanding nonobvious control solutions. To date there is no general approach that has been demonstrated to work on all systems without some input. However, very recent experiments and theoretical work have shown promise of gaining insight by developing the learning algorithms to not only find a solution, but also yield some information about the important aspects of the solution [41–45]. Thus, one can begin to separate the sufficient and necessary elements of the solution and understand why a solution is effective.

The experiments described in the following sections loosely follow the developments described above. Starting with a time domain approach to control of an atomic wavefunction, the experiments then close the loop on control, move to a more complicated molecular system where calculation of the the control fields is difficult and, finally, the learning algorithm is developed to gain insight into the control mechanism.

## 2. Laser system and pulse shaping

The laser used in our coherent control experiments is a chirped pulse amplifier system [46, 47], capable of producing 150 fs laser pulses with 10 mJ energy at a repetition rate of 10 Hz. Seed pulses for amplification originate from a commercial Kerr-lens mode-locked titanium sapphire oscillator, which produces 100 fs pulses with 3 nJ of energy at a repetition rate of about 100 MHz. These pulses are directed into a dispersive grating pulse expander where they are stretched in time to a duration of 120 ps. The stretched pulses are sent through a Pockel cell which selects pulses for amplification at 10 Hz. The selected pulses are amplified in a regenerative amplifier. The amplifier provides a gain of  $10^6$  in about 50 round trips, at which point the gain saturates and the pulses are switched out. The amplified pulses are split, with 10% of their energy directed into the pulse shaper and the remainder delayed for use as an unshaped reference pulse. The pulse shaper is inefficient, with a total throughput of about 10% including the diffraction efficiency of the gratings and the acousto-optic modulator (AOM). To compensate for the inefficiency of the pulse shaper,

the shaped pulses are re-amplified in a low gain multipass amplifier (three pass) to a final energy of up to 10 mJ. The re-amplified, shaped pulses are re-compressed to a minimum (transform-limited) pulse duration of 150 fs in a single-grating pulse compressor. The unshaped pulses are also re-compressed in a separate single-grating compressor after passing through a long delay stage so that they can be recombined spatially and temporally with the shaped pulses. The unshaped pulses can be used for spectral interferometry measurements as well as probing dynamics initiated by the shaped pulses.

The ability to shape ultrafast laser pulses is the most important technical aspect of this work. Programmable optical waveform control has enabled us to tailor the shape of electronic wavepackets and control the vibrational dynamics of molecules. The ultrafast optical pulses we want to shape are much faster than any current electronics and therefore shaping directly in the time domain is not feasible. Our approach to pulse shaping is a frequency filter, similar to others who pioneered the field [48, 49]. We employ a geometrical trick to access the frequency components that make up the pulse. By using a grating and lens in a spectrometer geometry, the different spectral components of the pulse can be mapped to different spatial locations in the image plane. A spatial mask placed at the focal plane of the lens modifies the amplitude and phase of individual frequency components:

$$E_{\text{shaped}}(\omega) = E_{\text{input}}(\omega) * A(\omega) * e^{i\phi(\omega)}.$$

Here  $A(\omega)$  and  $\phi(\omega)$  represent the amplitude and phase imposed on the pulse by the spatial mask with  $A(\omega) \leq 1$ . A second lens and grating recombine the frequency components, yielding a temporally shaped pulse. This device is essentially a zero-dispersion stretcher with a mask at the Fourier plane.

The spatial mask need not be fixed. If there is a way of controlling and changing the mask, then new pulse shapes can be programmed on the optical pulses by adjusting the phase and amplitude that the mask imposes on the pulse.

Our programmable mask is an AOM. The AOM filters and delays individual frequency components in the light through Bragg diffraction of the dispersed optical pulse from a periodic refractive index, formed by a travelling acoustic wave launched into a tellurium dioxide crystal. The light pulse is so short that during its transit time across the AOM, the acoustic wave is essentially frozen. The acoustic velocity in the crystal is  $4.2 \times 10^3 \text{ m s}^{-1}$ , and the light pulse spends less than 10 ps in the crystal; so the acoustic wave moves less than 42 nm which is  $0.002\lambda_{\text{acoustic}}$  during the transit of the light. Thus the temporal profile of the acoustic wave is mapped onto the optical pulse spectrum. In momentum space, the input optical  $\vec{k}$  vector adds with the acoustic  $\vec{k}$  vector to give the output optical  $\vec{k}$  vector. Because momentum is conserved in the scattering process, both the amplitude and phase of the acoustic wave are written onto the optical pulse. Modulating the phase and amplitude of the acoustic wave thus gives control over the phase and amplitude of the optical pulse.

The index variation across the optical beam wavefront due to the acoustic wave can be described by

$$n(x) = n_0 + A\kappa \cos\left(\frac{2\pi x}{\Lambda} + \delta\right).$$

Here,  $\Lambda$  is the acoustic wavelength,  $A$  is the amplitude of the acoustic wave,  $\kappa$  is a constant describing the index variation for a given acoustic wave amplitude and  $\delta$  is the phase of the acoustic wave. If the incident optical field is given by  $E = E_0 e^{i(k_z z + k_x x)} + \text{c.c.}$ , the optical wave, after passing through the AOM, can be written as

$$E(x, 0) = -E_0 e^{i k_x x} e^{i(n_0 + A\kappa \cos(2\pi x/\Lambda + \delta))} + \text{c.c.}$$

Expanding the exponential term in a power series gives the first-order diffracted field:

$$E_r(x, 0) = \frac{-iA\kappa e^{in_0}}{2} E_0 e^{i k_x x} \cos\left(\frac{2\pi x}{\Lambda} + \delta\right) + \text{c.c.}$$

Taking the +1 order we get a modulated optical field:

$$E_r(x, 0) = \frac{-iA\kappa e^{in_0}}{2} E_0 e^{i((k_x + 2\pi/\Lambda)x + \delta)} + \text{c.c.}$$

The amplitude of the acoustic wave,  $A$ , determines the amplitude of the optical wave while the phase of the acoustic wave,  $\delta$ , determines the optical phase.

AOM-based pulse shapers have several departures from ideal behaviour. Most of these can be compensated for by changing the acoustic pulse shape [49]. There are additional difficulties associated with amplification of shaped pulses for strong field experiments. Shaped pulses are distorted by saturation or nonlinear refraction in laser amplifiers. On the other hand, the pulse shaper has damage limits of its own. A solution is to insert the pulse shaper between the two amplification stages. This way one can produce shaped pulses with at least 1 mJ of energy, and avoid damage to either the pulse shaper or the amplifier while maintaining pulse shape fidelity. The pulse shaper follows the regenerative amplifier and its input is limited to 100  $\mu\text{J}$ . The output pulses are amplified in a low-gain multipass amplifier that produces the required pulse energy without saturation.

Spectral interferometry was employed to characterize the amplitude and phase of the shaped pulse [50]. This technique combines the shaped laser pulse with an unshaped reference on a beamsplitter and sends the combined pulses into a spectrometer. The spectrum is a frequency-resolved interference pattern which yields the phase of the shaped pulse if the phase of the reference pulse is known.

### 3. Atomic wavepackets

Rydberg atoms are an excellent testing ground for quantum classical correspondence and for quantum state preparation and measurement. The monograph by Gallagher describes Rydberg atoms and their properties [51]. Our ability to tailor ultrafast laser pulses allowed us to create shaped Rydberg wavepackets and develop new measurement techniques for dynamical quantum states. It also enabled us to start controlling dynamics in addition to observing them.

This section deals with ‘wavepacket sculpting’ and the development of a ‘quantum camera’. Wavepacket sculpting refers to the production of wavepackets whose phase and amplitude are tailored to produce a predetermined quantum state. Our quantum camera makes use of a holographic technique for wavepacket measurement, producing a complete

picture of the complex quantum wavefunction. This quantum photography might also be called quantum holography, because it uses interference between two wavepackets (one whose shape is known) to measure the phase as well as the amplitude of the constituent eigenstates.

Our experiments on Rydberg wavepackets take place inside a vacuum chamber where multiple laser beams interact with an effusive beam of caesium atoms. The wavepackets are measured by field ionizing the atoms and analysing the time of flight electron spectra. Prior to the arrival of the shaped laser pulse, we excite the ground state atoms to the 7s ‘launch’ state using a Raman shifted dye laser pulse centred at 1.079  $\mu\text{m}$ . A few ns later, the atoms are irradiated by the shaped laser pulse which excites a coherent superposition of p states with principal quantum number between 24 and 35. This is the Rydberg wavepacket.

The excitation of the Rydberg wavepacket from the launch state can be described by first-order perturbation theory using the usual dipole operator. In this approximation, the transition rate to a given eigenstate  $\psi_i$  from the 7s state with a broadband pulse is

$$w_i = \frac{\lambda_{laser}^3 \rho(\nu_0) A_e}{8\pi h}.$$

Here  $\rho(\nu_0)$  is the energy per unit bandwidth per unit volume.  $A_e$  is the Einstein A coefficient:

$$A_e = \frac{4\omega_o^3 \pi^2}{3\alpha^2} \left| \int dr \psi_i(r) \psi_{7s}(r) r \right|^2.$$

The wavefunctions necessary for calculating the dipole matrix element and also for reconstructing the wavepacket wavefunctions can be calculated numerically using quantum defect theory and a Numerov integration routine [51]. They are very close to hydrogenic wavefunctions, which can be described analytically, but they have a spatially dependent phase shift due to the deviation of the ion core potential from a pure Coulomb potential.

The dipole matrix elements for transitions between Rydberg states and the ground or low  $n$  states are generally very small because the Rydberg state wavefunctions have almost no overlap with the core, where all of the ground state wavefunction is concentrated. This is why they have such long lifetimes and also why the relatively high energy pulses that couple the 7s state to the Rydberg series still act only perturbatively.

For a pulse with a bandwidth of 3 THz and an energy of 1 mJ, focused to a spot size of 1  $\text{mm}^2$ , the total excitation probability for the  $n = 26$  Rydberg state is 0.016, assuming unit probability of being in the 7s state. Tighter focusing of the short pulse laser can however easily carry the interaction into the strong field regime where multiphoton ionization is observed.

#### 3.1. Making and measuring Rydberg wavepackets

The Rydberg wavepacket,  $\Psi_{signal}$  created by the shaped laser pulse at time  $t = 0$  can be written as

$$\Psi_{signal}(t = 0) = \sum_i a_i \psi_i(r, \theta, \phi) \quad a_i = |a_i| e^{i\delta_i} \quad (1)$$

where  $\psi_i(r, \theta, \phi)$  are the constituent eigenstates.

$\Psi$  is not an eigenstate. It evolves in time according to

$$\Psi(\mathbf{r}, t) = \sum_i a_i \psi_i(\mathbf{r}, \theta, \phi) e^{-i\omega_i t}. \quad (2)$$

The stationary Rydberg states  $\psi_{nlm}$  are real functions of  $\theta$  and  $r$  up to a global phase factor  $e^{i\delta}$ . Here  $(n, l, m)$  are the usual quantum numbers specifying the energy of the state, its angular momentum and  $z$ -component<sup>2</sup>. The azimuthal dependence of  $\psi_{nlm}$  is a complex exponential,  $e^{im\phi}$ , and will not concern us here. If the eigenstates and eigenvalues are known or can be measured, then the wavepacket is completely specified by measuring the amplitude and phase of the complex coefficients  $a_i$ . The phase information in these complex coefficients was not available using previous detection methods for the wavepacket.

The amplitude  $|a_i|$  can be controlled through the amplitude of the laser frequency component resonant with the  $i$ th state, and the relative phase of the state,  $\delta_i$ , can be adjusted using the phase of that same frequency component of the laser pulse. In other words, the Rydberg wavepacket is sculpted by transferring the coherence and control over the optical field to the atom.

The amplitude of the  $a_i$ 's can be analysed by state-selective field ionization (SSFI) [52], since each nondegenerate eigenstate in the wavepacket has a unique field ionization threshold. The critical field for classical field ionization is

$$F_{crit} = \frac{1}{16n^{*4}}.$$

Here,  $n^*$  is the principal quantum number for each of the states in the wavepacket including the quantum defect:

$$n^* = n - \delta_l.$$

By applying an adiabatically increasing electric field with time, the potential that the valence electron sees can be tilted so that each eigenstate in the coherent sum experiences  $F_{crit}$  at a separate time and therefore ionizes at a separate time.

The electron time of flight spectrum contains the information necessary to infer the population of each state in the wavepacket for an ensemble of identically prepared atoms. Working with an ensemble of identically prepared atoms is equivalent to performing the same experiment many times on a single atom, as would be required to determine the populations for each eigenstate with a single atom being excited each time.

A second approach to measuring the  $a_i$ 's is wavepacket interferometry, also called the optical Ramsey method [53]. Here two wavepackets are summed coherently by exciting the atom with two identical light pulses with a variable time delay  $\tau$ . If the usual weak field condition for time-dependent perturbation theory applies, then the wavefunction is given by

$$\Psi(\mathbf{r}, t, \tau) = \sum_i (a_i \psi_i e^{-i\omega_i t} + a_i \psi_i e^{-i\omega_i(t+\tau)} e^{i\omega_{gs}\tau}) \quad (3)$$

<sup>2</sup> The Schrödinger wavefunction shows states of pure  $l$ , neglecting the spin-orbit interaction. This is justified, because our measurements take place on a timescale (2 ps) that is short compared with the inverse frequency of the spin-orbit splitting in these states of the atom (about 60 ps in our case). Therefore we are both making and probing superposition states in the atom that correspond to pure  $l$ -states, not  $j$ -states. The same comment can be made to justify the neglect of the hyperfine interaction and the total angular momentum  $\mathbf{F} = \mathbf{I} + \mathbf{J}$ .

where  $e^{i\omega_{gs}\tau}$  represents the phase advance of the initial state during the delay time  $\tau$ . The total excited state population  $\langle \Psi(\mathbf{r}, t, \tau) | \Psi(\mathbf{r}, t, \tau) \rangle$  oscillates with the time delay  $\tau$ :

$$\langle \Psi(\tau) | \Psi(\tau) \rangle = \sum_i 2|a_i|^2 (1 + \cos(\omega_{gs} - \omega_i)\tau). \quad (4)$$

The spectral amplitudes can be extracted by Fourier analysis of this interference function, but the phase information is missing; therefore the shape of the wavepacket cannot be determined.

The phase problem in Rydberg wavepackets has an analogue in ultrafast optics. There are several techniques to determine the spectrum and phase of the electric field in a sub-picosecond laser pulse [54]. A recent paper suggested that *spectrally resolved cross-correlation*, the basic method used in spectral interferometry, was also directly applicable to the phase problem in wavepackets [55].

Cross-correlation is similar to optical interferometry, where the signal received by a detector is the coherent sum of two optical fields,  $E(t)$  and  $E'(t)$ , combined on a beamsplitter. In conventional Michelson interferometry the two fields are the same except for a time delay  $\tau$ , and so the detector sees an autocorrelation of the electric field superimposed on a background:

$$\begin{aligned} S_{auto}(\tau) &= \int |E(t) + E(t + \tau)|^2 dt \\ &= 2 \int E(t)E(t + \tau) dt + 2 \int |E(t)|^2 dt. \end{aligned} \quad (5)$$

Note that all phase information is lost, and that  $S_{auto}$  has a similar formal appearance to the wavepacket signal in equation (4). In a cross-correlation, the second pulse is replaced by a reference field whose properties are known:

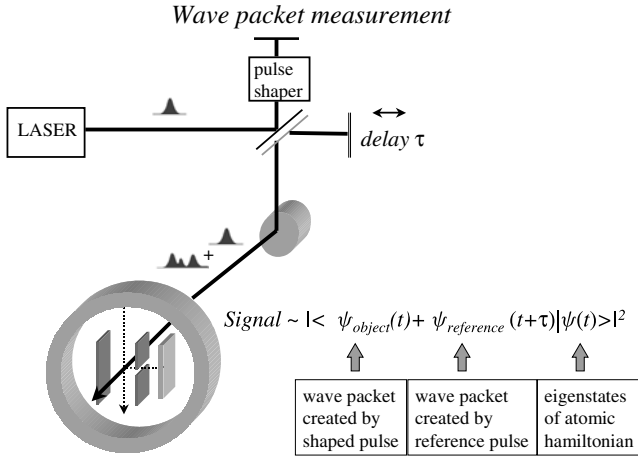
$$\begin{aligned} S_{cross}(\tau) &= 2 \int E(t)E_{ref}(t + \tau) dt \\ &+ \int [ |E(t)|^2 + |E_{ref}(t)|^2 ] dt. \end{aligned} \quad (6)$$

The first term is then a convolution integral, which can be Fourier decomposed to reveal the phase difference between each frequency component of  $E$  and  $E_{ref}$ . The phase of the reference light pulse can be analysed using methods such as frequency-resolved optical gating (FROG) [54].

When 'reference' and 'signal' quantum wavepackets interfere, the phase difference between each pair of eigenstates in the signal can be resolved through the temporal interference pattern of each state in the spectrum. Thus, spectral interferometry can be applied to quantum mechanics. This same measurement can be described as quantum holography as a recent paper by Leichtle *et al* [56] points out. The measurement of the phase of the quantum state of the Rydberg electron through interference with another known state is analogous to the measurement of the phase of an unknown light field through interference with a well characterized one. In our case, the interference is spectrally resolved, which is why the description in terms of spectral interferometry is useful.

After the creation of the shaped wavepacket with the initial pulse, a second reference wavepacket is created in each atom at time  $t = \tau$  with the unshaped reference laser pulse:

$$\Psi_{ref}(t = \tau) = e^{i\omega_{gs}\tau} \sum_i b_i \psi_i \quad b_i \in \Re. \quad (7)$$



**Figure 1.** Schematic diagram of wavepacket measurement scheme.

Figure 1 shows a diagram illustrating the creation of the signal and reference wavepackets.

The reference pulse contains a small amount of high-order dispersion; however, the phase of the spectrum over the bandwidth of the experiment is constant to within  $\pm 0.5$  rad, or about 0.05 rad between neighbouring  $n$ -states. In the absence of decoherence, the resulting total wavefunction can be written as

$$\Psi_{\text{total}}(t = \tau) = \sum_i (a_i e^{-i\omega_i \tau} + b_i e^{i\omega_{gs} \tau}) \psi_i. \quad (8)$$

The measured quantity for each laser shot is the total population ( $P_i$ ) in each Rydberg state as determined from the SSFI spectra:

$$P_i = |a_i|^2 + |b_i|^2 + 2|a_i||b_i| \cos((\omega_i - \omega_{gs})\tau - \delta_i). \quad (9)$$

From the  $P_i$ , we wish to extract the relative phases between the states in the wavepacket:  $\delta_{ij} = \delta_i - \delta_j$ .

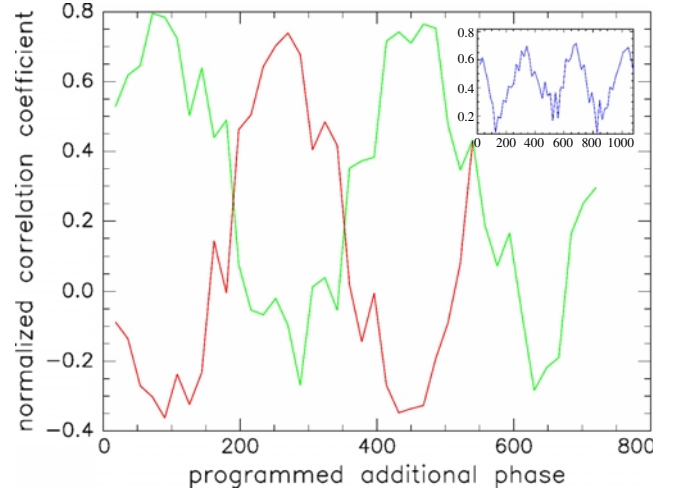
$P_i$  oscillates at the optical frequency. The time delay  $\tau$  should therefore be stable to much better than the optical period of about 2 fs, which is not practical in a system with amplifiers and pulse shapers in separate delay lines. Furthermore, the carrier frequency in the acoustic waveform in the AOM is not phase locked to the oscillator repetition rate, resulting in a random phase shift on the signal pulse relative to the reference pulse even in the absence of jitter in the time delay. Therefore, over several laser shots, the cosine term in  $P_i$  averages to zero, so the average values  $\langle P_i \rangle$  have no phase dependence:

$$\langle P_i \rangle = |a_i|^2 + |b_i|^2. \quad (10)$$

To extract the phase despite the time delay instability, we construct the normalized correlation function between each state  $i$  in the wavepacket and a chosen reference state  $j$  [57]:

$$r_{ij} = \frac{\langle P_i P_j \rangle - \langle P_i \rangle \langle P_j \rangle}{(\Delta P_i)(\Delta P_j)}. \quad (11)$$

We have also generalized this correlation function to a correlation map between all states in the wavepacket. This makes use of more information to give a more accurate phase measurement.



**Figure 2.** Measured correlation functions as a function of additional phase,  $\beta$ , programmed (in degrees) onto the reference state. The two curves in the main part of the figure are for two states which have been programmed to be  $\pi$  out of phase and the measurement was performed with a few ps delay between reference and shaped pulses. The inset shows the measured correlation function for a state in a wavepacket which was measured 1 ns after it was created.

Even after averaging over the optical frequency, the product  $\langle P_i P_j \rangle$  still has a dependence on the phase difference  $\delta_i - \delta_j$ :

$$\langle P_i P_j \rangle = (|a_i|^2 + |b_i|^2)(|a_j|^2 + |b_j|^2) + 4|a_i||b_i||a_j||b_j| \cos((\omega_i - \omega_j)\tau - \delta_{ij}). \quad (12)$$

Inserting this into the correlation function yields a simple expression for  $\delta_{ij}$ :

$$r_{ij} = \cos((\omega_i - \omega_j)\tau - \delta_{ij}). \quad (13)$$


This is a measure of second-order coherence of the wavefunction. Thus, our measurement can be viewed as a Hanbury-Brown–Twiss interferometric measurement of the wavepacket [58].

In order to make our second-order coherence measurements, the AOM is programmed to linearly increment the phase between each eigenstate  $n_i$  and a reference state  $n_j$  by an amount  $\beta = 0$  to  $4\pi$ , in 40 steps. Then  $r_{ij}(\delta_{ij}, \beta)$  is fitted to a cosine function to extract the initial phase  $\delta_{ij}$  with greater accuracy and precision than possible from a single measurement of the correlation function. Figure 2 shows the measured correlation functions for two states which are programmed to be  $\pi$  out of phase with each other. The inset shows the measured correlation function when the reference pulse is delayed by 1 ns, demonstrating that some coherence is still present after 1 ns.

The fitting routine is a nonlinear least squares fit with three parameters: the phase, amplitude and offset of the cosine wave. The positive vertical offset of the cosine curves from zero is partially associated with laser power fluctuations. The laser power fluctuations result in shot-to-shot population fluctuations that affect all states alike, therefore increasing the correlation between states. If we did not vary  $\beta$ , then we would not be able to distinguish this contribution to the correlation function from the quantum phase that we want to measure. The standard error in the fit for  $\delta_{ij}$  is the phase uncertainty listed in table 1.

**Table 1.** Measured amplitude and phase in degrees for each eigenstate in the wavepacket shown in figure 3. Amplitude errors are nonstatistical estimates of maximum systematic deviations during data collection. Phase errors are given by the standard deviation in fitting the phase to the acquired cosine curves.

| $n$ state | Amplitude        | Phase (degrees)  |
|-----------|------------------|------------------|
| 24        | $0.148 \pm 0.02$ | $34.37 \pm 6.0$  |
| 25        | $0.158 \pm 0.02$ | $115.2 \pm 1.0$  |
| 26        | $0.160 \pm 0.02$ | $78.2 \pm 0.7$   |
| 27        | $0.180 \pm 0.02$ | $-165.0 \pm 1.0$ |
| 28        | $0.118 \pm 0.02$ | $0.0 \pm 0.0$    |
| 29        | $0.158 \pm 0.02$ | $-70.5 \pm 0.7$  |
| 30        | $0.125 \pm 0.02$ | $-0.1 \pm 0.7$   |
| 31        | $0.010 \pm 0.02$ | $74.5 \pm 0.7$   |



**Figure 3.** Schrödinger wavefunction for measured wavepacket with phases and amplitudes for the constituent eigenstates given in table 1. Phase is described by shading, and amplitude by height. The bottom right axis is the atomic  $x$ -axis and the bottom left axis is the atomic  $z$ -axis. The laser was polarized along the  $x$ -axis. Only one quadrant of the  $x$ - $z$ -plane is shown.

Lifetime limitations on the coherence time of the wavepackets are far greater than the measured lower bound of 1 ns. For Na Rydberg states, the measured lifetime for the  $n = 17p$  state is  $11.4 \mu\text{s}$ , in the presence of blackbody radiation at 300 K [59]. Since the radiative lifetime  $\tau \propto n^3$  the states in the wavepackets that we created had radiative lifetimes greater than  $10 \mu\text{s}$ . Collisional dephasing even at pressures exceeding the pressure inside the oven is negligible on a ns timescale.

With knowledge of the amplitude and relative phase of each state in the wavepacket, the electronic wavefunction can be completely reconstructed. Figure 3 shows the Schrödinger wavefunction for one of the radial wavepackets we created and measured in the laboratory. The compositions of the measured wavepacket in figure 3 are given in table 1.

This measurement method can be extended to any wavepacket whose constituent eigenstates are nondegenerate, and whose eigenvalues are known or can be measured. Our wavepackets were made with optical pulses shaped by a spectrally resolved AOM, but the wavepackets could be made in any way that can be synchronized to a reference. Since the measurement is made by observing covariant fluctuations (COIN technique), the synchronization need only match the timescale of the motion of the wavepacket itself; interferometric stability down to an optical period is not necessary.

### 3.2. Implementing feedback

With complete information about the wavefunction, we can implement simple feedback to reshape the wavefunction to a desired target shape at any arbitrary delay between reference and shaped laser pulses. The difference between the measured wavefunction and the desired target state is computed and used to reprogram the phase of each state via the pulse shaper. Full reconstruction information eliminates the need for an iterative algorithm to implement feedback, so there is very rapid convergence to any target. In our simple feedback algorithm, the relative phase of each projection is adjusted by an amount equal to the difference between the measured phase and the target phase. The routine for extracting the relative phases between states from the measured correlation functions is taken from [60]. For any given delay  $\tau$ , the algorithm is able to control the shaped wavefunction to match the target within two iterations of the feedback loop. Nonlinearities in the pulse shaper response and other technical control problems are also corrected automatically by feedback.

This simple feedback system can also stabilize a quantum ensemble in the presence of changing conditions. For example, we showed that the wavepacket output could be fixed, even when slow changes were placed on the input by changing the delay  $\tau$  several times after the algorithm converged. The only limitation is the bandwidth of the feedback loop, which cannot be greater than the pulse repetition rate. Figure 4 shows the shaped wavepacket following each iteration.

The top frame shows the target wavefunction. All other frames show measured wavefunctions at the time the reference pulse was incident on the atoms. Each row starts with a different delay setting, and each new column shows the next iteration in the feedback loop.


### 3.3. Phase space wavefunction representations

The Schrödinger representation of the measured wavepackets in coordinate space is complete, i.e. all of the information about the quantum state is contained in the wavefunction. However, it does not provide information about all aspects of the wavepacket in a way that is easy to interpret. The position space wavefunction does not give direct insight into the wavepacket's motion at any given point in its orbit. The momentum space wavefunction is easily calculated from the position space representation, but it is not easy to glean information about the wavepacket's position by looking at the momentum space wavefunction. A particularly useful alternative representation is the Wigner function [61], and its close relative, the Husimi function [62]. The Husimi function is analogous to a phase space distribution function for a classical particle with wave mechanical limits on simultaneous position and momentum localization. It provides the viewer with an idea of the behaviour of the wavepacket in both position and momentum at the same time without violating the uncertainty principle. It can be constructed by transforming the position space wavefunction  $\Psi(x)$  according to

$$Q(x, p) = \int dp' \int dx' W(x', p') e^{-(x-x')^2 - (p-p')^2}$$

where

$$W(x, p) = \int dx' \Psi(x) \Psi(x - x') e^{-2\pi i p x'}$$



**Figure 4.** Schrödinger wavefunction for wavepackets reshaped with feedback to create target shape for arbitrary delays between shaped and reference pulses.

The Husimi function for the wavepacket in figure 3 is shown in figure 5. In exchange for being positive definite, the Husimi distribution loses the property of its marginals being the position and momentum space probability densities as they are for the Wigner function.

The Husimi distribution for the electronic wavepacket shown illustrates the close analogy between the quantum wavepacket we created and measured and a classical electron with a planetary orbit. The electron is localized and it is moving toward the atomic core, similar to the way one would imagine a classical electron with an uncertain energy heading toward the core.

Extensions of this work include measurement of wavepackets created by sub-single-cycle electromagnetic field pulses [63] and implementation of quantum algorithms [64]. The feedback loop that we established and tested in the experiments described in this section is the starting point for the work described in the next section.

#### 4. Learning control

The wavepacket sculpting and feedback experiments suggest the following question: under what circumstances can we find a pulse shape capable of achieving a particular state of a quantum system? For atomic Rydberg systems, the answer is known: the target state can be projected onto the Rydberg eigenstate basis to find the amplitude and phase required for each state. In principle, any state determined in this way can be prepared. In practice, there are limitations imposed by excitation selection rules and the range of optical spectrum available for shaping.

The problem is harder for molecules, or for systems that are not isolated from the environment. Then we may not even

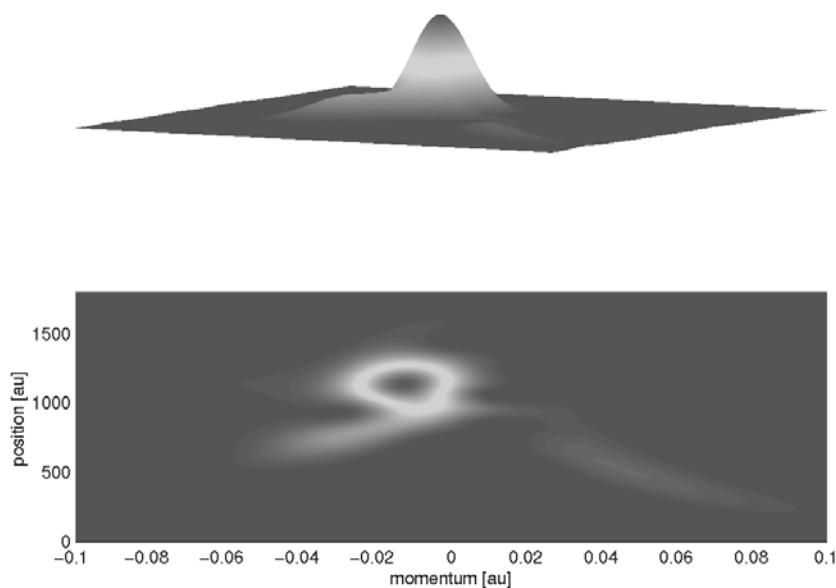
know the Hamiltonian well, let alone the eigenstates that span the relevant state space. Often in such cases, the best way to find the optimal pulse shape for a particular target is to search the results of many experimental trials. The process of refining the search based on results of previous trials is called ‘learning control’ [7].

The basic goal of this work was to construct and demonstrate a learning machine capable of controlling quantum systems, where the route to control is not known *a priori*. The learning algorithm must select new pulse shapes based upon the success of previous ones. The problem is to find an algorithm for the feedback that searches the enormous phase space the pulse shaper can access. Since the pulse shaper can control the amplitude and phase of well over 100 different frequency components, with eight-bit resolution per frequency component (in principle), there are over  $200^{100}$  different pulse shapes to try before one has exhaustively searched the phase space. The effect of each frequency component is not independent of the phase and amplitude of the others so they cannot each be optimized independently. One therefore requires an algorithm that is capable of handling many degrees of freedom which are inter-related and capable of optimizing them all simultaneously. Further requirements are that the algorithm be robust in the face of experimental noise and that it be capable of escaping local optima in a ‘rough’ potential energy landscape. A GA fulfills all of these requirements. It is a nonlocal search algorithm that can operate in a multi-dimensional phase space in the presence of noise. A brief description of the algorithm we employed is given below.

##### 4.1. Genetic algorithm

The GA starts with a random collection (population) of pulse shapes (individuals). (The community that studies search





**Figure 5.** Husimi distribution for the measured wavepacket. Two views are shown to emphasize that the distribution is positive definite. The bottom view is the same as for Wigner distributions shown above. The  $x$ -axis is radial momentum in atomic units and the  $y$ -axis is radial position in atomic units.

algorithms has developed its own terms of art, introduced here in parentheses.) Each individual is specified uniquely by a number string (genetic code). Our genetic code is the spectral phase  $\Phi(\omega)$  and amplitude  $A(\omega)$  of the pulse sampled at  $N$  and  $M$  different colours, respectively. The values of  $N$  and  $M$  can be different and can be set at the start of each experiment, with the phase linearly interpolated between these points and the amplitude set in discrete regions. Each individual is evaluated and given a score (fitness). The evaluation consists of performing the experiment: the laser pulse shines on the sample and the final state of the system is measured. This measure is converted into a single value which is the fitness. Each individual is ranked according to its fitness after evaluation. The next step in the GA is the combination of the fittest individuals to produce new ones (mating). The new set (generation) of individuals then goes through the same procedure of evaluation, ranking for fitness and mating.

There are several schemes for selecting individuals for mating; we use one called roulette wheel selection. In this scheme, each individual is selected for mating with a probability proportional to its fitness. The mating is traditionally a mixture of genetic mixing (crossover) and genetic alteration (mutation), both of which are described below. We have introduced an adaptive algorithm to determine the best mating mechanism. We evaluate the mating procedure following the fitness ranking in each generation. This helps speed up convergence, and may also help shed light on the control mechanism. Our GA uses a whole set of mating operators, which compete against each other for the chance to produce new pulse shapes. The combination of operators producing new pulse shapes evolves as the algorithm converges. Each operator is given an operator fitness which is the basis of operator selection: each operator is chosen to produce new pulse shapes with a frequency proportional to its own fitness. Operators are given credit toward their fitness only when they create pulse shapes that are better than the best of the last generation. Thus, operators that produce excellent

pulse shapes are given more opportunity to produce pulses than operators that do not improve the fitness of the pulses that they create by a substantial amount.

The operators separate into two classes: intuitive or physically motivated operators; and statistically motivated operators. Statistical operators generate new pulse shapes from old ones without any systematic physical insight into the generation process. Intuitive operators use insight into the problem to associate a physically relevant change to the pulse shape with the mating procedure. Intuitive operators include smoothing, time domain crossover and polynomial phase mutation. Statistical operators include mutation, two point crossover, average crossover and creep.

Crossover is an operator that creates two new individuals (children) by exchanging portions of the gene string for two old individuals (parents). We use two-point crossover, which snips the gene string at two random locations. Crossover is usually effective at the beginning of the algorithm, when there is ample genetic diversity. It becomes less effective as the GA converges. Average crossover works similarly, by averaging the gene values between randomly selected points in the gene string. This is useful near convergence when parents chosen for mating lie in the same region of phase space, so that averaging the gene values is more meaningful. Mutation creates a single child from a single parent by assigning random numbers for the values of selected genes. Initially, mutation is not effective since at the beginning there is already a large amount of genetic diversity if the population is randomly initialized. However, as the GA nears convergence, the mutation operator becomes important for maintaining genetic diversity. Creep is like mutation, but it reassigns selected gene values incrementally, by a small random amount. Creep is most useful near convergence where it is helpful in fine tuning the solutions that the GA has found.

The smoothing operator creates a new pulse shape from a single parent by performing a windowed smoothing over the gene string. This operator works very well for problems that

require smooth phase profiles across the bandwidth and it also aids in the interpretation of the results because it produces pulses that are not plagued as much by the entropically driven variations that arise from the GA. Smoothing the phase profile shortens the pulse and reduces structure at long times. Time domain crossover transforms the gene string into a time domain representation of the pulse by performing a fast Fourier transform (FFT). It then performs two-point crossover and transforms back to the frequency domain. This operator is extremely useful for problems that are sensitive to time domain correlations in the pulse, such as high-order processes that depend on  $I(t)$ . We have shown that the time domain crossover operator performs well in these situations [44]. Polynomial phase mutation produces children by selecting a portion of the gene string and replacing it with a polynomial plus a series of random numbers. This operator and smoothing work well in conjunction to produce pulses with polynomial phase. A more detailed mathematical description of the action of each operator on parents to produce children has appeared separately [44].

The fitness of the different operators can shed light on the dynamics of the GA and the problem itself. Operators that perform well produce genetic material that is valuable in solving the problem being addressed. By comparing the success of different operators, one can start to see what pulse shape characteristics are important in solving the problem. As one gains more and more insight into a given problem, one can start testing possible mechanisms for control by introducing intuitive operators and evaluating their success. Figure 9 illustrates the evolution of different operators in an experiment where we were optimizing second-harmonic generation. In this experiment, we understand the nature of the nonlinear interaction between the system and the light field so we could introduce physically relevant operators and watch their fitness evolve.

A multidimensional search is simplest if the search parameters are decoupled, i.e. they form an orthogonal basis. The problem then reduces to a series of one-dimensional searches. The GA power is best put to use in the opposite limit: when the experimental knobs are all coupled, so that different degrees of freedom are interdependent. The epistasis is a quantitative measure of the degree to which the different degrees of freedom are coupled. Problems in which there is a high degree of coupling have a high epistasis.

One way to extract more information about the system might be to search for a new basis that minimizes the epistasis, during or after the GA converges. We have not yet implemented this because of the low repetition rate of our laser system. This may be a viable approach with a kHz repetition rate laser system.

#### 4.2. Controlling second-harmonic generation in BBO

Frequency doubling in a noncentrosymmetric crystal with a large  $\chi^{(2)}$  provided many excellent opportunities for improving our understanding of the GA and its role in the learning process. Even though the interaction of the light field with this system can be described classically, this experiment illustrates features that are relevant to all the experiments described in this section.

This experiment explored two distinct regimes separated by the laser intensity. At low intensity the interaction is well

understood and easily modelled as a second-order nonlinear process. At higher intensity the system behaviour departs from this simple model. We explored the deviation from the simple model using the GA.

In the low intensity regime, the interaction of the laser pulses with the crystal can be described by a nonlinear polarization that is proportional to the square of the input light field:

$$P_{NL}(2\omega) = \chi^{(2)} E(\omega)^2.$$

This nonlinear polarization acts as a source or driving term in the wave equation for a field at  $2\omega$ . The experimental feedback signal was the integrated second-harmonic intensity:

$$\text{Signal} = \int dt E_{2\omega}^2(t) \propto \int dt E_{\omega}^4(t).$$

If the input field strength is not too large, this description of the interaction is adequate and gives accurate predictions for the second-harmonic generation (SHG) without including other nonlinear effects. This allowed us to simulate the experimental feedback signal and compare the GA runs on the model with the actual experiment. This is important for several reasons. One is that it allowed us to isolate features in the solutions that were a result of our GA encoding from experimental effects. Secondly, it allowed us to quickly test operator ideas and implementations, and finally, it allowed us to easily determine if our experimental model provided a sufficient description of the interaction between the laser and the system by comparing the experimental solutions with the simulation. For all of the frequency doubling experiments and simulations, we kept the laser pulse energy fixed and allowed the GA to only determine how to distribute the energy in time. The amplitude of each frequency component was kept fixed. Only the phase was allowed to vary. One key lesson we learned from comparing experiment and simulations is that whenever an experiment involves minimization of the laser intensity while maintaining a fixed energy, the resolution limits of the pulse shaper impose an artificial structure on the intensity of the optimal pulses. In order to look at the optimal pulse shapes in experiments with the GA, we constructed Husimi distributions as for the atomic wavefunctions. This allowed us to look at the energy distribution of the optimal pulses as a function of frequency and time simultaneously, which provided more insight than the electric field as a function of time or frequency. The Husimi distribution,  $Q(t, \nu)$  is calculated from the measured field  $E(\omega)$  in the frequency domain:

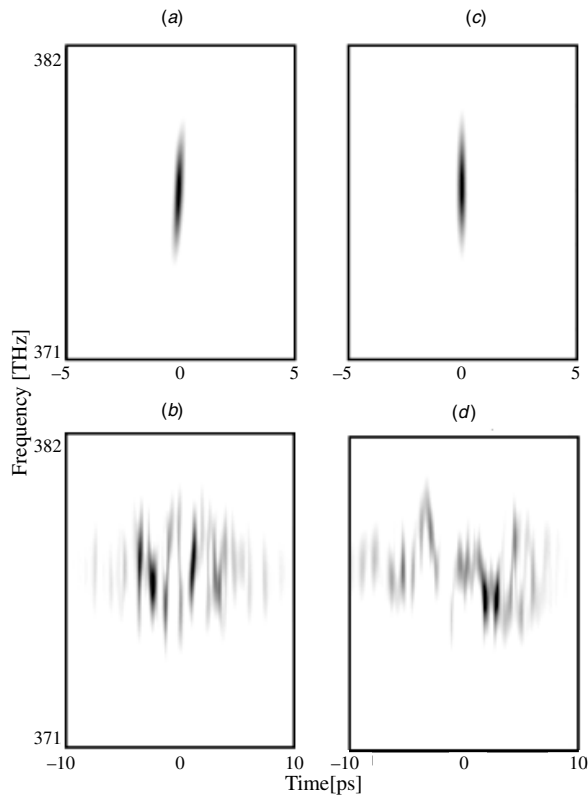
$$Q(t, \nu) = \int \int dt' d\nu' S(t', \nu') * e^{-(\nu-\nu')^2 - (t-t')^2}$$

$$S(t, \nu) = \int E(\nu) * E(\nu - \nu') * e^{2i\nu't} d\nu'.$$

$S(t, \nu)$  is the Wigner function whose marginals represent the power spectrum  $P(\nu)$  and the temporal intensity,  $I(t)$ :

$$\int d\nu' S(t, \nu') = I(t)$$

$$\int dt' S(t', \nu) = P(\nu).$$




**Figure 6.** Husimi distributions for optimal pulse shapes in experiments and simulations for frequency doubling in BBO. Panel (a) shows the optimal pulse for maximizing SHG experimentally and (c) shows the simulation. Panel (b) shows the optimal pulse for minimizing SHG experimentally and (d) shows the simulation. The value of the Husimi function is proportional to the darkness of the shading.

Figure 6 shows Husimi distributions for pulses optimized to maximize and minimize frequency doubling in BBO. The results of simulations and experiments are both shown for comparison. Comparison of panels (a) and (c) reveals that experiment and simulation arrive at the same result for maximizing SHG except for a small amount of quadratic phase (linear chirp) evident in the experimental result. The amount of chirp in the experimental solution varied within our measurement resolution limits, depending on laser alignment.

Both solutions from the simulation and experiment contain structure in  $I(t)$  that one not might expect if minimization of the integrated intensity is optimal as is the case in minimizing SHG. Furthermore, one may expect that the solutions for the spectral phase  $\phi(\omega)$  would contain only the lowest nontrivial order ( $\phi(\omega) = k\omega^2$ ) since this would be the most efficient way to programme rapid phase variations across the pulse which spread in time. A phase that varies linearly with frequency simply produces a global time delay of the pulse which has no physical significance. For a given maximum phase change between frequencies of the light pulse, quadratic phase is the most efficient of single-order phase variations because it allows for the greatest amount of phase variation across the spectrum. However, as we learned from the GA solutions, a single order of phase is not the most efficient way to a pulse that minimizes  $I(t)$ .

As the resolution of the pulse shaper is increased in both simulations and experiment,  $I(t)$  becomes smoother and



**Figure 7.** Fitness as a function of generation for a simulation of frequency doubling experiment. Three curves are shown, two with and one without the smoothing operator included in the pool of operators used for mating. The solid curve is without smoothing and the other two curves are with smoothing. The two curves without smoothing are shown to compare two GA runs which use the same operators but with different random initial populations.

spreads out over a longer time window. Also, when we evaluated pulses that were stretched in time by programming them with the maximum amount of quadratic phase allowed by the resolution of our pulse shaper, they did not perform as well as the solutions found by the GA which involved many orders of phase as a function of frequency:  $\phi(\omega) = \sum k_n \omega^n$ .


The simulations can evaluate candidates for GA operators if the simulation provides an accurate model of the system or at least of some aspect of the system. For instance, figure 7 shows the best fitness as a function of generation for two different runs of the simulation, with and without the smoothing operator.

As can be seen in the figure, the addition of the smoothing operator allows the GA to achieve a much higher fitness more rapidly, which is a demonstration of the relevance of smoothing.

The high intensity regime of the frequency doubling experiment illustrates a benefit of the learning algorithm. Figure 8 shows the changing optimal pulse shape for frequency doubling as the energy of the input pulse is increased. The high energy solutions are different from the low energy ones indicating a deviation from our simple model, which predicts the same solution for different pulse energies.

The optimal pulse changes from nearly transform limit to a pulse with large third- and fourth-order dispersion. Measurements of the shaped pulse spectrum at high intensities show significant evidence of self-phase modulation (SPM), whereas in the low intensity limit there is no evidence of SPM. Simulations are also consistent with the presence of SPM in the crystal, which contributes to the variation in the solutions as the intensity is increased [65]. Evidently, at higher pulse energies, SPM is no longer negligible and distorts the phase matching for the SHG process, thus making a transform-limited input pulse no longer optimal for maximum SHG.

The doubling experiment also highlights the interplay of multiple operators. Figure 9 shows the operator fitness for




**Figure 8.** Husimi plots for pulses optimized to maximize second-harmonic generation at different pulse energies. GA-optimized pulses are shown for three separate pulse energies with increasing pulse energy for higher plots.

several operators during an experimental run for the frequency doubling experiment.

The GA is initialized in a traditional configuration, with two-point crossover dominant and the other operators creating few children. For the first few generations, the operators create children in proportion to their initially assigned fitness, but after the fourth generation their fitnesses are allowed to change in proportion to the success of their children. Figure 9 demonstrates how the algorithm determines that two-point crossover and simple mutation alone are not always the best combination of operators. At different points during the evolution, different combinations of operators are optimal for producing the best children. Figure 9 is particularly interesting because the success of the different operators at different times can be explained in terms of the physical process and the dynamics of the GA. Throughout, crossover does well, which is consistent with it being the only operator in many implementations of the GA [8]. However, early on, polynomial phase becomes successful and then dies away. Initially it is valuable to acquire polynomial phase, as seen from figure 8, but as many of the pulse shapes start to have and spread this feature, the operator becomes useless. Because polynomial phase does not insert smooth phase but rather polynomial sections with random variations, smoothing becomes important around generation 15. At the end of the run, average crossover becomes useful because all of the pulse shapes are similar and this operation makes useful incremental changes to the pulse shapes.

It is important to note several points. Firstly, the operators cannot be evaluated in isolation, for they affect each other



**Figure 9.** Operator fitness as a function of generation for multiple operators during optimization of second-harmonic generation in the high energy regime.

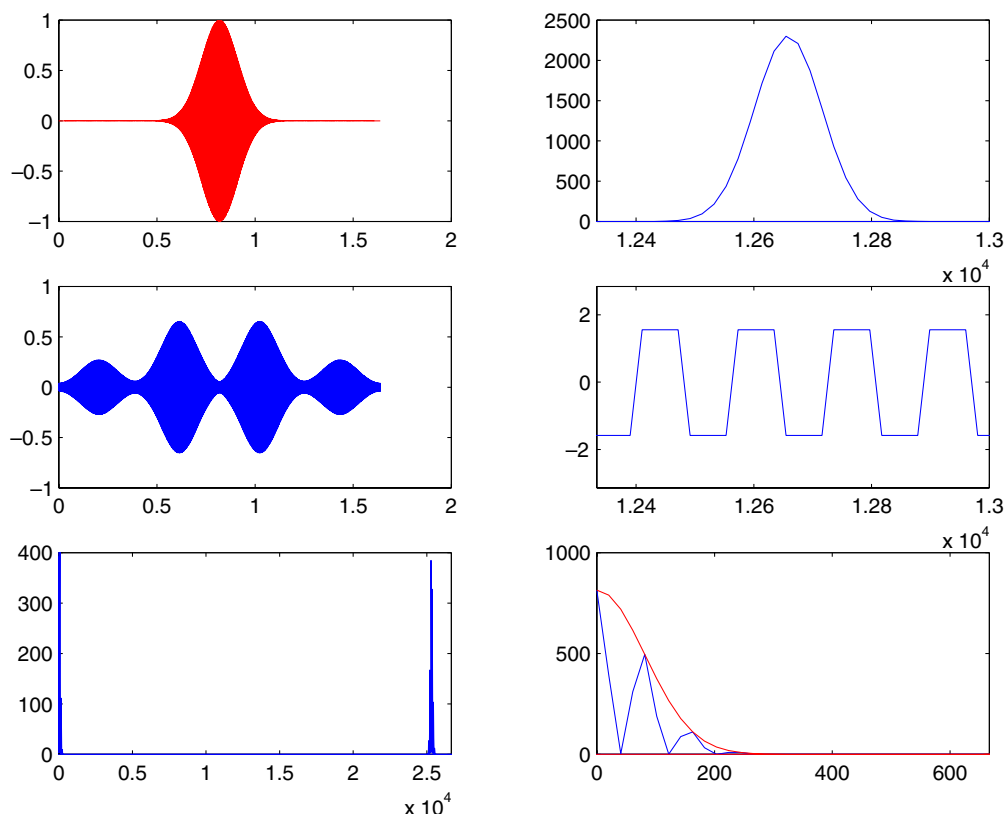
in the sense that certain operators work well in conjunction with others but not on their own. For example, smoothing is more important if other operators tend to introduce small random phase variations. Secondly, the performance of the operators cannot be evaluated instantaneously, but must be evaluated over the course of several generations. Finally, the performance of each operator is highly problem dependent. This is of course how operators can be used to learn about the problem since the performance of a given operator can help determine whether its action on an individual is physically relevant.

The frequency doubling experiment provides a valuable contrast with others in the interpretation of the feedback signal. The fitness assigned to each individual was simply the integrated blue light intensity as measured by a photodiode in a regime where the response of the diode was linear. In other experiments, it was not so clear how to assign a fitness to each individual given the nature of the measurement being made to evaluate the success of each pulse. The case of the methanol mode selection experiment discussed below is one where the feedback signal required careful analysis before being converted into a fitness value. In many cases, the goal that we wanted to achieve was not adequately expressed in terms of the fitness function as evidenced by trivial and meaningless solutions.

#### 4.3. Controlling molecular liquids

Our initial goal in molecular liquids was to control molecular vibrations through impulsive stimulated Raman scattering. Impulsive scattering occurs when the laser pulse is shorter than the vibrational period of the molecule. In the frequency domain, this means that the bandwidth of the laser is broad compared with the vibrational energy spacing of the mode in question. The result is that a stimulated Stokes wave can be seeded with light that is already present in the laser and does not have to build up from noise. This makes the process much more efficient than the nonimpulsive case.

Figure 10 shows the spectrum of the quadratic driving term  $E^2$  responsible for Raman scattering [66]. This *nonlinear*



**Figure 10.** Calculated pulse shapes (time domain) and spectra for phase-shaped and unshaped pulses. All times are in ps and all frequencies are in  $\text{cm}^{-1}$ . Top left shows unshaped pulse in time domain. Top right shows unshaped pulse in frequency domain. Middle panels show shaped pulse in time domain and phase of shaped pulse in frequency domain. Bottom panels show spectrum of  $E^2$  in full view (bottom left) and expanded view of low frequency portion (dashed curve is for unshaped pulse).


spectrum contains frequency components at the vibrational frequency when the pulse is in the impulsive limit. The two peaks centred around zero frequency and  $2\omega$  are important for two different kinds of two-photon transitions. The high frequency peak is relevant to frequency doubling, where the driving term is at the sum frequency; the low frequency peak is relevant to Raman scattering, where the nonlinear driving term is at a frequency difference between two photons in the pulse.

Figure 10 also shows simulations that demonstrate how a pulse with sufficient bandwidth to impulsively couple several modes in a multimode molecule can be shaped to excite only one or more of the modes. By shaping the pulse appropriately, the magnitude of  $E^2$  can be driven to zero for selected frequencies, suppressing modes at those frequencies, while maintaining high values of  $E^2$  for other modes. The bottom panels show the spectrum of  $E^2$  for shaped and unshaped pulses. The shaped laser pulse enhances Raman scattering for selective modes by controlling the phase of frequency pairs in the pulse spectrum. If all of the frequency pairs in the spectrum separated by the Stokes frequency for a given mode are phased to add constructively, then the amplitude of  $E^2$  at the Stokes frequency will be maximum. In the time domain, this corresponds to having a pulse that is longer than the vibrational period with the intensity modulated at the vibrational period of one mode and not the others, resonantly driving only one of many modes.

#### 4.4. Controlling self-phase modulation in liquids

Nonlinear propagation effects in liquid phase experiments can distort the shaped light pulse. This makes it more difficult to perform, or even observe, laser control over molecular modes. This feature was studied in experiments on  $\text{CCl}_4$ . We chose  $\text{CCl}_4$  to study multimode Raman control because it has several low frequency modes with relatively high cross-sections. It also has a large polarizability and therefore much of the light scattered near the laser bandwidth in the forward direction was a result of SPM.

We established feedback goals based upon small features that were barely visible in the spectra of the forward scattered radiation after transform-limited pulses illuminated the sample. One set of goals was established based on the observation of periodic spectral modulations in the measured spectra when a transform-limited pulse was incident on the cell. These modulations are typical spectral features for pulses that have undergone SPM under the influence of an intensity-dependent index of refraction [67]. We used the GA to enhance these modulations. We were able to control their frequency and phase by altering the shape of the driving pulse. For this set of experiments, we again used phase only modulation so that the pulse energy was kept fixed. Figure 11 shows the spectra of four different pulses after having propagated through the  $\text{CCl}_4$  sample. The first panel shows the spectrum for an unshaped pulse after propagation through the sample and the following three panels show spectra for pulses that were shaped to control the spectral modulations. The optimal shapes for generating these spectra were measured using spectral interferometry.




**Figure 11.** Power spectra for pulses propagating through 1 cm of  $\text{CCl}_4$ . The top panel shows the forward scattered spectrum for an unshaped incident pulse and the next three show the forward scattered spectra for pulses optimized to control spectral modulations.

Calculations employing the simplest description of SPM yield qualitatively similar spectra to our measurements.

#### 4.5. Controlling vibrations in multimode molecular liquids

Once we had evidence that we were able to control the intensity modulations in the forward scattered spectrum simply by changing the phase of the driving pulse through SPM, we attempted to control the interaction between the driving laser pulse and the vibrational modes of a multimode molecule without making use of impulsive scattering. The initial idea was that there might be an electronic mechanism such as SPM that would allow the driving laser pulse to selectively excite one of possibly many vibrational modes. The forward scattered radiation could once again serve as a diagnostic and be used as feedback for the GA. However, to avoid confusion between Stokes light and light generated by SPM alone, we decided to work with a molecule with a much larger Stokes shift. Further from the laser frequency, there is much less light generated through SPM. Methanol ( $\text{CH}_3\text{OH}$ ) was ideal because it had two closely spaced modes with large Stokes shifts and large cross-sections. It is also transparent and readily available.

One would not expect to see any backwards scattered Stokes radiation because of the short duration of the shaped pump pulse ( $\sim 1$  ps). The forward-backward symmetry of the scattering is broken for a short pump pulse because the backward travelling Stokes wave passes through the pump wave before any appreciable buildup [68]. We looked for



**Figure 12.** Short and long pulse stimulated Raman scattering in liquid methanol. Left two panels show Raman spectra for YAG laser pump and right two panels show Raman spectra for short pulse laser as pump. y-axes show spectral intensity and x-axes show frequency of the collected radiation.

backward-scattered Stokes light and we did not observe any. This was in direct contrast to the observation of backward-scattered Stokes light when we directed a 10 ns  $Q$ -switched doubled YAG laser pulse into the cell. Figure 12 provides a comparison between short pulse and long pulse Raman scattering.

The top left figure shows the Nd:YAG laser Stokes spectrum for both CH modes, although only one is visible in the plot. The symmetric mode has a slightly larger cross-section and therefore steals most of the gain in the stimulated process. The bottom left shows the same data plotted with a logarithmic scale for the intensity and shows the asymmetric mode as well, although the number of scattered photons is almost two orders of magnitude fewer. The two right plots show the Raman spectra for the same two modes with the short pulse laser at two different energies. Close to threshold (bottom right), two peaks are visible, but at higher pump laser energies, SPM broadens the peaks to the extent that they are no longer resolved (top right).

The timescales for the interaction between the molecules and the laser pulse are set by two molecular times. These are the vibrational period of the active modes and the coherence time of the vibrations. Stimulated scattering with pulses that are longer than the coherence time reaches steady state and exhibits strong pulse duration dependence. This is because it is a stimulated process and the more photons that interact with the molecules within the coherence time, the more likely the molecule will be stimulated to absorb a laser photon and emit a Stokes photon. Scattering with pulses that are shorter than the coherence time but longer than the vibrational period is transient scattering and exhibits little dependence on the duration of the laser pulse [69, 70]. Scattering with pulses that

are shorter than the vibrational period is impulsive scattering and results in very efficient stimulated scattering as discussed above [71, 72].


The vibrational period for the two modes is simply the inverse of the Raman shifts which, for the symmetric and asymmetric CH stretch modes in methanol, are 11.7 and 12.2 fs respectively.

With an unshaped laser pulse focused into a 10 cm long sample of methanol, the spectrum shown in figure 13(a) was obtained. The two peaks in the spectrum correspond to Stokes light for  $\delta\nu = 1$  for the symmetric and asymmetric C–H stretch modes. Using samples less than 10 cm in length, and without focusing, we were not able to see any evidence of SRS. The sample and lens focal lengths we used for all the data shown are 10 and 40 cm respectively. We also used a 30 cm lens with similar results. The first feedback goal was to maximize the contrast between the two Stokes peaks and the background light resulting from SPM. The forward scattered spectra in the spectral range of the Stokes radiation was collected for each laser shot. The interpretation of this spectrum and the translation into a single number was not as straightforward as in the frequency doubling experiment. The basic idea was that for each vibrationally excited molecule, there was a Stokes photon whose frequency labelled the mode of excitation. Thus, counting the number of Stokes photons at a particular frequency was a measure of the number of molecules excited in the mode labelled by the Stokes frequency. However, since SPM was present in the liquid and the bandwidth of the shaped laser pulses was large, the forward scattered spectra contained lots of misleading information. The feedback function had to filter out the fluctuating background, broadening of the Stokes peaks by SPM and shifting of the peaks because the shaped pulses had a large bandwidth compared with the linewidth of the vibrational levels. We tried many different fitness functions, with different emphases on peaks versus background and different spectral widths for the peaks. We found that depending on which peak(s) we were optimizing, different functions worked best. We found that differences always worked better than ratios, which tended to be much too sensitive to noise and also tended to drive the GA to turn off all of the frequency components yielding very large fitnesses through division by zero. A typical fitness function we used in optimizing one peak versus another plus background is

$$\sum_{\omega_r < \omega_i < \omega_b} \frac{N * C(\omega_i)}{\omega_b - \omega_r} - \sum_{\omega_i > \omega_b, \omega_i < \omega_r} \frac{C(\omega_i)}{\Delta\omega - (\omega_b - \omega_r)}$$

Here,  $C(\omega_i)$  is the number of spectrometer counts at  $\omega_i$ ,  $\Delta\omega$  is the bandwidth of the spectrometer,  $\omega_r$  is the low frequency limit for the desired peak,  $\omega_b$  is the high frequency limit for the desired peak and  $N$  is an empirically determined integer. We set the values of  $\omega_r$  and  $\omega_b$  by narrowing the bandwidth of the shaped laser pulse and measuring the width of the Stokes peaks in the forward scattered spectrum. We found that  $N = 2, 3$  worked well.

When we set out to optimize the contrast between the two Stokes peaks and the background, we measured the spectra shown in figure 13(b). The next goal that we established was to generate spectra with each peak separately. These spectra, shown in figure 13(c), correspond to exciting one or another



**Figure 13.** Control of Raman scattering in CH<sub>3</sub>OH. (a) shows the forward scattered spectrum when an unshaped laser pulse is incident. (b) shows the spectrum after the GA was set to optimize excitation of both modes while minimizing peak broadening due to other nonlinear effects. (c) shows spectra when the GA was set to optimize each mode independently. (d) shows the spectrum for a pulse optimized to minimize Raman scattering from both modes. y-axes are spectral intensity and the x-axes are the frequency in cm<sup>-1</sup>.

of the two modes alone. It is important to note that the Stokes shift for these two modes is large compared with the bandwidth of the driving laser pulse. This is equivalent to saying that the Raman excitation was nonimpulsive and therefore one would not expect to control mode selectivity through seeding of the Stokes radiation with the laser light directly. Our final feedback goal was to eliminate all forward scattered light at either of the two Stokes frequencies and the resulting spectrum is shown in figure 13(d). All of the spectra shown in figure 13 were the result of optimization with fixed spectral amplitudes (phase only modulation) but we also performed experiments with phase and amplitude modulation which yielded similar results.

#### 4.6. Discussion

Impulsive excitation of the two CH stretch modes would be one mechanism for control if the laser bandwidth were large enough for impulsive coupling.

Since the Stokes shift for the CH stretch is almost 3000 cm<sup>-1</sup> and the laser bandwidth is roughly 100 cm<sup>-1</sup>, the scattering is definitely nonimpulsive. Another possible mechanism that could account for our ability to selectively excite the symmetric or asymmetric stretch mode of methanol is a coupling between the electronic polarizability of the atoms and the vibrational modes. The results with CCl<sub>4</sub> have shown that the light generated from SPM of the pump beam is very sensitive to the input pulse shape. A large contribution to SPM, particularly for femtosecond pulses, is the atomic polarizability [73]. One might argue that the atomic polarizability provides a handle for vibrational mode selectivity by providing light generated by SPM to seed one of the two Raman modes but not the other.

#### 4.7. Further investigations


In an effort to determine whether the methanol result was made possible through an intermolecular or intramolecular effect, we decided to perform an experiment where the two modes of vibration that we aimed to select between were in two different molecules. If the energy deposited into vibration could be shared between molecules and not just redistributed within each molecule itself, then one should be able to drive mode selection where the competing modes are in different molecules. If SPM were responsible for the control described above, then mode-selective excitation should still be possible even if the modes are in separate molecules, since SPM light generated from one molecule could seed Raman Stokes generation in another.

We chose benzene and deuterated benzene as the two molecules for this experiment since the ring breathing mode of benzene ( $\nu = 992 \text{ cm}^{-1}$ ) has a large Raman cross-section, and the deuterated benzene has a frequency of  $\nu_D = 945 \text{ cm}^{-1}$  which is shifted by  $47 \text{ cm}^{-1}$ , similar to the mode splitting in the case of methanol. Initially, the laser was focused into the experimental cell with pure  $\text{C}_6\text{H}_6$ , and with a transform-limited pulse we measured no forward scattered Stokes light. We used the GA to find a pulse that generated forward scattered Stokes light and collected the spectrum shown in figure 14 (top panel dark curve). Then the benzene was replaced with  $\text{C}_6\text{D}_6$  and with the same pump pulse the spectrum shown in figure 14 (top panel, grey curve) was collected, demonstrating the shift of the mode frequency because of the six extra neutrons in  $\text{C}_6\text{D}_6$ . Finally, a 50/50 mixture of  $\text{C}_6\text{H}_6$  and  $\text{C}_6\text{D}_6$  was placed in the cell and we tried to use the GA to selectively drive each of the two modes. Figure 14 (bottom panel) shows that we were successful in driving the  $\text{C}_6\text{H}_6$  mode but not in driving the  $\text{C}_6\text{D}_6$  mode, which demonstrates a lack of control. The black curve shows the spectrum when we tried to optimize excitation of the  $\text{C}_6\text{H}_6$  and the light curve shows the spectrum when we tried to optimize excitation of the  $\text{C}_6\text{D}_6$ .

This result is consistent with the idea that the mode selection is an intramolecular effect that relies on coupling between the two modes inside each molecule rather than on a macroscopic buildup such as seeding of one of the modes with light from SPM.

Other experiments have also led us to conclude that SPM does not provide the control mechanism for the CH stretch modes in the methanol experiment. Looking for the CO mode in methanol, we noticed that the same pulse shape produced very different SPM spectra even closer to the laser frequency each time it was incident on the sample. Experiments in  $\text{CO}_2$  gas also showed that light generated by SPM further than a few  $100 \text{ cm}^{-1}$  away from the laser frequency was extremely noisy and not reproducible. The reproducibility of SPM spectra increases nearer the frequency of the driving laser. The control over SPM shown in figure 11 shows that SPM produces stable spectra very near the laser frequency. Far from the laser at the frequency of the Stokes light from the CH stretch mode, the light produced by SPM would be too noisy to reproducibly seed one of the two Raman modes and not the other.

We arrived at one more possible explanation for the control mechanism, which was suggested by analysing the pulse shapes that were optimal for exciting each of the two modes. This explanation is based on quasi-impulsive




**Figure 14.** Stimulated Raman scattering in  $\text{C}_6\text{H}_6$  and  $\text{C}_6\text{D}_6$ . The top panel shows spectra for each molecular liquid separately after optimization of the pulse shape to excite the breathing mode of the molecule. The bottom panel shows the results of trying to excite each molecule in a 50/50 mixture of the two using the GA to optimize the pulse shape.  $y$ -axes are spectral intensity and the  $x$ -axes are the frequency in  $\text{cm}^{-1}$ .

scattering. Although the laser bandwidth is narrow compared with the Stokes shift of each mode, it is wide in comparison with the spacing between the two modes. This is equivalent to saying that in the time domain the laser pulse is long compared with the vibrational period of the two modes, but it is short in comparison with the beat note between the two vibrations. The beat note period is 285 fs while the time duration of the unshaped laser pulse is 150 fs. Therefore energy could be transferred between the two modes by engineering temporal structure in the driving pulse at the coupling frequency between the two modes.


Figure 15 shows a Husimi distribution for a pulse that was optimized for excitation of the asymmetric stretch mode. Figure 16 shows a Husimi distribution for a pulse that was optimized for excitation of the symmetric stretch mode. In these GA runs, the spectral phases and amplitudes were both allowed to vary. Similar pulses were obtained for multiple runs of the GA with the same goal. The quasi-impulsive model is suggested by the structure in the optimal Husimi plots for the symmetric stretch mode where the frequency separation of the two bright subpulses is exactly the beat frequency between the symmetric and asymmetric modes. Once an initial vibrational population is established in some combination of the two modes, the population could be redistributed by the shaped pulse through impulsive coupling of the two levels.

Individual pulse shapes may provide insight into the physics of a given problem, but it is general trends in an entire population that tend to provide more valuable information because looking at a single individual, even if it is the very best pulse shape of all, does not give information about which features are necessary and which are merely sufficient. We decided to look at statistical variations among gene values of individuals in a population as a function of generation in an effort to see which genes were important for a given



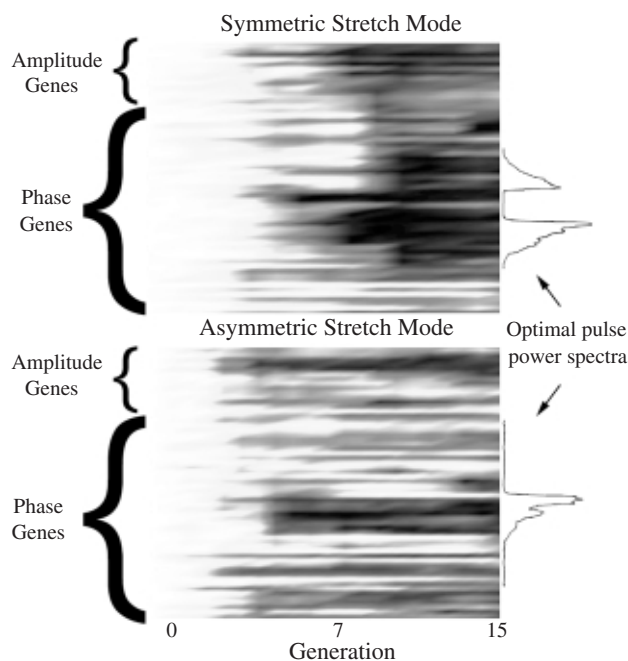


**Figure 15.** Husimi distribution for laser pulse optimized for excitation of the asymmetric stretch mode.



**Figure 16.** Husimi distribution for laser pulse optimized for excitation of the symmetric stretch mode.

problem and which were not. Figure 17 shows the genetic variation as a function of generation for a run of the GA for the methanol experiment. Light shading represents a large degree of variation among individuals for the value of a given gene and dark shading represents low variation. The variation was calculated as the normalized sum of the absolute values of the difference between each genes value and the average value for that gene in a given generation. Initially, all genes have light shading since they are randomly initialized. As the GA converges, all gene values tend to become similar through mating. However, the plots verify the intuitive idea that near convergence, frequency components whose amplitudes are large have smaller variation in their programmed phase values than frequency components whose amplitudes are zero. The power spectra shown on the right match with the darker regions (smaller variation) in the plots. This is a simple and crude demonstration of the fact that one can look at trends in populations and statistics among gene values to determine which gene values are important (i.e. the ones with darker shading in the plots) and which are not.



**Figure 17.** Genetic variation as a function of generation. Data are shown for two separate runs of the GA. The top shows genetic variation during optimization of the asymmetric stretch mode and the bottom shows genetic variation during optimization of the symmetric stretch mode.

## 5. Conclusion

The number of time domain coherent control experiments has grown enormously in the last few years. This is a result of many factors, including the development of laser technology, the implementation of feedback for learning control and the realization that coherent control is a powerful and general technique. Key aspects of future implementations will be demonstrating the utility of the final state achieved through control, as well as illustrating what has been learned in the control experiment.

## References

- [1] Zare R N 1998 *Science* **279** 1875
- [2] Rabitz H, de Vivie-Riedle R, Motzkus M and Kompa K 2000 *Science* **288** 824
- [3] Warren W S, Rabitz H and Dahleh M 1993 *Science* **259** 1581
- [4] Brumer P and Shapiro M 1986 *Chem. Phys. Lett.* **126** 541
- [5] Tannor D and Rice S 1985 *J. Chem. Phys.* **83** 5013
- [6] Weiner A M 1995 *Prog. Quantum Electron.* **19** 161
- [7] Rabitz H and Judson R 1992 *Phys. Rev. Lett.* **68** 1500
- [8] Davis L (ed) 1991 *Handbook of Genetic Algorithms* (Reinhold: Van Nostrand)
- [9] Zhu L, Kleiman V, Li X, Lu S, Trentelman K and Gordon R J 1995 *Science* **270** 77
- [10] Seideman T 1999 *J. Chem. Phys.* **111** 9168
- [11] Haché A, Kostoulas Y, Atanasov R, Hughes J L P, Sipe J E and van Driel H M 1997 *Phys. Rev. Lett.* **78** 306
- [12] Wang F, Chen C and Elliott D S 1996 *Phys. Rev. Lett.* **77** 2416
- [13] van Druen N J, Constantinescu R C, Schins J M, Nieuwenhuize H and Muller H G 1997 *Phys. Rev. A* **55** 622
- [14] Backus S, Durfee C G, Murnane M M and Kapteyn H C 1998 *Rev. Sci. Instrum.* **69** 1207
- [15] Dantus M 2001 *Rev. Phys. Chem.* **52** 639
- [16] Hornung T, Meier R, de Vivie-Riedle R and Motzkus M 2001 *J. Chem. Phys.* **267** 16

- [17] Knopp G, Pinkas I and Prior Y 2000 *J. Raman Spectrosc.* **31** 51
- [18] Bardeen C J, Che J W, Wilson K R, Yakovlev V V, Apkarian V A, Martens C C, Zadayan R and Kohler B 1997 *J. Chem. Phys.* **106** 8486
- [19] Blanchet V, Bouchene M A and Girard B 1998 *J. Chem. Phys.* **108** 4862
- [20] Amitay Z, Ballard J B, Stauffer H U and Leone S R 2001 *J. Chem. Phys.* **267** 141
- [21] DeCamp M F, Reis D A, Bucksbaum P H and Merlin R 2001 *Phys. Rev. B* **6409** 2301
- [22] Weiner A M, Leird D E, Wiederrecht G P and Nelson K A 1990 *Science* **247** 1317
- [23] Hase M, Mizoguchi K, Harima H, Nakashima S, Tani M, Sakaim K and Mangyo M 1996 *Appl. Phys. Lett.* **69** 2474
- [24] Ozgur U, Lee C W and Everitt H O 2001 *Phys. Rev. Lett.* **86** 5604
- [25] Wehner M U, Ulm M H, Chemla D S and Wegener M 1998 *Phys. Rev. Lett.* **80** 1992
- [26] Gershgoren E, Vala J, Kosloff R and Ruhman S 2001 *J. Phys. Chem. A* **105** 5081
- [27] Villeneuve D M, Aseyev S A, Dietrich P, Spanner M, Ivanov M Y and Corkum P B 2000 *Phys. Rev. Lett.* **85** 542
- [28] Kubasik M, Cebo A, Hertz E, Chauv R, Lavorel B and Faucher O 2001 *J. Phys. B: At. Mol. Opt. Phys.* **34** 2437
- [29] Bonadeo N H, Erland J, Gammon D, Park D, Katzer D S and Steel D G 1998 *Science* **282** 1473
- [30] Marie X, LeJeune P, Amand T, Brousseau M, Barrau J, Paillard M and Planel R 1997 *Phys. Rev. Lett.* **79** 3222
- [31] Baumberg J J, Heberle A P, Kohler K and Ploog K 1996 *J. Opt. Soc. Am. B* **13** 1246
- [32] Meshulach D and Silberberg Y 1998 *Nature* **396** 298
- [33] Bucksbaum P H 1998 *Nature* **396** 6708
- [34] Paye J 1992 *IEEE J. Quantum Electron.* **28** 2262
- [35] Weinacht T C 2000 *Dissertation* University of Michigan
- [36] Bardeen C J, Yakovlev V V, Wilson K R, Carpenter S D, Weber P and Warren W S 1997 *Chem. Phys. Lett.* **280** 151
- [37] Assion A, Baumert T, Bergt M, Brixner T, Kiefer B, Seyfried V, Strehle M and Gerber G 1998 *Science* **282** 919
- [38] Vajda S, Bartelt A, Kaposta E C, Leisner T, Minemoto S, Rosendo-Francisco P and Woste L 2001 *J. Chem. Phys.* **267** 231
- [39] Levis R J, Menkir G M and Rabitz H 2001 *Science* **292** 5
- [40] Bartels R A, Backus S, Kapteyn H C and Murnane M M 2000 *Nature* **406** 6792
- [41] Geremia J M, Zhu W S and Rabitz H 2001 *J. Chem. Phys.* **113** 8
- [42] Weinacht T C, Bartels R A, Backus S, Pearson P, Bucksbaum P H, Geremia J M, Rabitz H, Kapteyn H C and Murnane M M 2001 *Chem. Phys. Lett.* **344** 333
- [43] Brixner T, Kiefer B and Gerber G 2001 *J. Chem. Phys.* **267** 241
- [44] Pearson B J, White J L, Weinacht T C and Bucksbaum P H 2001 *Phys. Rev. A* **63** 063412-1
- [45] Zeidler D, Frey S, Kompa K L and Motzkus M 2001 *Phys. Rev. A* **6402** 3420
- [46] Strickland D and Mourou G 1985 *Opt. Commun.* **56** 219
- [47] Strickland D and Mourou G 1985 *Opt. Commun.* **55** 447
- [48] Weiner A M, Leird D E, Patel J S and Wullert J R 1992 *J. Quantum Electron.* **28** 908
- [49] Tull J X, Dugan M A and Warren W S 1996 *Adv. Magn. Opt. Res.* **20** 1
- [50] Fittinghoff D N, Bowie J L, Sweester J T, Jennings R T, Krumbugel M A, DeLong K W, Trebino R and Walmsley I A 1996 *Opt. Lett.* **21** 884
- [51] Gallagher T 1995 *Rydberg Atoms* (Cambridge, MA: Cambridge University Press)
- [52] Gallagher T F, Humphrey L M, Cooke W E, Hill R M and Edelstein S A 1977 *Phys. Rev. A* **16** 1098
- [53] Jones R R, Raman C S, Schumacher D W and Bucksbaum P H 1993 *Phys. Rev. Lett.* **71** 2575
- [54] Trebino R, DeLong K W, Fittinghoff D N, Sweetser J N, Krumbugel M A, Richman B A and Kane D 1997 *Rev. Sci. Instrum.* **68** 3277
- [55] Xin C and Yeazell J A 1990 *Phys. Rev. A* **56** 2316
- [56] Leichtle C, Schleich W P, Averbukh I Sh and Shapiro M 1998 *Phys. Rev. Lett.* **80** 1418
- [57] Kinrot O, Averbukh I S and Prior Y 1995 *Phys. Rev. Lett.* **75** 3822
- [58] Hanbury-Brown R and Twiss R Q 1996 *Nature* **175** 27
- [59] Gallagher T F and Cooke W E 1979 *Phys. Rev. Lett.* **42** 835
- [60] Press W H 1988 *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge: Cambridge University Press)
- [61] Leonhardt U 1997 *Measuring the Quantum State of Light* (Cambridge, MA: Cambridge University Press)
- [62] Lee H 1995 *Phys. Rep.* **259** 147
- [63] Raman C 1997 *Dissertation* University of Michigan
- [64] Ahn J, Weinacht T C and Bucksbaum P H 2000 *Science* **287** 463
- [65] Nibbering E J T, France M A, Prade B A, Grillon G, Le Blanc C and Mysryowicz A 1995 *Opt. Commun.* **119** 479
- [66] Boyd R 1992 *Nonlinear Optics* (London: Academic)
- [67] Corkum P and Rolland C 1989 *IEEE J. Quantum Electron.* **25** 2634
- [68] Shen Y R 1984 *The Principles of Nonlinear Optics* (New York: Wiley)
- [69] Carman R L, Shimizu F, Wang C S and Bloembergen N 1970 *Phys. Rev. A* **2** 60
- [70] Carman R L, Mack M E, Shimizu F and Bloembergen N 1969 *Phys. Rev. Lett.* **23** 1327
- [71] Korn G, Duhr O and Nazarkin A 1998 *Phys. Rev. Lett.* **81** 1215
- [72] Nazarkin A and Korn G 1998 *Phys. Rev. A* **58** R61
- [73] Hellwarth R W, Owyong A and George N 1971 *Phys. Rev. A* **4** 2342