Optimal Control of Molecular Motion: Design, Implementation, and Inversion

HERSCHEL RABITZ* AND WUSHENG ZHU
Department of Chemistry, Princeton University, Princeton, New Jersey 08544

Received October 18, 1999

ABSTRACT

This paper reviews recent theoretical and experimental developments aimed at controlling molecular motion using tailored laser fields. Emphasis is given to seeking optimal designs for the laser controls and optimal implementation of the controls in the laboratory. Optimization on both counts provides a rigorous, flexible, and physically attractive means for obtaining the best possible control over molecular motion under any specified conditions. The theoretical design and laboratory implementation of control are best effected by a closed-loop process that draws on observations of the evolving molecular sample to steer it toward the desired target. Going beyond control, similar closed-loop laboratory learning concepts may lead to automated molecular monitors for inversion to systematically identify details of molecular Hamiltonians.

I. Introduction

A central concern of chemistry is making and breaking chemical bonds. The development of lasers as tunable intense sources of radiation has stimulated efforts at employing them to favorably alter the course of molecular dynamics phenomena. Laser control over molecular dynamics is a long-standing dream in the chemical dynamics community,[1] with many frustrating years of research dating from the early 1960s until the late 1980s. Conceptual,[2−4] theoretical,[5,6] and experimental[7,8] developments starting in the late 1980s dramatically changed this state of affairs. There are now a number of successful experiments demonstrating the key principles of molecular control.[9−15]

This Account will focus on *optimal* control of molecular motion, as it is proving to be the most general means to treat both the design[4−6] and laboratory implementation[12−15] of control over molecular motion. Molecular control by nonoptimal means[3] also can be effective for simple systems or special circumstances where the Hamiltonian can be prescribed in detail. Many molecules of practical interest are complex in that the Hamiltonians are often not known well. Optimal control is especially attractive in such cases as it may be directly applied in the laboratory in a closed-loop fashion to achieve control with minimal or even no knowledge of the Hamiltonian.

A brief synopsis of the overall field is useful in order to appreciate the role of optimal control. The basic operating principle behind molecular control is laser-field manipulation of constructive and destructive interferences of the evolving molecular wave function. Control of intramolecular dynamical events can be achieved by creating constructive interferences in the desired product channel and destructive interferences in any unwanted product channels or undesirable states along the path toward the final goal. Successful control requires the laser field to cooperate with the dynamical capabilities of the molecule. A natural first step toward establishing this cooperation is the design[3−6] of an appropriate control laser field. As the control process involves manipulating wave interferences and molecular motion can often have a high degree of complexity, the use of simple intuition as a means for design is generally not expected to be satisfactory. Input from physical intuition remains essential, but it must be applied in a nontraditional manner within the optimal design process, as explained later. In keeping with the oscillatory nature of quantum wave packet motion, the typical controls involve phase- and amplitude-modulated laser pulses.[7,8,16] Many of the key elements of molecular control design and laboratory implementation are now becoming operational.

In molecular control, we always desire to achieve the best possible solution, and it is therefore natural to consider optimization procedures for both design and laboratory implementation to maximally approach the desired final state. These steps will be referred to as optimal control theory (OCT) and optimal control experiments (OCE), respectively. Both OCT and OCE can work together to attain the overall best molecular control results. For polyatomic molecules, there is considerable uncertainty in the Hamiltonians, and the ability to solve the quantum control design equations[4] to acceptable accuracy is a very challenging task. Thus, an OCT design will often just be an estimate of the optimal control, calling for refinement in the OCE process. This situation has led to the suggestion[17−20] of performing closed-loop OCE as an essential component of the laboratory studies, and an increasing number of experiments are of this nature.[12−16,21] The role of optimization is central, since an excursion from the initial state of a polyatomic molecule to the desired target could encounter many undesirable product channels or states along the way. These undesirable events need to be suppressed while maximizing the flux into the target and simultaneously carrying this out under constraints on the capabilities of the laser.[22]

Beyond considering the control of molecular motion for chemical or physical objectives, recent projections suggest that the same family of techniques may be

Herschel Rabitz graduated from Harvard University in 1970, with a Ph.D. degree in chemical physics, followed by postdoctoral work at the University of Wisconsin. Since 1971, Professor Rabitz has been on the faculty of the Department of Chemistry at Princeton University. Professor Rabitz's research interests lie at the interface of chemistry, physics, and engineering, with principal areas of focus including molecular dynamics, chemical kinetics, and optical interactions with matter.

Wusheng Zhu obtained his B.S. degree in 1990, and his Ph.D. degree in 1996 from Peking University. In 1996, he joined the group of Professor Herschel Rabitz at Princeton University as a postdoctoral associate. His research interests concern the development of methodology and algorithms for solving problems in chemical physics.

10.1021/ar990054h CCC: $19.00
Published on Web 05/17/2000

redirected to better understand molecules and materials.[23,24] The notion of "understanding" in this context refers to controlled manipulation for the purpose of inversion, to quantitatively extract information about Hamiltonians, especially intramolecular potentials. The remainder of this Account will review the state of OCT, OCE, and the promise for dynamical inversions.

II. Design by Optimal Control Theory (OCT)

A laser field may interact with a molecule through various optical coupling mechanisms. The general structure of OCT is the same, regardless of the coupling involved, and for illustrative purposes here, an electric dipole interaction $-\mu\epsilon(t)$ is assumed, where μ is the dipole moment and $\epsilon(t)$ is the sought-after optical field design. The full Hamiltonian is $H = H_0 - \mu\epsilon(t)$, with H_0 being the free molecular Hamiltonian without control. The evolution of the molecular system is described by the Schrödinger equation (or its density matrix analogue), $i\hbar(\partial|\psi(t)\rangle/\partial t) = H|\psi(t)\rangle$, with the objective of steering the system from an initial state $|\psi(0)\rangle = |\phi\rangle$ to an observable $\langle O(T)\rangle = \langle\psi(T)|O|\psi(\mathrm{T})\rangle$ in the final state $|\psi(T)$ at the target time T.[22] Here, O is an observable operator, such as a bond length or projection into a particular state, etc. The goal is to steer $\langle O(T)\rangle$ to a desired value O^*, but often there are other factors to consider. For example, it may be important to minimize the magnitude of the expectation value of a family of other operators, $\langle O_l(t)\rangle$, $l = 1, 2, ...$, over the full control interval $0 \le t \le T$, where these new operators correspond to molecular events which are undesirable (e.g., breaking the wrong chemical bonds). In addition, there may be a set of practical constraints on the control field $\epsilon(t)$, in keeping with realistic laboratory capabilities. Thus, the overall desire is to strike a balance between reaching the objective O^* and minimizing other undesirable processes, and doing both while keeping the control field within reachable bounds.

Although the tasks posed above are physically transparent, finding the control field $\epsilon(t)$ that meets the multiple criteria is a nontrivial matter, as the role of the field is subtly buried in the dynamics of the evolving molecular wave function. It is assumed that the objectives and penalties are physically compatible, such that at least some control field $\epsilon(t)$ exists[4] that will give an acceptable solution. An optimal field $\epsilon(t)$ is desired that does the best job possible under the various competing criteria. This situation naturally defines a problem for OCT,[4,6,22] expressed in terms of a positive definite variational cost functional, $J[\epsilon(t)]$. There is considerable flexibility in choosing the functional, and often, a quadratic form is chosen due to its simplicity:

$$J = (\langle\psi(T)|O|\psi(T)\rangle - O^*)^2 + \sum_l \int_0^T \omega_l(t)|\langle\psi(t)|O_l|\psi(t)\rangle|^2\,\mathrm{d}t + \int_0^T \omega_\epsilon(t)\epsilon^2(t)\,\mathrm{d}t + \int_0^T \mathrm{d}t\left[\left\langle\lambda(t)\middle|i\hbar\frac{\partial}{\partial t} - H\middle|\psi(t)\right\rangle + \mathrm{c.c.}\right] \quad (1)$$

where the positive weights $\omega_l(t)$, $l = 1, 2, ...$, and $\omega_\epsilon(t)$ are chosen to balance the physical significance of the various terms in the cost functional. The cost on the laser field, for illustration, is considered here as a fluence, but other costs in the time and/or frequency domain could also be introduced. The last term includes the Lagrange multiplier costate $|\lambda(t)\rangle$ to ensure that Schrödinger's equation is satisfied for any possible field designs under consideration. The overall form of the functional and its components are strictly guided on physical grounds. The designer's intuition is never lost in the process, and a change in any of the criteria can lead to different designs; design for any purpose is generally not unique, and quantum mechanical design is no exception. Consideration of the often complex physical issues involved, and their translation into an optimization problem, indicates why a priori choosing a control field is not likely to be successful. An intuitively chosen field is not expected to minimize a reasonable physical cost functional. However, in a well-posed OCT problem, there will likely be many good solutions, and the goal of OCT is to find at least one good solution.

The cost functional in eq 1 depends on the three unknown functions, $\epsilon(t)$, $|\psi(t)\rangle$, and $|\lambda(t)\rangle$. Setting to zero the first-order variations of the cost functional, with respect to these functions, leads to the OCT design equations[22]

$$i\hbar\frac{\partial|\psi(t)\rangle}{\partial t} = [H_0 - \mu\epsilon(t)]\psi(t)\rangle, \quad |\psi(0)\rangle = |\phi\rangle \quad (2a)$$

$$i\hbar\frac{\partial|\lambda(t)\rangle}{\partial t} = [H_0 - \mu\epsilon(t)]\lambda(t)\rangle + 2\sum_l \omega_l(t)\langle\psi(t)|O_l|\psi(t)\rangle O_l|\psi(t)\rangle \quad (2b)$$

$$|\lambda(T)\rangle = \frac{4}{i\hbar}[\langle\psi(T)|\mathrm{O}|\psi(T)\rangle - O^*]O|\psi(T)\rangle \quad (2c)$$

$$\epsilon(t) = \frac{1}{2\omega_\epsilon(t)}R\,\langle\lambda(t)|\mu|\psi(t)\rangle \quad (2d)$$

Equation 2d gives the sought-after field. The primary numerical complexity of solving these equations arises because eq 2a has an initial condition, while eq 2b has a final condition in eq 2c. Thus, these design equations form a nonlinear boundary value problem in *time*, and it is this character which gives rise to a discrete multiplicity of solutions.[25] Each solution produces a particular control design, offering the flexibility of finding at least one design that may be conveniently constructed in the laboratory.

In general, the OCT design equations for the field $\epsilon(t)$ must be solved iteratively due to their nonlinear nature. For this purpose, a variety of algorithms have been tried, and the topic is still one of emerging development. Under certain conditions, algorithms may be found which are guaranteed to be monotonically convergent, such that each iteration is assured to take the control a step closer to a local optimal solution.[26,27]

Many numerical calculations have been executed with OCT in recent years considering the control of rotational, vibrational, and electronic degrees of freedom. The com-

putational effort involved can push the limits of quantum dynamics simulation capabilities, and continued advances in the latter area may be directly carried over into improving OCT designs. A basic point to consider is that the Hamiltonian H for realistic systems is often uncertain to a significant degree. Thus, additional criteria ensuring that the control is robust to uncertainties may be included in the cost functional.[28] In practice, the field designs will always be approximate, at least for this reason. Fortunately, reasonable design estimates will generally suffice, as they may be refined in the laboratory by OCE, as explained in section III. It seems quite remarkable that so many good OCT solutions have been reported in the literature, given the fact that the optimization space is likely populated with solutions of diverse quality. The qualitative structure of control field designs ranges from simple to complex, with the diversity dictated by the physical demands placed in the design cost functional. The achievement of good-quality control requires that the field cooperate with the dynamical capabilities of the molecule, and postdesign analysis can reveal interesting mechanistic patterns and information on the dynamical pathways leading to control.

III. Implementation by Optimal Control Experiments (OCE)

The performance of successful control experiments is the anticipated outcome of performing OCT designs. The ability to create the necessary control fields is an essential first step in the laboratory. The fields can range in structure from having a few frequency components to broad-band phase- and amplitude-modulated pulses. Pulse shaping[7,8,16,29] is a rapidly developing technology which offers the ability to operate in cooperation with the system's dynamics, and the preparation of control fields is diminishing as a stumbling block to progress. Pulse shaping is generally performed by starting with an ultrafast laser pulse that is spectrally decomposed, followed by phase and amplitude modulation of the discrete frequency components. This modulation is usually carried out with liquid crystal[8] or acousto-optic[30] techniques. Laser pulse shaping capabilities are expected to further advance, especially into the longer wavelength infrared regime.

One possible approach to achieving molecular control consists of (1) performing an OCT laser field design, (2) implementing the design in the laboratory as a tailored laser pulse, and (3) applying it to the molecule, where it is hoped that satisfactory control is the outcome. In simple cases, this open-loop OCT execution effort should be successful.[10,31] However, in general, and especially for polyatomic molecules, a number of serious issues arise: (a) the Hamiltonian is only approximately known, (b) the design equations for complex systems will likely call for approximate solutions, and (c) the field designs produced in the laboratory may have systematic and random errors. This situation points to the need for closing the loop in the laboratory, thereby defining the OCE process.[17−20] Figure 1 presents a schematic combining OCT and OCE

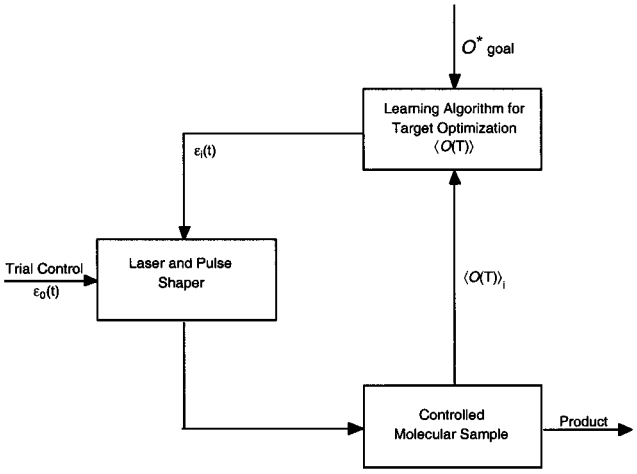

FIGURE 1. Closed-loop learning process for performing optimal control experiments (OCE). The process is initiated by specification of a product goal O^* for the expectation value $\langle O(T)\rangle$ of the target operator O at the final time T. A trial field $\epsilon_0(t)$ from optimal control theory (OCT) or another estimation procedure may be fed to the laser pulse shaper for the first control experiment. In a sequence of excursions $i = 1, 2, \ldots$ around the loop, the learning algorithm guides the shaped laser pulse to steer $\langle O(T)\rangle$ toward the product goal O^* by observing the patterns of behavior evident in the laser control settings for $\epsilon_i(t)$ and their molecular actions $\langle O(T)\rangle_i$.

into an overall framework. The initial OCT field design $\epsilon_0(t)$ enters for further refinement in the closed-loop OCE excursions. Beyond the initial OCT step, the OCE closed-loop process eliminates issues of Hamiltonian uncertainty and the need to solve Schrödinger's equation, as the actual molecules and their true dynamics are part of the loop. The only demand on OCT is that it produce a trial field $\epsilon_0(t)$ that gives a minimal signal in the target state for OCE refinement. Another important feature of OCE is that the field $\epsilon(t)$ need not be observed upon each excursion of the closed loop. It suffices to merely utilize the laser "knob" settings in the closed-loop learning algorithm. Furthermore, although the control and observation processes in the closed loop are carried out at the real ultrafast time scales of the dynamics, a full excursion of the loop can be performed at a leisurely rate, in keeping with the speed of the computers in the loop and the laser hardware operational capabilities.

The OCE process in Figure 1 is formally called learning control, in contrast to feedback control. The latter topic refers to circumstances where real-time control is being executed for a system whose model is generally known well.[32] Although feedback control is common in many engineering applications, it is problematic for controlling molecular dynamics due to the ultrafast dynamics involved, the lack of full knowledge of the Hamiltonian, and the fundamental question of whether the control feedback process may introduce inherent instabilities associated with the special nature of quantum mechanics (i.e., to observe a quantum system is also to disturb it). Quantum learning control in Figure 1 circumvents all of these difficulties by playing on the fact that we have readily available $\sim 10^{23}$ samples of the same system, and laser pulse shapers have very high duty cycles producing

distinct control pulses at the rate of possibly $\sim 10^2-10^6$ s^{-1}. Thus, a very large number of control experiments may be performed over manageable laboratory time scales on a sequence of samples.

The nature of the learning algorithm in the OCE loop is of basic importance. Although no definitive judgment has been made about the best algorithms for OCE in terms of stability and efficiency, computer simulations have shown that success may be attained with local gradient-based algorithms,[18,20] as well as more global genetic or evolutionary techniques.[17] The initial OCE simulations[17] and the current experiments[12−15] have all employed genetic-type algorithms. As the number of experiments that may be performed over realistic time scales is high, at this juncture, issues of algorithmic efficiency appear to be of secondary importance. Algorithmic stability for the OCE process has not been explored, but no negative results have been attributed to such difficulties. A small degree of experimental noise may actually be beneficial,[18] as a means to stimulate the control excursions in Figure 1 away from one local optimum to a better solution. Systematic errors in the control field cause no particular difficulty; however, noise beyond a certain level in the control field will become a significant issue. Simulations on simple systems[18,19] indicate a remarkable level of robustness in this regard, but the full extent of the matter is not known, especially for strong field nonlinear molecular control.

The closed-loop OCE process[17] operates with logic parallel to that employed in OCT design.[22] A laboratory cost functional $K[\tilde{\epsilon}(\omega)]$ is defined for OCE in analogy with $J[\epsilon(t)]$ of OCT. However, unlike theoretical design, the laboratory cost functional $K[\tilde{\epsilon}(\omega)]$ can contain only quantities which may be explicitly measured. At a minimum, this will include the laser control knob settings and an observation of the target objective. The laboratory cost functional $K[\tilde{\epsilon}(\omega)]$ is shown to depend on the Fourier components of the field $\tilde{\epsilon}(\omega)$, as this is coincident with current pulse shaping technology operating in the frequency domain.[7,8] The associated control knobs are the phases and/or amplitudes at a set of discrete field frequencies. Other observations may also be performed along the evolutionary path of the system and used to guide the molecule toward its objective. The latter extra data might be especially useful for seeking the control of complex molecular systems, where the wave packet evolution involves the excitation of a number of atoms during its excursion toward the target. This additional observational data could be employed to keep the wave packet evolution generally along a desired track toward the target goal. Pulling together all of these statements, a reasonable choice for the closed-loop OCE cost functional is

$$K = (\langle O(T)\rangle - O^*)^2 + \int \sigma(\omega) f(\tilde{\epsilon}(\omega))\, \mathrm{d}\omega + \sum_r \sum_l \sigma_{lr}[\langle O_l(t_r)\rangle - O_l^*(t_r)]^2 \quad (3)$$

The first term serves to steer the system to the target O^* at the final time T, while the second term contains the positive field cost function $f(\tilde{\epsilon}(\omega))$ that is typically dependent on the phases and/or amplitudes of the field, and the last term corresponds to intermediate observations at a sequence of times t_r, $r = 1, 2, ...$, with possibly several observable operators O_l, $l = 1, 2, ...$, all aiming to steer the dynamics toward a satisfactory minimization of the first term in eq 3. A specification of the final and intermediate times, T and t_r, $r = 1, 2, ...$, respectively, may be avoided, if desired, by using the total accumulated flux into the associated target states as the objectives. The positive weights $\sigma(\omega)$ and σ_{lr} in eq 3 balance the contributions of the various terms. In the current closed-loop experiments,[9,12−16] the OCE cost functional only explicitly contained the first target term, and any implicit costs on the field were borne by constraints inherent in the laser apparatus. It can be important to choose $f(\tilde{\epsilon}(\omega))$ such that the controller is guided to recognize that some phase and/or amplitude frequency control knobs need not be exercised, as they have little impact on the evolution to the target. Elimination of such extraneous control knob variations may accelerate the OCE convergence process, but most importantly, their elimination will leave a control field that contains only physically relevant features. This latter issue is important when attempting a physical interpretation of the controlled dynamics through an examination of the final optimal field structure.

An intriguing prospect is operation of the closed-loop OCE process without any prior OCT zeroth-order field guidance $\epsilon_0(t)$. Minimally, the input would originate from knowledge of appropriate portions of the system spectrum relevant for the posed objectives. The initial computer simulations[17] of OCE involved "going in blind" in this fashion, and current experiments,[9,12−16] including some on rather complex systems, have shown the approach to be successful. It remains to be seen if going in blind will be a broadly applicable technique, but the positive current results suggest that one future focus of OCT should be toward rapidly generated field design estimates capable of producing at least a minimal signal in the target for iterative OCE refinement. The closed-loop OCE process has the attractive feature of being robust in the sense that only those controls that produce stable manipulations will be identified in the laboratory.

As a forerunner to actual molecular control, a number of closed-loop experiments have been executed with the goal of generating a particular optical pulse shape.[16,29] In this context, the sample is replaced by an appropriate laser pulse shape detector. The experiments by Meshulach and Silberberg demonstrate that OCE can be very effective for learning two-photon atomic-level population control.[9] In their OCE molecular control experiments, Wilson et al. manipulated the emission of a laser dye using laser chirp control knobs.[12] Gerber et al. used laser phase modulation to control the ratios of products in the fragmentation and ionization of two organometallic compounds.[13] Sauerbrey et al. controlled the relative dynamics of two electronic states of CsCl.[14] Bucksbaum et al. controlled the wave packet structure of an H atom.[15] In a separate experiment,

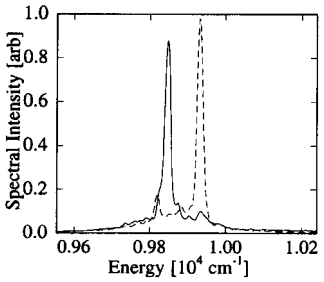

FIGURE 2. Stimulated Raman scattering spectra for the C−H stretch modes of CH_3OH excited by intense shaped ultrafast radiation using the closed-loop OCE process in Figure 1. By specification of distinct objectives, optimization was achieved for either the asymmetric (solid) or the symmetric (dashed) mode. The optimal pulses producing these two results are quite different.[21]

Bucksbaum et al. demonstrated OCE manipulation of the vibrational population of methyl alcohol.[21] An illustration from the latter work is shown in Figure 2, where closed-loop OCE was capable of finding distinct controls to steer the molecule clearly to either the symmetric or asymmetric C−H stretching modes of CH_3OH. The experiments by Wilson et al.[13] and Bucksbaum et al.[15] were carried out in the solution phase, suggesting a considerable degree of robustness of the closed-loop OCE process to environmental disturbances. Some of the cases may be considered easier in terms of their objectives than others, but collectively, these closed-loop OCE results[12−16,21,29] are most encouraging for the further growth of the field. A reasonable expectation is that *all* future molecular control experiments will involve closed-loop learning, except perhaps in special cases where high-quality OCT or other design techniques may be performed for open-loop guidance.[10,31]

IV. Controlled Inverse Dynamics for Learning about Molecules

The current emphasis in the molecular control field is on making desirable molecular transformations, and the same comment also applies to allied applications in manipulating electron motion in semiconductors,[33] quantum computing,[34] atom lasers,[35] etc. Going beyond the objective of making things, consideration has been initiated toward utilizing the ability to control molecular motion as a means for possibly learning more about molecules[23,24] and especially their intramolecular potential surfaces. The prospect for such a development follows from the subtle motion of evolving wave packets depending sensitively on details of the Hamiltonian.

Besides ab initio electronic structure calculations, the main source of information about intramolecular potentials has come from continuous wave spectral data. Currently, there is no dearth of such data, as spectrometers have evolved to become sophisticated instruments capable of generating massive amounts of high-quality data; the difficulty lies in identifying reliable algorithms for inverting the data to obtain the sought-after potential information. For diatomic molecules, algorithms exist, such as the classic RKR method,[36] to perform relatively reliable inversion for the potential. Also, for some special circumstances, such as inversion to extract the excited-state potential surface on the basis of a known ground-state reference potential surface, there has been some development of stable and accurate inversion algorithms.[37] However, there is no reliable algorithm for general purpose inversion. The inversion problem is inherently ill-posed, as a finite amount of data is never sufficient to fully prescribe the potential. The ill-posedness and resultant instability of the inversion is a major source of difficulty. A basic question is whether time-dependent dynamical observations offer an advantage for inversion purposes. In terms of the present quality of such data, considerable caution is called for, although the technology to perform the observations will surely improve. Most importantly, arguments suggest that appropriate time domain inversion techniques may possess special stabilizing features to more directly manage the ill-posedness of the inversion.[23,24] The ability to manipulate, and ideally focus, the wave packets in molecular configuration space can lead to a stabilization of the inversion process. In this fashion, although a given experiment may contain small incremental information about the potential, that information may be more readily obtained in a stable and unambiguous fashion.

Controlled inverse dynamics can involve various types of pump−probe data, but most intriguing is the prospect of obtaining ultrafast high-resolution imaging data, possibly from X-ray diffraction,[38] electron diffraction,[39] Coulomb explosions,[40] or other means. This latter data may be expressed in terms of the evolving probability density $\rho(x,t) = |\psi(x,t)|^2$, over the multidimensional configuration space x of the molecule. Although the overall complex phase of the wave function is irretrievably lost, such data are rich in information. Recent work has shown that an inverse algorithm[24] based on $\rho(x,t)$ may be developed with some special characteristics: (a) no knowledge is required of the pulse exciting the molecular dynamics, (b) the imaging process does not have to be synchronized with the excitation pulse, (c) it is not necessary to solve Schrödinger's equation, although the inversion is fully quantum mechanical, and (d) the desired potential is the solution to a rigorous linear integral equation with stable properties.

A most significant enabling development for OCE was the introduction[12−21,29] of the closed-loop learning process in Figure 1. The learning algorithm in the OCE loop guides the controller on the basis of specific input knowledge of the desired target state. For inversion, it would also be very attractive to introduce a closed-loop laboratory learning process. However, a significant difference with OCE arises, as in the case of inversion the sought-after potential surface target is not known beforehand. Nevertheless, a variant of the closed-loop process appears feasible for learning about molecular Hamiltonians,[41] and Figure 3 illustrates the components involved. The actual inversion would occur in a separate module containing software particular to the type of data involved. Upon the ith excursion around the loop, the latter module updates

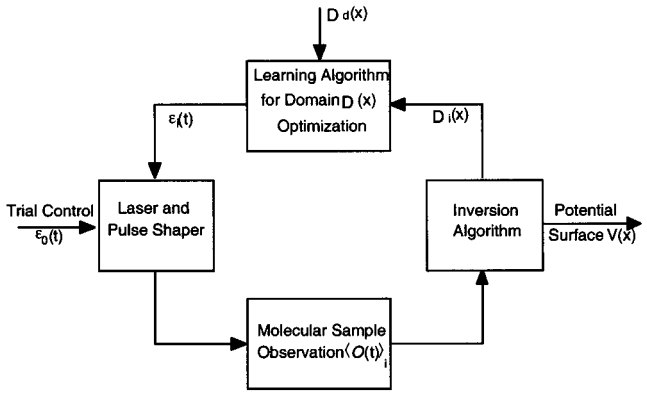

FIGURE 3. Closed-loop process for learning about molecular Hamiltonians (i.e., potential surfaces). The experiments are initiated by specifying a desired domain $D_d(x)$ to learn about the potential $V(x)$. The first experiment is started with a trial control $\epsilon_0(t)$, possibly based on some zeroth-order knowledge about the molecule. In contrast to the closed-loop control in Figure 1, here the learning algorithm serves the purpose of directing the laser to maintain the dynamics in a stable inversion domain $D(x)$, guided by knowledge of the desired domain $D_d(x)$ and the currently observed stable domain $D_i(x)$ for the ith excursion around the loop. The actual inversion is carried out as a separate operation by software that is particular to the type of data involved.

the inverted potential, as well as provides an estimate for the domain of stability $D_i(x)$ within which the potential may be reliably extracted. The learning algorithm guiding the generation of the laser control field $\epsilon(t)$ operates to ensure that the dynamics remain in the domain of stability by considering the current domain $D_i(x)$ and the input desired domain $D_d(x)$. A comparison of Figures 1 and 3 shows the similarities and significant differences between closed-loop control and inversion.

Loop closure for learning about molecules has many attractive features, and the full exploration and development of the concept await future research. Regardless of the nature of the experiments, two critical issues involve the generation of the appropriate quality data and the speed of the inversion step in the loop. Considering the latter point, it may be wise to employ a sequence of optimally controlled incremental experiments sampling the configuration space in a stepwise fashion (i.e., an overall combination of the processes in Figures 1 and 3). However, the inversion algorithm may still operate in a fully nonperturbative fashion, as suggested, for example, in the case of ultrafast high-resolution imaging data.[24,41] Although the implementation of the full loop in Figure 3 is not achieved to date, we optimistically expect to see its performance in the near future. Ultimately, the prospect of creating smart instruments capable of learning and adapting themselves in an evolutionary fashion opens up applications beyond those of molecular control alone.

V. A Look Ahead

Recent research has produced significant advances toward the goals presented approximately 40 years ago of employing lasers to alter the pathways of molecular dynamical events. The subject now has a firm conceptual foundation, and the proper theoretical tools and laboratory techniques are coming online. One of the most important steps ahead will be laboratory studies over a broad sampling of systems brought under dynamical control. A body of such OCT−OCE studies could begin to provide the information for identifying classes of quantum control mechanisms, as well as revealing the boundaries of what may be achieved. As yet, no characteristic mechanisms or rules of thumb exist for controlling quantum dynamics phenomena, other than the simple criteria of exploiting optical resonances and selection rules. Especially important will be further testing of the degree to which OCE may operate by going in blind,[17] or with minimal guidance from OCT. Regardless of the latter outcome, OCT will continue to be valuable, at least as an initial testing ground for finding potential objectives worthy of laboratory pursuit.

A special consideration for molecular control is whether manipulation in the coherent quantum mechanical regime has any advantages over operating in the more traditional incoherent manner. The full answer to this question must await evidence from further studies, but to some degree, the issue may be ephemeral. Quantum computing[34] appears to be the only current application of controlled atomic- and molecular-scale phenomena which critically depends on the operations being performed in a fully quantum mechanical fashion. Molecular motion, aside from that involving electronic excitation, typically operates in a gray-scale regime between classical and quantum dynamics. As effective control must work cooperatively with the dynamics, we may expect control to operate somewhere in the gray scale. This circumstance should not diminish the significance of laser-driven molecular dynamics phenomena, as even in the fully incoherent regime, an optimal tailored sequence of laser pulses still could prove to be quite effective in creating desired products.

Most of the emphasis in this Account and that of the quantum control community has been on *manipulation* of molecular dynamics. This type of goal will remain important, but equally significant will be the redirection of these same tools to learn more about molecular dynamics phenomena, and especially intramolecular potentials. Knowledge of potentials is at the heart of molecular dynamics, with or without control. Controlled dynamics phenomena sensitively depend on these potentials, and arguments suggest[23,24,41] that controlled temporal data may be an ideal source of information for inversion, to extract detailed knowledge about system Hamiltonians. Notions of optimization will likely play an essential role here, including the identification of optimal fields to more reliably and stably extract Hamiltonian information. An important enabling technology for control and learning about molecules is the emerging ability to perform massive numbers of high-throughput pump−probe experiments, which may be characterized as combinatorial laser chemistry.

Science projects lasting nearly 40 years are quite rare, and laser control over molecular motion is an unusual

example. The persistence of the scientific community involved attests to the significance of this endeavor. Tangible evidence of success is now finally in hand, and the overall field is expected to blossom in a variety of directions in the coming years.

The authors acknowledge support from the National Science Foundation and Department of Defense. We also thank Professor P. Bucksbaum for providing the results shown in Figure 2.

References

(1) Warren, W. S.; Rabitz, H.; Dahleh, M. Coherent Control of Quantum Dynamics: The Dream is Alive. *Science* **1993**, *259*, 1581−1589.

(2) Tannor, D. J.; Rice, S. A. Control of Selectivity of Chemical Reaction via Control of Wave Packet Evolution. *J. Chem. Phys.* **1985**, *83*, 5013−5018.

(3) Shapiro, M.; Brumer, P. Laser Control of Product Quantum State Populations in Unimolecular Reactions. *J. Chem. Phys.* **1986**, *84*, 4013−4104.

(4) Peirce, A. P.; Dahleh, M. A.; Rabitz, H. Optimal Control of Quantum Mechanical Systems−Existence, Numerical Approximation, and Applications. *Phys. Rev. A* **1988**, *37*, 4950−4964.

(5) Shi, S.; Rabitz, H. Selective Excitation in Harmonic Molecular Systems by Optimally Designed Fields. *Chem. Phys.* **1989**, *139*, 185−199.

(6) Kosloff, R.; Rice, S. A.; Gaspard, P.; Tersigni, S.; Tannor, D. J. Wavepacket Dancing−Achieving Chemical Selectivity by Shaping Light Pulses. *Chem. Phys.* **1989**, *139*, 201−220.

(7) Weiner, A.; Leaird, D.; Patel, J.;Wullert, J. Programmable Shaping of Femtosecond Optical Pulses by us of 128-Element Liquid-Crystal Phase Modulation. *IEEE J. Quantum Electron.* **1992**, *28* 908−920.

(8) Kawashima, H.; Wefers, M. M.; Nelson, K. A. Femtosecond Pulse Shaping, Multi-pulse Spectroscopy, and Optical Control. *Annu. Rev. Phys. Chem.* **1995**, *46*, 627−656.

(9) Meshulach, D.; Silberberg, Y. Coherent Quantum Control of Two-photon Transitions by a Femtosecond Laser Pulse. *Nature* **1998**, *396*, 239−242.

(10) Bardeen, C. J.; Che, J. W.; Wilson, K. R.; Yakovlev, V. V.; Apkarian, V. A.; Martens, C. C.; Zadoyan, R.; Kohler, B.; Messina, M. Quantum Control of I_2 in the Gas Phase and in Condensed Phase Solid Kr Matrix. *J. Chem. Phys.* **1997**, *106*, 8486−8503.

(11) Kleiman, V. D.; Zhu, L. C.; Allen, J.; Gordon, R. J. Coherent Control over the Photodissociation of CH_3I. *J. Chem. Phys.* **1995**, *103*, 10800−10803.

(12) Bardeen, C. J.; Yakovlev, V. V.; Wilson, K. R.; Carpenter, S. D.; Weber, P. M.; Warren, W. S. Feedback Quantum Control of Molecular Electronic Population Transfer. *Chem. Phys. Lett.* **1997**, *280*, 151−158.

(13) Assion, A.; Baumert, T.; Bergt, M.; Brixner, T.; Kiefer, B.; Seyfried, V.; Strehle, M.; Gerber, G. Control of Chemical Reactions by Feedback Optimized Phase Shaped Femtosecond Laser Pulses. *Science* **1998**, *282*, 919−922.

(14) Feurer, T.; Glass, A.; Rozgonyi, T.; Sauerbrey, R.; Szabo, G. Coherent Control of the Photodissociation Process of a Diatomic Molecule using a Feedback Controlled Self-Learning fs-Laser System. *Phys. Rev. Lett.* **1999**, submitted.

(15) Weinacht, T. C.; Ahn, J.; Bucksbaum, P. H. Controlling the Shape of a Quantum Wavefunction. *Nature* **1999**, *397*, 233−235.

(16) Meshulach, D.; Yelin, D.; Silberberg, Y. Adaptive Real-time Femtosecond Pulse Shaping. *J. Am. Opt. Soc. B* **1998**, *15*, 1615−1619.

(17) Judson, R. S.; Rabitz, H. Teaching Lasers to Control Molecules. *Phys. Rev. Lett.* **1992**, *68*, 1500−1503. Rabitz, H.; Shi, S. Optimal Control of Molecular Motion: Making Molecules Dance. In *Advances in Molecular Vibrations and Collision Dynamics*; Bowman, J., Ed.; (JAI Press: Greenwich, CT, 1991; Vol. 1, Part A, pp 187−214.

(18) Gross, P.; Neuhauser, D.; Rabitz, H. Teaching Lasers to Control Molecules in the Presence of Laboratory Field Uncertainty and Measurement Imprecision. *J. Chem. Phys.* **1993**, *98*, 4557−4566.

(19) Toth, G. J.; Lorincz, A.; Rabitz, H. The Effect of Control Field and Measurement Imprecision on Laboratory Feedback Control of Quantum Systems. *J. Chem. Phys.* **1994**, *101*, 3715−3722.

(20) Phan, M. Q.; Rabitz, H. A Self-Guided Algorithm for Learning Control of Quantum Mechanical Systems. *J. Chem. Phys.* **1999**, *110*, 34−41.

(21) Weinacht, T.; White, J.; Bucksbaum, P. Toward Strong Field Mode-Selective Chemistry. *J. Phys. Chem. A* **1999**, *103*, 10166−10168.

(22) Shi, S.; Rabitz, H. Quantum Mechanical Optimal Control of Physical Observables in Microsystems. *J. Chem. Phys.* **1990**, *92*, 364−376.

(23) Lu, Z.-M.; Rabitz, H. Unified formulation for control and inversion of molecular dynamics. *J. Phys. Chem.* **1995**, *99*, 13731−13735.

(24) Zhu, W.; Rabitz, H. Potential Surfaces from the Inversion of Time Dependent Probability Density Data. *J. Chem. Phys.* **1999**, *111*, 472−480.

(25) Demiralp, M.; Rabitz, H. Optimally controlled quantum molecular dynamics: A perturbation formulation and the existence of multiple solutions. *Phys. Rev. A* **1993**, *47*, 809−816.

(26) Zhu, W.; Botina, J.; Rabitz, H. Rapidly Convergent Iteration Methods for Quantum Optimal Control of Population. *J. Chem. Phys.* **1998**, *108*, 1953−1963.

(27) Ohtsuki, Y.; Zhu, W.; Rabitz, H. A Monotonically Convergent Algorithm for Quantum Optimal Control with Dissipation. *J. Chem. Phys.* **1999**, *110*, 9825−9832.

(28) Zhang, H.; Rabitz, H. Robust Optimal Control of Quantum Molecular Systems in the Presence of Disturbances and Uncertainties. *Phys. Rev. A* **1994**, *49*, 2241−2254.

(29) Baumert, T.; Brixner, T.; Seyfried, V.; Strehle, M.; Gerber, G. Femtosecond Pulse Shaping by an Evolutionary Algorithm with Feedback. *Appl. Phys. B* **1997**, *65*, 779−782.

(30) Tull, J. X.; Dugan, M. A.; Warren, W. S. High-Resolution Acousto-Optic Shaping of Unamplified and Amplified Femtosecond Laser Pulses. *J. Opt. Soc. Am. B* **1997**, 2348−2358.

(31) Uberna, R.; Amitay, Z.; Loomis, R.; Leone, S. Phase Control of Wave Packet Dynamics using Shaped Femtosecond Pulses. *Faraday Discuss. Chem. Soc.* **1999**, *113*, 385−400.

(32) Wiseman, H. M. Feedback in Open Quantum Systems. *Mod. Phys. Lett.* **1995**, *B9*, 629−654.

(33) Potz, W. Coherent Control of Terahertz Radiation from Semiconductor Nanostructures. *Appl. Phys. Lett.* **1998**, *72*, 3002−3004.

(34) Williams, C.; Clearwater, S. *Explorations in Quantum Computing*; Springer-Verlag: New York, 1997.

(35) Vitanov, N. V.; Suominen, K. A. Time-dependent Control of Ultracold Atoms in Magnetic Traps. *Phys. Rev. A* **1997**, *56*, 4377−4380.

(36) Schinke, R. *Photodissociation Dynamics: Spectroscopy and Fragmentation of Small Polyatomic Molecules*; Cambridge University Press: New York, 1993.

(37) Shapiro, M. Spectroscopic Wave Function Imaging and Potential Inversion. *J. Phys. Chem. A* **1996**, *100*, 7859−7866.

(38) Krause, J.; Schafer, K.; Ben-Nun, M.; Wilson, K. Creating and Detecting Shaped Rydberg Wave Packets. *Phys. Rev. Lett.* **1997**, *79*, 4978−4981.

(39) Williamson, J.; Cao, J.; Ihee, H.; Frey, H.; Zewail, A. Clocking Transient Chemical Changes by Ultrafast Electron Diffraction. *Nature* **1997**, *386*, 159−162.

(40) Stapelfeldt, H.; Constant, E.; Sakai, H.; Corkum, P. Time-resolved Coulomb Explosion Imaging: A Method to Measure Structure and Dynamics of Molecular Nuclear Wave Packets. *Phys. Rev. A* **1998**, *58*, 426−433.

(41) Zhu, W.; Rabitz, H. Molecular Dipole Function Inversion from Time-dependent Probability Density and Electric Field Data. *J. Phys. Chem. A* **1999**, *103*, 10187−10193.

AR990054H