Teaching Lasers to Control Molecules

Richard S. Judson[a]
Center for Computational Engineering, Sandia National Laboratories, Livermore, California 94551-0969

Herschel Rabitz
Department of Chemistry, Princeton University, Princeton, New Jersey 08544
(Received 26 August 1991)

We simulate a method to teach a laser pulse sequences to excite specified molecular states. We use a learning procedure to direct the production of pulses based on "fitness" information provided by a laboratory measurement device. Over a series of pulses the algorithm learns an optimal sequence. The experimental apparatus, which consists of a laser, a sample of molecules, and a measurement device, acts as an analog computer that solves Schrödinger's equation *exactly*, in real time. We simulate an apparatus that learns to excite specified rotational states in a diatomic molecule.

PACS numbers: 33.80.−b

In this Letter we suggest a method for designing laser pulses to control the motion of molecules, using an experimental apparatus as an analog computer that learns an optimal pulse sequence in real time. Much theoretical work [1–3] has gone into the design of laser sequences that can drive a reaction into a desired, thermally inaccessible state, but successful experimental implementation of these ideas has been an elusive goal. A major stumbling block is the complicated nature of molecular Hamiltonians which typically have many degrees of freedom tightly coupled together, all of which may have to be simultaneously controlled. It has been shown theoretically, under appropriate conditions, that molecules can be controlled, i.e., fields can be designed to drive them into desired final states [1]. However, translating these results from theory to experiment has not been possible until now. An underlying problem is that the methods for designing fields requires full knowledge of the molecular Hamiltonian which is known only approximately for systems with more than two or three atoms. In addition, laboratory uncertainties can arise due to optical pulse generation errors of various types. Fields designed theoretically on the basis of an approximate Hamiltonian may not be sufficiently robust to tolerate errors arising from the Hamiltonian as well as laboratory introduced uncertainties.

The method we suggest leapfrogs these difficulties by using an experimental apparatus as an analog computer that solves Schrödinger's equation exactly with the true laboratory field. Essentially the same system design concepts already introduced [1,3] can be adopted with laboratory experiments supplanting the need to solve any equations of motion on the computer. By working with the molecular sample of interest, the apparatus acts as an input-output device capable of reliably reporting the action of any introduced field upon the molecules. The apparatus is then coupled to a learning algorithm capable of recognizing patterns in the input-output measurement relationships and thus guiding an iterative sequence of new experiments. The iteration is facilitated by a cost functional as in current molecular control theory, but now only containing costs for the target state and laboratory considerations (e.g., constraints on the form of the field). The iteration and learning process would continue until satisfactory convergence is reached. The overall procedure is an example of an adaptive learning network, and a key element is the rapidity with which laser pulses can be created and the resultant effects probed (i.e., the pump-probe duty cycle). The latest laboratory tools indicate a conservative duty cycle of $\ll 1$ sec allowing for many iterations to be performed on a comfortable laboratory time scale. The remainder of this Letter will simulate this methodology using a genetic algorithm (GA) [4], although other methods might also be employed. To some degree, the ability of current control methodology [1] to successfully design laser pulses for particular applications also provides evidence for the ability to teach lasers to manipulate molecules.

Genetic algorithms are global optimization methods based on several metaphors from biological evolution. The first is the concept of a breeding population in which individuals who are more "fit" in some measurable sense will have a higher chance of producing offspring and passing their genetic information onto succeeding generations. The second is the concept of crossover in which a child's genetic material is a mixture of his parents. The third concept is that of mutation, meaning that genetic material is occasionally corrupted, leading to individuals who may or may not be more fit than they would have been otherwise, but always maintaining a certain level of genetic diversity in the population.

The apparatus we model needs to consist of a sample of the molecules of interest, a laser whose pulse sequence is supplied by a computer and a measurement device that feeds final population distributions or other observables back to the controlling computer. The genetic algorithm code runs on the controlling computer, supplying pulse sequences to the laser and receiving fitness values (some observable function of the molecular state) from the measurement device. Over several generations, the system as a whole will seek to optimize the fields.

The genetic algorithm is implemented as follows. An

individual (i.e., a single pulse sequence) is coded for by a "gene" which is a bit string of length N_{gene} that can be uniquely decoded to give the pulse sequence. A fitness function is defined that can discriminate between individuals. To drive molecules into state $\bar{j}$ this might be $\sum_j (\delta_{j,\bar{j}} - \rho_j)^2$, where ρ_j is the corresponding population of state j. Control over other observable properties of the molecule could also be formulated. An initial population of individuals (N_{pop}) is formed by choosing N_{pop} bit strings, often initially at random, and evaluating each individual's fitness. Children of these generation-1 parents are formed as follows. All the parents are ranked by fitness and the highest fitness individuals are placed directly into generation 2 with no change. From the remaining parents, pairs of individuals are chosen and their genes are crossed over to form genes of the remaining generation-2 individuals. The crossover is effected by taking some subset of the bits from parent 1 and the complementary set of bits from parent 2 and combining them to form the gene of child 1. The remaining bits from the two parents are combined to form the gene of child 2. Additionally, during replication there is a small probability of a bit flip or mutation in a gene. This serves primarily to prevent premature convergence in which a single very fit individual takes over the entire population. To bound the magnitude of the effect of mutations, the binary genes are usually gray coded. (A gray-coded binary number has the property that changing any bit between 1 and 0 changes the magnitude of the number by 1.) An interesting point about mutation is that it causes a GA to degenerate to a Metropolis algorithm in the situation that crossover is ineffective. The lowest fitness individuals in each generation may be discarded and replaced by children of more fit individuals. Many variants of the basic algorithm exist in the literature, but the basic model outlined here is the one we have used.

The experiment described above is designed to be executed with actual laboratory pump and observe-probe measurements. At this stage, a simulation of the algorithm could be carried out for virtually any molecular system capable of being modeled on the computer. For simplicity and ease of computation, here we show the results of a simulation where we design a field to derive a rigid diatomic molecule into a particular rotational state. The specific shape of the field we calculate is not as important as the recognition that the code calculating the time evolution of the molecules could be replaced with an actual experimental apparatus and the genetic algorithm would help the apparatus to find an optimal controlling field in exactly the same way it teaches our model code. The molecule we consider is KCl which we assume starts in the ground state. The target state we aim for is $j=3$, $m=0$ although the method works equally well for other final states. In the model code, we use states $j=0$ to 5. The laser is assumed to be linearly polarized in the z direction so only $m=0$ states are sampled. The field is referred to as a "laser" here, although it actually operates

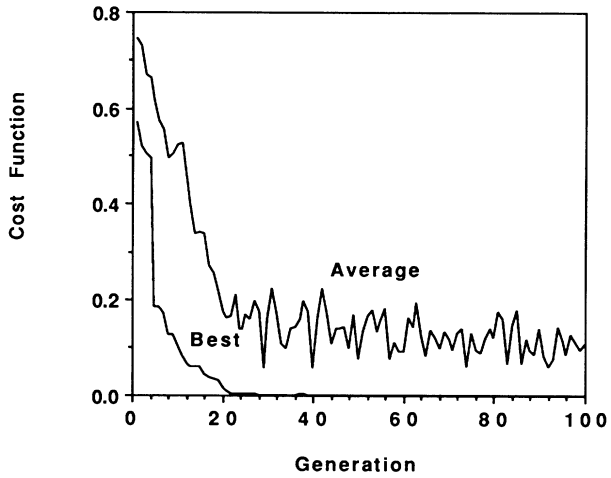

FIG. 1. The cost function vs GA generation for a case where no bias or filter is used to guide the determination of the optimal field. The cost is $\sum_j (\delta_{j,\bar{j}} - \rho_j)^2$, where $\bar{j}=3$ is the target state and ρ_j is the corresponding population of state j. The top curve gives the average value for the population and the bottom curve gives the value for the best individual in each generation. The gene specified the field amplitude at a series of 128 discrete times over 1 nsec.

in the microwave regime. In practice, currently available optical pulse shaping techniques [5] suggest that a Raman excitation experiment coupled to a learning algorithm might be practical for execution.

In any real problem the spectrum of the test molecule would surely be obtained to identify the appropriate regions (i.e., to find a filter) for the learning algorithm to operate in. Such a simulation will be considered below,

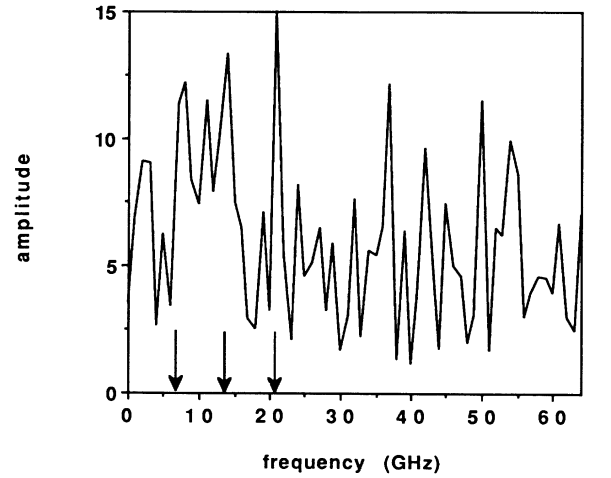

FIG. 2. The frequency spectrum for the optimal field with the gene described in Fig. 1. The arrows indicate the positions of the resonant transitions from $j=0$ to 1, 1 to 2, and 2 to 3. The lack of discriminating genetic pressure away from the region of spectral absorption allows for a congested, broad spectrum.

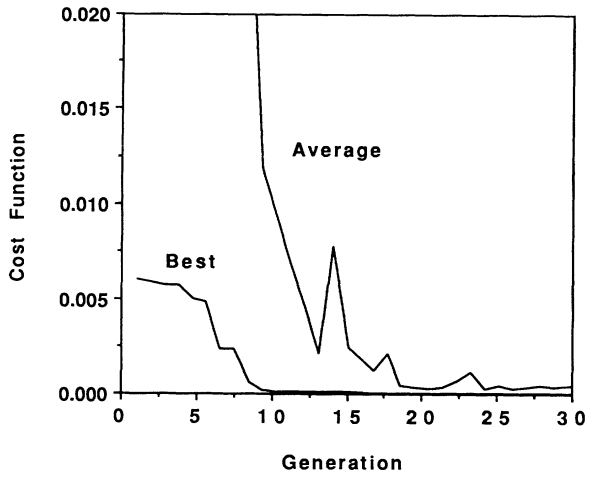

FIG. 3. Same as Fig. 1 except that the gene specified the real and imaginary parts of the spectrum. The raw gene was then passed through a filter (as described in the text) before being Fourier transformed to provide the time sequence.

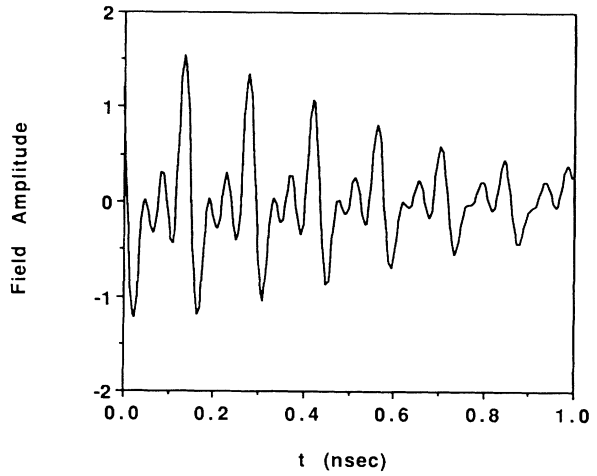

FIG. 4. The optimal field (in kV/cm) for driving the transition from $j=0$ to $j=3$ for the gene described in Fig. 3.

but first it is useful to show the power of the algorithm to teach the laser even when *no* spectral information is introduced. In the first example, the GA gene is a sequence of 128 points in the time series for the electric field over 1 nsec. The population contained 50 individuals with each initial gene consisting of 128 amplitudes chosen from a random uniform distribution in the range 0 to 1 and then scaled to a maximum of 5 kV/cm. All other parameters (mutation rates, crossover rates, etc.) were the defaults provided by GENESIS [6], the GA implementation we use. The GA was then run with the resulting cost as a function of generation number shown in Fig. 1. This exponential decrease in cost is typical of other GA applications. The algorithm performed very well with the best members of the population yielding $>99.9\%$ occupation in state $j=3$. The average population also performed well. The best field is very noisy as is indicated by its spectrum, shown in Fig. 2. It is not surprising that the spectrum is highly congested as this experiment contained no filtering, no fluence penalty, or any other genetic pressure to have spectral structure except where it is actually absorbed by the molecule. There are strong peaks near the resonant frequences (shown by the arrows) but there are equally strong peaks off resonance. This illustrates the power of the algorithm to recognize a small but significant control advantage in one of the initial random family members and then amplify this preference through successive generations to yield excellent results. Optimal control was literally lifted out of the noise.

In a second implementation, the GA gene was taken as a spectrum consisting of the real and imaginary parts of the field amplitude at a series of frequencies. A total of 200 amplitudes are uniformly spaced over the frequency range from 6 to 24 GHz (the KCl fundamental frequency ν_0 is 6.97 GHz). Each of the amplitudes can range from 0 to 1 in magnitude. Because prior spectral knowledge is available, the raw spectrum provided by the GA is then passed through a filter consisting of Gaussians of width 0.5 GHz centered on ν_0, $2\nu_0$, and $3\nu_0$ which are the transition frequencies needed to carry the system up the ladder from $j=0$ to $j=3$. The optimal field is essentially a repeating sequence of pulses with decreasing amplitude. The spectrum is then Fourier transformed and scaled to a maximum of 5 kV/cm to provide the actual pulse sequence. The pulses last a total of 1 nsec. The initial population of raw spectra was totally random; i.e., each spectrum consists of white noise. This insures that all frequencies that might be needed later on exist initially. Of course, the spectra are biased by the Gaussian filter. During the simulation, the GA essentially learns which

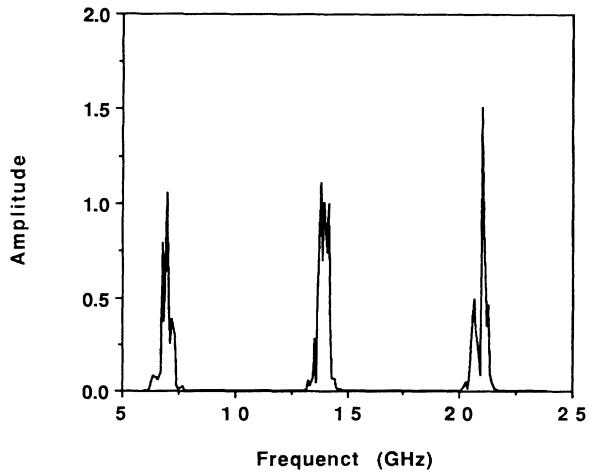

FIG. 5. The frequency spectrum for the field in Fig. 4. This figure should be contrasted with Fig. 2 where no filtering was imposed.

frequencies are needed and it discards the rest. The population size was 50. In Fig. 3, we show the cost for the population of the target state $j=3$ for the best individual in a generation and the average population cost versus generation. In Fig. 4, we show the final best field that transfers better than 99% of the population from $j=0$ to $j=3$. Figure 5 shows the spectrum corresponding to the field in Fig. 4. A total of about 30 generations was needed to reach this level of optimization. Some of the fine structure in Fig. 5 could possibly also be eliminated with little penalty for the target state by the introduction of additional constraints on the field. It is obvious from Fig. 4 that this field is completely nonintuitive which is one of the powers of learning methods such as genetic algorithms. They are able to intelligently search parameter space relatively unfettered by the limits of human intuition.

In conclusion, we have demonstrated that an adaptive learning procedure can teach a laser to selectively excite chosen states of a molecule. Constraints on the form or amplitude of the driving field can readily be included in the learning algorithm in accord with laboratory capabilities. Comparison of Figs. 2 and 5 indicates that the introduction of intelligent physical bias in the control field can yield forms more readily generated in the laboratory. Clearly, to implement this procedure experimentally, we must have available the necessary laser pulse tools. This is an evolving technology which should already be capable of being adapted to the learning technique for some current applications.

R.S.J. acknowledges support by the Department of Energy under Contract No. DE-ACO4-76DP00789. H.R. acknowledges AFOSR and ONR for support of this research.

[a]Electronic address: rsjuds@sandia.llnl.gov.

[1] S. Shi and H. Rabitz, J. Chem. Phys. **92**, 364 (1990).

[2] D. Tannor and S. A. Rice, J. Chem. Phys. **83**, 5013 (1985); P. Brummer and M. Shapiro, Chem. Phys. Lett. **126**, 541 (1986).

[3] R. S. Judson, K. K. Lehmann, H. Rabitz, and W. S. Warren, J. Mol. Struct. **223**, 425 (1990).

[4] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning* (Addison-Wesley, Reading, MA, 1989); L. Davis, *Handbook of Genetic Algorithms* (Van Nostrand Reinhold, New York, 1991); *Proceedings of the Third International Conference on Genetic Algorithms*, edited by J. D. Schaffer (Morgan Kaufman, San Mateo, 1989); in *Proceedings of the Fourth International Conference on Genetic Algorithms*, edited by R. K. Belew and L. B. Booker (Morgan Kaufman, San Mateo, 1991).

[5] R. Thurston, J. Heritage, A. Weiner, and W. Tomlinson, IEEE J. Quantum Electron. **22**, 682 (1986); M. Haner and W. Warren, Appl. Phys. Lett. **52**, 1459 (1988).

[6] J. J. Grefenstette and N. N. Schraudolph, GENESIS 1.2ucsd (available via anonymous ftp from the Artificial Life Archive server, iuvax.cs.indiana.edu [129.79.254.192] in the directory /pub/alife/software/unix/GAucsd).