Quant. Phys. Lett. 1 No. 1, 1-19 (2012)

Quantum Physics Letters
An International Journal

Control in the Sciences Over Vast Length and Time Scales

Herschel Rabitz
Department of Chemistry, Princeton University, Princeton, NJ 08544, USA
Email : hrabitz@Princeton.edu

Received March 02, 2012; Accepted March 25, 2012
Published online: 15 April 2012

Abstract: The desire to direct the outcome of chemical, physical and biological processes is pervasive in many areas of science. A set of protocols, rules and procedures is often followed in each domain frequently resulting in very favorable outcomes. These operations can be expressed in a control framework analogous to common practice in the engineering disciplines. As a foundation for assessing the value of taking a control perspective in the sciences, the paper first presents a summary of observations found when optimally manipulating quantum dynamics phenomena, maximizing the yield of chemical and material syntheses and properties, and enhancing the outcome of directed evolution. In addition, findings from natural evolution are considered where optimization is for the fitness of a species' population; in this case a control perspective provides a mathematical framework for assessing the behavior of naturally occurring evolutionary processes. Collectively the control of phenomena in these and other areas of science involve dynamics with distinctive characteristics spanning vast length and time scales. Notwithstanding the disparate dynamical behavior in each domain, the searches for optimal controls are strikingly efficient, especially considering that the available control resources are generally very extensive. The achieved high efficiency defies intuition, as the systems subjected to control are often quite complex by any reasonable measure. The basis for the surprising degree of efficiency in finding optimal solutions can be understood by considering the topology of the underlying control landscape defined as the objective in each case as a function of the controls. The general conclusion, upon satisfaction of some key physical assumptions, is that control landscapes are expected to be devoid of traps, which could hinder the search for the best outcome. In order to bolster this conclusion additional specific details are presented regarding control behavior found (i) while manipulating quantum dynamics and (ii) in the optimization of synthesis yields and properties in chemical and material science. The observed common **Opti**mal control behavior over vast length and time scales in the **Sci**ences and its foundations are referred to as OptiSci. Further research is needed to fully understand the basis of OptiSci and its implications. However, the current findings suggest that the principles of OptiSci may have wide ranging significance including for (i) enhancing the efficiency of searches for optimal controls, (ii) manipulating phenomena that transcend traditional domains in science, (iii) the early identification of flawed experimental designs and (iv) providing the basis to automate the discovery of systematic rules for finding effective controls.

Keywords: Quantum Control; common Optimal control.

I. Introduction

A common goal in many domains of science is to seek the optimal outcome from an experiment, which aims to manipulate specified phenomena. For example, in chemistry, the objective is often to optimize either the yield of a chemical synthesis or the property of a material.[1] In biology, it is understood that natural evolution is optimizing some measure of population fitness.[2] The field of directed evolution seeks to optimize or alter the outcome of an erstwhile naturally occurring biological process (e.g., an enzyme's activity).[3] A recent domain of optimization concerns altering quantum dynamical processes typically through the application of tailored electromagnetic fields.[4] Optimization in these contexts collectively entails events spanning vast length and time scales. Notwithstanding these distinctions and many other

differences, the goal of this paper is to show that optimization in the sciences share strikingly common characteristics. The observed behavior and its formulation will be referred to as OptiSci. Optimization in the sciences is typically not expressed in a control framework, although this is standard practice in the engineering disciplines.[5] Consideration of an optimal control perspective permits the introduction of a physical and mathematical structure unifying optimization of phenomena in seemingly diverse areas of science.

Optimization in the sciences normally involves an effort based on the prescription: (1) specify an objective, (2) arrange for a relevant set of control variables and (3) seek an optimal control that maximizes (or minimizes) the objective. This prescription can be readily understood to fit optimization of quantum dynamics[4], manipulation in chemical and material science[1], directed evolution[3] and optimization goals in various other scientific disciplines. However, consideration of optimization in natural evolution[2] from a control perspective calls for care. This paper uses the notion of control in natural evolution exclusively in a mathematical and physical analysis framework. In addition, laboratory experiments with micro-organisms provide the basis to include natural evolution in the context of observing controlled behavior analogous to like experiments in other domains of science[6], as will be explained later.

The remainder of the paper is organized as follows. Section II will succinctly summarize experience found upon performing optimization in various areas of science with an emphasis placed on the extensive control resources searched over, which is in surprising juxtaposition with the observed very high efficiency for identifying effective optimal controls. The collective observations will then be drawn together leading to several questions about the findings. A hypothesis will be presented suggesting that the similar highly efficient searches for optimal controls have a common foundation lying in the topology of the associated control landscapes, despite the vast length and time scales of the phenomena involved. In this context, a landscape is the optimization objective as a function of the controls[7]. Section III will take a closer look at control behavior in optimization over quantum dynamics phenomena (Section III. A) as well as optimization of synthesis yields and properties in the chemical and material sciences (Section **III. B)**. Upon satisfaction of a few key assumptions the control landscapes will be argued to have a generic trap free topology (i.e., a trap corresponds to a suboptimal extremum whose existence could hinder efforts at finding controls that produce the global maximum outcome). Finally, Section **IV** will return again to control prospects across the sciences to address the generality of the OptiSci concepts. Tentative conclusions about the significance of this common optimal control behavior will be presented.

II. Observations Regarding Optimization in the Sciences

As background for OptiSci, this section will summarize observations from performing optimization in several domains of science. A special focus will be on the general extent of the control resources available and the reported efficiency in searching over those resources to find an optimal solution. The following subsections will separately address control in the contexts of **II. A** quantum dynamics, **II. B** chemical and material science, **II. C** natural evolution and **II. D** directed evolution. The collective findings will be drawn together in Section **II. E** to pose a few questions for assessment.

II. A. Optimization of Quantum Dynamics

Until recently, the consideration of quantum phenomena in the laboratory was largely a matter of observation of its micro- and macroscopic manifestations. Quantum mechanical processes generally occur at time scales characterized by ω^{-1} where ω is a natural transition frequency of the system. Such time scales may range over many orders of magnitude, for example, from μsec with nuclear spins out to *atto*sec for electronic degrees of freedom. Concomitantly, the length scales vary accordingly over a wide range for particles (e.g., atoms, electrons, nucleons) that are spatially confined. Control in these diverse circumstances may be considered for a variety of purposes[8], but a basic goal is to direct the motion of the system

state, either $|\psi(0)\rangle \rightarrow |\psi(T)\rangle$ for the wave function or $\rho(0) \rightarrow \rho(T)$ for the density matrix through application of a tailored external field $\varepsilon(t)$, $0 \le t \le T$. In practice, the control objective often concerns optimizing the expectation value $\langle O \rangle$ of a physical observable operator O, where $\langle O \rangle$ is either $\langle \psi(T) | O | \psi(T) \rangle$ or $Tr[\rho(T)O]$. Regardless of the particular objective, the available controls are generally drawn from the electromagnetic spectrum, often in the form of shaped radiation pulses viewed as "photonic reagents" to steer about the quantum dynamics.[9] In principle, the entire electromagnetic spectrum (or at least its portions capable of interacting with a sample) may be drawn upon as a resource. In practice, current technology limits the ability to fully utilize these resources. Nevertheless, control of quantum dynamics with photonic reagents is a burgeoning field, which we will refer to as **Opti**mal control of **Q**uantum phenomena (OptiQ).

Consideration of control over quantum phenomena in this work is confined to experiments where *optimization* of the control outcome is sought. The same perspective is taken in Sections **II. B, C and D**. Thus, experiments that apply a single photonic reagent (i.e., perhaps even insightfully chosen) and then measure the dynamical outcome are put aside here, as they do not reveal information about the level of effort needed to identify an optimal control. The current literature reports ~ 150 successful OptiQ experiments[8]. The quality of the outcome in each case ranges from satisfactory to excellent with the difference likely associated with practical constraints present on the electromagnetic resources. Typically $10^3 \sim 10^4$ algorithmically guided experiments need to be performed in order to find an optimal photonic reagent.[10] An individual experiment consists of applying a trial control field and performing an associated observation of its dynamical impact on the system. Although executing $10^3 \sim 10^4$ experiments for any particular goal would constitute an extensive effort under normal conditions, the high-duty cycle of automated laser control apparatuses[11] permits reaching numbers of experiments of this magnitude often in a period of minutes to an hour. Importantly, the performance of $10^3 \sim 10^4$ experiments to identify an optimal control solution needs to be compared with the number of possible distinct accessible control fields, which is typically on the order of $\sim 10^{100}$. The latter estimate is even conservative, considering that a laser pulse shaper[12] may have ~ 500 pixels where each is digitized to ~ 30 settings. In addition, the number of experiments needed to reach optimal performance appears to be essentially invariant to the level of quantum system complexity, while also keeping in mind that the final outcome may be limited by the available resources.[8,13] The striking degree of efficiency for finding optimal controls is also evident in other diverse domains of science,[6,14,15] as summarized in the remainder of section **II**, and this observed behavior forms the basis of OptiSci.

A special feature of most OptiQ experiments is their automated operation[10,11] carried out in the adaptive fashion sketched in Figure 1. The automation of OptiQ enables the performance of thousands of experiments guided by suitable pattern recognition algorithms seeking[16] optimization of a quantum mechanical control objective. An effective algorithm operates by identifying features in the controls that are favorable for achieving a better objective outcome. The actual scientific goals for controlling quantum phenomena are wide ranging and ever increasing with a current list shown in Figure 2 along with additional comments about the nature of the experiments.[8] The OptiQ field is still quite young and is pushing the limits of available technology to find effective controls, especially in systems with dynamics having broad spectral character. Some of the comments on the state of OptiQ in Figure 2 have their own parallel in many other areas of science where complex phenomena are manipulated (see the topics in the remainder of Section **II**). A distinction needs to be made between finding a successful outcome from performing $10^3 \sim 10^4$ excursions around the loop in Figure 1 in approximately and hour compared to the often arduous 'overhead' in setting up the experiments that may take months or years. Additional extensive effort can be required to understand the mechanism of the control induced dynamics,[17] its quantum character and possibly identify rules which may be discerned from the

findings to inform the execution of new experiments in like systems.

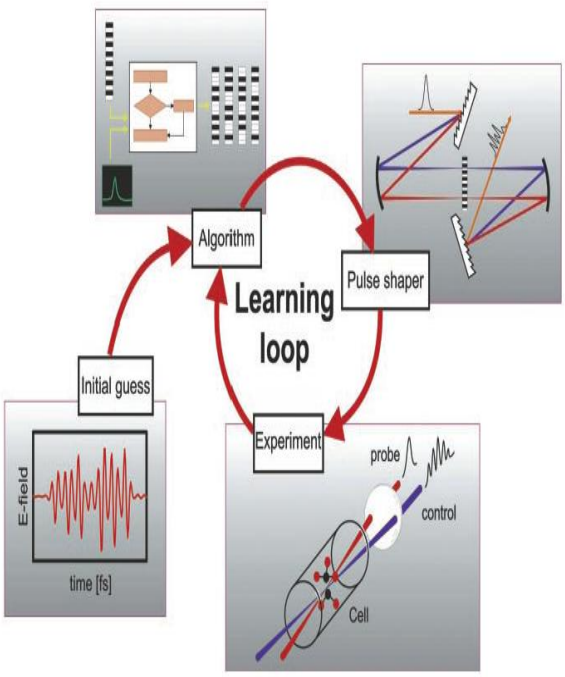

Figure 1. Automated optimization of quantum phenomena. *The sketched adaptive feedback control procedure[10] is widely employed in OptiQ to discover tailored light pulses (photonic reagents) for manipulating quantum dynamics phenomena.[18] An initial trial control field is introduced and digitized for creation by the pulse shaper. The trial shaped pulse is applied to the sample and a signal is recorded reflecting the degree of achieved control. The signal is then fed back to the algorithm which suggests a follow on photonic reagent aiming to enhance the control yield. The iterations are continued until convergence to an optimal result. Typically* $\sim 10^3 - 10^4$ *excursions around the loop are necessary to find an optimal photonic reagent, yet the number of potential distinct controls is on the order of* $\sim 10^{100}$. *The dramatic high efficiency of OptiQ experiments is one component of OptiSci.*

	Comments
1. Atomic excitation 2. Vibrational excitation 3. Vibronic excitation 4. Selective isotope excitation 5. Discrimination of very similar molecules	➢ 5~10 examples/domain
6. Ionization, single to multiple 7. Molecular isomerization 8. Molecular fragmentation 9. Molecular alignment 10. Nonlinear spectroscopy 11. Ultrafast semi-conductor optical switching	➢ Control mechanism often hard to discern
12. Biomolecular energy transfer 13. Filimentation 14. Pulse propagation is nonlinear media	➢ Rules to be discovered
15. Decelerating and trapping of molecules 16. Surface ablation 17. Subsurface imaging 18. Decoherence mitigation	➢ How quantum is the dynamics?

Figure 2: Domains subjected to optimal quantum control. *The orange highlighted table lists several current applications utilizing optimal photonic reagents deduced by the adaptive feedback procedure in Figure 1. There are presently* ~ 150 *experiments spread over the indicated domains of application.[8] The highlighted comments on the right refer to all of the collective experiments. The fact that there are only a modest number of experiments in each domain makes it difficult to discern quantum control mechanisms and possibly rules for guiding future experiments. A basic question regarding all the experiments is their degree of quantum character. However, a common goal is to maximize the objective regardless of whether the induced dynamics has a mix of quantum and classical character.*

II. B Optimization in Chemistry and Material Science

Optimization underlies many objectives in chemical and material science where finding either the highest synthesis yield or the best property value are frequent goals.[1] From a control perspective the variables employed to achieve optimization are generally drawn from chemicals, solvents, catalysts and processing conditions. It almost goes without saying that reaching high, if not optimal performance in the chemical sciences is widely attained in the laboratory and at an industrial scale. Building on that foundation, the efficiency of performing such optimization efforts is of interest in the context of this paper. Table 1 presents a cross section of results for the automated optimization of chemical and material synthesis and properties[14,15]. Each of the optimizations in Table 1 was carried out by a robotized apparatus similar in nature to that schematically indicated in Figure 1 where photonic reagents are used as controls. In the case of chemical and material science, current operations permit the employment of a modest number of variables on the order of ~ 10. Yet, with even relatively few variables, the total number of potential experiments can be quite large. The numbers (i.e., the product of the number of discrete choices for each of the control variables) indicated in the last column of Table 1 were gleaned from those references where they were included. In particular, the number of reported *accessible* control settings (i.e., experiments) range from $\sim 10^4$ out to $\sim 10^{22}$, while the number of experiments needed to *reach* optimality is often only $10^2 \sim 10^3$. Interestingly, the papers that present the number of experiments employed to reach optimality along with the total number of possible experiments typically provide no accompanying comment on the evident extreme disparity involved. Similar lack of comments also exists in the parallel photonic reagent control literature.

In summary, **Opti**mal control of synthesis yields and properties in the **Chem**ical and material sciences (OptiChem) is both widely successful and highly efficient.[14,15] Additionally, the automated performance of laboratory operations in this domain is often guided by rules riding on long experience drawn from prior studies in many laboratories. A fundamental question is why rules[19] even exist for the control of nominally highly complex chemical systems and whether their discovery may also be automated. This point will be returned to in Section **IV**. As with OptiQ using photonic reagents as controls in Section **II. A** , attempts at understanding chemical control mechanisms only follows after first optimizing the desired product or property.

Ref. Number	Number of Control Variables	Objective	Number of Experiments to Reach Optimal Outcomes	Number of Possible Control Sample Points
20	6	Binding to stromelysin	300	6.4×10^7
21	8	Propane → propene	328	NA
22	4	Inhibition of thombin	400	1.6×10^5
23	8	Propane → CO	150	NA
24	8	Propane → propene	280	NA
25	13	Propane → propene	60	NA
26	23	$NH_3 + CH_4 \rightarrow HCN$	644	NA
27	9	$CO \rightarrow CO_2$	189	NA
28	4	$CO + CO_2 + H_2 \rightarrow CH_3OH$	115	2.7×10^9
29	5	$3CO + 3H_2 \rightarrow C_2H_6O + CO_2$	160	2.4×10^{11}
30	6	$CO + CO_2 + H_2 \rightarrow CH_3OH$	235	4.7×10^9
31	10	*n*-Pentane isomerization	72	1.44×10^4
32	7	Propane → aldehydes	80	NA
33	8	Isobutane → methacrolein	90	10^9
34	8	Membrane permeability	192	9×10^{21}
35	4	Cyclohexane epoxidation	114	NA

36	3	Protein inhibition	160	10^{16}
37	6	Red luminescence	216	NA
38	7	Green luminescence	540	10^{14}
39	6	Color chromaticity	168	NA
40	8	Red luminescence	270	NA
41	7	Red luminescence	1080	NA

Table 1: Automated optimization of chemical and material synthesis and properties. *The objectives span a wide variety of applications where each experiment was performed in a robotized fashion similar to the schematic of OptiQ photonic reagent control in Figure 1. In the case of OptiChem the number of control variables (i.e., chemicals, solvents, processing conditions, etc.) is modest, as indicated in the table.*[14,15] *Typically only a few hundred iterated experiments are necessary to obtain an optimal outcome, and yet the number of possibilities is often many orders of magnitude larger as shown in the column on the right. This high degree of efficiency in OptiChem is one founding component of OptiSci.*

II. C. Optimization in Natural Evolution

The process of natural evolution garners great attention in the sciences and beyond for obvious reasons[2]. Darwin set out the principles of fitness optimization driving a species' population, and from a control perspective the variables at the most basic level are specified by the sequence of nucleic acids in a species' genome[42] whose size is $\sim 10^9$. The latter number of variables translates to $\sim 4^{10^9}$ distinct genomes considering the four possible bases. A number of this size is hard to fathom, and even if practical access to genome sequences during evolution is reduced by a large factor, the conclusion is the same that the genome of any species poses a vast space to stochastically search over for a sequence that optimizes fitness. The latter search is performed through mutation, crossover, insertion, deletion, etc., operations on the sequence of nucleic acids.[2]

Natural evolution may be viewed as an ongoing 'experiment' by nature, and certainly the existence of complex forms of life can be taken as adequate evidence that these evolutionary experiments by Nature are quite successful. From the perspective of this work, a primary focus is on the efficiency of the evolutionary optimization process. Given the sheer size of the genomic search space and considering an evolutionary time period of just $\sim 10^9$ years, one may conclude that evolution is startlingly efficient. Additional quantitative findings about natural evolution are contained in laboratory experiments with bacteria[43] and viruses[44]. These organisms turn over their populations very rapidly permitting access to large numbers of evolutionary generations on reasonable laboratory time scales. The experiments involved often can be characterized as disturb-and-observe, which is analogous to control-and-probe experiments in the OptiQ photonic reagent domain (Section **II. A**). In the case of evolution, the disturbances may be environmental or directly introduced as a mutation in the genome, with subsequent evolutionary results observed at the genotypic or phenotypic levels. This body of experimental data and its analysis has been referred to as global **Opti**mality of fitness landscapes in **Evo**lution (OptiEvo).[6] In particular, upon disturbing a population away from its optimal fitness (e.g., by changing the environment), the subsequent evolution often achieves global convergence to the original fitness value. These experiments also frequently show genotypic diversity at the latter global fitness maximum value corresponding to the existence of neutral mutations (i.e. the fitness landscape is flat at its absolute maximum). Finally, in these disturb-and-observe experiments local suboptimal traps on the fitness landscape are rarely found and appear to occur in either genomes of restricted size (e.g., in viruses) or in evolutionary experiments that have gone through relatively few generations.[6,45] In summary, Nature has access to enormous control resources and evidently can efficiently draw on those resources to readily optimize population fitness.

II. D. Optimization in Directed Evolution

Directed evolution is a laboratory effort aiming to take over some specific aspect of biological functionality and adapt it for practical purposes.[3] The goal often involves re-engineering an enzyme (i.e., changing a particular protein's sequence) so that it produces a desired chemical product. Bacteria are commonly chosen for this purpose, and the controls are the nucleic acids forming the genome of the organism. As with natural evolution, the available resources are vast. In practice, these resources are sparingly used in directed evolution for two reasons.[46] First, the organism generally must survive the re-engineering of its genome, and second, the experiments can be expensive and time consuming. Mutations are frequently performed one at a time with the results assessed for another round of variations. A major experimental step is the screening of the mutated organisms to identify those with more favorable performance at maximizing the desired product. Notwithstanding all of these circumstances, directed evolution is proving to be very successful with many notable applications, and its relation to natural evolution permits referring to it under the title of OptiEvo.

The actual number of mutations performed in directed evolution is a very small fraction of the total offered by the genome, and the process can be characterized as highly efficient. The subject is also reaching a stage where automated closed-loop operations[47] are involved much like analogous procedures in OptiQ and OptiChem, respectively summarized in Sections **II. A** and **B**. Very modest numbers of excursions around the closed-loop are typically sufficient to achieve significant optimization of a directed evolutionary product, although consideration should be given to the fact that the starting point is frequently a naturally occurring enzyme or biological process that has been 'pre-optimized' by Nature for a distinct, but like, purpose.

II. E. Questions Posed by Observations of Control Performance in the Sciences

At this juncture, we may draw together the collected findings from control in the areas of science summarized above. All of the information is succinctly presented in Table 2. The most striking observation from the performance of optimization across these domains of science is the high efficiency of the control search efforts. In all cases, the resources are very extensive, yet only a modest number of experiments are generally needed to reach optimality. The reason for this common behavior is the focus of OptiSci. Control in these diverse areas of science occurs for phenomena with distinctive dynamics across vast length and time scales. The latter clear distinctions amongst the dynamics of chemical, physical and biological phenomena is *not* the focus of this work, but rather the attention of OptiSci is on the evident high efficiency in finding optimal solutions to controlling such phenomena. Table 2 also presents a number of subjective comments about the nature of control experiments in these domains. Importantly, the execution of the experiments can be a complex venture, and this practical overhead does not figure into considering the efficiency of the control discovery process. OptiSci is concerned with the number of experiments needed to reach optimality verses the number of possible experiments. Another consideration is the desire to operate under the guidance of rules given the complexity of performing individual experiments on the way to identifying an optimal control. The most advanced examples of rule based operations are evident in the chemical and material sciences,[19] but the desire for revealing utilitarian rules exists in all areas. Additionally, the level of mechanistic understanding of the controlled phenomena greatly differs from one application domain to another.

The following questions arise from observations of control in the sciences

- Are the common features greater than the differences?
- Is the common high search efficiency a coincidence?
- Under optimization, is a unified picture operative?
- What are the practical control consequences?

One could set aside these questions and proceed as usual to seek optimal performance in any particular scientific application. The tenet of OptiSci is that continuing in this fashion may overlook fundamental and practical issues extending beyond

optimization considerations in any single application. The material in Section **III** will address these questions by expanding on the findings in two scientific areas followed by the presentation of some general conclusions on the potential significance of OptiSci in Section **IV**.

Science	Control Variables	Control Resources	What Is Optimized	Outcome	Control Search Efficiency	Mechanism Understanding	Systematics	Laboratory Implementation
Quantum	Photonic	Vast, but constrained at present	Yield, Fidelity	Good	High	Initial stages	Rules emerging	Automated
Chemical & Material	Chemicals, Processing Conditions	Vast	Yield, Property	Good-Excellent	High	Good	Rule driven	Automated
Natural Evolution	Genomic	Vast	Fitness	Excellent	High	Poor	Not clear	Micro-organisms (laboratory)
Directed Evolution	Genomic	Vast, but sparingly utilized	Enzymatic activity	Good	High	Fair	Vague rules	Semi-automated

Table 2: Common features of optimization in the sciences. *The table draws together the findings from optimization in the sciences listed in the left column. Subjective comments garnered from the collective literature are made characterizing the performance of control in these domains. In each scenario extensive control resources are available, yet searching amongst them for an optimal control is highly efficient with the sampling only drawing on an infinitesimal fraction of the possibilities. This common behavior across the domains of OptiQ, OptiChem and OptiEvo is the puzzle addressed by OptiSci.*

III. Taking a Closer Look at OptiSci

The questions raised in Section **II. E** can be addressed by (i) taking a closer examination of controlled phenomena in particular domains of science and then (ii) assessing whether a common foundation exists to collectively explain the observed highly efficient control optimizations. This section will consider the basis of OptiSci by examining the control of quantum dynamics,[7,8] OptiQ in **III. A,** and allied control in the chemical and material sciences,[14,15] OptiChem in **III. B.** There are extensive mathematical and physical details involved in both analyses beyond those presented here (see the cited references). In addition, similar assessments may be performed for the domains of natural and directed evolution.[6]

III. A. Optimal Control of Quantum Dynamics (OptiQ)

The control of quantum dynamical processes is an active area of research[8] with many potential applications summarized in Figure 2. Figure 3 presents a sketch of a quantum system, illustrated as a molecule, subjected to control of its wave function by means of a photonic reagent $\varepsilon(t)$ consisting of a shaped electric field pulse.[9,12] The photonic reagent $\varepsilon(t)$ in Figure 3 has broad bandwidth character in keeping with the general need for simultaneously addressing multiple transitions or degrees of freedom in a complex, evolving quantum system. Regardless of the particular application, a general principle underlies OptiQ for the control of all quantum phenomena. The quantum system starts out in initial state $|\psi_i\rangle$ (or its density matrix analog ρ_i, as appropriate) with the aim of arriving at a desired final state $|\psi_f\rangle$ at a later time *T*. Achieving a particular final state $|\psi(T)\rangle$ often is not the actual goal, but rather the objective is to maximize the expectation value of an observable operator O in the created state $J = \langle\psi(T)|O|\psi(T)\rangle$. The underlying assumption is that evolution with the free Hamiltonian H_0 alone does not satisfactorily reach the objective. Thus, the photonic reagent $\varepsilon(t)$ is introduced with

the overall dynamics described by Schrödinger's equation

$$i\hbar\frac{\partial}{\partial t}|\psi(t)\rangle = [H_0 - \mu\varepsilon(t)]\,|\psi(t)\rangle \quad (1)$$

where for illustration the field coupling operator is taken as the dipole μ. The goal is then to seek an optimal photonic reagent $\varepsilon(t)$ that maximizes the objective $\langle\psi(T)|O|\psi(T)\rangle$ at a suitable time *T*. A general means for achieving this goal is through the creation of constructive interferences in the dynamics to arrive at the state $|\psi(T)\rangle \sim |\psi_f\rangle$ while simultaneously creating destructive interferences in the other accessible final states $|\psi_{f'}\rangle$, $f' \neq f$. The practical execution of control in the laboratory is frequently carried out through the adaptive feedback process[10,11] depicted in Figure 1, and many types of applications are shown in Figure 2. The growing number of experiments may be characterized as quite successful, especially considering the practical constraints often present in the available controls. Thus, a basic question is *why* it appears to be easy to find effective, or even truly optimal, photonic reagents in the laboratory.[48]

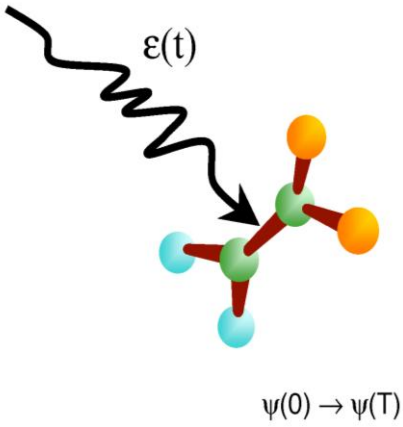

Figure 3: Optimal control of quantum dynamics. *The sketch shows a photonic reagent control* $\varepsilon(t)$ *impinging on a quantum system considered as a molecule for illustration. The goal is for the photonic reagent to direct the molecular dynamics reflected in the wave function evolution* $\psi(O) \rightarrow \psi(T)$. *The delicate shaping of the photonic reagent achieves this goal by cooperating with the often complex dynamical capabilities of the quantum system.*

The context of the word "easy" above needs explanation. First, the actual experiments can be arduous, even taking years to set up with great attention to detail required. The notion of easy does not refer to this heavy technological and humanly intensive overhead, but rather to the effort required in finding an optimal control relative to the frequently vast extent of controls searched over. As explained in Section **II. A**, a mere infinitesimal fraction of the accessible controls is typically sampled during a sequence of experiments to find an optimal photonic reagent.[13] The essence of the process may be captured by considering, for example, that the collective effort may take a period of four years, *plus one additional hour*. The laboratory set up overhead occurs during the four years while the latter additional hour is when the actual experiment is finally run to discover an optimal photonic reagent. With the typical extensive overhead put aside, the emphasis in OptiQ is on analyzing *why* it appears easy to find optimal controls over quantum dynamics phenomena during a brief experimental search (e.g., an hour) when $10^3 \sim 10^4$ experiments are done, while the possible number of distinct controls is $\sim 10^{100}$.

Nominally, the answer to the latter question and all matters about control of quantum phenomena lie in a detailed understanding of the Hamiltonian operators H_0 and μ in Equation 1 and the dynamics induced by $\varepsilon(t)$. Many of the applications in Figure 2 involve highly complex systems, which in turn have equally rich Hamiltonian structure that is frequently not known in quantitative detail. Most important in this regard is the *potential energy landscape V(**r**)* residing in the Hamiltonian term H_0, which can depend on many variables ***r*** when multiple atoms or electrons are simultaneously controlled by an applied field. The adaptive feedback control experiments[10,11] of Figure 1 inherently fold in all of the relevant potential energy landscape details for determining an optimal field. However, neither the potential energy landscape nor consideration of the detailed quantum dynamics (i.e., even if it was available) addresses *why* it is operationally easy to find optimal controls $\varepsilon(t)$ over quantum

phenomena. The answer to the posed question appears to reside in the topology of the *quantum control landscape* $J=[\varepsilon(t)]$ which is the observable as a function of the control.[7,8,13,48] Despite the fact that each physical system has its own *unique* Hamiltonian, theoretical analysis shows that quantum control landscapes should have a *generic* topology whose character only depends on the system's initial state and the nature of the observable operator. The origin of the latter conclusion rests on the satisfaction of three assumptions specified below. Before considering the details, it is reasonable to expect that a broad foundation must be behind the evident universality of easily finding effective quantum controls, which occurs regardless of the particular system's Hamiltonian complexity.

Extensive research is underway to understand quantum control landscapes[7,8,13,48] and only a basic summary of the concepts will be presented here for manipulating closed systems of finite dimension (i.e., *N* states) whose dynamics are described by

$$i\hbar\frac{\partial}{\partial t}U(t)=\left[H_0-\mu\varepsilon(t)\right]U(t) \qquad (2)$$

where $U(t)$ is the time evolution operator such that $|\psi(t)\rangle=U(t)|\psi_i\rangle$. OptiQ quantum control landscape analysis rests on three assumptions:

(a) The system is controllable

(b) The control $\rightarrow$ state map is full rank

(c) No (significant) constraints are placed on the controls

Each of these assumptions requires special consideration. Satisfaction of assumption (a) implies that some control exists to reach the target state $|\psi_f\rangle$ at a suitable time *T* from $|\psi_i\rangle$ at $t=0$. Although exceptions to fulfillment of this assumption can be found, mathematical and physical analysis suggests that most finite dimensional quantum systems are likely to be controllable.[49] Assumption (b) refers to the rank of the matrix $\frac{\delta U(T)}{\delta\varepsilon(t)}=\frac{-i}{\hbar}U(T)U^{\dagger}(t)\mu U(t)$ being full, which is equivalent to stating that a differential change in the state at the final time, $\delta U(T)$, has a corresponding differential control $\delta\varepsilon(t)$ producing it. This assumption could be violated corresponding to a so-called singular control, although these circumstances seem to be rare.[50] The last assumption (c) on the ready availability of photonic reagents of arbitrary shape will surely never be fully met due to limited laboratory resources. However, as assumption (c) states, the practical requirement is that no "significant" constraints be placed on the controls, which is a much milder condition (i.e., the controls only need to be available for addressing all the necessary transitions to create the desired final state). Collectively, satisfaction of these assumptions provides the conditions for assessing the topological nature of control landscapes.[7,8,13,48] The conclusion of this analysis is that suboptimal local maxima forming traps should generally *not* exist on quantum control landscapes. A trap is a local submaximal value of the objective $J[\varepsilon(t)]=\langle O\rangle$, from which it is not possible to climb further on the landscape to a higher value through variation of $\varepsilon(t)$ guided by a myopic algorithm (e.g., a gradient based procedure).

Rigorously establishing that the three assumptions above are satisfied in any particular case is a difficult task, but it is plausible to expect that these assumptions should be fulfilled, likely including (c), in many applications, especially as control resources become richer. Accepting satisfaction of the assumptions, the resultant conclusion about the trap free nature of the quantum control landscape topology can be tested in simulations and in the laboratory. Simulations have a special role to play, as they can be carried out with great care while putting aside laboratory issues including noise, technologically limited control resources and other factors. The recent literature contains two works carefully exploring this matter with simulations for the control of (i) the state-to-state transition probability[51] $P_{i\rightarrow f}$ and (ii) the fidelity[52] $\|W-U\|$ of creating a target unitary transformation W. Collectively, these studies involved over 50,000 computer simulations with a broad variety of model multi-state Hamiltonians, and all tests reached perfection for the objective to at least three decimal places upon paying due attention to numerical details. More studies are warranted, but

these results are consistent with the three assumptions above being satisfied and the consequent conclusion that quantum control landscapes are generally expected to have a trap-free landscape topology. Care is needed in performing such computational tests, as the nature of the numerical discretization and other approximations can become issues, especially when very high fidelity results are sought. A key limiting factor is likely satisfaction of assumption (c) with regard to free access to all essential control resources. In computer simulations this should not be an issue, but in the laboratory there will always be a need for further control resources, especially when treating complex systems. In some cases more bandwidth and energy may suffice to meet the control needs of assumption (c). But, even subtle resource limitations can become important in some circumstances. For example, the target time *T* must be sufficiently large to avoid undue constraints on the intervening dynamics.

Direct exploration of control landscapes in the laboratory has only just begun.[53,54] A particularly interesting prediction[7,8,48] from the landscape analysis concerns the rank of the Hessian $\frac{\delta^2 \mathrm{P}_{i\to f}}{\delta\varepsilon(t)\delta\varepsilon(t')}$ at the top and bottom of the transition probability landscape for $J = \mathrm{P}_{i\to f}$. The analysis predicts that at the top of the landscape (i.e., $\mathrm{P}_{if} = 1.0$) for a system with N states there will be at most $2N-2$ routes off, and at the bottom (i.e., $\mathrm{P}_{if} = 0.0$), there will be at most two routes up. At both landscape extremes the routes are specified by the associated eigenvectors of the Hessian with non-zero eigenvalues. Importantly, the numbers of non-zero eigenvalues of the Hessian, at the top or bottom of the landscape, does not depend on the quantitative details of the Hamiltonian (i.e., besides the need to satisfy the assumptions (a) and (b) of OptiQ). The Hessian is infinite dimensional (i.e., time or frequency is continuous), but in practice it would be of finite size through digitization of the controls in the laboratory. An experimental test of the Hessian spectral predictions at the top and bottom of the $\mathrm{P}_{i\to f}$ landscape was performed in atomic Rb vapor, and the results were found to be fully consistent with the predictions.[54] In addition, many simulations also confirm these predictions about the Hessian rank.[51] An ancillary issue of practical importance is that the remaining null Hessian spectrum at the top of the landscape implies an inherent degree of robustness to noise when controlling quantum phenomena.[48]

The general conclusion from these OptiQ analyses is that quantum control landscapes should be devoid of traps upon satisfaction of the three assumptions, thereby providing very favorable circumstances for maximizing physical observables. Many additional issues remain to be explored, including the nature of structural (i.e., non-topological) features on the control landscapes, as these can influence the efficiency of optimization. The observed practical ease of performing optimization suggests that tortured features on the landscapes must be rare, but careful studies remain to be done to directly affirm this conclusion. Another important direction for OptiQ is consideration of how constraints on the controls enter to limit the attainable yields and introduce landscape traps. Such traps will be 'artificial' due to their constrained control resource origin, although they will appear to be real in the associated experiment or simulation.

The perspective taken in OptiQ is to choose a sample from the chemical or material stockroom and then manipulate the quantum dynamics of the sample by searching through the photonic reagent 'stockroom' to find an effective control $\varepsilon(t)$. This process seeks control of the system's dynamics by performing tailored variations of the time-dependent portion of the Hamiltonian. Samples can, of course, be freely chosen from the material stockroom, and some demanding control applications may benefit from dual material-photonic reagent optimization as schematically indicated in Figure 4. Simulations have already considered the converse of the analysis discussed above where now the control field $\varepsilon(t)$ is fixed in form, and the time independent structural components H_0 and/or μ of the Hamiltonian (i.e., accessed through the chemical or material stockroom) are treated as the controls.[55] In the latter situation, the same trap-free landscape topology was found for $\langle O \rangle$ as a function of the

control matrix elements of H_0 and μ, upon satisfaction of the analogous three assumptions stated earlier for photonic reagents. An extreme circumstance would entail turning off the external field and treating the controls as strictly drawn from the chemical or material stockroom along with processing conditions. This latter view naturally leads to considering control within chemistry as addressed below.

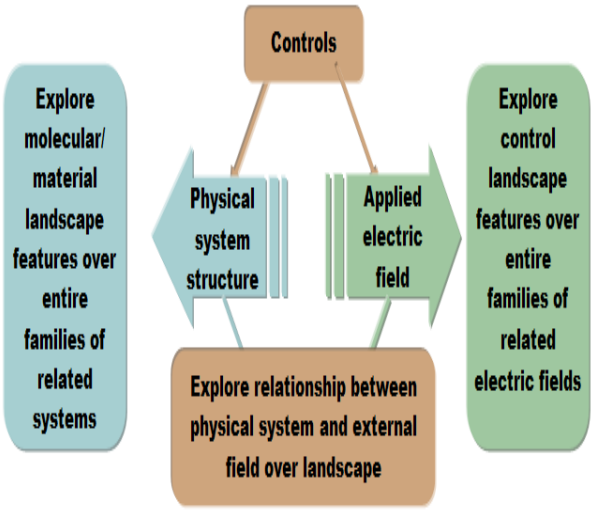

Figure 4: The full prospects for control resources. *The traditional perspective for controlling quantum dynamics phenomena in OptiQ is to start with a molecular or material sample and search through the photonic reagent 'stockroom' to find an optimal control field by the operations in Figure 1, and a schematic of this application is presented in Figure 3. The search process is observed to be very efficient, which may be understood from the favorable topology of quantum control landscapes. An enhanced scenario involves simultaneously searching over the chemical/material and photonic reagent stockrooms to meet demanding objectives, as indicated above. A fully automated machine of this type, analogous to that sketched out in Figure 1, may be envisioned for this purpose. A final prospect is to turn off the photonic reagent to then only utilize the chemical and material stockroom as well as processing conditions for resources to optimize synthesis yields and properties, which forms the basis of OptiChem.*

III. B Optimization of Chemical Synthesis and Properties (OptiChem)

Two chief goals in the chemical and material sciences are synthesis and property optimization.[1] These objectives are widely met to a practical degree in a broad spectrum of applications, and as a result one can simply declare that chemistry works! But, should this success be expected upon consideration of the underlying complexities inherent in meeting these goals? A commonly circulated quote[56] is that "Chemistry is all about getting lucky". The analogous comment could have been made in the context of controlling quantum dynamics. But, as explained above, attaining successful control of quantum phenomena is not about getting lucky, and a basis exists through OptiQ to understand the ease of finding effective photonic reagents. OptiChem contends that a similar analysis may be applied to the observed ease of optimization in chemistry.[14,15] In the following, for simplicity we will use the word 'chemistry' to include like circumstances in material science.

It is attractive to express the optimization objectives in chemistry within an optimal control framework. In this case, the control variables include (a) reagents, catalysts, solvents and their concentrations, (b) reaction time, temperature and other processing conditions, (c) the mole fractions of components in a material and (d) the substituents on a molecular scaffold. Any particular chemical application will draw on necessary aspects of these controls in keeping with the objective. Lurking in the background upon optimization in chemistry is the so-called *curse of dimensionality*, which states that the effort at finding an optimal solution should scale exponentially with the number of variables. As an example, consider a reaction between two classes of molecules with one referred to as R and the second being a molecular scaffold S with two sites for chemical functionalization having variables referred to as x_1 and x_2. A scaffold is a molecular framework (e.g. a benzene ring) to be built upon by functionalization at the specified sites. The variables x_1 and x_2 take on 'values' drawn from two respective lists of substituents capable of

bonding to the scaffold sites. In addition to these variables, in this illustration the temperature, reaction time, solvent and the concentration of R are all part of the controls subject to optimization, and the objective is to maximize a particular reaction product. For this example, there are six variables, and a reasonable supposition is that each can take on ~ 10 values. Even in this modest circumstance of reacting species R with a companion molecule (i.e., a member of the functionalized scaffold set) there are a total of $\sim 10^6$ possibilities available for consideration in the laboratory. Such an extensive set of control conditions is not possible to thoroughly explore in practice. In the context of traditionally performed chemistry experiments, scanning over the potential control choices is typically guided by insight and experience, which often produces very successful outcomes.

Importantly, Table 1 provides experimental evidence that automated optimization over the vast available resources in chemistry is far easier than might be expected and can beat the curse of dimensionality. The search for optimal photonic reagents discussed in Section **III. A** should also be plagued by the curse of dimensionality, but evidence in that domain also shows the ability to circumvent this perceived bottleneck. In addition, the same class of algorithms[16] used to systematically discover optimal photonic reagents in Figure 1 is also employed in the automated optimization of chemical objectives exemplified in Table 1. The ease of finding optimal solutions in chemistry is beyond consideration of the overhead involved with setting up the experiments and their careful execution. The notion of "ease" refers to the number of experiments required in practice to reach an optimal outcome versus the total number of possibilities, which is consistent with the same assessment of efficiency in considering photonic reagent control and that found in biological evolution.

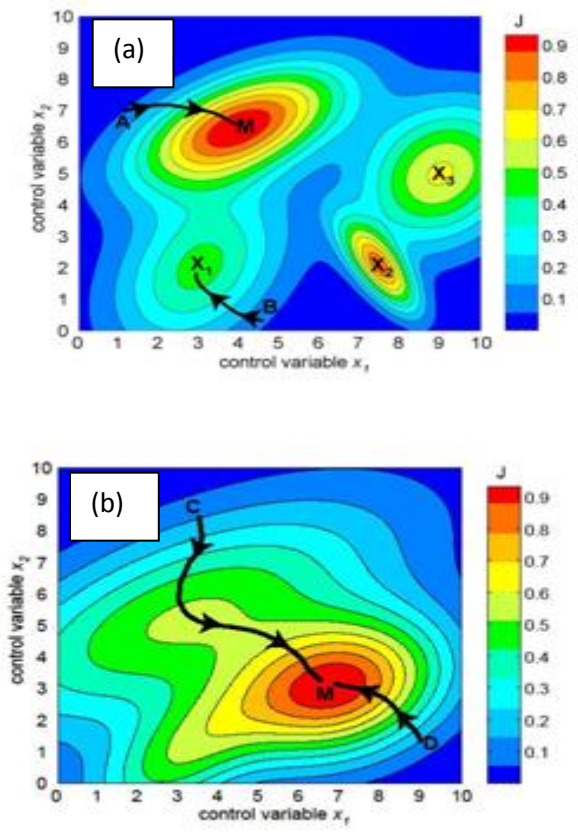

Figure 5: The yield/property control landscape in OptiChem. *The variables x_1 and x_2 (likely amongst a larger number in practice) are shown controlling a yield or property value.[14,15] The hypothetical associated landscape in the left panel (a) has multiple local suboptimal extrema acting as traps for searches seeking the best possible outcome. A search initiated at an arbitrary point on such a landscape would likely terminate at a less than optimal yield. In contrast, the landscape on the right (b) has rich structure but only a single extremum at the absolute maximum value. The foundations of OptiChem imply that chemical and material landscapes should have the latter attractive features upon satisfaction of some assumptions explained in the text. (a) Contains "Traps Hard to Optimize (b) Trap-Free Easier to Optimize*

Given the circumstances above, it is natural to explore the character of chemical control landscapes[14,15] in analogy to their photonic reagent counterparts in OptiQ. Figure 5 depicts two distinct prospective types of chemistry yield/property landscapes illustrated as well with two control variables x_1 and x_2. In practice,

realistic applications can have several control variables, $x_i, i = 1,2$ as indicated in Table 1. Figure 5(a) schematically shows a landscape with many local suboptimal extrema; such a landscape might be expected to be the norm upon considering the complexities generally evident in controlling chemical phenomena. Since a full picture of the landscape will not be available *a priori*, starting from any point on the landscape in Figure 5(a) would likely lead to a search arriving at a suboptimal local maximum acting as a trap. Without insight into what may lie elsewhere on the landscape, encountering such a trap could produce an unsatisfactory result and termination of the search to seek a better solution. Even stochastic algorithms[16] having global search capabilities could become lost with large numbers of control variables when many local traps are present. In contrast, Figure 5(b) depicts a landscape having only a single maximum at the global optimal value. In this case, experiments initiated from anywhere on the landscape should be able to readily climb to the absolute maximum, even using simple algorithms for guidance. OptiChem refers to the analysis of chemical transformations or properties and their associated landscapes for optimization. The assessment of the landscape topology in OptiChem applications rests on three assumptions[14,15]:

(a) The objective is well posed

(b) Freedom exists to move over the landscape

(c) No significant constraints are placed on the controls

Chemical and material systems under control generally are open to the environment, and in many cases the environment (e.g., a solvent) is an integral part of the controls. Thus, the analysis of OptiChem landscape topology needs to take into account the control of open chemical/physical systems. This analysis may be performed from either a quantum mechanical or a classical perspective (see Section **IV**) to reach the same fundamental conclusion.[14,15] In particular, upon satisfaction of the three assumptions above, it follows that yield/property landscapes appear as sketched in Figure 5(b) with no inherent traps. As for the case of photonic reagents in Section **III. B**, the satisfaction of the three assumptions needs consideration. First, the objective being well posed in assumption (a) refers to seeking a reachable physical or chemical objective; a counter example would be the targeted synthesis of a molecule where an atom is in an unattainable valence state. Chemical insight and experience provide the prime means for posing reasonable objectives. Satisfaction of the second assumption (b) is difficult to prove in general, but it is likely that a small change in the objective value has at least one associated small change in the controls thereby permitting free movement over the landscape. The last assumption (c) that no significant constraints are placed on the controls is the main concern when considering yield/property landscapes. In practice, an initial choice would be made for the controls, followed by their iterative improvement based on the measured chemical yield/property outcome and even guided by observations of the resultant landscape (see Section **IV**). Extensive data are available on the nature of landscapes in chemistry, [14,15] and there is more evidence in this regard than in the case of control in OptiQ with photonic reagents due to the large size of the chemistry community and the lengthy period of intense activity. Landscapes within OptiChem are available for synthesis yields and for a broad variety of properties, including spectral, electrical, mechanical, thermodynamic, etc. The reported landscapes are generally only up to four dimensions (i.e., the number of control variables) due to the arduous sampling effort required as the dimension rises and for graphical reasons. In summary, $\sim 90\%$ of the reported chemical and material landscapes have no evident traps. [14,15] Interestingly, regardless of whether the landscapes are reported to have or not have traps, the papers presenting these landscapes usually provide no discussion of the topological significance of the findings. This body of landscape work also needs to be drawn together with the results from the automated optimization efforts summarized in Table 1, where the algorithm employed seeks to guide a climb of each landscape without thoroughly sampling the full control possibilities. The latter reported optimizations were carried out in dimensions up to ~ 20, with the evident ease of finding good solutions defying the curse of dimensionality. Two examples of landscapes will

be given here, and many more are summarized in refs. 14 and 15.

Optimization of solid-state catalyst composition is commonly performed aiming to enhance the yield and selectivity of chemical reactions. A particular example is the controlled oxidation of isobutane to form a set of products, including methacrolein, isobutene and finally combustion to produce carbon dioxide.[57] Here the controls consisted of the molybdenum, vanadium and antimony fractional catalyst components. Thus, a landscape for each product depends on these three variables, taking into account that their mole fractions add to 1.0. In each case, the reported landscapes were trap free with particular interest focused on the case of methacrolein as the desired product. In addition, these landscapes generally exhibited broad, flat optimal regions corresponding to a level set of successful controls, as similarly found when dealing with photonic reagents. The authors also found that landscape traps could occur with catalysts made from other particular atomic components.[57] This circumstance points to the significance of assumption (c) on having a proper set of resources when considering optimal control and the resultant impact on the observable landscape.

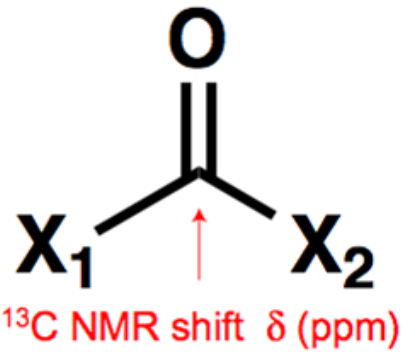

Figure 6 (a): An NMR application of OptiChem with discrete variables. *A family of ketone molecules is characterized by the chemical substituents* x_1 *and* x_2*, where the physical property of interest is the* ^{13}C *NMR shift at the indicated carbon atom.*[14,15,58] *The shift is a "function" of the particular (discrete) substituents* x_1 *and* x_2*.*

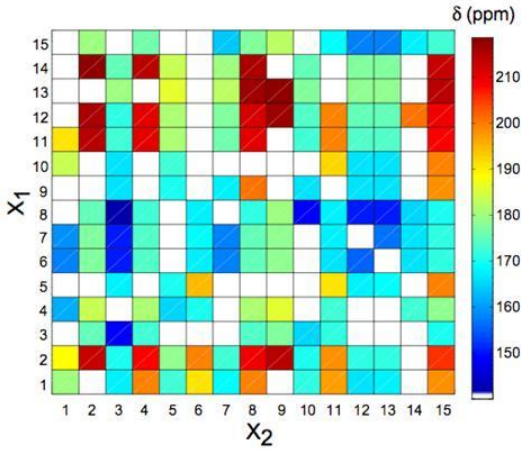

Figure 6 (b): The landscape with arbitrary assignment of substituent ordering. *A family of molecules of the form in 6(a) was chosen with both* x_1 *and* x_2 *separately drawn from respective sets of fifteen substituents. Although the* ^{13}C *NMR shift depends on* x_1 *and* x_2*, there is no a prioi way to plot such data, since the sequence of chemical substituents labeling the axes of such a plot would be arbitrary. As an illustration, the figure was made using a random, but unique, integer assigned to each of the fifteen chemical substituents* x_1 *and* x_2 *on the scaffold in Figure 6(a). The resultant chemical shift landscape appears random (the white squares correspond to molecules whose data was not available).*

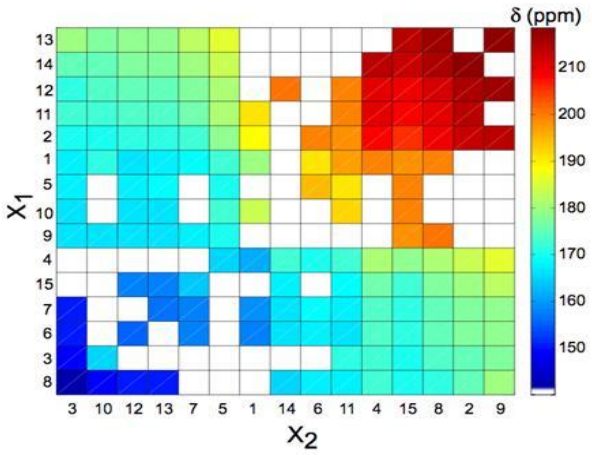

Figure 6(c): Uncovering trap-free landscapes with discrete controls. *Following the concepts of OptiChem, a combinatorial reordering algorithm*[58]

was applied to the data shown in Figure 6(b). Accordingly, Figure 6(c) is produced by reordering the rows and columns in Figure 6(b) resulting in a regular and essentially monotonic landscape. The particular ordering along the two axes in Figure 6(c) reflected the rules of NMR as well as provided additional physical insights. The fact that OptiChem principles could identify physical rules in this illustration opens up the prospect for automated rule discovery in broader domains of OptiSci applications[14,15,58]

The second example refers to the common goal of optimizing the property of a molecule drawn from a 'library' of possibilities where each member of the library has a common underlying scaffold with particular sites where chemical substituents may be attached. In a circumstance with two variables x_1 and x_2 and a scaffold each of the controls would take on a discrete set of values (i.e. denoted by the labels methyl, bromo, ethyl, etc.) corresponding to the chemical substituents respectively considered for bonding to the two sites. Upon selecting (synthesizing) a member of the library, the property of interest would be measured. The property $\mathrm{P}(x_1, x_2)$ depends on the chemical substituents bonded to the two sites, but plotting the associated landscape poses an inherent problem due to the discrete nature of the control variables x_1 and x_2 That is, the appearance of the landscape is influenced by the initially ambiguous ordering of the discrete values for x_1 and x_2. In addition, the property landscape $\mathrm{P}(x_1, x_2)$ will have a degree of grainy resolution due to the discrete variables. This type of molecular control (discovery) problem occurs in many chemistry applications, and a special substituent reordering algorithm[59] was established for the purpose of revealing the features of such OptiChem landscapes. As an example, Figure 6 considers the case of measuring the $\mathrm{P}(x_1, x_2) = {}^{13}C$ chemical shift of the carbon atom in the carbonyl group with x_1 and x_2 drawn from a set of prescribed chemical substituents. The library of compounds may be codified by assigning a unique, but otherwise arbitrary, integer label for each substituent on both sites. An initial arbitrary labeling of the substituents will generally produce a landscape that appears random, as found to be the case in Figure 6(b). The underlying property landscape $\mathrm{P}(x_1, x_2)$ should be trap free, (i.e., while taking into consideration the grainy nature of such discrete landscapes), provided that the three assumptions underlying OptiChem are satisfied. With that in mind, the reordering algorithm operates by performing a combinatorial optimization of the substituent integer labels on both sites seeking to identify if a regular landscape is hidden in the data. Starting with Figure 6(b), the result after reordering the integer labels of the substituents (i.e., the rows and columns in 6(b)) is shown in Figure 6(c), which clearly reveals a smooth and essentially monotonic landscape. Many similar results have been found with other molecular libraries for a wide variety of properties considered as observables including from IR spectroscopy, thermodynamics, chemical binding affinity, etc.[14,15] The outcome in each case was similar to that shown in Figure 6 producing a monotonic trap free landscape upon reordering the variables. The significance of these results, and in particular those in Figure 6, can be easily appreciated. First, the regular structure in Figure 6 readily permits interpolation over the reported $^{13}\mathrm{C}$ chemical shift values to estimate those that have not already been measured (i.e., the values for the white squares of the Figure 6). This capability has significance in NMR spectroscopy when trying to identify the products from a complex chemical synthesis. Second, the optimal ordering of the chemical substituents on the two sites in Figure 6 corresponded to known rules of NMR as well as provided additional chemical insights.[58] The capability of landscape principles to reveal rules associated with the controls has potential importance beyond NMR, as rules ultimately often drive the operations in many domains of science. The identification of the rules lying hidden in the original ordering of the substituents in Figure 6(b) to finally produce Figure 6(c) utilized the reordering algorithm,[59] which rests on the expectation of finding a regular trap free landscape upon satisfaction of the three assumptions underlying OptiChem. Turning around this capability also has led to the identification of data

errors and other aberrations that stood out from an otherwise regular landscape in applications of this kind.[60] This prospect has significance for potentially providing the means to identify flawed experimental designs at an early stage in the laboratory.

IV. Conclusions

Each area of science has its own objectives and operational differences, with the collective observed properties and variables spanning vast length and time scales as well as often exhibiting highly distinct subtleties. Thus, these circumstances would seem to suggest that little commonality should arise upon considering optimization in different domains of science. Counter to this expectation, OptiSci is based on the observed strikingly similar findings of extraordinary high efficiency upon seeking optimal control solutions across wide domains of science. This situation is summarized in Table 2, and Figure 7 addresses these points by showing two typical control laboratories with one dealing with photonic reagents and the other chemical and material reagents. The latter laboratory is modern with heavy instrumentation and the more traditional chemical apparatus is relegated to the hood in the background. A comparison of the two photographs in Figure 7 would suggest that the activities underway in the laboratories do not share much in common. However, the primary focus in many laboratories is optimization of a desired outcome by a tailored choice of the available controls. The basic contention of OptiSci is that regardless of the distinctions at a detailed operational level in such diverse experimental settings, seeking optimal performance provides a unifying concept and mathematical foundation to link together the common behavior found in seemingly diverse applications of control in science. Importantly, the notion of common behavior in OptiSci goes beyond just considering the implementation of an optimal control algorithm, which can be utilized in virtually any application. OptiSci is concerned with the universal control behavior found upon performing such implementations.

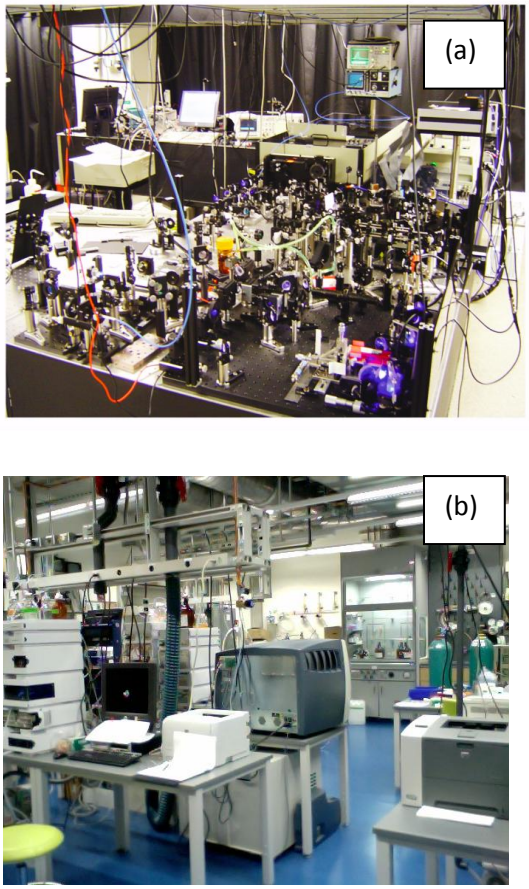

Figure 7: Laboratories controlling events in two distinct areas of science. *The photograph on the left is of a photonic reagent quantum control laboratory with lasers, pulse shapers, optics, computers, etc. The photograph on the right is of a modern chemical or material laboratory which is highly instrumented with the traditional apparatus residing in the background hood. A simple glance at these two photographs would nominally lead to the conclusion that little if anything in common exists between the experiments in these laboratories. The detailed operations in both laboratories are quite distinct with the associated phenomena typically involving very different length and time scales. Nevertheless, a frequent goal in both cases is to optimize a product or outcome using available control resources. The evidence shows that under optimization the behavior found in both laboratories is strikingly similar, especially in terms of the high efficiency of finding optimal controls, despite the extensive possibilities to search over. The basic premise of OptiSci is that a unified origin for this behavior*

exists which may be understood upon considering the associated objective control landscapes, and this behavior is exhibited across the sciences. The high efficiency of searches in each area of science is important in its own right. But, the observed similar behavior across the sciences opens up many opportunities for the control and understanding of what might hitherto have appeared as complex objectives possibly transcending the sciences. (a) Control Laboratory Employing Photonic Reagents and (b) Control Laboratory Employing Molecular/Material Reagents.

The unification of the control operations in OptiQ, OptiChem and OptiEvo to finally form OptiSci is based on satisfying three assumptions reiterated here as: (a) the objective is well posed, (b) free movement is available on the objective landscape and (c) no significant constraints are placed on the control resources. Upon satisfaction of these assumptions, the objective may be either expressed quantum mechanically $\langle O_{qm} \rangle = Tr(\rho_{qm} O)$ or classically $\langle O_{cl} \rangle = \int d\omega \; \rho_{cl}(\omega) O(\omega)$ where the goal is corresponding optimization of $\langle O_{qm} \rangle$ or $\langle O_{cl} \rangle$ over the control variables. In both cases, ρ is either the quantum or classical density, which depends on the controls. There are various ways of establishing the trap free nature of either the classical or quantum control landscape,[7,8,61] with some procedures revealing distinct levels of detail. However, the essence of the analysis rests on either the quantum mechanical expectation value or the classical average value being a convex function of its appropriate density. In addition, upon satisfaction of the three key assumptions, these respective densities form convex sets. Thus, the landscape is itself convex. The generality and simplicity of this analysis speaks to the observed generic landscape behavior found in a wide variety of scientific disciplines.

There are still many issues to be resolved in OptiSci calling for additional mathematical and experimental analyses, and the fundamental and practical implications for the findings of common control behavior in the sciences remains to be determined. However, at this point some speculative assessments can be made based on the evidence and analysis that a universal control principle is operative in the sciences. This behavior captured by OptiSci suggests that optimal controls may be efficiently discovered for phenomena spanning vast length and time scales. The systematics of OptiSci implies that it may be possible to develop control based 'machines' whose goal is to identify rules rather than just seek an optimal outcome.[14,15] In this case a rule may be defined as a minimal set of variables that produce a trap free landscape for a particular objective. Turning this around, failure to find an expected trap free landscape in a sequence of control experiments should indicate the presence of an ill-defined goal that violates one or another of the assumptions of OptiSci.[14,15] Hints at these prospects are already evident,[60] and the early discovery of flawed experimental designs could be of high value in complex optimization efforts. Finally, a deeper understanding of the topology and general structure of control landscapes may enable more efficient control algorithms, which could be especially important in cases where the experiments are expensive to perform. At this juncture the main conclusion from OptiSci is that the strikingly common control behavior found within specific areas of science, and between them, deserves attention for its full implications. In the context of optimal control in the sciences, perhaps we may adopt G. Liebnitz's adage that "It is the best of all possible worlds".

Acknowledgements: The author acknowledges support from NSF, DOE and ARO.

References

[1] G. Fleming and M. Ratner, Directing Matter and Energy: Five Challenges for Science and the Imagination, US Department of Energy, Washington, DC (2007).
[2] S. Wright, The roles of mutation, inbreeding, crossbreeding and selection in evolution. (1932).
3 F. Arnold and G. Georgiou, Directed enzyme evolution: screening and selection methods. (Humana Press, New Jersey, 2003).
4 S.A. Rice and M. Zhao, Optical control of molecular dynamics. (John Wiley New York, 2000).
5 R. Stengel, Optimal control and estimation. (Dover Publications, 1994).
6 X. Feng, A. Pechen, A. Jha, R. Wu, and H. Rabitz, Chem. Sci. (2011).

7 R. Chakrabarti and H. Rabitz, Int. Rev. Phys. Chem. 26 (4), 671 (2007).
8 C. Brif, R. Chakrabarti, and H. Rabitz, Adv. Chem.Phys. 148 (2012).
9 H. Rabitz, Science 299 (5606), 525 (2003).
10 R.S. Judson and H. Rabitz, Phys. Rev. Lett. 68 (10), 1500 (1992).
11 P. Nuernberger, G. Vogt, T. Brixner, and G. Gerber, Phys. Chem. Chem. Phys. 9 (20), 2470 (2007).
12 AM Weiner, Rev. Sci. Instrum. 71, 1929 (2000).
13 H.A. Rabitz, M.M. Hsieh, and C.M. Rosenthal, Science 303 (5666), 1998 (2004).
14 K.W. Moore, A. Pechen, X.J. Feng, J. Dominy, V.J. Beltrani, and H. Rabitz, Phys. Chem. Chem. Phys. 13, 10048 (2011).
15 K.W. Moore, A. Pechen, X.J. Feng, J. Dominy, V. Beltrani, and H. Rabitz, Chem. Sci. 2, 417 (2011).
16 C.M. Bishop, Pattern recognition and machine learning. (Springer, Berlin, 2006).
17 R. Rey-de-Castro and H. Rabitz, Phys. Rev. A 81 (6), 063422 (2010).
18 H. Rabitz, R. de Vivie-Riedle, M. Motzkus, and K. Kompa, Science 288 (5467), 824 (2000).
19 William B. Jensen, The Lewis acid-base concepts : an overview. (Wiley, New York, 1980).
20 J. Singh, M.A. Ator, E.P. Jaeger, M.P. Allen, D.A. Whipple, J.E. Soloweij, S. Chowdhary, and A.M. Treasurywala, J. Am. Chem. Soc 118 (7), 1669 (1996).
21 U. Rodemerck, M. Baerns, M. Holena, and D. Wolf, Appl. Surf. Sci. 223 (1-3), 168 (2004).
22 L. Weber, S. Wallbaum, C. Broger, and K. Gubernator, Angew 34 (20), 2280 (1995).
23 U. Rodemerck, D. Wolf, OV Buyevskaya, P. Claus, S. Senkan, and M. Baerns, Chem. Eng. J. 82 (1-3), 3 (2001).
24 OV Buyevskaya, A. Brückner, EV Kondratenko, D. Wolf, and M. Baerns, Catal. 67 (4), 369 (2001).
25 OV Buyevskaya, D. Wolf, and M. Baerns, Catal. 62 (1), 91 (2000).
26 S. Moehmel, N. Steinfeldt, S. Engelschalt, M. Holena, S. Kolf, M. Baerns, U. Dingerdissen, D. Wolf, R. Weber, and M. Bewersdorf, Appl. Catal. A. 334 (1-2), 73 (2008).
27 F. Clerc, M. Lengliz, D. Farrusseng, C. Mirodatos, S.R.M. Pereira, and R. Rakotomalala, Rev. Sci. Instrum. 76, 062208 (2005).
28 T. Umegaki, Y. Watanabe, N. Nukui, K. Omata, and M. Yamada, Energ. Fuel. 17 (4), 850 (2003).
29 K. Omata, T. Ozaki, T. Umegaki, Y. Watanabe, N. Nukui, and M. Yamada, Energ. Fuel. 17 (4), 836 (2003).
30 Y. Watanabe, T. Umegaki, M. Hashimoto, K. Omata, and M. Yamada, Catal. 89 (4), 455 (2004).
31 J.M. Serra, A. Chica, and A. Corma, Appl. Catal. A. 239, 35 (2003).
32 Y. Yamada, A. Ueda, K. Nakagawa, and T. Kobayashi, Res. Chem. Intermed. 28 (5), 397 (2002).
33 J.S. Paul, R. Janssens, J.F.M. Denayer, G.V. Baron, and P.A. Jacobs, J. Comb. Chem. 7, 407 (2005).
34 M. Bulut, L.E.M. Gevers, J.S. Paul, I.F.J. Vankelecom, and P.A. Jacobs, J. Comb. Chem. 8 (2), 168 (2006).
35 A. Corma, JM Serra, P. Serna, S. Valero, E. Argente, and V. Botti, J. Catal. 229 (2), 513 (2005).
36 A. Rusinko, S.S. Young, D.H. Drewry, and S.W. Gerritz, Comb. Chem. High Throuput Screen. 5, 125 (2002).
37 C. Kulshreshtha, A.K. Sharma, and K.S. Sohn, J. Comb. Chem. 10 (3), 421 (2008).
38 K.S. Sohn, D.H. Park, S.H. Cho, B.I. Kim, and S.I. Woo, J. Comb. Chem. 8 (1), 44 (2006).
39 Y.S. Jung, C. Kulshreshtha, J.S. Kim, N. Shin, and K.S. Sohn, Chem. Mat. 19 (22), 5309 (2007).
40 K.S. Sohn, D.H. Park, S.H. Cho, J.S. Kwak, and J.S. Kim, Chem. Mat. 18 (7), 1768 (2006).
41 K.S. Sohn, B.I. Kim, and N. Shin, J. Electrochem. Soc 151, H243 (2004).
42 S. Gavrilets, Fitness landscapes and the origin of species. (Princeton University Press, Princeton, New Jersey, 2004), pp.610.
43 Z.D. Blount, C.Z. Borland, and R.E. Lenski, Proceedings of the National Academy of Sciences USA 105 (23), 7899 (2008).
44 C.L. Burch and L. Chao, Nature 406 (6796), 625 (2000).
45 D.M. Weinreich, R.A. Watson, and L. Chao, Evol. Int. J. Org. Evol. 59 (6), 1165 (2005).
46 E.L. Haseltine and F.H. Arnold, Annu. Rev. Biophys. Biomol. Struct. 36, 1 (2007).
47 K.M. Esvelt, J.C. Carlson, and D.R. Liu, Nature 72 (7344), 499 (2011).
48 T.S. Ho and H. Rabitz, J. Photochem. Photobiol. A 180 (3), 226 (2006).
49 C. Altafini, J. Math. Phys. 43, 2051 (2002).
50 R. Wu et al, to be published (2012).
51 K.W. Moore and H. Rabitz, Phys. Rev. A 84 (1), 012109 (2011).
52 K.W. Moore, R. Chakrabarti, G. Riviello, and H. Rabitz, Phys. Rev. A 83 (1), 012326 (2011).
53 T. Bayer, M. Wollenhaupt, and T. Baumert, J. Phys. B: At., Mol. Opt. Phys. 41, 074007 (2008).
54 J. Roslund and H. Rabitz, to be published (2012).
55 A. Donovan, V. Beltrani, and H. Rabitz, Phys. Chem. Chem. Phys. 13, 7348 (2011).
56 Attributed to Robert Curl.
57 J.S. Paul, P.A. Jacobs, P.W. Weiss, and W.F. Maier, Appl. Catal. A. 265 (2), 185 (2004).
58 K. Moore, R. Li, I. Pelczer, and H. Rabitz, to be published (2012).
59 N. Shenvi, J.M. Geremia, and H. Rabitz, J. Phys. Chem. A 107 (12), 2066 (2003).
60 X. Feng, J. Sanchis, M.T. Reetz, and H. Rabitz, Eur. J. Chem. in press (2012).
61 A. Pechen and H. Rabitz, European Phys. Lett. 91, 60005 (2010)