Robust control of quantum gates via sequential convex programming

Robert L. Kosut,[1,*] Matthew D. Grace,[2,†] and Constantin Brif[2,‡]
[1]*SC Solutions, Inc., 1261 Oakmead Parkway, Sunnyvale, California 94085, USA*
[2]*Department of Scalable & Secure Systems Research, Sandia National Laboratories, Livermore, California 94550, USA*
(Received 21 June 2013; published 21 November 2013)

Resource trade-offs can often be established by solving an appropriate robust optimization problem for a variety of scenarios involving constraints on optimization variables and uncertainties. Using an approach based on sequential convex programming, we demonstrate that quantum gate transformations can be made substantially robust against uncertainties while simultaneously using limited resources of control amplitude and bandwidth. Achieving such a high degree of robustness requires a quantitative model that specifies the range and character of the uncertainties. Using a model of a controlled one-qubit system for illustrative simulations, we identify robust control fields for a universal gate set and explore the trade-off between the worst-case gate fidelity and the field fluence. Our results demonstrate that, even for this simple model, there exists a rich variety of control design possibilities. In addition, we study the effect of noise represented by a stochastic uncertainty model.

DOI: 10.1103/PhysRevA.88.052326 PACS number(s): 03.67.−a, 02.30.Yy, 02.60.Pn

I. INTRODUCTION

Robust control and robust optimization of uncertain systems are essential in many areas of science and engineering [1–8]. Recently, there has been much interest in achieving robust control of quantum information systems in the presence of uncertainty [9–40]. An important property of quantum information processing that distinguishes it from most other applications is the requirement for an unprecedented degree of precision in controlling the system dynamics. Also, due to the very fast time scale of physical processes in the quantum realm, implementing closed-loop feedback control is extremely difficult and thus open-loop control arises as the most feasible option in most circumstances.

For quantum information systems, a robust optimization problem can be formulated as a search for *design variables* $\theta \in \Theta$ (where Θ is the *design set*) that maximize a measure of *quantum gate fidelity* $\mathcal{F}$ over a range of *uncertain parameters* $\delta \in \Delta$ (where Δ is the *uncertainty set*). Fidelity compares a target unitary transformation with the actual quantum channel, which depends on both θ and δ. Fidelity is typically normalized, $\mathcal{F} \in [0,1]$, and the maximum value $\mathcal{F} = 1$ corresponds to a perfect generation of the target transformation. The design variables θ can include time-dependent control fields (for both open-loop and closed-loop control), measurement configurations (for closed-loop feedback control), constants associated with physical implementation, the circuit layout, and so on. The uncertainties δ can affect any element of the system Hamiltonian (including the design variables), with specific manifestations and ranges depending on the details of the physical implementation and external hardware. For example, uncertainties can represent dispersion and/or slow time variation of parameters such as coupling strengths, exchange interactions, and applied electromagnetic fields, as well as additive and/or multiplicative noise in control fields. The uncertainty set Δ can thus, in general, contain deterministic and random variables. Whatever the case, we assume that θ and δ are constrained to known sets Θ and Δ, respectively.

One common approach to robust control of quantum gates (e.g., see Refs. [21] and [40]) is based on maximizing the *average fidelity*, given by

$$\mathcal{F}_{\rm avg}(\theta) = \mathbb{E}_{\delta\in\Delta}\{\mathcal{F}(\theta,\delta)\}, \tag{1}$$

where $\mathcal{F}(\theta,\delta)$ denotes the fidelity as a function of design and uncertain variables, and $\mathbb{E}_{\delta\in\Delta}\{\cdot\}$ is the expectation with respect to the underlying distribution in Δ. Often the average fidelity is well approximated as the sum over a discrete sample with associated probabilities $\{\delta_i \in \Delta, p_i \in [0,1]\}$, i.e., $\mathcal{F}_{\rm avg}(\theta) = \sum_i p_i \mathcal{F}(\theta,\delta_i)$. While the use of the average fidelity is applicable in some cases (e.g., when the uncertainty represents weak random noise), the stringent performance requirements of quantum information processing make it more appropriate, in general, to estimate gate errors by using the *worst-case fidelity* with respect to all uncertainties $\delta \in \Delta$:

$$\mathcal{F}_{\rm wc}(\theta) = \min_{\delta\in\Delta} \mathcal{F}(\theta,\delta). \tag{2}$$

Also, *worst-case robust optimization* (or *minimax optimization*) is a well-known approach employed in many classical problems [7,41–57], and some of the methods developed for these applications can be adapted for robust control of quantum gates. The worst-case robust optimization problem for quantum gate fidelity is formulated as

$$\begin{aligned} &\text{maximize} && \min_{\delta} \mathcal{F}(\theta,\delta) \\ &\text{subject to} && \theta \in \Theta, \quad \delta \in \Delta. \end{aligned} \tag{3}$$

The goal reflected in problem (3) is to find the design variables $\theta \in \Theta$ that maximize the worst-case fidelity of Eq. (2).

In control applications, the design set Θ represents the set of control constraints and is most often convex or sufficiently well approximated by a convex set. In some cases so is the uncertainty set Δ, although this is not necessary for solving (3). What makes the problem difficult is that the fidelity is not a convex function of θ for any sample $\delta \in \Delta$. Nonconvex optimization problems are common in all of science and engineering and have engendered numerous

*kosut@scsolutions.com
†mgrace@sandia.gov
‡cnbrif@sandia.gov

numerical approaches to finding local optimal solutions. In particular, effective methods have been developed in recent years for worst-case robust optimization with nonconvex cost functions [51–57].

In optimal control applications, the functional dependence of the objective (e.g., fidelity) on the control variables is referred to as the *optimal control landscape* [58–61]. For an ideal model of a closed quantum system with no uncertainties, the optimal control landscape for the generation of unitary transformations has a very favorable topology [62–64]. Specifically, provided that a number of physically reasonable conditions are satisfied [65], the landscape is free of local optima, i.e., there exist one manifold of global minimum solutions (resulting in $\mathcal{F} = 0$) and one manifold of global maximum solutions (resulting in $\mathcal{F} = 1$), while all other critical points reside on saddle-point manifolds [62–64]. Such a favorable landscape topology facilitates easy optimization, as any gradient-based search (various types of which are popular in quantum optimal control [21,66–75]) is guaranteed to reach the global maximum [76]. Unfortunately, when uncertainties are present, this landscape topology is not preserved. Typically, uncertainties cause a decrease and fragmentation of the global maximum manifold, resulting in the emergence of multiple local maxima [40] (the landscape also undergoes a similar transformation when control fields are severely constrained [77]). Provided that the range of uncertainty is not too large, many of these local optimal solutions will have fidelities close to 1.

For quantum information systems, there is considerable ongoing effort to develop efficient methods for obtaining a good solution to the problem of robust control, for either average or worst-case fidelity. The majority of existing approaches rely on a numerical optimization procedure, mostly involving a gradient-based search for maximizing the average fidelity of Eq. (1). In some cases, a randomized search such as a genetic algorithm is employed [40]. The results demonstrate the existence of many solutions with high fidelities, consistent with the control landscape picture discussed above. Additionally, the optimal controls are often similar to the corresponding initial controls, provided the latter are reasonably good. This phenomenon, also observed in many engineering and design applications employing local search algorithms, supports the need for developing tools to efficiently calculate a good initial control. In particular, empirical evidence and simulations suggest that robust controls for an uncertain quantum system can be found by searches that start from solutions generated by applying *optimal control theory* or *dynamical decoupling* to the ideal (zero-uncertainty) counterpart system (see, e.g., [39,61], and references therein).

In this paper, we propose the use of sequential convex programming (SCP), which is one of several methods available for numerically solving optimization problems like (3). (See [78] for a collection of earlier SCP varieties and uses and [79] for a recent informative overview.) SCP provides a general framework for finding local optimal solutions to the worst-case robust optimization problem (3). The specific SCP algorithm used here, delineated in Algorithm 1, below, follows directly from [52] and [55]. It was used previously for robust design of slow-light tapers in photonic-crystal waveguides [55,56] and quantum potential profiles for electron transmission in semiconductor nanodevices [57]. In this paper, we apply this SCP algorithm to identify robust control fields for the generation of quantum gates in an uncertain one-qubit system.

II. SEQUENTIAL CONVEX PROGRAMMING

The SCP algorithm used here is shown in abstract form in Algorithm 1. The algorithm is initialized with (i) a control in the feasible set Θ, which is assumed to be convex; (ii) samples δ_i, $i = 1, \ldots, L$, taken from the uncertainty set Δ, which need not be convex; and (iii) a convex trust region $\tilde{\Theta}_{\text{trust}}$. The trust region is selected so that the linearized fidelity $\mathcal{F}(\theta,\delta_i) + \tilde{\theta}^{\mathsf{T}}\nabla_\theta \mathcal{F}(\theta,\delta_i)$, where $\tilde{\theta} \in \tilde{\Theta}_{\text{trust}}$, used in the optimization step retains sufficient accuracy. In each iteration the SCP algorithm returns the optimal increment $\tilde{\theta}$ and the associated worst-case linearized fidelity. To compute the actual worst-case fidelity requires simulating the system's evolution with the control variables $\theta + \tilde{\theta}$ as indicated in step 3 of Algorithm 1. The centerpiece is the optimization step, which, in the version shown in Algorithm 1, is gradient based, thereby resulting in L affine constraints in $\tilde{\theta}$ and, hence, is a convex optimization. The Hessian, perhaps not so easily computed, can be easily incorporated as shown in Appendix A. In some cases the number of samples, L, can be very large. Fortunately, however, computational complexity grows gracefully with the number of constraints and thus does not grossly affect the convex optimization efficiency [8].

Algorithm 1. Robust control via SCP.

Initialize:

control $\theta \in \Theta \subseteq \mathbb{R}^N$;
uncertainty/noise sample $\{\delta_i \in \Delta,\ i = 1, \ldots, L\}$;
trust region $\tilde{\Theta}_{\text{trust}} \subseteq \mathbb{R}^N$.

Repeat:

(1) Calculate fidelities and gradients:

$$\mathcal{F}(\theta,\delta_i), \quad \nabla_\theta \mathcal{F}(\theta,\delta_i) \in \mathbb{R}^N, \quad i = 1, \ldots, L.$$

(2) Using the linearized fidelity, solve for the increment $\tilde{\theta}$ from the convex optimization:

$$\text{maximize} \quad \min_i [\mathcal{F}(\theta,\delta_i) + \tilde{\theta}^{\mathsf{T}}\nabla_\theta \mathcal{F}(\theta,\delta_i)]$$
$$\text{subject to} \quad \theta + \tilde{\theta} \in \Theta,\ \tilde{\theta} \in \tilde{\Theta}_{\text{trust}}.$$

(3) Update:

if $\min_i \mathcal{F}(\theta + \tilde{\theta},\delta_i) > \min_i \mathcal{F}(\theta,\delta_i)$, **then**
replace θ by $\theta + \tilde{\theta}$ and increase $\tilde{\Theta}_{\text{trust}}$
else
decrease $\tilde{\Theta}_{\text{trust}}$
endif

Until stopping criteria satisfied.

In Appendix A we show how the gradient and Hessian can be cast in standard forms compatible with freely available software specifically designed to solve such convex optimization problems. In general, solving the convex optimization is not the most time-consuming step in the SCP algorithm. The time burden in each iteration falls more often on simulations required to compute the fidelities and gradients (and the Hessian if used) at each uncertainty sample. Of course, as is the case with numerical simulations of any quantum information

system, there always lurks the exponential scaling with the number of qubits.

Despite its many advantages, SCP is a local optimization method. As such, there is no way to verify that a globally optimal solution has been found. Since the fidelity by construction cannot exceed 1, it would seem that at least the maximum is known, so if $\mathcal{F} = 1$ is achieved, it is an optimal solution. However, even as we often obtain fidelities that are extremely close to 1, for example, $\log_{10}(1-\mathcal{F}) \in [-6,-4]$, this does not guarantee that the algorithm did not miss a better solution. Although a fidelity value with four to six 9's following the decimal point is effectively 1 for most engineering problems, for quantum computing every additional improvement in fidelity is important, since it can greatly decrease the physical resources required for fault-tolerant operation.

III. SEQUENTIAL CONVEX PROGRAMMING FOR AN UNCERTAIN QUBIT

In this section, we show how to use SCP for robust control of quantum gates in the presence of common types of uncertainties and constraints. We consider a one-qubit system modeled by the time-dependent Hamiltonian ($\hbar = 1$),

$$H(t) = c(t)\omega_x X + \omega_z Z, \tag{4}$$

where $c(t)$ is the external control field (a real-valued function of time defined on the interval $[0,T]$), and X and Z are the respective Pauli matrices. The real parameters ω_x and ω_z are constant but uncertain over the time interval $[0,T]$. Correspondingly, the uncertain parameters δ in Eq. (3) are specified by the parameter vector $\omega = [\omega_x,\omega_z]^{\mathsf{T}}$.

A. Control generation and constraints

The control field $c(t)$ is typically the output of a signal generation device whose dynamics impose constraints on magnitudes, bandwidth, and so on. To illustrate the use of SCP we make the simplifying assumption that the control is piecewise constant over N uniform time intervals of width $h = T/N$,

$$c(t,\theta) = \theta_k \quad \text{for} \quad t \in (t_{k-1},t_k], \quad k = 1,\ldots,N, \tag{5}$$

where $t_k = kh$. Correspondingly, the design variables θ in the optimization problem (3) are specified by the vector of field values $\theta = [\theta_1,\ldots,\theta_N]^{\mathsf{T}}$. The set Θ reflects control constraints, typical examples of which are listed in Table I. The appearance of control constraints due to signal generation dynamics is discussed in Appendix B.

TABLE I. Typical control constraints. The bounding parameters α, β, γ are positive constants and $c^{\min}$, $c^{\max}$ are real constants. Also, a is a real $N \times N$ matrix and b is a real vector of length N.

Constraint	Set Θ
None	$\mathbb{R}^N$
Fluence	$\Phi(\theta) \leqslant \gamma$
Magnitude	$c^{\min} \leqslant c(t,\theta) \leqslant c^{\max}$, $t \in [0,T]$
Slew rate	$\lvert\dot{c}(t,\theta)\rvert \leqslant \beta$, $t \in [0,T]$
Area	$A(\theta) \leqslant \alpha$
Linear	$a\theta = b$

A couple of important characteristics of the control field, used in Table I, are the fluence (a measure of the field energy),

$$\Phi(\theta) = \int_0^T c^2(t,\theta)dt = \|\theta\|_2^2 h, \tag{6}$$

and the area (a measure of the field strength),

$$A(\theta) = \int_0^T |c(t,\theta)|dt = \|\theta\|_1 h, \tag{7}$$

where $\|\theta\|_p = (\sum_{k=1}^N |\theta_k|^p)^{1/p}$ is the vector L^p norm.

The list of control constraints in Table I is certainly not exhaustive. However, since $c(t,\theta)$ is a linear function of θ, each of these constraints or any combination thereof forms a convex set in $\mathbb{R}^N$. The bounding parameters in Table I can also be used as design variables to establish control resource trade-offs via SCP. In particular, the trade-off between the gate fidelity and the field fluence is explored in Sec. V.

B. Evolution operator and fidelity

For a given realization of Hamiltonian (4) (i.e., for given values of ω_x and ω_z), the system undergoes a unitary evolution, governed by the Schrödinger equation,

$$i\frac{d}{dt}U(t) = H(t)U(t), \quad U(0) = I, \tag{8}$$

where $U(t) \equiv U(t,0)$ is the time-evolution operator (propagator) from time $t = 0$ to t, and I is the identity operator. For the piecewise-constant control, (5), the evolution operator $U(t_k)$ is given by a product of incremental propagators:

$$U(t_k) = U(t_k,t_{k-1})\cdots U(t_2,t_1)U(t_1,t_0), \tag{9}$$

$$U(t_k,t_{k-1}) = \exp\left[-ih(\theta_k\omega_x X + \omega_z Z)\right]. \tag{10}$$

In particular, the evolution operator attained at the final time T is $U_T \equiv U(T) = U(t_N)$. This evolution operator is a function of θ and ω.

The fidelity of a quantum gate is a measure of alignment between the target unitary transformation W and the actual final-time evolution operator U_T. Specifically, for the one-qubit system, we use the fidelity defined as

$$\mathcal{F}(\theta,\omega) = \tfrac{1}{4}|\mathrm{Tr}(W^\dagger U_T)|^2. \tag{11}$$

This fidelity, normalized to $[0,1]$, is independent of the phase of either W or U_T. Along with fidelity, we also use the normalized distance between W and U_T, which is defined as

$$\mathcal{D}(\theta,\omega) = 1 - \mathcal{F}(\theta,\omega). \tag{12}$$

In accordance with Eq. (12), $\mathcal{D}_{\mathrm{avg}}(\theta) = 1 - \mathcal{F}_{\mathrm{avg}}(\theta)$ and $\mathcal{D}_{\mathrm{wc}}(\theta) = 1 - \mathcal{F}_{\mathrm{wc}}(\theta)$.

C. Uncertainty modeling

One general approach to modeling the uncertainty in the Hamiltonian parameters ω is via a *deterministic* (or set-membership) model,

$$\Delta = \{\|\Omega^{-1}(\omega - \bar{\omega})\|_p \leqslant 1\}, \tag{13}$$

where $\bar{\omega} = [\bar{\omega}_x,\bar{\omega}_z]^{\mathsf{T}}$ is the vector of nominal values, Ω is a positive-definite weighting matrix (here 2×2), and p is

typically 2 or ∞. If $p = \infty$ and Ω is diagonal, then ω_x and ω_z are not correlated, in which case Eq. (13) reduces to

$$\Delta = \{|\omega_x - \bar{\omega}_x| \leqslant \tilde{\omega}_x,\ |\omega_z - \bar{\omega}_z| \leqslant \tilde{\omega}_z\}, \tag{14}$$

where $\Omega = \mathrm{diag}(\tilde{\omega}_x,\tilde{\omega}_z)$. If Ω is not diagonal, then ω_x and ω_z are correlated, possibly arising, respectively, from an approximation of a joint Gaussian or uniform distribution, with Ω, typically, being the covariance matrix associated with a specified confidence region for the parameters.

The uncertainty in the parameters can often be best described via a *probabilistic* model, for example,

$$\Delta = \{\mathbb{E}\{\omega\} = \bar{\omega},\ \mathbb{E}\{(\omega - \bar{\omega})(\omega - \bar{\omega})^{\mathsf{T}}\} = C\}, \tag{15}$$

where $\mathbb{E}\{\cdot\}$ is the expectation with respect to the underlying probability distribution of ω. If this distribution is Gaussian, then $\Delta = \{\omega \in \mathcal{N}(\bar{\omega},C)\}$.

Random uncertainty also arises from noise in the control field and/or environment, best represented by a *stochastic* model. In this case, the uncertainty set Δ can have the same form as in Eq. (15), but here the elements of $\omega(t) = [\omega_x(t),\omega_z(t)]^{\mathsf{T}}$ are stochastic variables with the moments $\mathbb{E}\{\omega(t)\} = \bar{\omega}$ and $\mathbb{E}\{[\omega(t) - \bar{\omega}][\omega(t') - \bar{\omega}]^{\mathsf{T}}\} = C(t,t')$.

As mentioned above, the uncertainty set Δ for SCP need not be convex; for example, the set of Eq. (13) is convex, but that of Eq. (15) is not. Step 2 in Algorithm 1 only requires that the uncertain parameters be *sampled* from the set Δ. In a numerical example studied below, we use a simple uniform sampling from an uncertainty set of the form (14). More sophisticated methods cycle through a sampling in the optimization step followed by validation on a different sampled set; bad parameters revealed in the validation step can be used in a new sampling for a repeat of the optimization step (e.g., see Ref. [52]).

D. Robust optimization

Now we can formulate a specific instance of the optimization problem, (3), corresponding to finding a robust control field for generating a target quantum gate in an uncertain one-qubit system. Specifically, the goal is to solve for the field values $\theta \in \mathbb{R}^N$ from the optimization problem:

$$\begin{array}{lll} \text{maximize} & \min\limits_{\omega\in\Delta} \mathcal{F}(\theta,\omega) & \\ \text{subject to} & U_T & \text{obtained from Eq. (9),} \\ & \theta \in \Theta & \text{from a combination of sets in Table I,} \\ & \omega \in \Delta & \text{from Eq. (13) or Eq. (15).} \end{array} \tag{16}$$

Since Θ is a convex set and samples are taken from Δ to compute gradients of $\mathcal{F}$ with respect to θ, then step 2 of Algorithm 1 will be a convex optimization.

IV. ROBUST ONE-QUBIT GATES

We use the SCP routine to find robust control fields corresponding to the following target unitary transformations:

$$W_{\mathrm{I}} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad W_{\mathrm{H}} = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}, \quad W_{\mathrm{P}} = \begin{bmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{bmatrix}. \tag{17}$$

Here, W_{I}, W_{H}, and W_{P} represent the identity, Hadamard, and phase ($\pi/8$) gates, respectively. Note that W_{H} and W_{P} comprise a universal gate set for one-qubit operations.

The uncertainty set used for all optimizations presented in this section is

$$\Delta = \{|\omega_x - 1| \leqslant 0.01, \quad |\omega_z - 2| \leqslant 0.20\}, \tag{18}$$

corresponding to a deterministic model with 1% control amplitude uncertainty and 10% drift magnitude uncertainty. For each target gate, SCP is used to solve for $\theta \in \mathbb{R}^N$ from

$$\begin{array}{lll} \text{maximize} & \min\limits_{\omega\in\Delta} \left\{\mathcal{F}(\theta,\omega) = \frac{1}{4}|\mathrm{Tr}(W^{\dagger}U_T)|^2\right\} & \\ \text{subject to} & U_T & \text{obtained from Eq. (9),} \\ & \theta \in \mathbb{R}^N & \text{(unconstrained),} \\ & \omega \in \Delta & \text{from Eq. (18).} \end{array} \tag{19}$$

We obtain solutions of Eq. (19) for all combinations of $W \in \{W_{\mathrm{I}},W_{\mathrm{H}},W_{\mathrm{P}}\}$ and $T \in \{1,2,4\}$, along with selected values of $N \in \{5,10,20,80\}$. For each SCP optimization presented in this section, we first used a gradient-based search [76,80] to find a control-field vector $\theta^{(0)}$ that achieves $\mathcal{F}\left(\theta^{(0)},\bar{\omega}\right) \simeq 0.999$ for the nominal parameter values $\bar{\omega}_x = 1$ and $\bar{\omega}_z = 2$. For a fixed ω (i.e., in the absence of uncertainty), it is easy to achieve unit fidelity to a desired numerical accuracy, so $\theta^{(0)}$ is a solution which is close to the top of the landscape but not fully optimal. Then $\theta^{(0)}$ was used as the initial field to start the SCP search for the uncertain system.

Figure 1 shows control fields that are solutions of the worst-case robust optimization problem, (19), for various choices of W, T, and N, along with corresponding distances $\mathcal{D}(\theta,\omega)$ that are plotted on a logarithmic scale as functions of the parameters ω_x and ω_z. Properties of these robust controls, including logarithms of the corresponding worst-case and average distances, the field fluence, and the maximum field value, are listed in Table II. For each target gate, we present results for eight combinations of N and T. With the one-qubit system of Eq. (4) and the uncertainty set Δ of Eq. (18), controls with worst-case fidelities $\mathcal{F}_{\mathrm{wc}}(\theta) \geqslant 0.9999$ are obtained for $N \geqslant 10$ and $T \geqslant 2$ for all three target gates. These results demonstrate that robust, high-fidelity control is possible with a relatively small number of control variables N, provided that the final time T is chosen properly.

Interestingly, the worst-case fidelity $\mathcal{F}_{\mathrm{wc}}$ can decrease as the number of field values N increases; this behavior is seen in the results in Table II for the identity gate with $T = 2$ when N increases from 10 to 20 and for the Hadamard gate with $T = 1$ when N increases from 5 to 20. Since the set of controls with N_1 field values is a proper subset of controls with $N_2 > N_1$ field values, these results suggest that the control landscape for the optimization problem, (19), possesses local optima that can trap SCP searches. Thus, even more robust solutions are, in principle, achievable by combining SCP with a nonlocal algorithm capable of exploring multiple optima.

V. TRADEOFF BETWEEN GATE FIDELITY AND CONTROL FIELD FLUENCE

The success of the optimization depends on available control resources, and it is expected that constraints on the control field will, generally, decrease the attainable fidelity [77,80]. As a further illustration of the utility of SCP, we use it

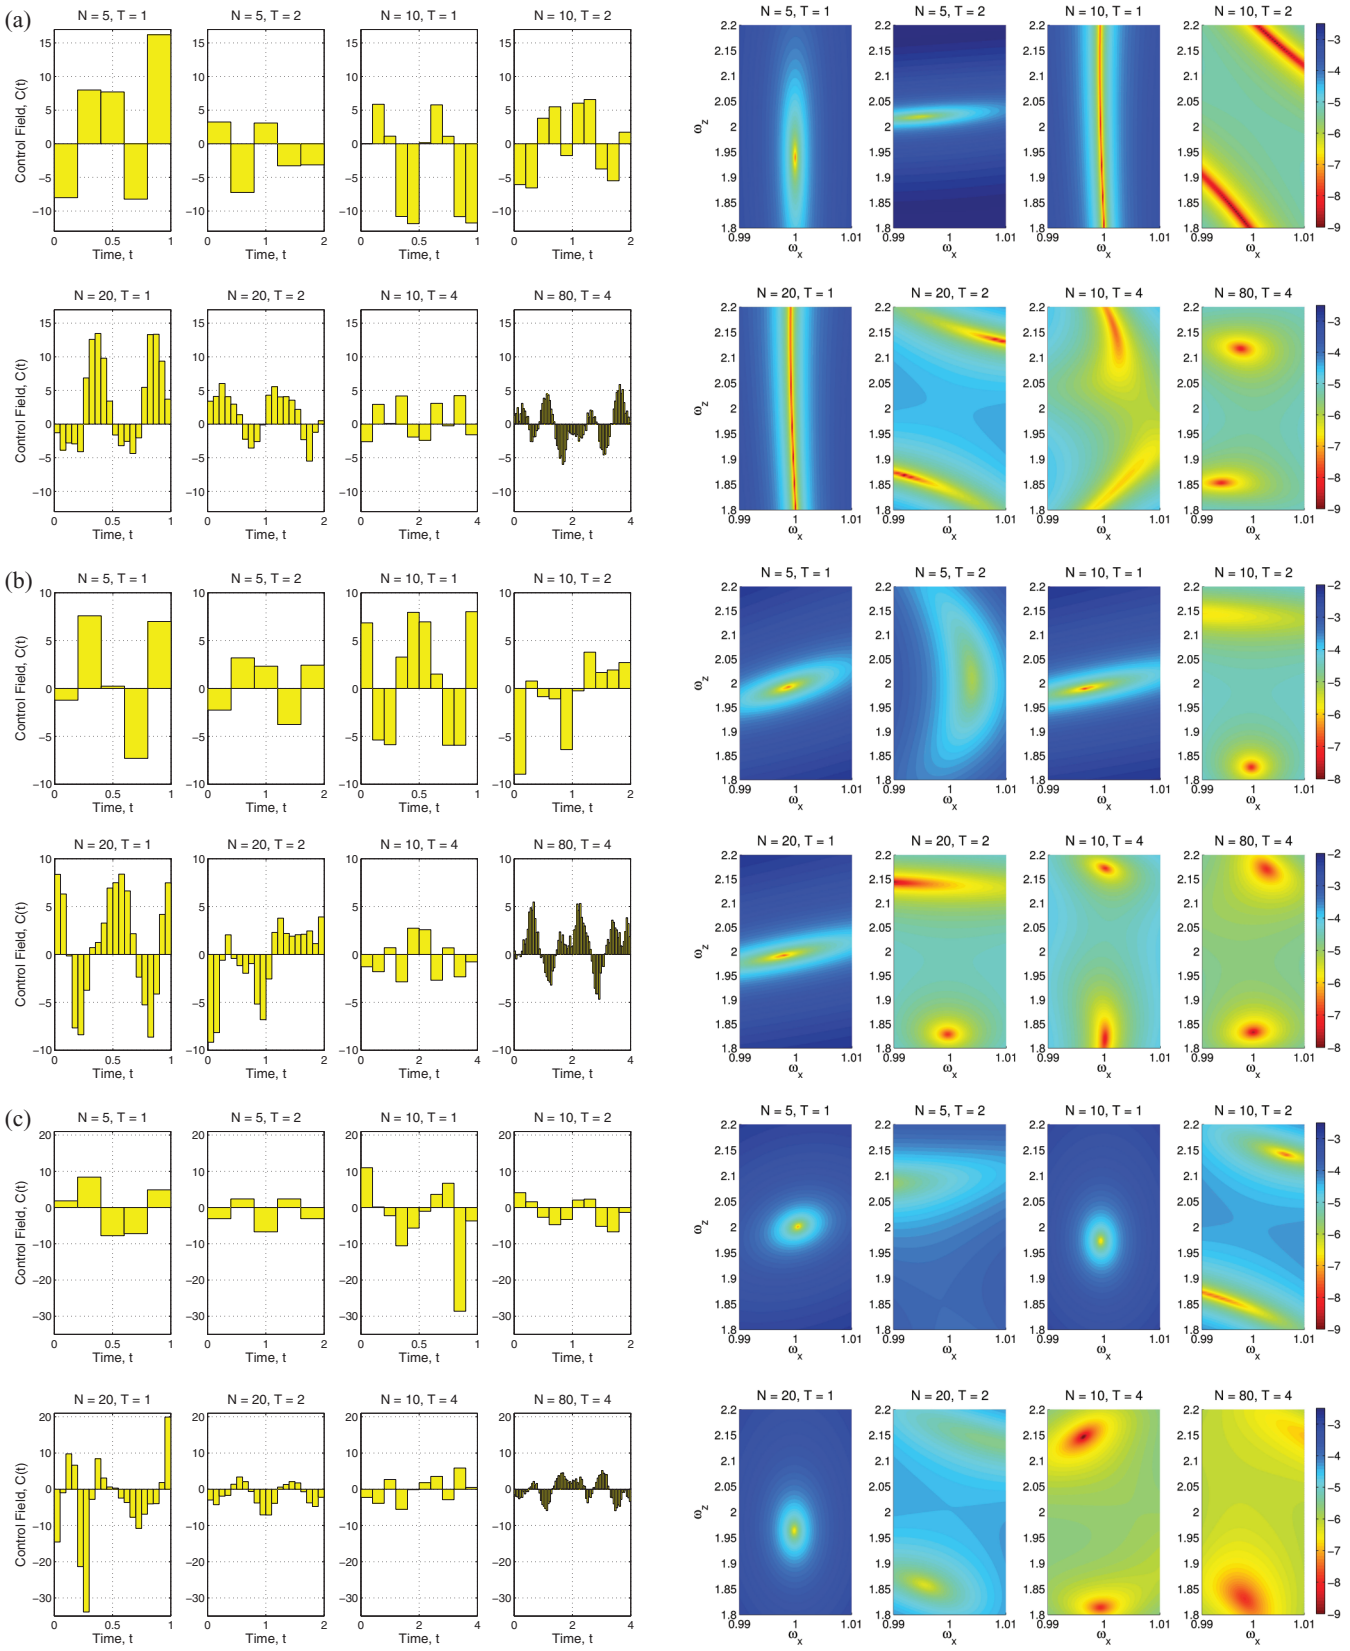

FIG. 1. (Color online) Left column: Control fields $c(t,\theta)$ that are solutions of the worst-case robust optimization problem, (19). Right column: Logarithms of corresponding distances, $\log_{10}\mathcal{D}(\theta,\omega)$, as functions of $\omega_x \in [0.99,1.01]$ and $\omega_z \in [1.8,2.2]$. Results are shown for eight combinations of N and T for each of the three target gates of Eq. (17): (a) the identity gate, (b) the Hadamard gate, and (c) the phase gate.

TABLE II. Properties of control fields that are solutions of the worst-case robust optimization problem, (19), for the three target gates of Eq. (17) and various combinations of N (the number of field values) and T (the final time).

N	T	$\log_{10}\mathcal{D}_{\rm wc}(\theta)$	$\log_{10}\mathcal{D}_{\rm avg}(\theta)$	$\Phi(\theta)$	$\max(\theta)$
		Target gate: Identity			
5	1	−3.13	−3.82	103.63	16.21
5	2	−2.35	−3.16	37.24	7.25
10	1	−3.28	−4.20	58.45	11.88
10	2	−5.23	−5.79	51.00	6.59
20	1	−3.31	−4.24	53.66	13.47
20	2	−4.35	−4.98	25.34	6.03
10	4	−4.62	−5.66	28.96	4.22
80	4	−5.08	−5.60	31.74	6.00
		Target gate: Hadamard			
5	1	−2.20	−3.08	32.27	7.59
5	2	−3.02	−3.74	16.35	3.77
10	1	−2.17	−3.05	36.93	8.02
10	2	−4.33	−4.80	30.33	8.95
20	1	−2.17	−3.06	34.38	8.64
20	2	−4.34	−4.86	30.07	9.17
10	4	−4.06	−4.63	16.60	2.86
80	4	−4.69	−5.12	25.61	5.48
		Target gate: Phase			
5	1	−2.77	−3.51	41.98	8.39
5	1	−3.71	−4.19	29.99	6.70
10	1	−2.96	−3.55	116.17	28.62
10	2	−4.34	−4.88	25.40	6.80
20	1	−3.02	−3.61	136.06	33.88
20	2	−4.30	−4.77	23.39	7.12
10	4	−5.57	−6.02	46.82	5.87
80	4	−6.00	−6.34	33.51	5.91

to explore the trade-off between the gate's worst-case fidelity and the control field's fluence. Specifically, we consider five uncertainty sets for ω:

$$\Delta_1 = \{|\omega_x - 1| \leqslant 0.001,\ |\omega_z - 2| \leqslant 0.02\}, \quad (20a)$$

$$\Delta_2 = \{|\omega_x - 1| \leqslant 0.010,\ |\omega_z - 2| \leqslant 0.02\}, \quad (20b)$$

$$\Delta_3 = \{|\omega_x - 1| \leqslant 0.001,\ |\omega_z - 2| \leqslant 0.10\}, \quad (20c)$$

$$\Delta_4 = \{|\omega_x - 1| \leqslant 0.010,\ |\omega_z - 2| \leqslant 0.10\}, \quad (20d)$$

$$\Delta_5 = \{|\omega_x - 1| \leqslant 0.010,\ |\omega_z - 2| \leqslant 0.20\}. \quad (20e)$$

These sets correspond to deterministic models with relative variations ranging from 0.1% to 1% in ω_x and from 1% to 10% in ω_z. Note that Δ_5 is the uncertainty set in Eq. (18) with relative variations in ω_x and ω_z at 1% and 10%, respectively. For each of the uncertainty sets in Eqs. (20), we use SCP to solve for $\theta \in \mathbb{R}^N$ from

$$\begin{aligned} &\text{maximize} && \min_{\omega\in\Delta}\left\{\mathcal{F}(\theta,\omega) = \tfrac{1}{4}|\mathrm{Tr}(W^\dagger U_T)|^2\right\} \\ &\text{subject to} && U_T \text{ obtained from Eq. (9)}, \\ & && \Phi(\theta) \leqslant \gamma, \\ & && \omega \in \Delta_m \text{ from Eqs. (20)}. \end{aligned} \quad (21)$$

The solutions of Eq. (21) are obtained for the target identity gate $W_{\rm I}$, final time $T = 2$, number of field values $N = 10$, and varying values of the fluence bound γ. For each uncertainty set Δ_m ($m = 1,\ldots,5$), we perform a series of SCP searches with decreasing γ. In the first SCP search in the series, the fluence bound is set to $\gamma = \infty$ (i.e., the fluence is unconstrained) and the solution of the optimization problem, (19), with uncertainty set (18) is used as the initial field. In each subsequent search in the series, γ is set to 0.95 of the fluence of the control field found in the previous search, and all values of the initial field are reduced proportionally so as to match the new fluence constraint. This process is repeated until the SCP routine fails to achieve $\mathcal{F}_{\rm wc} \geqslant 0.9$ due to the severity of the fluence constraint.

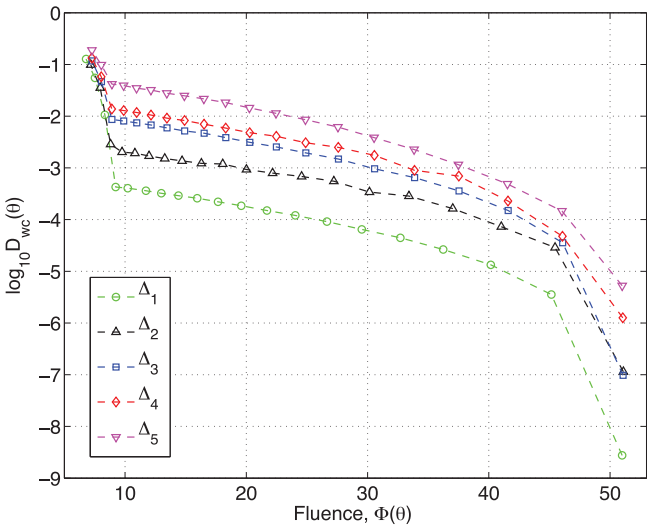

FIG. 2. (Color online) The logarithm of the worst-case distance, $\log_{10}\mathcal{D}_{\rm wc}(\theta)$, versus the fluence $\Phi(\theta)$ for control fields that are solutions of the optimization problem, (21), with $W = W_{\rm I}$, $T = 2$, and $N = 10$. Five data series, denoted by shape and color, correspond to five uncertainty sets Δ_m of Eqs. (20), as indicated in the legend.

Figure 2 shows the resulting trade-offs between the logarithm of the worst-case distance, $\log_{10}\mathcal{D}_{\rm wc}$, and the *achieved* field fluence, $\Phi(\theta)$, for each of the uncertainty sets Δ_m in Eq. (20). The rightmost point in each series corresponds to unconstrained fluence ($\gamma = \infty$). The rate of increase in the distance as the fluence bound is decreased is seen to be essentially the same for all sets Δ_m. Additionally, the fluence value where the distance abruptly changes for the worse is also about the same: $\Phi \approx 10$. It is important to note that it is not known whether any of the trade-off curves in Fig. 2 represent a true *Pareto front* for distance versus field fluence.

The trade-off curves in Fig. 2 show that the gate error, on average, exhibits a similar sensitivity of about one order of magnitude to either 1% variation in ω_x or 5% variation in ω_z. The greater sensitivity to variations in ω_x is due in part to the fact that in our model system, (4), ω_x is a direct *uncertain multiplicative gain* on the control signal. In other words, a perturbation $\tilde{\omega}_x$ around $\bar{\omega}_x = 1$ is equivalent to a perturbation $\tilde{\omega}_x\theta$ in the control field. Following the procedure presented in Refs. [40] and [80], a Taylor series approximation of the fidelity up to the second order in $\tilde{\omega}_x$ gives

$$\begin{aligned} \mathcal{F}(\theta,\bar{\omega}_x+\tilde{\omega}_x,\omega_z) &= \mathcal{F}(\theta+\tilde{\omega}_x\theta,\bar{\omega}_x,\omega_z) \\ &\approx \mathcal{F}(\theta,\bar{\omega}_x,\omega_z) + \tilde{\omega}_x\theta^{\mathsf{T}}\nabla_\theta\mathcal{F}(\theta,\bar{\omega}_x,\omega_z) \\ &\quad + \tfrac{1}{2}\tilde{\omega}_x^2\theta^{\mathsf{T}}\nabla_\theta^2\mathcal{F}(\theta,\bar{\omega}_x,\omega_z)\theta. \end{aligned} \quad (22)$$

For any control field that is a solution of the optimization problem, (19) or (21), the Hessian $\nabla_\theta^2 \mathcal{F}(\theta,\bar{\omega}_x,\omega_z)$ is negative semidefinite and the Hessian term dominates the gradient term. Assuming $|\tilde{\omega}_x| \leqslant \epsilon$, we use Eq. (22) to obtain a lower bound on the fidelity:

$$\mathcal{F}(\theta,\bar{\omega}_x+\tilde{\omega}_x,\omega_z) \geqslant \mathcal{F}(\theta,\bar{\omega}_x,\omega_z) - \epsilon|\theta^{\mathsf{T}}\nabla_\theta \mathcal{F}(\theta,\bar{\omega}_x,\omega_z)| - \tfrac{1}{2}\epsilon^2\left|\theta^{\mathsf{T}}\nabla_\theta^2 \mathcal{F}(\theta,\bar{\omega}_x,\omega_z)\theta\right|. \quad (23)$$

We evaluated the lower bound $\underline{\mathcal{F}}(\theta,\bar{\omega}_x,\epsilon,\omega_z)$, given by the right-hand side of Eq. (23), for control fields whose worst-case distances are shown in Fig. 2, using $\bar{\omega}_x = 1$ and ϵ values ($\epsilon = 0.001$ and $\epsilon = 0.01$) and ω_z ranges ($\omega_z \in [1.98,2.02]$, $\omega_z \in [1.9,2.1]$, $\omega_z \in [1.8,2.2]$) that correspond to the five uncertainty sets Δ_m of Eqs. (20). Then, for each field, we minimized $\underline{\mathcal{F}}(\theta,\bar{\omega}_x,\epsilon,\omega_z)$ over the respective ω_z range and found that the resulting distance values coincide almost exactly with points on the corresponding trade-off curves in Fig. 2. This coincidence indicates that the lower bound $\underline{\mathcal{F}}(\theta,\bar{\omega}_x,\epsilon,\omega_z)$ approximates well the minimum of the fidelity $\mathcal{F}(\theta,\omega)$ over the ω_x variation.

The trade-off analysis is valuable for understanding the interplay between constraints in control and system designs. In particular, a limitation on the maximum field fluence can reflect not only signal generator constraints, but also system design considerations such as thermal loads on the system. For example, if variations in ω_x and ω_z are small, such as in the uncertainty set Δ_1, we observe from the corresponding curve in Fig. 2 that a distance value $\mathcal{D}_{\text{wc}} \sim 10^{-4}$ is possible with a fairly low fluence ($\Phi \approx 25$), and it can even be as low as $\mathcal{D}_{\text{wc}} \sim 10^{-8}$ with a higher fluence ($\Phi \approx 50$). Attaining parameter uncertainties of the order of Δ_1 could be accomplished via material and hardware improvements and better manufacturing and testing procedures. Certainly, the possibility of achieving such high fidelities is a motivation to explore these options. Thus, establishing the trade-off between the gate fidelity and the field fluence provides important information about possibilities for enhancing the robust gate performance.

VI. EFFECT OF NOISE

When uncertainty is due to noise, performing SCP generally requires some form of sampling from the noise distribution. If the noise is sufficiently weak, then, based on ideas from Ref. [40], we show in Appendix C that an approximation can be utilized which avoids expensive sampling. We explore this approach in more detail in a future paper. Here, we consider a simplified scenario which captures some of the salient features of robust control in the presence of noise.

Consider the Hamiltonian of Eq. (4), where the parameter ω_x in the control term is constant, $\omega_x = 1$, and the parameter ω_z in the drift term is a noisy time series, i.e.,

$$\omega_z(t) = \bar{\omega}_z + \tilde{\omega}_z(t), \quad t \in [0,T], \quad (24)$$

where $\bar{\omega}_z$ is the average value of $\omega_z(t)$ and $\tilde{\omega}_z(t)$ is a stochastic variable obtained as the output of a linear filter G driven by stationary, Gaussian white noise $u(t)$ with variance σ^2. Specifically, $u(t)$ satisfies

$$\mathbb{E}\{u(t)\} = 0, \quad \mathbb{E}\{u(t)u(t')\} = \sigma^2\delta(t-t'), \quad (25)$$

the filter action is

$$\tilde{\omega}_z(t) = (G * u)(t), \quad t \in (-\infty,T], \quad (26)$$

G is a linear first-order filter with the transfer function

$$G(s) = 1/(s\tau + 1), \quad (27)$$

and τ is the filter time constant.

The average gate fidelity $\mathcal{F}_{\text{avg}}(\theta) = \mathbb{E}\{\mathcal{F}(\theta,\omega_z)\}$ under a noise process affecting ω_z can be evaluated using random sampling from the noise distribution:

$$\mathcal{F}_{\text{avg}}(\theta) \approx \sum_{l=1}^{L} \mathcal{F}\big(\theta,\omega_z^{(l)}\big). \quad (28)$$

Here, $\omega_z \in \mathbb{R}^M$ is the vector whose elements represent a piecewise-constant approximation of the time series $\omega_z(t)$ with a uniform time step $\tilde{h} = T/M$, $\omega_z^{(l)} = \bar{\omega}_z + \tilde{\omega}_z^{(l)}$ is the vector corresponding to the lth realization of the noise process, and L is the number of noise realizations in the sample. Another method for evaluating $\mathcal{F}_{\text{avg}}(\theta)$ is the weak noise approximation described in Appendix C. Specifically, using the Taylor series expansion of the fidelity about $\bar{\omega}_z$ up to second order in $\tilde{\omega}_z$ and assuming that $\tilde{\omega}_z$ has zero mean and covariance matrix $C = \mathbb{E}\{\tilde{\omega}_z\tilde{\omega}_z^{\mathsf{T}}\}$, we obtain [cf. Eq. (C9)]

$$\mathcal{F}_{\text{avg}}(\theta) \approx \mathcal{F}(\theta,\bar{\omega}_z) - \tfrac{1}{2}\text{Tr}\big(C R_{\omega_z\omega_z}\big), \quad (29)$$

where $R_{\omega_z\omega_z} = -\nabla_{\omega_z}^2 \mathcal{F}(\theta,\bar{\omega}_z)$ is the negative Hessian matrix. For the filtered noise process of Eqs. (25)–(27), elements of the covariance matrix C are given by

$$C_{mm'} = \tilde{\sigma}^2 \frac{1-\alpha}{1+\alpha}\alpha^{|m-m'|}, \quad m,m' = 1,\ldots,M, \quad (30)$$

where $\tilde{\sigma}^2 = \sigma^2/\tilde{h}$ and $\alpha = \exp(-\tilde{h}/\tau)$.

It is interesting to analyze how a control field designed to be robust for a deterministic uncertainty model performs in the presence of noise. For example, consider the control field that is a solution of the optimization problem, (21), with $W = W_{\text{I}}$, $T = 2$, $N = 10$, Δ_1 of Eq. (20a), and no fluence constraint ($\gamma = \infty$); this field corresponds to the rightmost point on the bottom curve in Fig. 2. For this field, we use both the random sampling method of Eq. (28) and the weak noise approximation of Eq. (29) to evaluate the average fidelity $\mathcal{F}_{\text{avg}}(\theta)$ under the noise process of Eqs. (25)–(27) with $\bar{\omega}_z = 2$, $\sigma \in \{0.001,0.02\}$ and various values of τ. Figure 3 shows the corresponding values of $\log_{10}\mathcal{D}_{\text{avg}}(\theta)$, for a range of filter time constants relative to the control time, $\tau/T \in [10^{-4},10^4]$. We observe an excellent agreement between the weak noise approximation (solid lines) and simulated data from random sampling (circles).

Equations (29) and (30) can be further used to investigate the asymptotic behavior in the limits of low-bandwidth and high-bandwidth filters. For $\tau/T \gg 1$, the filter bandwidth is very low, and the noise is effectively blocked. In this limit, all elements of C are approximately 0, and $\mathcal{D}_{\text{avg}}(\theta) \approx \mathcal{D}(\theta,\bar{\omega}_z) \approx 10^{-8.88}$ is independent of σ. For $\tau/T \ll 1$, the filter bandwidth is very high, which allows for the white noise to pass through unaltered. In this limit, C is proportional to the identity matrix, $C_{mm'} \approx \tilde{\sigma}^2\delta_{mm'}$, and $\mathcal{F}_{\text{avg}}(\theta) \approx \mathcal{F}(\theta,\bar{\omega}_z) - \frac{1}{2}\tilde{\sigma}^2\text{Tr}(R_{\omega_z\omega_z})$. For a control field $c(t,\theta^\star)$ which is globally optimal for the

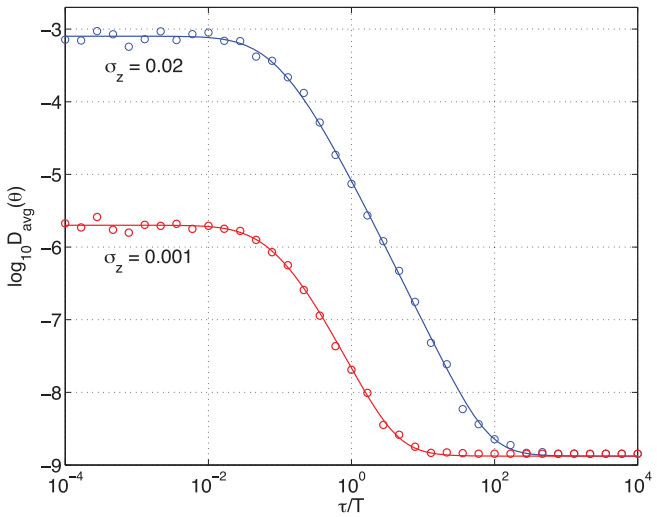

FIG. 3. (Color online) The logarithm of the average distance, $\log_{10}\mathcal{D}_{\rm avg}(\theta)$, versus τ/T, where τ is the noise filter timeconstant from Eq. (27). The control field used here is a solution of the optimization problem, (21), with $W = W_{\rm I}$, $T = 2$, $N = 10$, Δ_1 of Eq. (20a), and no fluence constraint ($\gamma = \infty$). The distance is averaged under the noise process described by Eqs. (24)–(27) with $\bar{\omega}_z = 2$ and two σ values: $\sigma = 0.001$ [lower (red) line and circles] and $\sigma = 0.02$ [upper (blue) line and circles]. The weak noise approximation (solid lines) computed using Eqs. (29) and (30) is in a very good agreement with simulated data (circles) obtained using random sampling from the noise distribution according to Eq. (28).

objective of maximizing $\mathcal{F}(\theta,\bar{\omega}_z)$, each diagonal matrix element of $R_{\omega_z\omega_z}$ equals $2\tilde{h}^2$, and we obtain a simple analytic result:

$$\mathcal{D}_{\rm avg}(\theta^\star) \approx \sigma^2 T. \tag{31}$$

The field that we use here is not exactly $\theta^\star$, but the value $\mathcal{D}(\theta,\bar{\omega}_z) \approx 10^{-8.88}$ is sufficiently close to the optimum for the result of Eq. (31) to be a very good approximation. Then, with $T = 2$, we obtain $\mathcal{D}_{\rm avg} \approx 2\times 10^{-6} \approx 10^{-5.70}$ for $\sigma = 0.001$ and $\mathcal{D}_{\rm avg} \approx 8\times 10^{-4} \approx 10^{-3.10}$ for $\sigma = 0.02$. The asymptotic results in both limits are very well confirmed by the data shown in Fig. 3.

The results in Fig. 3 also show that the control field, though not designed for stochastic uncertainty, nonetheless performs admirably. In fact, we used SCP to find control fields that are specifically robust against noise in ω_z but did not obtain a significant improvement. While in this example a robust control designed for a deterministic uncertainty model also works well against a stationary stochastic process, we do not know whether this behavior holds in general.

VII. SUMMARY

Using SCP we have demonstrated that it is possible to generate high-fidelity quantum gates with a substantial robustness against uncertainties, while simultaneously using limited control resources such as field amplitude, bandwidth, and fluence. Designing such robust control fields requires a specific knowledge of the range and character of the uncertainties, a process referred to in the control theory literature as "uncertainty modeling." Although we have focused on a one-qubit system, even this simple example clearly shows the strong effect of control constraints on the attainable degree of robustness. Our analysis of this system has also revealed that a control field designed for a deterministic (set-membership) uncertainty model can be quite effective against stochastic uncertainty (noise).

This work shows that SCP is useful for exploring possible improvements in the robust gate performance for different values and ranges available for both control and system designs. Specifically, SCP makes it possible to quantify a variety of trade-offs between constraints on control and system parameters. For example, one can determine how many control variables are required to achieve a desired worst-case fidelity for a given uncertainty range or, alternatively, how tight the uncertainty range should be for a given limitation on the maximum field fluence. Such a trade-off analysis could reveal a combination of physical design and robust control design resulting in a "sweet spot" among the possibilities.

Of course, SCP is not the only approach to finding locally optimal solutions to nonconvex problems. An important advantage of SCP is the ease with which various uncertainty models and constraints on design variables can be directly incorporated in the local convex optimization step of the algorithm. It would be desirable to develop a hybrid approach, integrating SCP with a nonlocal optimization method, in order to make it possible to search among multiple solutions.

The array of results presented here hopefully heralds what would be seen in more complex systems, involving multiple qubits, controlled ancillae, coupling to a bath, and so on. In addition, the results also begin to provide an insight into unanticipated control-field structures. Many of these potentialities are under consideration at present and will be forthcoming.

ACKNOWLEDGMENTS

We gratefully acknowledge helpful discussions with Kevin Young (SNL-CA), Kaveh Khodjasteh, and Lorenza Viola (Dartmouth College). This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories. Sandia is a multiprogram laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the United States Department of Energy's National Nuclear Security Administration under Contract No. DE-AC04-94AL85000. R.L.K. acknowledges support from ARO MURI Grant No. W911NF-11-1-0268 to USC and the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center Contract No. D11PC20165. The US Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the US Government.

APPENDIX A: CONVEX OPTIMIZATION

The convex optimization step in the SCP algorithm can be equivalently expressed as

$$\begin{aligned} &\text{maximize} && f_0 \\ &\text{subject to} && f_i + \tilde{\theta}^{\mathsf{T}} g_i \geqslant f_0, \quad i = 1,\ldots,L, \\ &&& \theta + \tilde{\theta} \in \Theta,\ \tilde{\theta} \in \tilde{\Theta}_{\text{trust}}, \end{aligned} \tag{A1}$$

where $f_i = \mathcal{F}(\theta,\delta_i)$ and $g_i = \nabla_\theta \mathcal{F}(\theta,\delta_i)$. Both $\tilde{\theta}$ and f_0 are now the optimization variables. If both Θ and $\tilde{\Theta}_{\text{trust}}$ bound their respective elements in a "box" in $\mathbb{R}^N$, then (A1) is a *linear program*.

The Hessian can be employed in the optimization step by using its negative-definite part, $R_i = -[\nabla^2_\theta \mathcal{F}(\theta,\delta_i)]_-$ where $[\cdot]_-$ retains only the negative eigenvalues of the Hessian; specifically, $R_i = V_i V_i^{\mathsf{T}}$ with $V_i \in \mathbb{R}^{N\times r}$, where r is the number of strictly negative eigenvalues of the Hessian less than or equal to a chosen threshold. Then the worst-case fidelity constraint can be formulated as

$$f_i + \tilde{\theta}^{\mathsf{T}} g_i - \tfrac{1}{2}\tilde{\theta}^{\mathsf{T}} V_i V_i^{\mathsf{T}} \tilde{\theta} \geqslant f_0, \quad i = 1,\ldots,L. \tag{A2}$$

Each of the inequalities in Eq. (A2) is equivalent to a linear matrix inequality in the variables $\{\tilde{\theta}, f_0\}$ [8]: $Q_i(\tilde{\theta}, f_0) \geqslant 0$, where

$$Q_i(\tilde{\theta}, f_0) = \begin{bmatrix} f_i - f_0 - \tilde{\theta}^{\mathsf{T}} g_i & \tilde{\theta}^{\mathsf{T}} V_i/\sqrt{2} \\ V_i^{\mathsf{T}}\tilde{\theta}/\sqrt{2} & I_r \end{bmatrix}, \tag{A3}$$

and I_r is the $r \times r$ identity matrix. The optimization step in SCP is now given by the *semidefinite program*:

$$\begin{aligned} &\text{maximize} && f_0 \\ &\text{subject to} && Q_i(\tilde{\theta}, f_0) \geqslant 0, \quad i = 1,\ldots,L, \\ &&& \theta + \tilde{\theta} \in \Theta,\ \tilde{\theta} \in \tilde{\Theta}_{\text{trust}}. \end{aligned} \tag{A4}$$

Optimization problems (A1) and (A4) are now expressed in standard forms suitable for use with existing software specially developed for these classes of convex optimization. In particular, YALMIP [81] and CVX [82,83] are convex compilers compatible with MATLAB. Using these software tools makes it very easy to code the convex optimization problems almost exactly as expressed mathematically. These compilers call for convex solvers such as SDPT-3 [84] and SeDuMi [85], which have been developed and in use for many years and, as a result, are generally efficient and reliable. There are limits imposed by both memory and speed for a particular problem instance and computer platform. In these cases it could be necessary to use or develop specialized versions with modifications that take into account the specific underlying structure of the problem.

APPENDIX B: SIGNAL GENERATION

In general, the control field $c(t,\theta)$ is the output of a signal generation device. As an example, consider a field generated by a device with rate ν and piecewise-constant commands θ,

$$\dot{c}(t,\theta) = \nu[\bar{c}(t,\theta) - c(t,\theta)], \quad c(0) = 0, \tag{B1a}$$

$$\bar{c}(t,\theta) = \theta_k \quad \text{for} \quad t \in (t_{k-1},t_k], \quad k = 1,\ldots,N, \tag{B1b}$$

where $t_k = kh$ and $h = T/N$. The field in this example can be expressed in a general form:

$$c(t,\theta) = \sum_{k=1}^{N} s_k(t)\theta_k = s(t)^{\mathsf{T}}\theta, \quad t \in [0,T]. \tag{B2}$$

This expression holds for *any* signal generation well represented by known linear dynamics whose input is a finite sequence of control commands $\{\theta_k\}$ at a uniform sampling rate. The linear dynamics are captured in the shape-function vector $s(t) \in \mathbb{R}^N$. For example, for the field of Eqs. (B1), the elements of $s(t)$ are given by

$$s_k(t) = \begin{cases} 0 & \text{for} \quad t \leqslant t_k, \\ 1 - e^{-\nu(t-t_{k-1})} & \text{for} \quad t \in (t_{k-1},t_k], \\ (1-e^{-\nu h})e^{-\nu(t-t_k)} & \text{for} \quad t > t_k. \end{cases} \tag{B3}$$

In the limit of very fast dynamics ($\nu \to \infty$), the element $s_k(t)$ of Eq. (B3) becomes the indicator function of the interval $(t_{k-1},t_k]$, and the control field $c(t,\theta)$ is piecewise constant over N uniform time intervals of width h, as given by Eq. (5).

Generally, when the dynamics of the signal generation device have an appreciable effect on the shapes of $\{s_k(t)\}$, the numerical integration of the Schrödinger equation, (8), would require using a time step over which the field $c(t,\theta)$ does not change much, i.e., finer than the command interval h.

Note that for any field of the form (B2), the control constraint sets in Table I are convex. For example, the constraint on the field fluence can be expressed as $\Phi(\theta) = \theta^{\mathsf{T}} B \theta \leqslant \gamma$ with $B = \int_0^T s(t)s(t)^{\mathsf{T}} dt$.

The field form of Eq. (B2) can be further generalized by considering multiple command vectors $\{\theta_i\}$ and shape-function vectors $\{s_i(t)\}$, i.e.,

$$c(t,\theta) = \sum_{i=1}^{K} s_i(t)^{\mathsf{T}}\theta_i. \tag{B4}$$

For example, laser pulse shaping in a liquid crystal modulator generates a control field of the form

$$c(t) = A_0(t)\sum_{i=1}^{K} a_i \sin(\omega_i t + \phi_i), \tag{B5}$$

where the envelope function $A_0(t)$ and frequencies $\{\omega_i\}$ are fixed, while the amplitudes $\{a_i\}$ and phases $\{\phi_i\}$ are the control variables. The field, (B5), can be equivalently expressed in the form (B4), where $s_i(t) = A_0(t)[\sin(\omega_i t), \cos(\omega_i t)]^{\mathsf{T}}$ and $\theta_i = a_i[\cos\phi_i, \sin\phi_i]^{\mathsf{T}}$.

For a control field of the form (B5), constraints are typically imposed on the amplitudes $\{a_i\}$, and since $\|\theta_i\|_2 = a_i$, they can be equivalently expressed as constraints on θ. For example, the magnitude constraint $a_i \leqslant a_{\max}$ is equivalent to the convex set $\|\theta_i\|_2 \leqslant a_{\max}$. However, the constraint that all the amplitudes are the same, i.e., $a_i = a_0$ is equivalent to the nonconvex set $\|\theta_i\|_2 = a_0$. This problem can be circumvented by using the constraint set $\|\theta_i\|_2 \leqslant a_0$, which is a *convex relaxation* [8] of the actual nonconvex one; then SCP will return a local solution to the relaxed problem. Some relaxations can be proven to be optimal, but that is not known here and hence a postoptimization analysis is required.

APPENDIX C: WEAK NOISE APPROXIMATION

When the noise variance is small, it is possible to avoid the expensive simulation of noise realizations drawn from the underlying distribution. The approach for evaluating the effect of weak noise, the basics of which are presented here, was introduced in Refs. [40] and [80] and will be explored in depth in a subsequent paper.

Consider an n-level quantum system with the Hamiltonian

$$H(t) = c(t)H_c + w(t)H_w, \tag{C1}$$

where $c(t)$ and $w(t)$ are, respectively, the control field and the noisy field (real-valued functions of time defined on the interval $[0,T]$). We assume that $c(t) = c(t,\theta)$ is a piecewise-constant function of the form (5). Let the elements $\{w_m\}$ of the vector $w \in \mathbb{R}^M$ represent a piecewise-constant approximation of $w(t)$, i.e.,

$$w(t) = w_m \quad \text{for} \quad t \in (\tilde{t}_{m-1},\tilde{t}_m], \quad m = 1,\ldots,M, \tag{C2}$$

where $\tilde{t}_m = m\tilde{h}$ and $\tilde{h} = T/M$. Assuming that $M \geqslant N$ and $p = M/N$ is an integer, the control can be represented as

$$c(t) = c_m \quad \text{for} \quad t \in (\tilde{t}_{m-1},\tilde{t}_m], \quad m = 1,\ldots,M, \tag{C3}$$

where $\{c_m\}$ are the elements of the vector $c = \theta \otimes e_p \in \mathbb{R}^M$ and e_p denotes the vector of 1's of length p. Analogous to Eqs. (9) and (10), the time-evolution operator is given by

$$U(\tilde{t}_m) = U(\tilde{t}_m,\tilde{t}_{m-1})\cdots U(\tilde{t}_2,\tilde{t}_1)U(\tilde{t}_1,\tilde{t}_0), \tag{C4}$$

$$U(\tilde{t}_m,\tilde{t}_{m-1}) = \exp[-i\tilde{h}(c_m H_c + w_m H_w)], \tag{C5}$$

and, in particular, $U_T = U(\tilde{t}_M)$. The gate fidelity is

$$\mathcal{F}(\theta,w) = \frac{1}{n^2}|\mathrm{Tr}(W^\dagger U_T)|^2. \tag{C6}$$

Assume that the noisy field has the form $w = \bar{w} + \tilde{w}$, where $\bar{w} \in \mathbb{R}^M$ is a deterministic mean and $\tilde{w} \in \mathbb{R}^M$ is a stochastic variable that represents a stationary noise process. For a specified control θ, the Taylor series expansion of the fidelity about $\bar{w}$ up to second order in $\tilde{w}$ gives the approximation,

$$\mathcal{F}(\theta,w) \approx \mathcal{F}(\theta,\bar{w}) + \tilde{w}^\mathsf{T} g_w - \tfrac{1}{2}\tilde{w}^\mathsf{T} R_{ww}\tilde{w}, \tag{C7}$$

where $g_w = \nabla_w \mathcal{F}(\theta,\bar{w}) \in \mathbb{R}^M$ is the gradient vector and $R_{ww} = -\nabla_w^2 \mathcal{F}(\theta,\bar{w}) \in \mathbb{R}^{M\times M}$ is the *negative* Hessian matrix. Assume that the stochastic variable $\tilde{w}$ has zero mean and covariance matrix $C \in \mathbb{R}^{M\times M}$, i.e.,

$$\mathbb{E}\{\tilde{w}\} = 0, \quad \mathbb{E}\{\tilde{w}\tilde{w}^\mathsf{T}\} = C. \tag{C8}$$

The fidelity averaged over all noise realizations is given by the statistical expectation: $\mathcal{F}_{\mathrm{avg}}(\theta) = \mathbb{E}\{\mathcal{F}(\theta,w)\}$. Using Eqs. (C7) and (C8), we obtain the weak noise approximation for the average fidelity:

$$\mathcal{F}_{\mathrm{avg}}(\theta) \approx \mathcal{F}(\theta,\bar{w}) - \tfrac{1}{2}\mathrm{Tr}(C R_{ww}). \tag{C9}$$

Since the dependence on noise in Eq. (C9) is only through the covariance matrix, the evaluation of $\mathcal{F}_{\mathrm{avg}}(\theta)$ via this approximation does not require random sampling from the noise distribution, providing a huge advantage in numerical efficiency.

For Gaussian white noise with variance σ^2, the covariance matrix is given by $C = (\sigma^2/\tilde{h})I_M$, and Eq. (C9) yields

$$\mathcal{F}_{\mathrm{avg}}^{(\mathrm{wn})}(\theta) \approx \mathcal{F}(\theta,\bar{w}) - \frac{\sigma^2}{2\tilde{h}}\mathrm{Tr}(R_{ww}). \tag{C10}$$

For a control $\theta^\star$ which is globally optimal for the objective of maximizing $\mathcal{F}(\theta,\bar{w})$, all diagonal matrix element of R_{ww} are equal to each other:

$$(R_{ww})_{mm} = \frac{2\tilde{h}^2}{n}\mathrm{Tr}\big(H_w^2\big) - \frac{2\tilde{h}^2}{n^2}[\mathrm{Tr}(H_w)]^2, \quad \forall m. \tag{C11}$$

If the operator H_w is traceless, substituting Eq. (C11) into Eq. (C10) leads to a simple analytical expression:

$$\mathcal{F}_{\mathrm{avg}}^{(\mathrm{wn})}(\theta^\star) \approx 1 - \frac{1}{n}\mathrm{Tr}\big(H_w^2\big)\sigma^2 T. \tag{C12}$$

This result shows that robustness against additive white noise can be improved only by reducing the control duration T; however, this can be done only as long as $T \geqslant T^*$, where T^* is a critical value below which the nominal objective is not reachable [80]. In a quantum information system, H_w is typically given by a tensor product of Pauli matrices and identity operators for individual qubits. In this case, $H_w^2 = I_n$, and Eq. (C12) is further simplified:

$$\mathcal{F}_{\mathrm{avg}}^{(\mathrm{wn})}(\theta^\star) \approx 1 - \sigma^2 T. \tag{C13}$$

The weak noise approximation together with a similar expansion for a small control change (from θ to $\theta + \tilde{\theta}$) can be used in the optimization step of SCP for designing controls robust to a stochastic uncertainty model. Expanding the fidelity about $\{\theta,\bar{w}\}$ up to second order in $\{\tilde{\theta},\tilde{w}\}$ gives

$$\mathcal{F}(\theta+\tilde{\theta},w) \approx f + \tilde{x}^\mathsf{T} g - \tfrac{1}{2}\tilde{x}^\mathsf{T} R\tilde{x}, \tag{C14}$$

where

$$\tilde{x} = \begin{bmatrix}\tilde{\theta}\\ \tilde{w}\end{bmatrix}, \quad g = \begin{bmatrix}g_\theta\\ g_w\end{bmatrix}, \quad R = \begin{bmatrix}R_{\theta\theta} & R_{\theta w}\\ R_{w\theta} & R_{ww}\end{bmatrix}, \tag{C15}$$

$f = \mathcal{F}(\theta,\bar{w})$ is the fidelity, $g_a = \nabla_a \mathcal{F}(\theta,\bar{w})$ are gradient vectors, and $R_{ab} = -\nabla_a\nabla_b \mathcal{F}(\theta,\bar{w})$ are negative Hessian matrices ($a,b \in \{\theta,w\}$). Given a model of the noise distribution, we can then pose the robust optimization problem:

$$\begin{array}{ll}\text{maximize} & \gamma\\ \text{subject to} & \mathrm{Prob}\{\mathcal{F}(\theta+\tilde{\theta},w) \geqslant \gamma\} \geqslant \eta, \quad \tilde{\theta} \in \Theta.\end{array} \tag{C16}$$

Assume further that the stochastic variable $\tilde{w}$ has a zero-mean Gaussian distribution with covariance matrix C, i.e., satisfies Eq. (C8), with $\|C\| = O(\sigma^2)$. Following the approach to robust optimization described in Chapter 4 of [8], the problem, (C16), is equivalent, up to $O(\sigma^2)$, to the *second-order cone program* (SOCP) with optimization variables $\tilde{\theta}$ and γ:

$$\begin{array}{ll}\text{maximize} & \gamma\\ \text{subject to} & \bar{\mathcal{F}}(\theta+\tilde{\theta}) \geqslant \gamma + \Phi^{-1}(\eta)V^{1/2}, \quad \tilde{\theta} \in \Theta,\end{array} \tag{C17}$$

where

$$\bar{\mathcal{F}}(\theta+\tilde{\theta}) = f + \tilde{\theta}^\mathsf{T} g_\theta - \tfrac{1}{2}\tilde{\theta}^\mathsf{T} R_{\theta\theta}\tilde{\theta} - \tfrac{1}{2}\mathrm{Tr}(C R_{ww}), \tag{C18}$$

$$V = (R_{w\theta}\tilde{\theta} - g_w)^\mathsf{T} C (R_{w\theta}\tilde{\theta} - g_w), \tag{C19}$$

and $\Phi(\eta)$ is the cumulative distribution function for the normal Gaussian density. The SOCP of Eq. (C17) can be used in the

optimization step in Algorithm 1. The use of the weak noise approximation makes this approach very numerically efficient. Indeed, calculations at each new control θ require only the knowledge of the noise covariance matrix, thus eliminating the need for random sampling from the noise distribution. A full exploration of this approach is forthcoming.

[1] A. Weinmann, *Uncertain Models and Robust Control* (Springer, Wien, 1991).
[2] K. Zhou and J. C. Doyle, *Essentials of Robust Control* (Prentice Hall, Englewood Cliffs, NJ, 1998).
[3] G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach* (Springer, New York, 2000).
[4] F. Lin, *Robust Control Design: An Optimal Control Approach*, rev. ed. (Wiley, Chichester, UK, 2007).
[5] A. Belmiloudi, *Stabilization, Optimal and Robust Control: Theory and Applications in Biological and Physical Sciences* (Springer, London, 2008).
[6] G. Taguchi, R. Jugulum, and S. Taguchi, *Computer-Based Robust Engineering: Essential for DFSS* (ASQ Quality Press, Milwaukee, WI, 2004).
[7] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization* (Princeton University Press, Princeton, NJ, 2009).
[8] S. Boyd and L. Vandenberghe, *Convex Optimization* (Cambridge University Press, Cambridge, UK, 2004).
[9] E. M. Fortunato, M. A. Pravia, N. Boulant, G. Teklemariam, T. F. Havel, and D. G. Cory, J. Chem. Phys. **116**, 7599 (2002).
[10] M. A. Pravia, N. Boulant, J. Emerson, A. Farid, E. M. Fortunato, T. F. Havel, R. Martinez, and D. G. Cory, J. Chem. Phys. **119**, 9993 (2003).
[11] N. Boulant, J. Emerson, T. F. Havel, D. G. Cory, and S. Furuta, J. Chem. Phys. **121**, 2955 (2004).
[12] M. Henry, A. Gorshkov, Y. Weinstein, P. Cappellaro, J. Emerson, N. Boulant, J. Hodges, C. Ramanathan, T. Havel, R. Martinez, and D. Cory, Quantum Inf. Process. **6**, 431 (2007).
[13] T. W. Borneman, M. D. Hürlimann, and D. G. Cory, J. Magn. Reson. **207**, 220 (2010).
[14] J. Wesenberg and K. Mølmer, Phys. Rev. A **68**, 012320 (2003).
[15] J. H. Wesenberg, Phys. Rev. A **69**, 042323 (2004).
[16] I. Roos and K. Mølmer, Phys. Rev. A **69**, 022321 (2004).
[17] T. E. Skinner, T. O. Reiss, B. Luy, N. Khaneja, and S. J. Glaser, J. Magn. Reson. **163**, 8 (2003).
[18] K. Kobzar, T. E. Skinner, N. Khaneja, S. J. Glaser, and B. Luy, J. Magn. Reson. **170**, 236 (2004).
[19] K. Kobzar, B. Luy, N. Khaneja, and S. J. Glaser, J. Magn. Reson. **173**, 229 (2005).
[20] B. Luy, K. Kobzar, T. E. Skinner, N. Khaneja, and S. J. Glaser, J. Magn. Reson. **176**, 179 (2005).
[21] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, J. Magn. Reson. **172**, 296 (2005).
[22] T. E. Skinner, K. Kobzar, B. Luy, M. R. Bendall, W. Bermel, N. Khaneja, and S. J. Glaser, J. Magn. Reson. **179**, 241 (2006).
[23] N. Timoney, V. Elman, S. Glaser, C. Weiss, M. Johanning, W. Neuhauser, and C. Wunderlich, Phys. Rev. A **77**, 052334 (2008).
[24] N. Khaneja, C. Kehlet, S. J. Glaser, and N. C. Nielsen, J. Chem. Phys. **124**, 114503 (2006).
[25] B. Pryor and N. Khaneja, J. Chem. Phys. **125**, 194111 (2006).
[26] J.-S. Li and N. Khaneja, Phys. Rev. A **73**, 030302 (2006).
[27] P. Owrutsky and N. Khaneja, Phys. Rev. A **86**, 022315 (2012).
[28] J. Ruths and J.-S. Li, J. Chem. Phys. **134**, 044128 (2011).
[29] J. Ruths and J.-S. Li, arXiv:1102.3713 [math.OC].
[30] M. Steffen and R. H. Koch, Phys. Rev. A **75**, 062326 (2007).
[31] M. J. Testolin, C. D. Hill, C. J. Wellard, and L. C. L. Hollenberg, Phys. Rev. A **76**, 012302 (2007).
[32] B. Mischuck, I. H. Deutsch, and P. S. Jessen, Phys. Rev. A **81**, 023403 (2010).
[33] B. E. Mischuck, S. T. Merkel, and I. H. Deutsch, Phys. Rev. A **85**, 022302 (2012).
[34] B. Khani, S. T. Merkel, F. Motzoi, J. M. Gambetta, and F. K. Wilhelm, Phys. Rev. A **85**, 022306 (2012).
[35] M. D. Grace, J. M. Dominy, W. M. Witzel, and M. S. Carroll, Phys. Rev. A **85**, 052313 (2012).
[36] X. Wang, L. S. Bishop, J. P. Kestner, E. Barnes, K. Sun, and S. D. Sarma, Nat. Commun. **3**, 997 (2012).
[37] C. Stihl, B. Fauseweh, S. Pasini, and G. S. Uhrig, arXiv:1210.4311 [quant-ph].
[38] T. J. Green, J. Sastrawan, H. Uys, and M. J. Biercuk, New J. Phys. **15**, 095004 (2013).
[39] G. Quiroz and D. A. Lidar, Phys. Rev. A **88**, 052306 (2013).
[40] C. Brif, M. D. Grace, K. C. Young, D. L. Hocker, K. W. Moore, T.-S. Ho, and H. Rabitz, in QEC11: Second International Conference on Quantum Error Correction, Los Angeles, CA, 2011, available at: http://qserver.usc.edu/qec11/program.html (unpublished).
[41] A. Ben-Tal and A. Nemirovski, Math. Oper. Res. **23**, 769 (1998).
[42] A. Ben-Tal and A. Nemirovski, Math. Program. Ser. B **92**, 453 (2002).
[43] L. El Ghaoui, F. Oustry, and H. Lebret, SIAM J. Optim. **9**, 33 (1998).
[44] D. Bertsimas and M. Sim, Math. Program. Ser. B **107**, 5 (2006).
[45] G. C. Calafiore and M. C. Campi, IEEE Trans. Autom. Control **51**, 742 (2006).
[46] S. A. Vorobyov, A. B. Gershman, and Z.-Q. Luo, IEEE Trans. Signal Process. **51**, 313 (2003).
[47] R. G. Lorenz and S. P. Boyd, IEEE Trans. Signal Process. **53**, 1684 (2005).
[48] B. Rustem and M. Howe, *Algorithms for Worst-Case Design and Applications to Risk Management* (Princeton University Press, Princeton, NJ, 2002).
[49] L. El Ghaoui, M. Oks, and F. Oustry, Oper. Res. **51**, 543 (2003).
[50] G. R. Lanckriet, L. El Ghaoui, C. Bhattacharyya, and M. I. Jordan, J. Mach. Learn. Res. **3**, 555 (2003).
[51] Y. Zhang, J. Optim. Theory Appl. **132**, 111 (2007).
[52] A. Mutapcic and S. Boyd, Optim. Methods Softw. **24**, 381 (2009).
[53] D. Bertsimas, O. Nohadani, and K. M. Teo, Oper. Res. **58**, 161 (2010).
[54] D. Bertsimas, O. Nohadani, and K. M. Teo, INFORMS J. Comput. **22**, 44 (2010).
[55] A. Mutapcic, S. Boyd, A. Farjadpour, S. G. Johnson, and Y. Avniel, Eng. Optim. **41**, 365 (2009).
[56] A. Oskooi, A. Mutapcic, S. Noda, J. D. Joannopoulos, S. P. Boyd, and S. G. Johnson, Opt. Express. **20**, 21558 (2012).
[57] J. Zhang and R. L. Kosut, IEEE Trans. Control Syst. Technol. **21**, 869 (2013).

[58] H. A. Rabitz, M. M. Hsieh, and C. M. Rosenthal, Science **303**, 1998 (2004).

[59] R. Chakrabarti and H. Rabitz, Int. Rev. Phys. Chem. **26**, 671 (2007).

[60] C. Brif, R. Chakrabarti, and H. Rabitz, New J. Phys. **12**, 075008 (2010).

[61] C. Brif, R. Chakrabarti, and H. Rabitz, in *Advances in Chemical Physics*, edited by S. A. Rice and A. R. Dinner, Vol. 148 (Wiley, New York, 2012), pp. 1–76.

[62] H. Rabitz, M. Hsieh, and C. Rosenthal, Phys. Rev. A **72**, 052337 (2005).

[63] M. Hsieh and H. Rabitz, Phys. Rev. A **77**, 042306 (2008).

[64] T.-S. Ho, J. Dominy, and H. Rabitz, Phys. Rev. A **79**, 013422 (2009).

[65] A. N. Pechen and D. J. Tannor, Phys. Rev. Lett. **106**, 120402 (2011); H. Rabitz, T.-S. Ho, R. Long, R. Wu, and C. Brif, *ibid.* **108**, 198901 (2012); A. N. Pechen and D. J. Tannor, *ibid.* **108**, 198902 (2012); A. Pechen and N. Il'in, Phys. Rev. A **86**, 052117 (2012); P. de Fouquieres and S. G. Schirmer, arXiv:1004.3492 [quant-ph].

[66] V. F. Krotov, *Global Methods in Optimal Control Theory* (Marcel Dekker, New York, 1996); D. J. Tannor, V. Kazakov, and V. Orlov, in *Time Dependent Quantum Molecular Dynamics*, edited by J. Broeckhove and L. Lathouwers (Plenum, New York, 1992), pp. 347–360; S. G. Schirmer and P. de Fouquieres, New J. Phys. **13**, 073029 (2011); D. M. Reich, M. Ndong, and C. P. Koch, J. Chem. Phys. **136**, 104103 (2012); R. Eitan, M. Mundt, and D. J. Tannor, Phys. Rev. A **83**, 053426 (2011).

[67] W. S. Zhu and H. Rabitz, J. Chem. Phys. **109**, 385 (1998); G. Maday and G. Turinici, *ibid.* **118**, 8191 (2003); Y. Ohtsuki, G. Turinici, and H. Rabitz, *ibid.* **120**, 5509 (2004).

[68] M. Ndong, M. Lapert, C. P. Koch, and D. Sugny, Phys. Rev. A **87**, 043416 (2013).

[69] U. Hohenester, Phys. Rev. B **74**, 161307 (2006); M. Grace, C. Brif, H. Rabitz, I. A. Walmsley, R. L. Kosut, and D. A. Lidar, J. Phys. B: At. Mol. Opt. Phys. **40**, S103 (2007); M. D. Grace, C. Brif, H. Rabitz, D. A. Lidar, I. A. Walmsley, and R. L. Kosut, J. Mod. Opt. **54**, 2339 (2007); S. Montangero, T. Calarco, and R. Fazio, Phys. Rev. Lett. **99**, 170501 (2007); M. Wenin and W. Pötz, Phys. Rev. A **78**, 012358 (2008); Phys. Rev. B **78**, 165118 (2008); R. Roloff, M. Wenin, and W. Pötz, J. Comput. Theor. Nanosci. **6**, 1837 (2009); S. G. Schirmer and P. J. Pemberton-Ross, Phys. Rev. A **80**, 030301 (2009); P. Rebentrost, I. Serban, T. Schulte-Herbrüggen, and F. K. Wilhelm, Phys. Rev. Lett. **102**, 090401 (2009); T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, and S. J. Glaser, J. Phys. B: At. Mol. Opt. Phys. **44**, 154013 (2011); F. F. Floether, P. de Fouquieres, and S. G. Schirmer, New J. Phys. **14**, 073023 (2012).

[70] A. Rothman, T.-S. Ho, and H. Rabitz, J. Chem. Phys. **123**, 134104 (2005); Phys. Rev. A **73**, 053401 (2006); J. Dominy and H. Rabitz, J. Phys. A: Math. Theor. **41**, 205305 (2008); K. W. Moore and H. Rabitz, Phys. Rev. A **84**, 012109 (2011).

[71] T. Caneva, T. Calarco, and S. Montangero, Phys. Rev. A **84**, 022326 (2011).

[72] P. de Fouquieres, S. Schirmer, S. Glaser, and I. Kuprov, J. Magn. Reson. **212**, 412 (2011); S. Machnes, U. Sander, S. J. Glaser, P. de Fouquières, A. Gruslys, S. Schirmer, and T. Schulte-Herbrüggen, Phys. Rev. A **84**, 022305 (2011).

[73] G. von Winckel and A. Borzì, Comput. Phys. Commun. **181**, 2158 (2010); G. von Winckel, A. Borzì, and S. Volkwein, SIAM J. Sci. Comput. **31**, 4176 (2010).

[74] P. de Fouquieres, Phys. Rev. Lett. **108**, 110504 (2012).

[75] I. Degani and A. Zanna, SIAM J. Sci. Comput. **34**, A1488 (2012).

[76] K. W. Moore, R. Chakrabarti, G. Riviello, and H. Rabitz, Phys. Rev. A **83**, 012326 (2011).

[77] K. W. Moore and H. Rabitz, J. Chem. Phys. **137**, 134113 (2012).

[78] K. Schittkowski and C. Zillober, in *Stochastic Programming*, Lecture Notes in Economics and Mathematical Systems, edited by K. Marti and P. Kall, Vol. 423 (Springer, Berlin, 1995), pp. 123–141.

[79] S. P. Boyd, EE364b: Convex Optimization II, Lecture Notes, 2011, available at: http://www.stanford.edu/~boyd (unpublished).

[80] K. W. Moore Tibbetts, C. Brif, M. D. Grace, A. Donovan, D. L. Hocker, T.-S. Ho, R.-B. Wu, and H. Rabitz, Phys. Rev. A **86**, 062309 (2012).

[81] J. Löfberg, *2004 IEEE International Symposium on Computer Aided Control Systems Design, Taipei, Taiwan* (IEEE, Piscataway, NJ, 2004), pp. 284–289, http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=1393890.

[82] CVX: Matlab Software for Disciplined Convex Programming, Version 2.0 (beta) (CVX Research Inc., 2012), http://cvxr.com/cvx.

[83] M. Grant and S. Boyd, in *Recent Advances in Learning and Control*, Lecture Notes in Control and Information Sciences, edited by V. Blondel, S. Boyd, and H. Kimura (Springer-Verlag, Berlin, 2008), pp. 95–110, http://stanford.edu/~boyd/graph_dcp.html.

[84] K.-C. Toh, M. J. Todd, and R. H. Tütüncü, Optim. Methods Softw. **11**, 545 (1999).

[85] J. F. Sturm, Optim. Methods Softw. **11**, 625 (1999).