Automated synthesis of dynamically corrected quantum gates

Kaveh Khodjasteh,[1] Hendrik Bluhm,[2] and Lorenza Viola[1,*]

[1]*Department of Physics and Astronomy, Dartmouth College, 6127 Wilder Laboratory, Hanover, New Hampshire 03755, USA*

[2]*2nd Institute of Physics C, RWTH Aachen University, 52074 Aachen, Germany, and Jülich-Aachen Research Alliance, Fundamentals of Future Information Technology*

Dynamically corrected gates are extended to non-Markovian open quantum systems where limitations on the available controls and/or the presence of control noise make existing analytical approaches unfeasible. A computational framework for the synthesis of dynamically corrected gates is formalized that allows sensitivity against non-Markovian decoherence and control errors to be perturbatively minimized via numerical search, resulting in robust gate implementations. Explicit sequences for achieving universal high-fidelity control in a singlet-triplet spin qubit subject to realistic system and control constraint are provided, which simultaneously cancel to the leading order the dephasing due to non-Markovian nuclear-bath dynamics and voltage noise affecting the control fields. Substantially improved gate fidelities are predicted for current laboratory devices.

I. INTRODUCTION

Achieving high-precision control over quantum dynamics in the presence of decoherence and operational errors is a fundamental goal across coherence-enabled quantum sciences and technologies. In particular, realizing a universal set of quantum gates with sufficiently low error rate is a prerequisite for fault-tolerant quantum computation [1]. Open-loop control based on time-dependent modulation of the system dynamics has been extensively explored as a physical-layer error-control strategy to meet this challenge. Two main approaches have been pursued to date: On the one hand, if the underlying open-system relaxation dynamics is fully known, then powerful variational techniques and/or numerical algorithms from optimal quantum control theory (OCT) may be invoked to optimize the target gate fidelity; see, e.g., [2–5] for representative contributions. On the other hand, dynamically corrected gates (DCGs) [6] have been introduced having maximum design simplicity and *portability* in mind: Close in spirit to well-established dynamical decoupling techniques for quantum state preservation in non-Markovian environments [7], DCG sequences can achieve a substantially smaller net decoherence error than individual "primitive" gates by making minimal reference to the details of the system and control specifications. In principle, the use of recursive control design makes it possible for the final accuracy to be solely limited by the shortest achievable control time scale [8]. Remarkably, simple DCG constructions underly the fidelity improvement reported for spin-motional entangling gates in recent trapped-ion experiments [9].

While obtaining a detailed quantitative characterization of the noise mechanisms to overcome is imperative to guarantee truly *optimal* control performance, this remains practically challenging for many open quantum systems of interest. In addition, current approaches for applying OCT methods to *non-Markovian* environments rely on obtaining suitable simplifications of the open-system equations of motion (e.g., via identification of a finite-dimensional Markovian embedding [2] or approximation through time-local coupled linear equations [5])—which may be technically challenging and/or involve nongeneric assumptions. Since in DCG schemes the error cancellation is engineered at the level of the *full* system-plus-environment Hamiltonian evolution, two significant advantages arise: (1) environment operators may be treated *symbolically*, avoiding the need for an explicit equation of motion for the reduced dynamics to be derived; and (2) in contrast to error-control approaches designed in terms of gate propagators (notably, fully compensating composite pulses for systematic control errors [10,11]), working at the Hamiltonian level allows one to more directly relate to physical error mechanisms and operational constraints. Despite incorporating realistic requirements of finite maximum control rates and amplitudes, analytic DCG constructions nonetheless rely on the assumption that complete control over the target system can be afforded through a universal set of *stretchable* control Hamiltonians [6]. This requirement is typically too strong for laboratory settings where only a *limited* set of control Hamiltonians can be turned on and off with sufficient precision and speed, and universality also relies on internal always-on Hamiltonians. Furthermore, portability comes at the expense of longer sequence durations, making DCGs more vulnerable to uncompensated Markovian decoherence mechanisms.

In this work, we introduce a control methodology that results in an *automated* recipe for synthesizing DCGs via numerical search. This is accomplished by relaxing the portability requirement and utilizing the full details of the control. While the resulting "automated DCGs" (aDCGs) are still synthesized without *quantitative* knowledge of the underlying error sources, they overcome the restrictive assumptions of analytical schemes and lead to drastically shorter sequences. As an additional key advantage, our Hamiltonian-engineering formulation lends itself naturally to incorporating robustness against multiple error sources, which can enter the controlled open-system Hamiltonian in either additive or multiplicative fashion. This allows for aDCGs to *simultaneously* cancel non-Markovian decoherence *and* control errors, as long as the combined effects remain perturbatively small.

We quantitatively demonstrate these advantages by focusing on a highly constrained control scenario—the two-

———————
*lorenza.viola@dartmouth.edu

electron singlet-triplet (S-T$_0$) spin qubit in GaAs quantum dots (QDs) [12]. In spite of groundbreaking experimental advances [13–16], boosting single-qubit gate fidelities is imperative for further progress towards scalable quantum computation and is attracting intense theoretical effort [17–19]. Recently introduced SUPCODE composite-pulse sequences [19], for instance, are (analytically) designed to achieve insensitivity against decoherence induced by coupling to the surrounding nuclear-spin bath; however, they do not incorporate robustness against voltage noise, which is an important limitation in experiments [16]. Here, we provide explicit aDCG sequences for high-fidelity universal control in S-T$_0$ qubits, which cancel the dominant decoherence and exchange-control errors, while respecting the stringent timing and pulsing constraints of realistic S-T$_0$ devices. The resulting sequences use a very *small* number of control variables and a *fixed* base pulse profile, which streamlines their experimental implementation. Up to two orders of magnitude improvement in gate fidelities are predicted for parameter regimes appropriate for current experimental conditions.

II. CONTROL-THEORETIC SETTING

We consider in general a *d*-dimensional target quantum system *S* coupled to an environment (bath) *B*, whose total Hamiltonian *H* on $\mathcal{H}_S \otimes \mathcal{H}_B$ reads

$$H = [H_S + H_{\text{ctrl},0}(t)] \otimes \mathbb{I}_B + H_e,$$
$$H_e \equiv H_{e,\text{int}} + H_{e,\text{ctrl}}(t), \quad (1)$$

where $\mathbb{I}_{B(S)}$ denotes the identity operator on $B(S)$, H_S accounts for the internal ("drift") system's evolution in the absence of control, and the time-dependent $H_{\text{ctrl},0}(t)$ represents the intended control Hamiltonian on *S*. The total "error Hamiltonian" H_e encompasses the bath Hamiltonian, unwanted interactions with the bath, as well as deviations of the applied control Hamiltonian from $H_{\text{ctrl},0}(t)$, subject to the requirement that the underlying correlation times are sufficiently long. Formally, we require that $\|H_e\| \ll \|H(t)\|$, where $\|X\|$ is the operator norm of $X = X^\dagger$ (maximum absolute eigenvalue of *X*) [6,7]. In order to "mark" the error sources, we characterize the strength of each independent contribution to $H_{e,\text{int}}$ in terms of dimensionless parameters $\{\delta_\alpha\}$, in such a way that, without loss of generality, we may express

$$H_{e,\text{int}} = \mathbb{I}_S \otimes H_B + H_{SB} \equiv \sum_\alpha \delta_\alpha S_\alpha \otimes B_\alpha,$$

with S_α being a Hermitian operator basis on $\mathcal{H}_S$ and B_α acting on $\mathcal{H}_B$, respectively, and the bath internal Hamiltonian $H_B \equiv \delta_0^e B_0$. We assume that the B_α are norm bounded but otherwise quantitatively unspecified. In particular, if B_α are treated as scalars ($B_\alpha = \ell_\alpha \mathbb{I}_B$), then we may formally recover the limit of a classical bath, whereby $B_0 = 0$ and the system Hamiltonian is effectively modified in a random (yet slowly time-dependent) fashion. Note that, as long as we are interested in canceling effects that are first order in the error sources, there is no distinction between the B_α being actual operators or scalars. Similarly, we characterize the independent error

sources in $H_{\text{ctrl},0}(t)$ by letting

$$H_{e,\text{ctrl}}(t) = \sum_\beta \delta_\beta H_\beta(t) \otimes \mathbb{I}_B,$$

where $H_\beta(t)$ are known system operators, while the parameters δ_β remain unspecified. For notational convenience, we shall label all of the unknown parameters symbolically and collectively by $\delta \equiv \{\delta_\alpha, \delta_\beta\}$.

In an ideal error-free scenario, $\delta = 0$, the system evolves directly under the action of the control, in the presence of its internal drift Hamiltonian. We assume that in this limit, *S* is *completely controllable*, that is, arbitrary unitary transformations on *S* can be synthesized as "primitive gates" by suitably designing $H_{\text{ctrl},0}(t)$ in conjunction with H_S. As mentioned, we are particularly interested in the situation where the latter is essential for controllability to be achieved [20]. The available control resources may be specified by describing

$$H_{\text{ctrl},0}(t) = \sum_a c_a(t) H_a \otimes \mathbb{I}_B,$$

in terms of the admissible (nominal) control inputs and Hamiltonians. Beside restrictions on the set of tunable Hamiltonians H_a, limited "pulse-shaping" capabilities will typically constrain the control inputs $c_a(t)$ as system-dependent features of the control hardware. For concreteness, we assume here that $H_{\text{ctrl},0}(t)$ is decomposed as a sequence of shape-constrained pulses applied back to back and also constrains pulse amplitudes $\{h_i\}$ and durations $\{\tau_i\}$ to technological limitations such as $h_{\min} \leqslant h_i \leqslant h_{\max}, \tau_{\min} \leqslant \tau_i \leqslant \tau_{\max}$.

Ideally, if the target unitary gate is *Q*, then the objective for gate synthesis is to devise a control Hamiltonian $H_{\text{ctrl},0}(t)$ such that (up to a phase)

$$Q = \mathcal{T} \exp\left\{ -i \int_0^\tau [H_S + H_{\text{ctrl},0}(s)] ds \right\}, \quad (2)$$

where $\mathcal{T}$ denotes time ordering and τ is the running time of the control. The ideal evolution naturally defines a toggling-frame unitary propagator given by

$$U_Q(t) = \mathcal{T} \exp\left\{ -i \int_0^t [H_S + H_{\text{ctrl},0}(s)] ds \right\}, \quad (3)$$

which traces a path from $\mathbb{I}_S$ to *Q* over τ. If $H_e \neq 0$, then application of $H_{\text{ctrl},0}(t)$ over the same time interval results in a total propagator of the form

$$U_{Q[\tau]}^{(\delta)} \equiv U_{Q[\tau]}^{(0)} \exp\left(-i E_{Q[\tau]}^{(\delta)} \right),$$

where $U_{Q[\tau]}^{(0)} = Q$, and $E_{Q[\tau]}^{(\delta)}$ is an "error action" operator on $\mathcal{H}_S \otimes \mathcal{H}_B$ that isolates the effects of undesired terms in the evolution [6]:

$$\exp\left(-i E_{Q[\tau]}^{(\delta)} \right) = \mathcal{T} \exp\left[-i \int_0^\tau U_Q(s)^\dagger H_e U_Q(s) ds \right]. \quad (4)$$

The norm of the error action can be taken to quantify the error amplitude per gate (EPG) in the presence of δ. The EPG in turn upper bounds the fidelity loss between the ideal and actual evolution on *S* once its "pure-bath" components are removed. More concretely, define

$$\text{mod}_B E \equiv E - \mathbb{I}_S \otimes \text{Tr}_S(E)/d,$$

that is, a projector that removes the pure-bath terms in E (note that $\mathrm{mod}_B E = E$ if E is a pure-system operator of the form $A \otimes \mathbb{I}_B$, as for a classical bath). Then the following (not tight) upper bound for the (Uhlman) fidelity loss holds *independently* of the initial states [8,21]:

$$1 - f_U \leqslant \big\| \mathrm{mod}_B E_{Q[\tau]}^{(\delta)} \big\|.$$

Thus, reducing the EPG can be used as a proxy for reducing gate fidelity loss. While for a primitive gate implementation the EPG scales linearly with δ, the goal of DCG synthesis is to *perturbatively cancel* the dependence on δ in $E_{Q[\tau]}^{(\delta)}$, up to a desired order of accuracy, to realize the gate in a manner that is as error free as possible as long as δ is small. For simplicity and immediate application, we focus here on *first-order* aDCG constructions, for which $\| \mathrm{mod}_B E_{Q[\tau]}^{(\delta)} \| = O(\delta^2)$.

III. SYNTHESIZING DYNAMICALLY CORRECTED GATES

A. Existential approach

Recall that there are two main requirements in first-order analytical DCG constructions [6,8]: (i) primitive gate implementations of the generators of a "decoupling group" associated with the algebraic structure of EPGs and (ii) particular implementations of the target gate Q (as Q^*) and the identity gate (as I_Q) as sequences of primitive gates such that Q^* and I_Q share the *same* first-order EPG, making them a "balance pair." While (i) is provided by controllability and leads directly to a constructive procedure for correcting to the first order the identity evolution, (ii) is essential for modifying this procedure in such a way that the net first-order error cancellation is maintained, but Q is effected instead.

Generating balance pairs requires further adjustment of gate control parameters to form a *controllable relationship* between EPGs of gate implementations, holding as an identity regardless of the value of δ (or B_α). For example, in the absence of drift dynamics and control errors, such a controllable relationship can be engineered by "stretching" pulse profiles in time while the amplitudes are reduced proportionally, resulting in different realizations of the *same* target, with EPGs that scale linearly with the gate duration. Similarly, in the presence of a multiplicative control error, primitive gates with physically equivalent (modulo 2π) angles of rotation result in different EPGs (note that similar geometric ideas are used in designing composite pulses [11]). We argue next that knowing the control description and marking the error sources $\{\delta_j\}$ does still lead to (ii), as long as control constraints allow us to tap into a continuum of different gate implementations.

The multitude of pathways for realizing a primitive gate Q increases with gate duration and subsegments as a result of availability of more control choices and ultimately a simpler control landscape [22]. Assume (A1) that such primitive implementations may be parametrized as $Q[\tau]$. We aim to show that a balance pair or, alternatively, a direct cancellation of the EPG of Q may be found. The gist of our argument is most easily given for a single qubit, with the Pauli operators chosen as the operator basis $\{S_\alpha\}$ for error expansion. Using the fact that the interactions among different error sources can be ignored up to the first order, the basic idea is to start

with the first error source δ_1 and then use the resulting gates recursively for the next error source until all error sources are exhausted. Ignoring error sources other than δ_1, let us thus expand $E_{Q[\tau]}^{(\delta_1)} = \delta_1 \sum_\alpha e_{Q[\tau],\alpha} S_\alpha \otimes B_1$.

In addition, assume (A2) that as a function of the parameter τ, the range of the real-valued functions $e_{Q[\tau],\alpha}$ extends to infinity in positive or negative directions. Consider now "projection blocks" composed of two Pauli gates applied back to back, that is,

$$P_\alpha[\tau_\alpha] \equiv S_\alpha[\tau_\alpha] S_\alpha[\tau_\alpha],$$

with a corresponding EPG given by $2\delta_1 e_{S_\alpha[\tau_\alpha],\alpha} S_\alpha \otimes B_1$, which is *purely along* S_α. By virtue of assumption (A2), we can find a continuum of (τ, τ_α) pairs such that $2 e_{S_\alpha[\tau_\alpha],\alpha} = \pm e_{Q[\tau],\alpha}$ for all Pauli directions α, meaning that we may reproduce each error component in $E_{Q[\tau]}^{(\delta_1)}$ *up to a sign*. Those Pauli components α_- that reproduce error with a negative sign are combined in sequence with $Q[\tau]$ to form a longer gate,

$$Q^* = Q[\tau] \prod_{\alpha_-} P_{\alpha_-}[\tau_{\alpha_-}].$$

If all Pauli components can be matched with negative signs, then the resulting gate will cancel all error components and a DCG construction is provided by Q^*. Otherwise, the Pauli components α_+ that are matched with positive sign are combined to produce an identity gate,

$$I_Q = \prod_{\alpha_+} P_{\alpha_+}[\tau_{\alpha_+}],$$

which matches the error of Q^*. Hence, (Q^*, I_Q) form a balance pair and can be used to produce a continuum of constructions of a DCG $Q^{(1)}[\tau]$ that cancels the error source δ_1. Provided that the assumption (A2) remains valid for this new composite construction, we can repeat the procedure to remove the other error sources.

We remark that assumption (A2) essentially implies that the domain of the errors as a function of implementation parameters for a fixed unitary gate is not compact, so that arbitrary magnitudes of each error component can be sampled by appropriately choosing the implementation parameters. Such arbitrary large domains need not exist in the primitive gate implementations (naturally or due to constraints), or only discrete error values may be reachable. Nonetheless, we may still enlarge the accessible range of errors for the target gate Q by attaching a continuously parametrized family of identity gates. Universal controllability of the system implies that not only any gate U but also its inverse U^{-1} may be reached. Implementing U, followed by its inverse U^{-1}, produces an implementation of the identity I_U that is "parametrized" by the original gate U. In the absence of degeneracies (relationships between the errors that could be used separately to provide a balance pair), the EPG associated with E_U has then a continuous domain. Clearly, applying I_U followed by the target gate Q still realizes the gate Q, but the resulting EPG is now given by $E_Q + E_{I_U}$, which is parametrized by U. By applying sufficiently many copies of I_U before applying Q, the error can be extended to arbitrary large domains as desired.

B. aDCGs: Computational approach

The aDCG sequences generated by following the above existential argument tend to be far too long and complex to be useful in realistic control scenarios. Also note that, in principle, the construction of single- or two-qubit DCGs in n-qubit registers may be handled similarly by using multiqubit Pauli operators as a basis for EPG expansion and for building projection blocks. However, the sequence complexity tends in this case to also grow exponentially with n [8], making the need for more efficient synthesis procedures even more essential. Just as complete controllability provides an existential foundation to numerical OCT approaches for unitary gate synthesis when $H_e = 0$ [18], our argument legitimates a numerical search for aDCGs in the presence of H_e. Since the objectives of gate realization and perturbative error cancellation are not inherently competing, the numerical search can be described as a *multiobjective minimization* problem, as we detail next.

Let the nominal control Hamiltonian $H_{\text{ctrl},0}(t)$ be parametrized in terms of control variables $\{x_i\}$ and define objective functions as follows:

$$F(\{x_i\}) = \text{dist}\big(U_{Q[\tau]}^{(\delta=0)}, Q\big), \qquad (5)$$

$$B_j G_j(\{x_i\}) = \big\| \partial \, \text{mod}_B E_{Q[\tau]}^{(\delta)} \big/ \partial \delta_j \big\|_{\delta=0}, \qquad (6)$$

where j labels independent error sources and B_j *symbolically* denotes bath operators that mark error sources in $H_{e,\text{int}}$ (recall that $B_j = \mathbb{I}_B$ for control error sources) to ensure that F and G_j only depend on the known quantities x_i. Minimizing only the first objective, $F = 0$, corresponds to achieving exact ideal primitive gate synthesis, given by Eq. (2). As an appropriate distance measure for unitary operators in Eq. (5), we use

$$\text{dist}(U,V) = [1 - |\text{Tr}(U^\dagger V)|/d]^{1/2}, \qquad (7)$$

which is a standard phase-invariant choice [18]. Minimizing the objectives in Eq. (6) corresponds to first-order sensitivity minimization. Thus, solving for $F = 0 = G_j, \forall j$, results in an implementation of Q that is *insensitive to the perturbative parameters* δ_j, yielding a *robust control solution* as long as δ_j is small.

Evaluating G_j apparently requires solving the full time-dependent system-plus-bath Schrödinger equation parametrized by the controllable pulse shapes. In fact, once the error sources (including bath operators) are treated as first-order *symbolic variables*, G_j can be evaluated by effectively solving the Schrödinger equation on the system only, in order to determine the appropriate toggling-frame propagator, given by Eq. (3), and then evaluating the required error action, given by Eq. (4), by invoking a Magnus expansion [6]. Specifically, if the control variables $x_i \equiv \{(\tau_i, h_i)\}$, then the sequence propagator reads

$$U_{Q[\tau]}^{(\delta)} \equiv U_{x_n}^{(\delta)}(\tau_n) \cdots U_{x_1}^{(\delta)}(\tau_1), $$

where $\tau = \sum_i \tau_i$ and $U_{x_i}^{(\delta)}(s)$ is the ith pulse propagator corresponding to the variable x_i at time s, with its associated first-order error action,

$$E_{U_{x_i}}^{(\delta)} = \int_0^{\tau_i} U_{x_i}^{(0)}(s)^\dagger H_e U_{x_i}^{(0)}(s) ds. $$

To the first order in δ, the total EPG is in turn given by

$$E_{Q[\tau]}^{(\delta)} = \sum_i V_i^\dagger E_{U_{x_i}}^{(\delta)} V_i, $$

where $V_i \equiv U_{x_{i-1}}^{(0)} \cdots U_{x_1}^{(0)}$ denote the "partial" product of gate propagators up to and excluding the ith gate [6]. While, as noted, for a first-order aDCG the resulting accuracy $\|\text{mod}_B E_{Q[\tau]}^{(\delta)}\| = O(\delta^2)$, a more quantitative estimate of the actual conditions of applicability requires estimating the dominant uncorrected second-order errors. Technically, this can be carried out by means of standard algebraic techniques; however, it is not straightforward [23] and is beyond our present scope. Instead, we focus in what follows on addressing the construction and performance of first-order aDCGs in concrete illustrative settings.

IV. APPLICATION TO SINGLET-TRIPLET QUBITS

Consider first the following single-qubit specialization of Eq. (1):

$$H = \tfrac{1}{2}[B\sigma_x + J_0(t)\sigma_z] \otimes \mathbb{I}_B + H_e, \qquad (8)$$
$$H_e = \mathbb{I}_S \otimes \delta_0 B_0 + \sigma_x \otimes \delta_1 B_x + J_0(t)\delta_2\sigma_z,$$

where the operator-valued $\delta_1 B_x$ and the system drift B couple to the system along σ_x, and the nominal control $J_0(t)$ and a multiplicative error δ_2 couple along σ_z. Although explicitly included, the bath internal Hamiltonian does not play a role in the first-order removal of decoherence and is automatically accounted for in mod_B. On the other hand, the drift term B is essential for complete controllability and analytical DCG constructions are not viable even in the limit $\delta_2 \to 0$. Thus, the need to effectively address *both* noise sources δ_1, δ_2 for a *generic* operating point B mandates the use of aDCGs.

While useful as a template for single-axis control scenarios in the presence of internal drift and dephasing, a semiclassical version of the above model is relevant, in particular, to describe a universally controllable S-T$_0$ qubit. In this case, the logical qubit subspace is spanned by $\{|S\rangle, |T_0\rangle\}$, i.e., the singlet and triplet states of two electrons on separate QDs, [14,15,19,24] and, provided that the number of bath nuclear spins is sufficiently large [24], the following simpler Hamiltonian is appropriate and widely used for this system [18,19]:

$$H = \tfrac{1}{2}[B + \delta B(t)]\sigma_x + \tfrac{1}{2}J(t)\sigma_z. \qquad (9)$$

Physically, the drift term B is a known static magnetic field gradient between the two QDs that includes an Overhauser field from the nuclear spin bath, $\delta B(t)$ (corresponding to $2\delta_1 B_x$) accounts for random fluctuations of B due to coupling to nuclear flip-flop processes [25], and $J(t)$ is the exchange splitting. In practice, $J(t)$ is tuned by control of an electrostatic gate voltage [14], and voltage fluctuations due to charge noise result in a noisy control Hamiltonian of the form $J(t) = J_0(t)[1 + \delta J(t)]$, where $\delta J(t)$ thus corresponds to $2\delta_2$. We assume that both noise sources may be treated as Gaussian quasistatic processes, with their "run-to-run" distribution being characterized by standard deviations $\sigma_{\delta B}$ and $\sigma_{\delta J}$. While in practice the noise is not completely static, we expect our considerations to remain valid as long as high-frequency noise components decay sufficiently fast and the resulting aDCGs

are short relative to time scales over which white charge noise may become important. Phenomenologically, the dephasing induced by the fluctuating Overhauser field is consistent with a power-law noise spectrum of the form $S(\omega) \sim \omega^{-2}$ over a wide spectral range [26]. Likewise, recent experiments indicate that voltage noise also arises overwhelmingly due to low-frequency components with an approximate $1/f$ decay at low operating temperatures [16]. From experimentally measured values of $T_2^\star$, we use here $\sigma_{\delta B/(2\pi)} \lesssim 0.15$ MHz [13] and $\sigma_{\delta J} \lesssim 1/50$ [27].

In constructing aDCGs, we shall choose values of the internal drift (B) and the nominal control field $J_0(t)$ that are appropriate for the QD setting of Eq. (9). We stress, however, that the *same* solution is found from (and applies to) the fully quantum model Hamiltonian of Eq. (8). In practice, the drift term B can be set to a fixed value, $B/(2\pi) \in [0.03, 0.2]$ GHz, which we choose at 0.1 GHz [28]. The control field $J_0(t)/(2\pi)$ is taken to be positive and smaller than $J_{max}/(2\pi) = 0.3$ GHz. We recognize the finite rise, delay, and drop times associated with pulse generators by *fixing* a pulse profile. Thus, during each pulse, with time t' measured from the pulse start, the control function $J_0(t')$ is given by $h_i c(t'/\tau_i)$, where $c(x)$ is the pulse-shape function. We digitize the pulse-shape function for numerical evaluation. In contrast to merely bounding the pulse times and allowing pulse durations as extra control variables, we enforce the pulse durations to be *fixed* at $\tau_i \equiv \tau = 3$ ns, which is compatible with the currently most widespread pulse generators' temporal resolution of 0.83 ns. The search space of the pulse amplitude control variables is thus given by $x_i = h_i$. While removing τ_i from the control variables results in more severe constraints, it also corresponds to a reduction of the search space. We verified that all of our results were reproduced with variable but lower-bounded pulse widths as well. The objective functions F, G_1, and G_2 are computed explicitly in terms of each constituting pulse parameter h_i according to the general procedure described in Sec. III B.

For the resulting multiobjective minimization, we introduce numerical weight factors λ_1 and λ_2 and form a single objective function

$$O(\{h_i\}) \equiv F + \lambda_1 G_1 + \lambda_2 G_2.$$

Choosing small values of λ_i ($= 10^{-5}$) work best in directing the search from solutions that first synthesized the target gate only ($F = 0$), to an error-corrected solution minimizing all objectives, $F = 0, G_1 = 0 = G_2$. Motivated by our existential argument, the intuition is to avoid the local minima associated with multiple objectives and focus on a single objective which, once realized, will give weight to the other objectives iteratively. We solve each aDCG search problem using an off-the-shelf (Matlab's FMINCON function) search routine (within minutes), with the default choice for solving constrained nonlinear optimization without specifying a precalculated gradient or Hessian. We start the search with a small number of pulses, n, which is then incremented until the minimal value of the objective function comes close to the machine precision ($\approx 10^{-16}$). Figure 1 (top) depicts the synthesized control profiles for a universal set of single-qubit aDCGs.

Once aDCG sequences are found, evaluating their effectiveness for the S-T_0 qubit can take direct advantage of the effectively closed-system nature of the model Hamiltonian
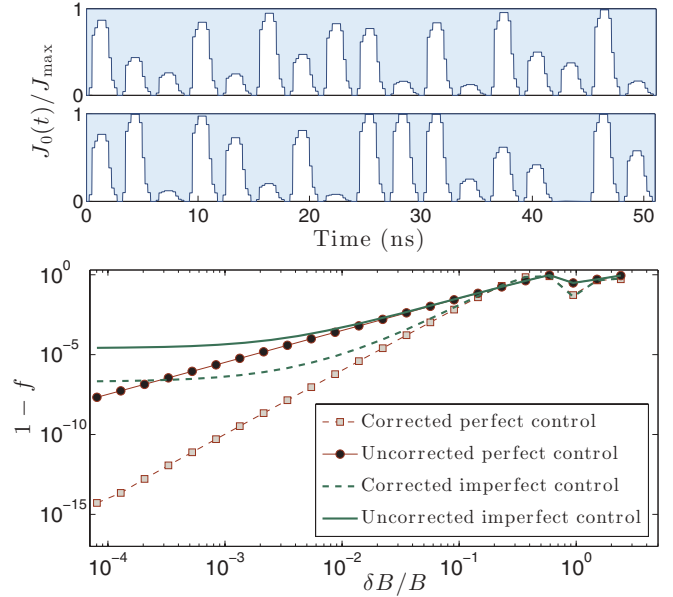

FIG. 1. (Color online) Upper panel: aDCG control profiles for $\exp(-i\pi\sigma_z/8)$ (top) and the Hadamard gate (bottom). Lower panel: Fidelity loss for uncorrected vs corrected gates, evaluated as $1 - f = \text{dist}(U_Q^{(\delta B, \delta J)}, Q)^2$ [Eq. (7)], for $Q = \exp(-i\pi\sigma_z/8)$ as a function of relative magnetic-field-gradient error $\delta B/B$ [$B/(2\pi) = 0.1$ GHz]. $\delta J = 0$ corresponds to perfect exchange control, while $\delta J = 0.01$ is close to current experimental values. The fact that $\delta J = 0.01$ is *fixed* is responsible for the eventual performance plateau, where δJ noise dominates over δB effects. Nevertheless, the aDCG advantage *is always maintained*.

in Eq. (9), thus avoiding the need for explicit spin-bath simulations and quantum process tomography. Figure 1 (bottom) depicts the fidelity loss for an uncorrected ($n = 3$ pulses, obtained through the same numerical procedure with $\lambda_1 = \lambda_2 = 0$) vs corrected implementation ($n = 17$ pulses). The higher slope of the fidelity loss as a function of δB when
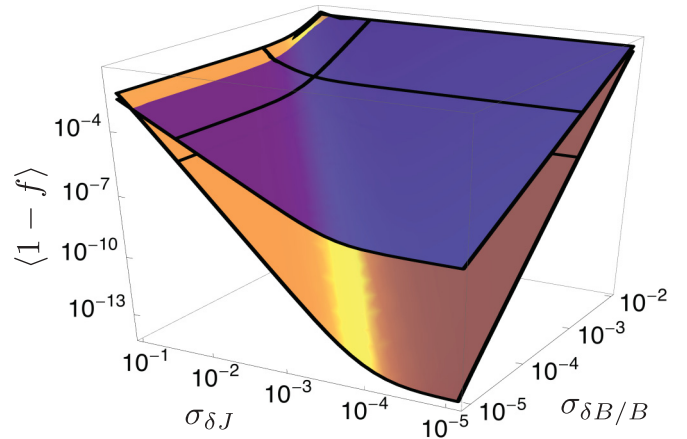

FIG. 2. (Color online) Ensemble-averaged fidelity loss as a function of relative magnetic-field-gradient and exchange-control noise for uncorrected (top, dark) vs corrected (bottom, light) implementation of $Q = \exp(-i\pi\sigma_z/8)$. The solid (black) lines on each surface correspond to typical values $\sigma_{\delta J} \approx 0.02$ and $\sigma_{\delta B/B} \approx 10^{-3}$ for current S-T_0 devices.

$\delta J = 0$ is the signature of a perturbative error cancellation and the aDCG advantage is maintained even with $\delta J > 0$, implying robustness with respect to both error types.

In order to make contact with experimentally relevant ensemble-averaged fidelities, we further evaluate the average of single-run fidelities $f(\delta B, \delta J)$ with respect to noise realizations, by assuming that δB and δJ are independent and normally distributed random variables with variance $\sigma_{\delta B}^2$ and $\sigma_{\delta J}^2$. The results are summarized in Fig. 2. Both noise sources adversely impact the expected gate fidelity, but aDCGs are far less affected, resulting in robust gates roughly as long as $\sigma_{\delta B/B} + \sigma_{\delta J} \lesssim 0.1$.

V. CONCLUSION

Our procedure can be interpreted as an *automated gate compiler*, which incorporates detailed information about the controllable parameters and their range of operations, along with qualitative information about the error sources affecting the evolution. Compared to mere (primitive) gate synthesis, the resulting increase in complexity scales proportionally to the number of error sources. Our approach applies to any Hamiltonian control setting, and for weak-enough error sources, even higher-order cancellation can be achieved in principle.

Thanks to the slow dynamics of the nuclear spin bath and fast control pulses available, electron spin qubits provide an ideal experimental testbed for validating our approach. While additional experimental details may be captured in more sophisticated ways, we believe that our framework is general and flexible enough for its effectiveness not to be compromised. In particular, further analysis is needed to quantify the effect of white electrical noise on aDCG sequences, as well as to possibly minimize its influence by penalizing large values of the exchange splitting in the numerical search. It is thus our hope that significantly improved single-gate fidelities will be achievable in S-T_0 qubits by aDCG sequences that operate under realistic noise levels and control limitations.

[1] E. Knill, Nature (London) **434**, 39 (2005).

[2] N. Khaneja, T. Reiss, C. Kehlet, T. Schulte-Herbrüggen, and S. J. Glaser, J. Magn. Reson. **172**, 296 (2005); T. Schulte-Herbrüggen, A. Spörl, N. Khaneja, and S. J. Glaser, J. Phys. B **44**, 154013 (2011).

[3] M. Möttönen, R. deSousa, J. Zhang, and K. B. Whaley, Phys. Rev. A **73**, 022332 (2006).

[4] J. Clausen, G. Bensky, and G. Kurizki, Phys. Rev. Lett. **104**, 040401 (2010).

[5] B. Hwang and H.-S. Goan, Phys. Rev. A **85**, 032321 (2012).

[6] K. Khodjasteh and L. Viola, Phys. Rev. Lett. **102**, 080501 (2009); Phys. Rev. A **80**, 032314 (2009).

[7] L. Viola, E. Knill, and S. Lloyd, Phys. Rev. Lett. **82**, 2417 (1999); L. Viola and E. Knill, *ibid.* **90**, 037901 (2003).

[8] K. Khodjasteh, D. A. Lidar, and L. Viola, Phys. Rev. Lett. **104**, 090501 (2010).

[9] D. Hayes, S. M. Clark, S. Debnath, D. Hucul, I. V. Inlek, K. W. Lee, Q. Quraishi, and C. Monroe, Phys. Rev. Lett. **109**, 020503 (2012); see also D. Hayes, K. Khodjasteh, L. Viola, and M. J. Biercuk, Phys. Rev. A **84**, 062323 (2011).

[10] M. Levitt, Prog. Nucl. Magn. Reson. Spectrosc. **18**, 61 (1986).

[11] J. T. Merrill and K. R. Brown, arXiv:1203.6392.

[12] J. Levy, Phys. Rev. Lett. **89**, 147902 (2002).

[13] J. R. Petta, A. C. Johnson, J. M. Taylor, E. A. Laird, A. Yacoby, M. D. Lukin, C. M. Marcus, M. P. Hanson, and A. C. Gossard, Science **309**, 2180 (2005); H. Bluhm, S. Foletti, D. Mahalu, V. Umansky, and A. Yacoby, Phys. Rev. Lett. **105**, 216803 (2010).

[14] S. Foletti, H. Bluhm, D. Mahalu, V. Umansky, and A. Yacoby, Nature Phys. **5**, 903 (2009).

[15] M. D. Shulman, O. E. Dial, S. P. Harvey, H. Bluhm, V. Umansky, and A. Yacoby, Science **336**, 202 (2012).

[16] O. E. Dial, M. D. Shulman, S. P. Harvey, H. Bluhm, V. Umansky, and A. Yacoby, arXiv:1208.2023.

[17] Preliminary results were reported by L. Viola, "Towards optimal constructions of dynamically corrected gates," Invited talk at QEC 2011, qserver.usc.edu/qec11/slides/Viola_QEC11.pdf

[18] M. D. Grace, J. M. Dominy, W. M. Witzel, and M. S. Carroll, Phys. Rev. A **85**, 052313 (2012).

[19] X. Wang, L. S. Bishop, J. P. Kestner, E. Barnes, K. Sun, and S. Das Sarma, Nature Commun. **3**, 997 (2012).

[20] If universal control and sufficiently fast dynamical decoupling pulses are available, alternative strategies are possible for protecting quantum gates, most simply by embedding the desired gate into the initial and/or final free evolution of a decoupling cycle, following the ideas of L. Viola, S. Lloyd, and E. Knill, Phys. Rev. Lett. **83**, 4888 (1999). A recent experimental implementation for a single solid-state qubit that also incorporates robustness against amplitude control errors was reported by A. M. Souza, G. A. Álvarez, and D. Suter, arXiv:1206.2933.

[21] D. A. Lidar, P. Zanardi, and K. Khodjasteh, Phys. Rev. A **78**, 012308 (2008).

[22] K. W. Moore, C. Brif, M. D. Grace, A. Donovan, D. L. Hocker, T.-S. Ho, R. Wu, and H. Rabitz, arXiv:1112.0333.

[23] Higher-order terms need more work as the (commutatively built) products of $\delta_j B_j$ do not form a basis for expansion of the error action operator used to define G_j. A so-called Hall basis can be used for expanding the latter in terms of algebraically

free operator elements of a given perturbation order. See, e.g., C. Reutenauer, *Free Lie Algebras* (Oxford University Press, New York, 1993).

[24] J. M. Taylor, J. R. Petta, A. C. Johnson, A. Yacoby, C. M. Marcus, and M. D. Lukin, Phys. Rev. B **76**, 035315 (2007).

[25] Ł. Cywiński, W. M. Witzel, and S. Das Sarma, Phys. Rev. Lett. **102**, 057601 (2009).

[26] M. J. Biercuk and H. Bluhm, Phys. Rev. B **83**, 235316 (2011).

[27] This corresponds to $T_2^* = 150$ ns for δB noise and about 11 coherent oscillations within T_2^* for δJ.

[28] Changing the choice of the drift may change the solutions considerably; for instance, having $B = 0$ removes controllability altogether. Note also that instead of fixing B, we could in principle consider optimizing its value to be fixed across a set of desired gates.