SAND2013-7871C

Adaptive quantum control via direct fidelity estimation and indirect model-based parametric process tomography

Robert L. Kosut* Hersch Rabitz† Matthew D. Grace‡

Abstract

The single and two-qubit logic gates which are universal for building a quantum computer are not, as yet, produced "naturally" – error correction and fault tolerant constructions are required, and making these requires control. To meet the requisite stringent performance goals places resource demands both spatially (ancilla qubits for error correction) and temporally (complex well timed control signals). On-line adaptive tuning of initially good controls offers a possible means to significantly reduce these overhead requirements. Two methods are proposed for control tuning: (i) direct estimation of fidelity between the actual system and the desired (unitary) logic gate, and (ii) estimating model parameters via compressive sensing. Both methods are evaluated numerically for a single qubit system with Hamiltonian parameter uncertainty.

1. Introduction

Exploiting quantum superposition and/or entanglement for information processing is a delicate process: quantum information is fragile. A ”bit” in our present computers utilizes an overwhelming number of electrons which makes for an inherent robustness. A quantum bit (qubit) may use just one electron which will clearly be extremely sensitive to many sources of error: thermal fluctuations, other qubits, unwanted states induced by material imperfections, etc. If the inherent errors can be driven below a threshold value, which depends on the physical implementation *and* the algorithm being coded, then the specific quantum computation can proceed unabated [1, 2]. To drive these inherent errors down requires control.

*SC Solutions, 1261 Oakmead Parkway, Sunnyvale, CA kosut@scsolutions.com

†Department of Chemistry, Princeton University, Princeton, NJ hrabitz@princeton.edu

‡Department of Scalable & Secure Systems Research, Sandia National Laboratories, Livermore, CA 94550 mgrace@sandia.gov

As outlined in [3], and numerous other surveys, control for quantum systems can be broadly categorized as *open-loop control*, *learning control* [4, 5], *measurement-based feedback control* [6], and *fully quantum control* sometimes referred to as *coherent-feedback quantum control* [7, 8].Coherent quantum feedback might be considered the Holy Graille of quantum control, *i.e.*, one quantum system is used to control another, a empquantum flyball governor. Currently the predominant form of quantum control for quantum computation is open-loop control, and as might be expected, a major problem is robustness. As such, there has been considerable recent effort to develop robust open-loop control procedures, all of which rely on numerical optimization. Even a partial listing would be too long for here so we refer the reader to the more complete list in [9].

If the parameters do not change significantly over long periods of time, then it would certainly be expected that online control tuning could improve performance. Even the promise of modest improvements may prove significant because of the ensuing overhead reduction in resources required for fault-tolerant computation. To accomplish this requires adjusting the control to increase fidelity. So the question becomes first of how to estimate fidelity from measurements, and considering the large number of quantum circuits, how to do that efficiently and with a minimum amount of data collection. In this paper two methods are proposed for adaptive control tuning: (i) a direct approximation of fidelity between the actual system and the desired (unitary) logic gate, and (ii) estimating mode parameters via a compressive sensing algorithm to obtain the process matrix. A numerical example for a single-qubit system with an uncertain Hamiltonian parameter is presented to illustrate the ideas.

2. Fidelity and robust control design

A quantum system under open-loop control over the time interval $0 \leq t \leq T$ is depicted schematically in Fig. 1. The part labeled *system* (S) is of dimension n

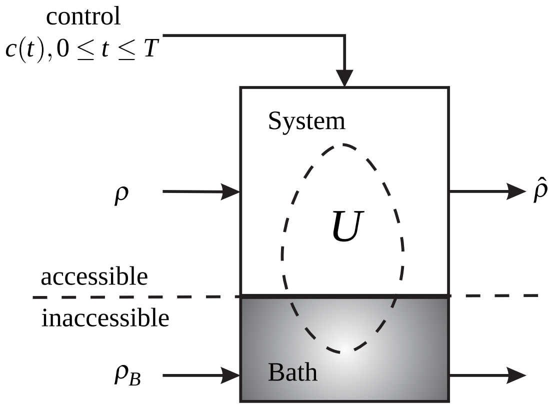

Figure 1. Quantum system under control.

and is accessible, whereas the *bath* (B) of dimension n_B is not. The total system-bath dynamics are described by an $nn_B \times nn_B$ unitary U which maps the initial (assumed decoupled) state (density matrix) $\rho \otimes \rho_B$ at time $t=0$ to the final state ρ at $t=T$, not necessarily decoupled, *i.e.*, not a tensor product of a system and bath state. The accessible system state $\hat{\rho}$ at $t=T$ is given by by the *partial trace* over the bath [1],

$$\hat{\rho} = \mathbf{Tr}_B \left(U(\rho \otimes \rho_B) U^\dagger \right) = \sum_{\mu=1}^{n_B} S_\mu \rho S_\mu^\dagger \qquad (1)$$

where $\{S_\mu \in \mathbf{C}^{n\times n}\}_{\mu=1}^{n_B}$ are the matrix elements of the *operator-sum-representation* (OSR) which under the stated conditions are trace preserving, that is, $\mathbf{Tr}\,\hat{\rho}=1$ if and only if $\sum_\mu S_\mu^\dagger S_\mu = I_S$.

The control goal is to make the map $\rho \rightarrow \hat{\rho}$ as close as possible to a desired unitary W acting solely on the system states, *i.e.*, $\hat{\rho} \approx W\rho W^\dagger$. A common measure of this performance goal is to select the control to maximize the *channel fidelity* [10, 11],

$$F = \frac{1}{n^2} \sum_{\mu=1}^{n_B} \left| \mathbf{Tr}\left(W^\dagger S_\mu \right) \right|^2 \qquad (2)$$

It is known that $F \in [0,1]$ with $F=1$ only when each OSR element is proportional to the desired unitary, specifically when $S_\mu = \alpha_\mu W$ with $\sum_\mu |\alpha_\mu|^2 = 1$ [12].

A robust quantum control design problem is typically formulated as maximizing the channel fidelity (**??**) for either a worst-case or average-case with respect to uncertainties in the system. In an abstract form, the problem we would like to solve is:

$$\begin{array}{ll} \text{maximize} & \left\{ \begin{array}{ll} \min_\delta F(c,\delta) & \text{(worst-case)} \\ \text{or} & \\ \text{avg}_\delta F(c,\delta) & \text{(average-case)} \end{array} \right. \\ \text{subject to} & c = \{c(t),\ o \le t \le T\} \in \mathscr{C},\ \delta \in \Delta \end{array} \qquad (3)$$

The optimization variable is the external control denoted by c representing the time varying control field $c(t), 0 \le t \le T$. The set $\mathscr{C}$ represents the control constraint set which is typically a convex set [9]. Similarly δ represents uncertainties which can be both deterministic (parameter uncertainty) as well as stochastic (noise from the control hardware and/or bath fluctuations). The set Δ represents the uncertainty set which may not be convex. Though (3) is not convex, as mentioned in the introduction, numerical techniques have been used to find very good local solutions. For example, in [9], we used *sequential convex programming* (SCP) to establish performance tradeoffs in level of uncertainty vs. control constraints for a single qubit system. We will use that example system here as well, tuning robust controls from [9].

3. Quantum process tomography via compressive sensing

We will assume that data is collected from a standard off-line procedure [1, Ch.8] using repeated identical experiments for each of ℓ system configurations of state and observable pairs $\{\rho_j, O_j \in \mathbf{C}^{n\times n},\ j=1,\ldots,\ell\}$. Each observable is represented by an Hermitian matrix $O_j \in \mathbf{C}^{n\times n}$ which has the spectral decomposition $O_j = \sum_{i=1}^m \lambda_{ij} O_{ij}$ with λ_{ij} real, and with $O_{ij} \geq 0$ and $\sum_i O_{ij} = I_S$. Using Born's rule and (1), the probability that the measurement outcome is λ_{ij} with configuration (ρ_j, O_j) is,

$$p_{ij} = \sum_\mu \mathbf{Tr}(O_{ij} S_\mu \rho_j S_\mu^\dagger),\ i=1,\ldots,m,\ j=1,\ldots,\ell \quad (4)$$

Let $\{\Gamma_\alpha \in \mathbf{C}^{n\times n},\ \alpha=1,\ldots,n^2\}$ be a basis for matrices in $\mathbf{C}^{n\times n}$. Expressing the OSR elements in this basis via $S_\mu = \sum_{\alpha=1}^{n^2} x_{\mu\alpha}\Gamma_\alpha$ gives the outcome probabilities and fidelity as,

$$p_{ij} = \mathbf{Tr}(M_{ij} X), \quad F = \mathbf{Tr}(GX) \qquad (5)$$

where $X, M_{ij}, G \in \mathbf{C}^{n^2\times n^2}$ with elements given as follows for $\alpha, \beta = 1, \ldots, n^2$,

$$\begin{array}{rcl} X_{\alpha\beta} & = & \sum_\mu x_{\mu\alpha} x_{\mu\beta}^* \\ (M_{ij})_{\beta\alpha} & = & \mathbf{Tr}(O_{ij}\Gamma_\alpha \rho_j \Gamma_\beta^\dagger) \\ G_{\beta\alpha} & = & \mathbf{Tr}(W^\dagger \Gamma_\alpha)\mathbf{Tr}(W\Gamma_\beta^\dagger)/n^2 \end{array} \qquad (6)$$

X is referred to as the *process matrix* and is in the convex set,

$$X \geq 0, \quad \sum_{\alpha,\beta=1}^{n^2} X_{\alpha\beta}\Gamma_\beta^\dagger \Gamma_\alpha = I_S \qquad (7)$$

The positivity constraint follows from construction of X, the linear constraint reflects the trace preserving constraint $\sum_\mu S_\mu^\dagger S_\mu = I_S$.

Clearly estimating the process matrix X from the probability outcomes would then provide an estimate of fidelity F. There are two issues with this approach. First, the probability outcomes p_{ij} are not known and can only be estimated from finite data, resulting in the empirical probabilities $\hat{p}_{ij} = N_{ij}/N_j$ where N_{ij} is the number of times outcome i occurred with configuration (ρ_j, O_j) from N_j trials; the total number of trials is $N = \sum_j N_j$. Thus the standard approach to estimating the process matrix, referred to as *Quantum Process Tomography* (QPT), is to estimate X from the empirical probability outcomes. This can be formulated as a least-squares objective together with the convex constraint that X satisfies (7), *e.g.*, [3]. Secondly, and actually of more significance, the number of real parameters in X is $n^4 - n^2$ where $n = 2^q$ for q system qubits. This gives rise to an exponential scaling of the number of parameters to estimate, *e.g.*, $q = [1,2,3,4] \;\Rightarrow\; n^4 - n^2 = [12, 240, 4032, 65280]$. As shown in [13] this $\mathscr{O}\left(n^4\right)$ scaling can be alleviated by selecting a basis corresponding to the ideal unitary W, and under the assumption that the system is designed (by control) to be close to the ideal, then methods of *compressive sensing* (CS) can be applied, because in this basis the process matrix will be almost sparse. The result is a scaling of $\mathscr{O}\left(k \log n\right)$ where k is the approximate sparsity level [14, 15]. To see this, and because this basis will be used subsequently for approximate fidelity estimation, observe that the vectorized version of W can be decomposed via a singular value decomposition as,

$$\vec{W} = \left[\vec{\Gamma}_1 \cdots \vec{\Gamma}_{n^2}\right] \begin{bmatrix} \sqrt{n} \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{8}$$

This basis is orthonormal ($\mathbf{Tr}(\Gamma_\alpha^\dagger \Gamma_\beta) = \delta_{\alpha\beta}$) with $\Gamma_1 = W/\sqrt{n}$ which gives the fidelity as,

$$F = X_{11}/n \tag{9}$$

In this basis if the channel were the ideal unitary, then $X_{11} = n$ (equivalently $F = 1$) with all the other elements of X equal to zero. If fidelity is close to one, then except for the large X_{11} element, the rest are small. Specifically, if $F = 1 - \varepsilon$, then all other elements of X are small, *i.e.*, $|X_{\alpha\beta \neq 11}/X_{11}| = \mathscr{O}\left(\varepsilon\right)$. Obtaining the process matrix via compressive sensing in this basis, referred to in [13] as *Compressed Quantum Process Tomography* (CQPT), is accomplished by solving the convex optimization,

$$\begin{array}{ll} \text{minimize} & \|\vec{X}\|_1 \\ \text{subject to} & \sum_{i,j} \left(\hat{p}_{ij} - \mathbf{Tr}(M_{ij} X)\right)^2 \leq \gamma \\ & X \text{ satisfies (7)} \end{array} \tag{10}$$

Here $\|\vec{X}\|_1$ is the usual heuristic for sparsity. Since fidelity is a single number, and in fact a simple scaling of the-11-element of the process matrix X, one would think that it ought to be considerably easier to estimate compared with having to estimate the entire process matrix even by CQPT (10) just to obtain the 11-element. Efforts in this direction are reported in [16, 17, 18] for randomized benchmarking and related statistical approaches. These approaches have a strong theoretical base in statistical analysis and are not strongly dependent (if at all) on the system order. Nonetheless, these approaches may still require a significant amount of data collection to provide statistically meaningful averages. If, however, the intended purpose of estimating fidelity is not just to estimate fidelity, but rather, to adjust the control to increase fidelity, then the simpler approach described in what follows may suffice. In particular, it is proposed here that it is more important that when the approximate fidelity increases so does the actual fidelity, rather than trying to make the error between the two small. Of course it would also be good if they are not too far apart.

4. Adaptive control via fidelity estimation

Ii is assumed that the initial control provides good fidelity over the anticipated range of model parameters. Moreover, if the parameter range were smaller, then it is also assumed known that another control would provide even better performance. Rather than try to estimate model parameters, a very difficult task even when if the Hamiltonian structure is known [3, 19], the control will be tuned directly from the estimated fidelity.

Towards this end, in the SVD-basis of $\vec{W}$ (8), where $F = X_{11}/n$ (9) and $\Gamma_1 = W/\sqrt{n}$, isolating the X_{11} term in (6) gives,

$$p_{ij} = F\,\mathbf{Tr}(O_{ij} W \rho_j W^\dagger) + \tilde{p}_{ij} \tag{11}$$

Here $\tilde{p}_{ij}$ contains all elements of the process matrix except X_{11}. Note that the coefficient of F above is *exactly* the probability outcome corresponding to the ideal unitary W. If the assumption holds that the initial fidelity is good, then as $F \to 1$ $\tilde{p}_{ij} \to 0$. Using the empirical estimates $\hat{p}_{ij}$, an approximate fidelity estimate can be obtained by solving for F from,

$$\begin{array}{ll} \text{minimize} & \sum_{ij} \left(\hat{p}_{ij} - F\,\mathbf{Tr}(O_{ij} W \rho_j W^\dagger)\right)^2 \\ \text{subject to} & 0 \leq F \leq 1 \end{array} \tag{12}$$

For a given control $c = \{c(t), 0 \le t \le T\}$ let $\hat{p}(c)_{ij}$ and $\hat{F}(c)$ denote, respectively, the empirical probabilities and estimated fidelity. A variety of gradient and/or randomized algorithms can now be used to adjust the control so that the estimated fidelity increases. An open question remains: under what conditions will increasing the estimated fidelity also increase the actual fidelity? In the next section we explore this with a numerical example.

5. Numerical example

To illustrate the procedure we concentrate on the single qubit system with Hamiltonian and control given by,

$$\begin{aligned} H(t) &= c(t)\omega_x X + \omega_z Z = \begin{bmatrix} \omega_z & c(t)\omega_x \\ c(t)\omega_x & -\omega_z \end{bmatrix} \\ c(t) &= \theta_k,\ (k-1)T/N \le t \le kT/N,\ k = 1,\ldots,N \end{aligned} \quad (13)$$

Here the control is in the Pauli-X term with a drift in Pauli-Z. In addition, the control is piece-wise-constant over uniform time intervals T/N with the design variable the control magnitude vector $\theta = [\theta_1 \cdots \theta_N]^T \in \mathbf{R}^N$. The robust control problem (3) now takes the form:

$$\begin{array}{ll} \text{maximize} & \min_{\omega_z,\omega_x} \{F = |\mathbf{Tr}(W^\dagger U)/2|^2\} \\ \text{subject to} & U = \prod_{k=1}^N \exp\{-iT/N(\theta_k \omega_x X + \omega_z Z)\} \\ & \omega_z \in [\omega_z^{\min}, \omega_z^{\max}],\ \omega_x \in [\omega_x^{\min}, \omega_x^{\max}], \\ & |\theta_k| \le \theta^{\max},\ k = 1,\ldots,N \end{array} \quad (14)$$

The optimization variables are $\theta_k, k = 1,\ldots,N$. Using the SCP routine in [9], we find a very good initial robust control for the following parameters:

$$T = 2, N = 4, \omega_z \in [2.5, 3.0], \omega_x = 1, \theta^{\max} = 4 \quad (15)$$

As seen in Fig. 2, the worst-case fidelity error in the assumed range for ω_z (shown in thick blue) is $\log_{10}(1 - F_{\text{initial}}) = -3.6$ or $F_{\text{initial}} = 0.99975$. However, suppose the true value has drifted out of the assumed range $[2.5, 3.0]$ and is actually $\omega_x^{\text{true}} = 2$. The actual fidelity is then $\log_{10}(1 - F_{\text{true}}) = -1.2223$ or $F_{\text{true}} = 0.9401$. For quantum computing this heralds impending failure, or would require serious overhead for fault tolerance mitigation.

To adapt the control to the actual system we use (12) with a minimal measurement scheme consisting of the single Pauli-Z observable $O = Z$ and a single pure input state $\psi = [0\ 1]^T$. The probability outcomes for a given control pulse sequence θ are averaged to produce the single number,

$$< Z(\theta) > = (U(\theta)\psi)^\dagger Z(U(\theta)\psi) \quad (16)$$

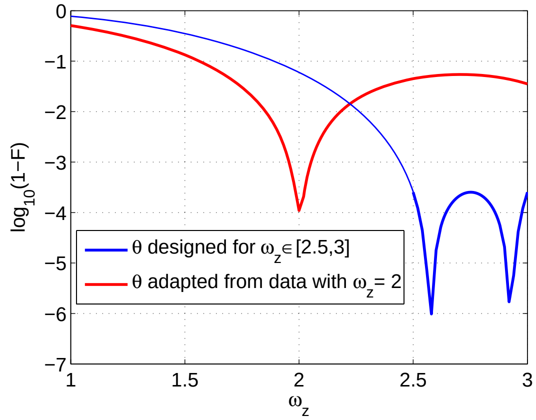

Figure 2. Fidelity errors for initial robust and adapted control pulses, respectively, for $\omega_z \in [2.5, 3.0]$ and for $\omega_z^{\text{true}} = 2.0$.

The fidelity estimate (12) then becomes for a given control pulse sequence θ,

$$\hat{F}(\theta) = \frac{< Z(\theta) >}{(W\psi)^\dagger Z(W\psi)} = \frac{(U(\theta)\psi)^\dagger Z(U(\theta)\psi)}{(W\psi)^\dagger Z(W\psi)} \quad (17)$$

In this case the fidelity estimate can be interpreted as the ratio of the actual response to the desired response. We implement a direct Newton algorithm which obtains the gradient and Hessian by perturbing the current control pulse magnitudes by small amounts and then moving in the direction of steepest ascent. The algorithm converges in this case in 3-4 steps. The fidelity error associated with the resulting adapted pulse is shown by the red curve in Fig.2. At the true value $\omega_z^{\text{true}} = 2.0$, $\log_{10}(1 - \hat{F}(\theta_{\text{adapt}})) = -3.96$ or $\hat{F}(\theta_{\text{adapt}} = 0.99989$; a considerable improvement over the initial control for which $F_{\text{true}} = 0.9401$. Fig. 3 shows a comparison of the initial and adapted control pulses. Although close in magnitude, these small differences are known to make an enormous difference in performance, *e.g.*, [20, 21, 9].

The adaptation scheme is repeated for the case where the true value is within the anticipated range, *i.e.*, here we set $\omega_z^{\text{true}} = 2.75$ at the center of the range for which the initial control is designed. Figs. 4-5 show, respectively, the fidelity errors and pulse magnitude comparison. Again, the fidelity error associated with the resulting adapted pulse is shown by the red curve in Fig.4. At the true value $\omega_z^{\text{true}} = 2.75$, $\log_{10}(1 - \hat{F}(\theta_{\text{adapt}})) = -6.968$ or $\hat{F}(\theta_{\text{adapt}} = 0.999999892$; even more of an improvement then in previous case over the initial control within the region $F_{\text{true}} = 0.99975$. Fig. 5 reveals that the pulse magnitude difference are very small.

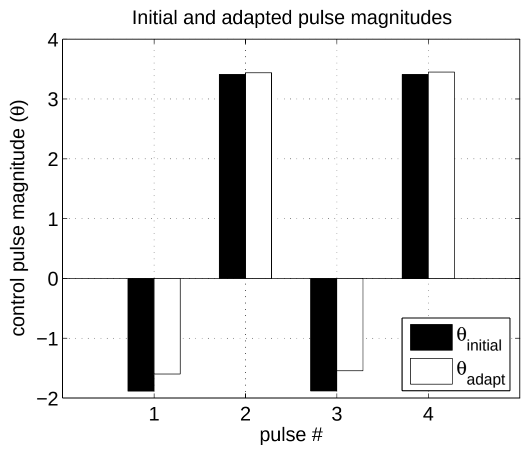

Figure 3. Initial robust and adapted control pulse magnitudes, respectively, for $\omega_z \in [2.5, 3.0]$ and for $\omega_z^{\text{true}} = 2.0$.

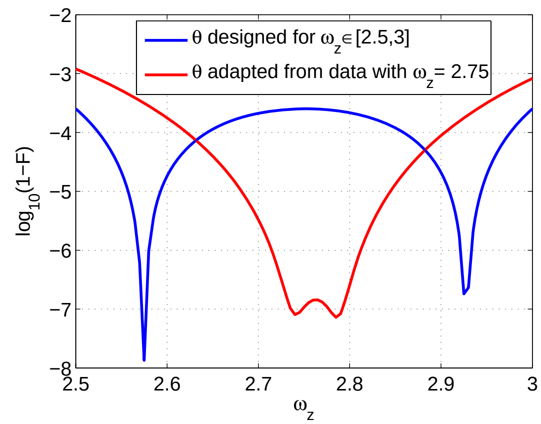

Figure 4. Fidelity errors for initial and adapted control pulses, respectively, for $\omega_z \in [2.5, 3.0]$ and for $\omega_z^{\text{true}} = 2.75$.

6. Adaptive control via parametric tomography

An alternative to direct fidelity estimation is to utilize the stronger prior that a good parametric system model is available. To illustrate the idea suppose in (13) we seek ω_z. If we take a sample in a known range, here $\omega_z \in [1.5, 2.5]$. Then to each sample $\omega_z^{()} k$ there corresponds a unitary $U^{(k)}$ (at $t = T$) *and* an associated basis set $\{\Gamma_\alpha^{(k)}\}$ and maximally sparse process matrix $X^{(k)}$, that is, if the model is perfect and and the actual parameter is the sampled value, then it has a *single* nonzero element equal to the Hilbert space dimension n. Figure 6 shows the results. The measurements are again from the single Pauli Z-observable. In each of the plots in Figure 6 we use only 10 samples of ω_z and in each keep refining the range. It is important to emphasize that the data set remains the same, only the off-line processing changes as we change the basis set for CS. The final plot shows the estimate converging to very near the true value of $\omega_z = 1.95$. A new control can then be designed to enhance performance akin to what is seen in Fig. 4.

7. Concluding remarks

Adaptation may offer a means to significantly reduce the spatial and temporal overhead to achieve continual quantum computing. One obstacle to such as approach is the time to do the adaptation. Though the speed of data collection for quantum systems though, is very high (10^4 samples can be done in a few milliseconds. *e.g.*, [13]), the specific demands are still to be determined. The results from the numerical examples shown here are encouraging. The procedure heralds the use of a minimal data collection method. The examples shown support using only a single observable and a single input state. An important caveat in our examples is that we have used the *exact* probability outcomes rather than the estimated, or empirical probabilities. This will make a difference since measurement errors will cause the adapted control to converge, at best, to a neighborhood of the perfectly tuned control, whose size depends on the level of measurement error.

In summary, given the stringent performance demands for quantum computing, the potential application which drives this work, the enticement of significant fidelity improvement through the use of a minimal data processing procedure would be valuable. Given the delicate nature of engineering quantum information, brings to mind what Richard Feynman said, “Experiment is the sole judge of scientific ’truth’.”

Acknowledgments

The authors acknowledge support from the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center contract number D11PC20165 (for RLK, HR), the Laboratory Directed Research and Development program at Sandia National Laboratories under contract DE-AC04-94AL85000 (for RLK, MDG), and the ARO MURI grant W911NF-11-1-0268 to USC (for RLK).

References

[1] M.A. Nielsen and I.L. Chuang. *Quantum Computation*

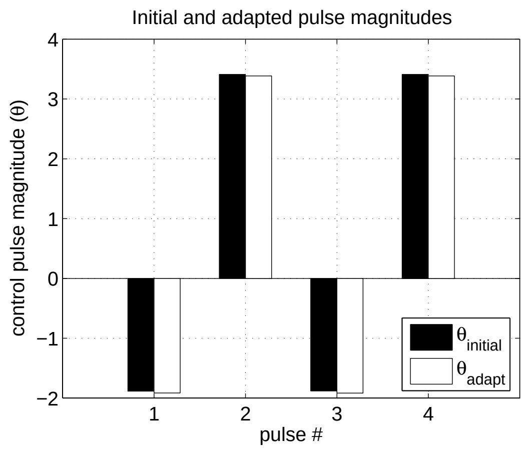

Figure 5. Initial robust and adapted control pulse magnitudes, respectively, for $\omega_z \in [2.5, 3.0]$ and for $\omega_z^{\text{true}} = 2.75$.

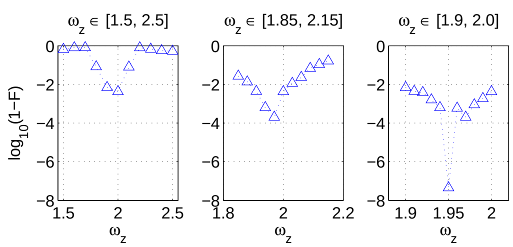

Figure 6. A sequence of CS estimates from a single data set using samples of the uncertain Hamiltonian model parameter ω_z. Associated with each sample is a basis set which would yield the maximally sparse process matrix if the true value had been selected.

and Quantum Information. Cambridge University Press, Cambridge, UK, 2000.

[2] D. Aharonov, A. Kitaev, and J. Preskill. Fault-tolerant quantum computation with long-range correlated noise. *Phys. Rev. Lett.*, 96:050504, 2006. eprint quant-ph/0510231.

[3] M.R. James and R.L. Kosut. Quantum estimation and control. In William S. Levine, editor, *The Control Handbook, Second Edition*, volume 31-1-31-42. CRC Press, 2010.

[4] I. Walmsley and H. Rabitz. Quantum physics under control. *Physics Today*, pages 43–49, August 2003.

[5] H.A. Rabitz, M.M. Hsieh, and C.M. Rosenthal. Quantum optically controlled transition landscapes. *Science*, 303:1998–2001, 2004.

[6] M.R. James, H. Nurdin, and I.R. Petersen. H^∞ control of linear quantum systems. *IEEE Trans Auto. Control*, 53(8):1787–1803, quant-ph/0703150 2008.

[7] S. Lloyd. Coherent quantum feedback. *Phys. Rev. A*, 62:022108, 2000.

[8] Kurt Jacobs and Xiaoting Wang. Coherent feedback that beats all measurement-based feedback protocols. *quant-ph/arXiv:1211.1724*, 2012.

[9] Robert L. Kosut, Matthew D. Grace, and Constantin Brif. Robust control of quantum gates via sequential convex programming. *quant-ph/arXiv:1301.3235*, 2012.

[10] A. Gilchrist, N. K. Langford, and M. A. Nielsen. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A*, 71(062310), June 2005. quant-ph/0408063.

[11] M. D. Grace, J. Dominy, R. L. Kosut, C. Brif, and H. Rabitz. Environment-invariant measure of distance between evolutions of an open quantum system. *NEW J.PHYS*, 12:015001, 2010.

[12] E. Knill and R. Laflamme. Theory of quantum error-correcting codes. *Phys. Rev. A*, 55(2):900, February 1997.

[13] A. Shabani, R. L. Kosut, M. Mohseni, H. Rabitz, M. A. Broome, M. P. Almeida, A. Fedrizzi, and A. G. White. Efficient measurement of quantum dynamics via compressive sensing. *Phys. Rev. Lett.*, 106:100401, Mar 2011.

[14] E. J. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information. *IEEE Trans. Inform. Theory*, 52(2):489–509, February 2006.

[15] R. G. Baraniuk, M. Davenport, R. A. DeVore, and M. B. Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28:253–263, December 2008.

[16] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland. Randomized benchmarking of quantum gates. *Phys. Rev. A*, 77:012307, Jan 2008.

[17] Marcus P. da Silva, Olivier Landon-Cardinal, and David Poulin. Practical characterization of quantum devices without tomography. *Phys. Rev. Lett.*, 107:210404, Nov 2011.

[18] Steven T. Flammia and Yi-Kai Liu. Direct fidelity estimation from few pauli measurements. *Phys. Rev. Lett.*, 106:230501, Jun 2011.

[19] A. Shabani, M. Mohseni, S. Lloyd, R. L. Kosut, and H. Rabitz. Estimation of many-body quantum hamiltonians via compressive sensing. *Phys. Rev. A*, 84(1):012107, Jul 2011.

[20] Gregory Quiroz and Daniel A. Lidar. Optimized dynamical decoupling via genetic algorithms. *quant-ph/arXiv:1210.5538*, 2012.

[21] Constantin Brif, Matthew D. Grace, Kevin C. Young, David L. Hocker, Katharine W. Moore, Tak-San Ho, and Herschel Rabitz. Protecting quantum gates from control noise. In *QEC11: Second International Conference on Quantum Error Correction*, Los Angeles, CA, December 2011. Available online at url-http://qserver.usc.edu/qec11/program.html.