

Necessary condition for the quantum adiabatic approximation

S. Boixo^{1,*} and R. D. Somma^{2,†}¹*Institute for Quantum Information, California Institute of Technology, Pasadena, California 91125, USA*²*Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*

(Received 4 December 2009; published 9 March 2010)

A gapped quantum system that is adiabatically perturbed remains approximately in its eigenstate after the evolution. We prove that, for constant gap, general quantum processes that approximately prepare the final eigenstate require a minimum time proportional to the ratio of the length of the eigenstate path to the gap. Thus, no rigorous adiabatic condition can yield a smaller cost. We also give a necessary condition for the adiabatic approximation that depends on local properties of the path, which is appropriate when the gap varies.

DOI: [10.1103/PhysRevA.81.032308](https://doi.org/10.1103/PhysRevA.81.032308)

PACS number(s): 03.67.Lx, 03.65.Xp, 03.67.Ac, 64.70.Tg

I. INTRODUCTION

The quantum adiabatic theorem asserts that a continuously perturbed and gapped quantum system remains in its instantaneous eigenstate in the limit where the rate of change of the perturbation vanishes [1]. This assertion is quantified via the adiabatic approximation [2], which provides a relation between the rate of change of the perturbation and the fidelity of the evolved state with the final eigenstate. The adiabatic approximation is a key part of quantum computing as it determines the complexity of several quantum algorithms [3–6]. In fact, adiabatic quantum computation [3], in which the result of a problem is encoded in the ground state of a (final) Hamiltonian, is equivalent to standard quantum computation [7]. Further, adiabatic approximations play an important role in areas like Born-Oppenheimer theory, the quantum Hall effect, and stimulated Raman adiabatic passage (STIRAP) [8].

It is important to remark that some familiar adiabatic approximations are known to be insufficient [9] and sometimes unnecessary [10,11]. The growing interest in the adiabatic approximation has spurred work on corresponding rigorous conditions [10,12]. In this article we give a rigorous lower bound for the evolution time (or cost) of adiabatic processes that prepare the final state. This bound is also valid for more general quantum evolutions [5,6,11,13].

Let $\{H(r)\}$, with $r \in [0, 1]$, be a given continuous Hamiltonian path and $\{|\psi(r)\rangle\}$ the corresponding nondegenerate eigenstate path (eigenpath). Adiabatic evolutions aim to prepare $|\psi(1)\rangle$ at bounded precision from $|\psi(0)\rangle$ by choosing a proper schedule $r(t)$. We recently argued [11] that the relevant quantities for the adiabatic approximation are not only the minimum eigenvalue gap of the Hamiltonians, Δ , but also the length of the path to be traversed, L . We presented a method that adiabatically prepares the final state by evolving with the Hamiltonians for suitable random times [11]. The average cost of the randomization method is $O(L^2/\Delta)$ when the rate of change of the eigenstates along the path and the corresponding eigenvalues are known. A more efficient method to traverse the eigenpath for this case was introduced in Ref. [6]. The method in Ref. [6] uses Grover's fixed point search and the schedule $r(t)$ is nonmonotonic. It results in a cost $\tilde{O}(L(\log L)^2/\Delta)$.

(We use the soft order notation \tilde{O} to hide doubly logarithmic factors.) Finally, we recently derived a nonmonotonic quantum algorithm or process that dynamically estimates the rate of change of the eigenstates and results in a cost $O(L \log L/\Delta)$ under broader assumptions [13].

Here, we also consider general quantum processes that prepare the final eigenstate from the initial one, at bounded precision, by evolving with the Hamiltonians. Thus, we do not exploit the unknown structure of $\{H(r)\}$, rather we work in the so-called black-box model where the only assumption is to be able to evolve with $H(r(t))$ for some schedule $r(t)$. We then prove that the cost of such processes is, at least, $O(L/\Delta)$.

To prove such a lower bound on the cost, we introduce particular instances of Hamiltonian paths $\{H(r)\}$ and reduce them to problems for which a query-complexity bound is known or can be obtained. For example, to show the scaling with Δ , we can simply consider the adiabatic version of Grover's search [3,11,14,15]. If N is the problem size, the minimum gap is $1/\sqrt{N}$ and $L \leq \pi$ in this case. The lower bound $\sqrt{N} \in O(1/\Delta)$ for the search problem is a celebrated result in quantum computation [16]. To show the dependence of the cost on the path length and the minimum gap requires a different analysis that constitutes our main contribution.

For the instances considered below, the relevant Hamiltonian eigenvalues, gaps, and rates of change of the eigenstates are known and remain constant for all r . We also clarify that a better bound may be obtained if these quantities vary along the path. Still, we show that any rigorous adiabatic approximation based on local properties of the path cannot yield a schedule that satisfies $\dot{r}(t) > c\Delta(r)/\|\partial_r \psi(r)\|$ for all r and specific $c > 0$ given below.

Before obtaining the necessary condition for the adiabatic approximation, we give a precise definition of L and comment on the resulting cost for the worst case. In the instances considered below, the path length is

$$L = \int_0^1 \|\partial_r \psi(r)\| dr. \quad (1)$$

(See Ref. [11] for a general definition of L .) With no loss of generality we assume $\langle \partial_r \psi(r) | \psi(r) \rangle = 0$. L is the only natural length in projective Hilbert space (up to irrelevant normalization factors). An upper bound on L is $\|\dot{H}\|/\Delta$, with $\|\dot{H}\| = \max \|\partial_r H\|$, and $\|\cdot\|$ is the operator norm. In the worst case, this bound is tight and $L/\Delta \sim \|\dot{H}\|/\Delta^2$. However, in

*boixo@caltech.edu

†somma@lanl.gov

many cases of interest, L can be bounded independently of Δ and the algorithms in Refs. [6,11,13] result in much smaller implementation costs than those determined by other rigorous adiabatic approximations [10,12].

The remainder of this article is organized as follows. In Sec. II, we give a simple proof of the necessary condition for the adiabatic approximation when the eigenstates are degenerate. This proof will motivate the desired result in the nondegenerate case, which is established in Sec. III. We establish the concluding remarks in Sec. IV.

II. EIGENPATH WITHIN A DEGENERATE SUBSPACE

We use constant-gap, piece-wise adiabatic universal computation [17]. Let $U = U_n \cdots U_1$ be a unitary quantum circuit acting on a Hilbert space \mathcal{H} , initially in $|\psi\rangle$. For $l \geq 1$, we define the parametrized Hamiltonians (see Fig. 1)

$$H_l(s) = -\frac{\Delta}{2} [\cos(\pi s) \mathbb{1} \otimes (|l-1\rangle\langle l-1| - |l\rangle\langle l|) + \sin(\pi s)(U_l \otimes |l\rangle\langle l-1| + U_l^\dagger \otimes |l-1\rangle\langle l|)], \quad (2)$$

acting on $\mathcal{H} \otimes \mathcal{I}$, with $s \in [0, 1]$ and Δ the relevant gap. The Hilbert space \mathcal{I} encodes the step of the circuit using the so-called clock states [7,18]. For arbitrary $|\psi\rangle$, we define $|\gamma_l\rangle = (U_l \cdots U_1 |\psi\rangle) |l\rangle$ and $|\gamma_0\rangle = |\psi\rangle |0\rangle$. Each $H_l(s)$ leaves invariant the subspace spanned by $\{|\gamma_{l-1}\rangle, |\gamma_l\rangle\}$. The eigenstates of $H_l(s)$ are degenerate; however, we are interested in the continuous path determined by the eigenstates $|\phi^l(s)\rangle = \cos(\pi s/2) |\gamma_{l-1}\rangle + \sin(\pi s/2) |\gamma_l\rangle$.

The continuous Hamiltonian path $\{H(r)\}_{0 \leq r \leq 1}$ is constructed by concatenation of n path segments; each segment linking $|\gamma_{l-1}\rangle$ to $|\gamma_l\rangle$. Then, $|\psi(1)\rangle = (U_n \cdots U_1 |\psi\rangle) |n\rangle$ and the final eigenstate of the path contains the state prepared by the circuit. The l th path segment, where $(l-1)/n \leq r < l/n$, is determined by $H(r) = H_l(rn - l + 1)$ (Fig. 1). The relevant eigenstate for this path is $|\psi(r)\rangle = |\phi^l(rn - l + 1)\rangle$. Because the rate of change of the eigenstate is constant in the intervals $(l-1)/n \leq r < l/n$ and each path segment occurs in a two-dimensional subspace, the path length is $L = \pi n$. The adiabatic procedure considered is universal for quantum computation.

No general unitary quantum process that evolves with the Hamiltonians $\{H(r)\}$ for some schedule $r(t)$ and that interleaves these evolutions with other (known) operations can prepare the final state at bounded precision with cost less than $O(L/\Delta) \in O(n/\Delta)$. This lower bound easily follows from considering those quantum circuits U built from particular instances of Grover's search. Each operation U_j that composes the circuit is then a so-called Grover iteration [15], which uses one search query. The size of the circuit is $n = \sqrt{N}$, where N is the size of the problem. In Ref. [19] it was shown that the

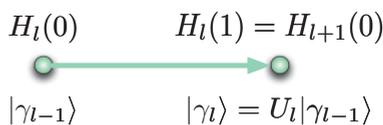


FIG. 1. (Color online) A quantum process that traverses the l th path segment applies U_l .

cost of any continuous-time quantum algorithm that uses the Hamiltonians $H(r)$ and outputs the desired marked state is, at least, $T \in O(\sqrt{N}/\Delta)$; otherwise the search problem could be solved using less than \sqrt{N} queries. Since $L \in O(\sqrt{N})$, this proves the result. A basis for the eigenspace of the final Hamiltonian is $\{|\chi\rangle |n\rangle\}_\chi$, and thus reaching the final eigenspace is easier than drawing up the right state.

III. NONDEGENERATE EIGENPATH

We now consider the case where the eigenstates are nondegenerate. We show that, for any given $\{H(r)\}$ and Δ , no generic evolution induced by a Hamiltonian $H(r(t)) + H_D(t)$, where $H_D(t)$ is a driving Hamiltonian, can prepare an approximation to the final eigenstate in time less than $O(L/\Delta)$. The result is valid for any $r(t)$ and $H_D(t)$ that do not depend on the unknown structure of $H(r)$, but may depend on the known gap and rate of change of the eigenstates [20]. This information is usually not available when deriving adiabatic approximations; our lower bound on the cost clearly encompasses those cases as well. Our setup is general, because it may include controlled-Hamiltonian evolutions as well as intermediate measurements. To prove the result, we consider instances that reduce to instances of the ordered search problem [21]. A discrete-query complexity lower bound, based on the adversary method [16,22,23], is known for this problem. We extend this result to the continuous-time query setting.

Let $x = [x_0, \dots, x_{n-1}]$, $x_j \in \{0, 1\}$, be a secret word (input) and define the Hamiltonians

$$H_x^l = \Delta |x(l), +, \dots, +\rangle \langle x(l), +, \dots, +| \quad (3)$$

acting on n qubits, where $x(l) = [x_0, \dots, x_{l-1}]$, and $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. The evolution induced by H_x^l for time π/Δ is equivalent to a phase query Q_x^l that, for input state $|a_0, \dots, a_{l-1}\rangle$, outputs $-|a_0, \dots, a_{l-1}\rangle$ when $a_j = x_j \forall 0 \leq j \leq l-1$ or does nothing otherwise. Basically, with one query Q_x^l , we answer the question: are the first l bits of x equal to $[a_0, \dots, a_{l-1}]$? The Hamiltonian path we will consider is one that has $|x(l), +, \dots, +\rangle$ as intermediate eigenstates. Then $L \in O(n)$. As in the degenerate case, this path is built upon n segments, each interpolating H_x^l with H_x^{l+1} (see below). A measurement on the final state allows us to learn the input.

We now reduce our problem to the ordered search problem. The goal of ordered search is to find the input x that represents a marked item in an ordered list of $N = 2^n$ elements. A query in this case, acting on input $a = [a_0, \dots, a_{n-1}]$, indicates whether the marked element x is before position a or not. The corresponding phase query R_x puts a \pm sign on the input state depending on the answer (see below). It requires $O(n)$ queries of type R_x to solve the problem [22]. We remark that a query of type Q_x^l can be implemented using two queries of type R_x ; to decide if the first l bits of the secret word coincide with $a(l) = [a_0, \dots, a_{l-1}]$, it suffices to call R_x with inputs $[a_0, \dots, a_{l-1}, 0, \dots, 0] - 1$ (binary subtraction) and $[a_0, \dots, a_{l-1}, 1, \dots, 1]$, respectively. If the corresponding outputs of R_x are different, then $a(l) = x(l)$. Then,

$$Q_x^l = U^l R_x V^l R_x W^l, \quad (4)$$

where U^l , V^l , and W^l are the x -independent unitaries used to build a circuit that simulates Q_x^l using R_x .

A simple argument based on the above analysis intuitively explains our result. It is now clear that no general quantum process that uses type- Q_x^l queries can find the secret word with less than $O(n)$ queries. Otherwise, by replacing the queries, we would solve ordered search with less than $O(n)$ type- R_x queries. Since each type- Q_x^l query can be implemented using the Hamiltonians H_x^l for time π/Δ , it is plausible that the total cost of a continuous-time quantum process that uses H_x^l is at least $O(n/\Delta) \in O(L/\Delta)$. In fact, using the equivalence between continuous- and discrete-time query models in Ref. [24] yields a lower bound for the cost $\tilde{O}(L/[\Delta \log(L/\Delta)])$. Our formal proof below, which uses a version of the adversary method in the continuous-time setting, will avoid the logarithmic correction in the cost.

Let $H_x^l(s) = \cos(\pi s/2)H_x^{l-1} + \sin(\pi s/2)H_x^l$. We build a particular Hamiltonian path as in the degenerate case: for $(l-1)/n \leq r \leq l/n$ we set $H_x(r) = H_x^l(rn - l + 1)$. The eigenstates $|\psi(r)\rangle$ of $H_x(r)$ with the lowest eigenvalue are nondegenerate. Further, the gap of $H_x(r)$ is Δ for all r . A quantum algorithm that traverses the eigenpath will aim to prepare $|x(l), +, \dots, +\rangle$ from $|x(l-1), +, \dots, +\rangle$, learning a bit of information about the secret word at each path segment. The path length is $L = \pi n/2$.

We consider evolutions with Hamiltonian $H_x(r(t)) + H_D(t)$, where $r(t)$ and $H_D(t)$ do not depend on x , the only unknown quantity in this case. Ideally, after some time $T > 0$, the initial state approximately evolves to the desired eigenstate $|\psi(1)\rangle = |x_0, \dots, x_{n-1}\rangle$. Evolutions of this type include continuous-time processes based on eigenpath traversal, such as adiabatic evolutions where $H_D = 0$. If $|\phi_x(t)\rangle$ is the evolved state,

$$i\partial_t|\phi_x(t)\rangle = [H_x(r(t)) + H_D(t)]|\phi_x(t)\rangle. \quad (5)$$

The initial state $|\phi_x(0)\rangle$ is also independent of x .

We use the adversary method to show that T is at least $O(n/\Delta)$. At its core, the adversary method provides a bound for the rate of change of the overlap between evolutions corresponding to different inputs. To distinguish between these inputs, the evolved states must satisfy $|\langle\phi_x(T)|\phi_y(T)\rangle| \leq \epsilon$, for some $T > 0$ and small ϵ .

Let Γ be an adversary matrix—an irreducible symmetric matrix with nonnegative entries and zeros in the diagonal. Denote by $\|\Gamma\|$ its operator norm, and by v the principal unit eigenvector; $\Gamma v = \|\Gamma\|v$. The following function serves as a measure of the distinguishability between evolutions with different inputs:

$$W(t) = \sum_{x,y} \Gamma_{x,y} v_x v_y \langle\phi_x(t)|\phi_y(t)\rangle. \quad (6)$$

Since the initial state is independent of the input, $W(0) = \|\Gamma\|$. Moreover, $W(T) \leq \epsilon\|\Gamma\|$.

An upper bound on $|\partial_t W(t)|$ can be obtained if we extend known results for the ordered search problem to our case. We write the Hamiltonians H_x^l as [see Eq. (4)]

$$H_x^l = \Delta(\mathbb{1} - Q_x^l) = \Delta(\mathbb{1} - U^l R_x V^l R_x W^l). \quad (7)$$

Then, using Eq. (5) and the triangle inequality,

$$\begin{aligned} |\partial_t \langle\phi_x(t)|\phi_y(t)\rangle| &\leq \max_l |\langle\phi_x(t)|H_x^l - H_y^l|\phi_y(t)\rangle| \\ &\leq 2\Delta \max_l |\langle\phi_x(t)|U^l R_x V^l R_x W^l \\ &\quad - U^l R_y V^l R_y W^l|\phi_y(t)\rangle|. \end{aligned} \quad (8)$$

For each x , we define a 2^n -bit string α_x with entries $\alpha_x^i = 0$ if $i < x$ and $\alpha_x^i = 1$ otherwise. Let $P_i = |i\rangle\langle i|$ and write $\langle\phi_x^l(t)\rangle = \langle\phi_x(t)|U^l R_x V^l$ and $|\tilde{\phi}_y^l(t)\rangle = V^l R_y W^l|\phi_y(t)\rangle$. The triangle inequality and Eq. (8) yield

$$\begin{aligned} |\partial_t \langle\phi_x(t)|\phi_y(t)\rangle| &\leq 2\Delta \max_l \sum_{i:\alpha_x^i \neq \alpha_y^i} [|\langle\phi_x^l(t)|P_i W^l|\phi_y(t)\rangle| \\ &\quad + |\langle\phi_x(t)|U^l P_i|\tilde{\phi}_y^l(t)\rangle|]. \end{aligned}$$

For $i \in \{0, \dots, 2^n - 1\}$, we introduce the matrices Γ^i as

$$\Gamma_{x,y}^i = \begin{cases} \Gamma_{x,y} & \text{if } \alpha_x^i \neq \alpha_y^i, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

Then,

$$\begin{aligned} |\partial_t W(t)| &\leq 2\Delta \sum_{i,x,y} \Gamma_{x,y}^i v_x v_y \max_l [|\langle\phi_x^l(t)|P_i W^l|\phi_y(t)\rangle| \\ &\quad + |\langle\phi_x(t)|U^l P_i|\tilde{\phi}_y^l(t)\rangle|] \\ &\leq 4\Delta \max_i \|\Gamma^i\|. \end{aligned} \quad (10)$$

In addition, $(1 - \epsilon)\|\Gamma\| \leq W(0) - W(T)$ and

$$\begin{aligned} W(0) - W(T) &\leq \int_0^T |\partial_t W(t)| dt \\ &\leq T 4\Delta \max_i \|\Gamma^i\|. \end{aligned} \quad (11)$$

The spectral lower bound for the cost of the process that approximates the final eigenstate is $T \in O[\|\Gamma\|/(4\Delta \max_i \|\Gamma^i\|)]$. We use the adversary matrix for the ordered search problem [22]:

$$\Gamma_{x,y} = \begin{cases} \frac{1}{\text{Hd}(\alpha_x, \alpha_y)} & \text{if } \alpha_x \neq \alpha_y, \\ 0 & \text{otherwise,} \end{cases} \quad (12)$$

where $\text{Hd}(\alpha_x, \alpha_y)$ is the Hamming distance between its arguments. This choice yields $\|\Gamma\| \geq n$ and $\max_i \|\Gamma^i\| \leq \pi$. Thus, $T \in \Omega[n/(4\pi\Delta)]$. Since $L = \pi n/2$, this proves the result $T \in O(L/\Delta)$.

With a similar construction we can prove a lower bound on the cost of continuous-time query algorithms that solve the ordered search problem, obtaining $T \in \Omega[n/(2\pi)]$ in this case.

In the derivation of our result we used the fact that the gaps of $\{H_x(r)\}$ are constant along the path. Nevertheless, the Hamiltonians in the path may have different (known) gaps and one could be interested in designing algorithms for eigenpath traversal with a schedule $r(t)$ that depends on the local gap $\Delta(r)$. Consider again the instances above and redefine $H_x(r) \leftarrow q(r)H_x(r)$, for known $q(r) > 0$. The new Hamiltonians have gaps $\Delta(r) = q(r)\Delta$. We can replace Δ by $\Delta(r(t))$ in Eq. (8). Following the steps above, and using the adversary matrix of Eq. (12), we obtain

$$(1 - \epsilon)n \leq 4\pi \int_0^T dt \Delta(r(t)). \quad (13)$$

We prove by contradiction that adiabatic approximations based on local properties of the path cannot yield a schedule satisfying the local condition $\dot{r}(t) > c\Delta(r)/\|\partial_r\psi(r)\|$, for some $c > 0$ and all r . If such a condition is satisfied, the inverse function $t(r)$ exists and Eq. (13) yields

$$(1 - \epsilon)n < \frac{4\pi}{c} \int_0^1 dr \|\partial_r\psi(r)\| = \frac{4\pi L}{c}. \quad (14)$$

Since $L = \pi n/2$ in this case, it is clear that $c < 2\pi^2/(1 - \epsilon)$ or otherwise the above inequality is inconsistent.

IV. CONCLUSIONS

We proved that no general quantum process that approximately prepares the eigenstate of a Hamiltonian $H(1)$, by evolving with the path of Hamiltonians $\{H(r)\}_{0 \leq r \leq 1}$, for any schedule $r(t)$, can achieve its goal in time less than $O(L/\Delta)$. The same bound on the time applies also for more general evolutions with additional driving Hamiltonians. Interestingly, some quantum processes for eigenpath traversal

almost achieve the bound under some assumptions [13]. We also gave a necessary local condition, valid even when the gaps of the Hamiltonians and rates of change of the eigenstates are known along the path. In this case we proved that if the schedule satisfies $\dot{r}(t) > c\Delta(r)/\|\partial_r\psi(r)\|$, for a specific $c > 0$ and all r , the quantum process will not succeed in the state preparation. If only $\Delta \leq \min_r \Delta(r)$ is known, our result suggests that no general monotonic schedule yields a cost better than $O(L/\Delta)$.

ACKNOWLEDGMENTS

We thank H. Barnum, R. Cleve, A. Childs, and E. Knill for discussions, and the Perimeter Institute where part of this work was done. SB thanks the National Science Foundation for support under grant PHY-0803371 through the Institute for Quantum Information at the California Institute of Technology. RS thanks the Laboratory Directed Research and Development Program at Los Alamos National Laboratory for support.

-
- [1] A. Messiah, *Quantum Mechanics* (Dover, New York, 1999).
 - [2] M. Born and V. Fock, *Z. Phys.* **51**, 165 (1928).
 - [3] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, e-print arXiv:quant-ph/0001106 (2000).
 - [4] D. Aharonov and Ta-Shma, in *Proceedings of the 35th Annual ACM Symposium on Theory of Computing, San Diego, CA, 2003*, p. 20; *Quantum Annealing and Related Optimization Methods*, edited by A. Das and B. K. Chakrabarti (Springer, Berlin, 2005); K. Eckert, O. Romero-Isart, and A. Sampera, *New J. Phys.* **9**, 155 (2007).
 - [5] R. D. Somma, S. Boixo, H. Barnum, and E. Knill, *Phys. Rev. Lett.* **101**, 130504 (2008).
 - [6] P. Wocjan, C. F. Chiang, D. Nagaj, and A. Abeyesinghe, *Phys. Rev. A* **80**, 022340 (2009).
 - [7] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, *SIAM J. Comput.* **37**, 166 (2007).
 - [8] S. Teufel, *Adiabatic Perturbation Theory in Quantum Dynamics* (Springer, Berlin, 2003).
 - [9] J. Schwinger, *Phys. Rev.* **51**, 648 (1937); K. P. Marzlin and B. C. Sanders, *Phys. Rev. Lett.* **93**, 160408 (2004); D. M. Tong, K. Singh, L. C. Kwek, and C. H. Oh, *ibid.* **95**, 110407 (2005).
 - [10] G. Schaller, S. Mostame, and R. Schutzhold, *Phys. Rev. A* **73**, 062307 (2006); S. Jansen, M. Ruskai, and R. Seiler, *J. Math. Phys.* **48**, 102111 (2007).
 - [11] S. Boixo, E. Knill, and R. Somma, *Quantum Inf. Comput.* **9**, 833 (2009).
 - [12] J. E. Avron, R. Seiler, and L. G. Yaffe, *Commun. Math. Phys.* **110**, 33 (1987); G. A. Hagedorn and A. Joye, *J. Math. Anal. & App.* **267**, 235 (2002); D. A. Lidar, A. T. Rezakhani, and A. Hamma, *J. Math. Phys.* **50**, 102106 (2009).
 - [13] S. Boixo, E. Knill, and R. D. Somma (in preparation, 2010).
 - [14] J. Roland and N. J. Cerf, *Phys. Rev. A* **65**, 042308 (2002).
 - [15] L. K. Grover, in *Proceedings of the 28th Annual ACM Symposium on Theory of Computing, Philadelphia, PA 1996*, p. 212.
 - [16] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, *SIAM J. Comput.* **26**, 1510 (1997).
 - [17] A. Y. Kitaev, *Ann. Phys. (NY)* **303**, 2 (2003); P. Zanardi and M. Rasetti, *Phys. Lett.* **A264**, 94 (1999); D. Bacon and S. T. Flammia, *Phys. Rev. Lett.* **103**, 120504 (2009).
 - [18] R. P. Feynman, *Found. Phys.* **16**, 507 (1986); A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and Quantum Computation* (American Mathematical Society, 2002).
 - [19] E. Farhi and S. Gutmann, e-print arXiv:quant-ph/9612026 (1996).
 - [20] Margolus-Levitin theorem would give a bound in terms of the total energy, including that of $H_D(t)$, but not in terms of the gap. N. Margolus and L. B. Levitin, *Physica D* **120**, 188 (1998).
 - [21] H. Buhrman and R. de Wolf, e-print arXiv:quant-ph/9811046 (1998); A. Ambainis, in *Proceedings of the 40th Annual Symposium on Foundations of Computing Science, New York, NY, 1999*, p. 352; P. Høyer and R. Spalek, *Bulletin of the European Association for Theoretical Computer Science* **87**, 78 (2005); A. Childs and T. Lee, *Automata, Languages, and Programming* (Springer, Berlin, 2008).
 - [22] P. Høyer, J. Neerbek, and Y. Shi, *Algorithmica* **34**, 429 (2002).
 - [23] A. Ambainis, in *Proceedings of the 32nd Annual ACM Symposium on Theory of Computing, Portland, OR, 2000*, p. 636; H. Barnum, M. Saks, and M. Szegedy, in *Proceedings of the 18th IEEE Annual Conference, Washington, DC, 2003*, p. 179.
 - [24] R. Cleve, D. Gottesman, M. Mosca, R. D. Somma, and D. Yonge-Mallo, in *Proceedings of the 41st Annual Symposium on Theory of Computing, Bethesda, MD, 2009*, p. 409.