Generalized quantum search with parallelism

Robert M. Gingrich,[1] Colin P. Williams,[2] and Nicolas J. Cerf[1,2,3]
[1]*W. K. Kellogg Radiation Laboratory, California Institute of Technology, Pasadena, California 91125*
[2]*Information and Computing Technologies Research Section, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109*
[3]*Ecole Polytechnique, Université Libre de Bruxelles, B-1050 Brussels, Belgium*
(Received 3 November 1999; published 19 April 2000)

We generalize Grover's unstructured quantum search algorithm to enable it to work with arbitrary starting superpositions and arbitrary unitary operators. We show that the generalized quantum search algorithm, when cast in a special orthonormal basis, can be understood as performing an *exact* rotation of a starting superposition into a target superposition. We derive a formula for the success probability of the generalized quantum search algorithm after n rounds of amplitude amplification. We then use this formula to determine the optimal strategy for a *punctuated* quantum search algorithm, i.e., one in which the amplitude amplified state is observed before the point of maximum success probability. On average, the optimal strategy is about 12% better than the naive use of Grover's algorithm. The speedup obtained is not dramatic but it illustrates that a hybrid use of quantum computing and classical computing techniques can yield a performance that is better than either alone. In addition, we show that a punctuated quantum algorithm that takes the same average computation time as Grover's standard algorithm only requires half the coherence time. We then extend the analysis to the case of a society of k quantum searches acting in parallel. We derive an analytic formula that connects the degree of parallelism with the expected computation time for k-parallel quantum search. The resulting parallel speedup scales as $O(\sqrt{k})$, while the minimum number of agents needed to ensure success, k, decreases as the inverse of the *square* of the achievable coherence time. This result has practical significance for the design of rudimentary quantum computers that are likely to have a limited coherence time.

PACS number(s): 03.67.Lx, 89.70.+c

I. INTRODUCTION

The field of quantum computing has undergone a rapid growth over the past few years. Simple quantum computations have already been performed using nuclear magnetic resonance [1–6] and nonlinear optics technologies [7,8]. Recently, proposals for specialized devices that rely on quantum computing have also been made [9]. Such devices are far from being general-purpose computers; nevertheless, they constitute significant milestones along the road to practical quantum computing.

In tandem with these hardware developments, there has been a parallel development of new quantum algorithms. Several important quantum algorithms are now known [10–15]. Of particular importance is the quantum algorithm for performing an unstructured quantum search discovered by Lov Grover in 1996 [12]. Further analysis of this algorithm is given by Jozsa [16] and an optical implimentation is given by Kwiat [17]. Grover's algorithm is able to find a marked item in a virtual ''database'' containing N items in $O(\sqrt{N})$ computational steps. In contrast, the best classical algorithm requires $O(N/2)$ steps on average, and $O(N)$ steps in the worst case. Thus Grover's algorithm exhibits a quadratic speedup over the best classical counterpart.

Although Grover's algorithm exhibits only a polynomial speedup, it appears to be much more versatile than the other quantum algorithms. Indeed, Grover has shown how his algorithm can be used to speed up almost any other quantum algorithm [18]. More surprisingly, even search problems that contain ''structure'' in the form of correlations between the items searched over, often reduce to an exhaustive search amongst a reduced set of possibilities. Recently, it was shown how Grover's algorithm can be nested to exploit such a problem structure [13]. This is significant because NP-hard problems, which are among the most challenging computational problems that arise in practice, possess exactly this kind of problem structure.

In order to appreciate the full versatility of Grover's algorithm it is important to examine all the ways in which it might be generalized. For example, whereas the original Grover algorithm was started from an equally weighted superposition of eigenstates representing all the indices of the items in the database, a natural generalization would be to consider how it performs when started from an arbitrary initial superposition instead. This refinement is important, because if Grover's algorithm is used within some larger quantum computation, it is likely to have to work on an arbitrary starting superposition rather than a specific starting eigenstate. Similarly, the original Grover algorithm uses a particular unitary operator, the Walsh-Hadamard operator, as the basis for a sequence of unitary operations that systematically amplifies the amplitude in the target state at the expense of the amplitude in the nontarget states. However, it is now known that this is not the best choice if there is partial information as to the likely location of the target item in the database. In such a situation a different unitary operator is desirable [19]. Hence, it is equally important to understand how Grover's algorithm performs when using an arbitrary unitary operator instead of the Walsh-Hadamard operator.

Each of these refinements have been analyzed in detail *separately*: Biham *et al.* have considered the case of an arbitrary starting superposition [20], while Grover considered the case of an arbitrary unitary operator [19]. In this paper, we present the analysis of the fully generalized Grover algo-

rithm in which we incorporate both of these effects simultaneously. Our goal is to determine the exact analytic formula for the probability of the fully generalized Grover algorithm succeeding after n iterations when there are r targets amongst N candidates. Having obtained this formula, we will recover the Biham *et al.* and Grover results as special cases. We will then show that the optimal strategy, on average, for using the fully generalized Grover algorithm consists of measuring the memory register after about 12% fewer iterations than are needed to obtain the maximum probability of success. This result confirms a more restricted case reported in [21]. Finally, we show how to boost the success probability and reduce the required coherence time by running a society of k quantum searches independently in parallel. In particular, we derive an explicit formula connecting the degree of parallelism, i.e., k, to the optimal number of iterations (for each agent in the society) that minimizes the expected search cost overall. We then derive the expected cost of an optimal k-parallel quantum search.

II. GROVER'S ALGORITHM

The problem we have to solve is the following. Given a function $f(x_i)$ on a set $\mathcal{X}$ of input states such that

$$f(x_i)=\begin{cases}1 & \text{if } x_i \text{ is a target element}\\ 0 & \text{otherwise.}\end{cases} \tag{1}$$

How do we find a target element by using the least number of calls to the function $f(x_i)$? In general, there might be r target elements, in which case any one will suffice as the answer.

To solve the problem using Grover's algorithm we first form a Hilbert space with an orthonormal basis element for each input $x_i \in \mathcal{X}$. In this paper, we refer to the basis of input eigenstates as the measurement basis. Let $N=|\mathcal{X}|$ be the cardinality of $\mathcal{X}$. Without loss of generality, we will write the target states as $|t_i\rangle$ (with $i=1,\ldots,r$), and the nontarget states as $|l_i\rangle$ (with $i=1,\ldots,N-r$). The function call is to be implemented by a unitary operator that acts as follows:

$$|x_i\rangle|y\rangle \rightarrow |x_i\rangle|y\oplus f(x_i)\rangle, \tag{2}$$

where $|y\rangle$ is either $|0\rangle$ or $|1\rangle$. By acting on

$$\left(\sum_{i=1}^{N-r} l_i|l_i\rangle+\sum_{j=1}^{r} k_j|t_j\rangle\right)\frac{1}{\sqrt{2}}(|0\rangle-|1\rangle) \tag{3}$$

with this operator we construct the state

$$\left(\sum_{i=1}^{N-r} l_i|l_i\rangle-\sum_{j=1}^{r} k_j|t_j\rangle\right)\frac{1}{\sqrt{2}}(|0\rangle-|1\rangle), \tag{4}$$

where the r measurement basis states $|t_i\rangle$ are the target states and the $N-r$ measurement basis states $|l_i\rangle$ are the nontarget states. If we now disregard the state $(1/\sqrt{2})(|0\rangle-|1\rangle)$ then all we have done is to *invert* the phase of the target states. Hence, the operator we have achieved is equivalent to the inversion operator

$$1-2\sum_{i=1}^{r}|t_i\rangle\langle t_i|, \tag{5}$$

although it must be emphasized that this operation can be performed without knowing the target states $|t_i\rangle$ *explicitly* but only through the knowledge of $f(x)$.

Next, we construct the operator Q defined as the product of the above operator and an inversion operator with respect to the state $|a\rangle$, that is

$$Q=-(1-2|a\rangle\langle a|)\left(1-2\sum_{i=1}^{r}|t_i\rangle\langle t_i|\right). \tag{6}$$

Different choices of $|a\rangle$ give rise to different unitary operators for performing amplitude amplification. In the original Grover algorithm, the state $|a\rangle$ was chosen to be

$$|a\rangle=\frac{1}{\sqrt{N}}\sum_{x\in\mathcal{X}}|x\rangle \tag{7}$$

and was obtained by applying the Walsh-Hadamard operator U to a starting state $|s\rangle=|00\cdots\rangle$, i.e., $|a\rangle=U|s\rangle$. Hence, the operation $2|a\rangle\langle a|-1$, which Grover referred to as ''inversion about the average,'' is equivalent to $-UI_sU^{\dagger}$, with U being the Walsh-Hadamard operator and I_s being $1-2|s\rangle\langle s|$. As we shall see, Grover's algorithm is based on applying Q for a certain number of iterations. By knowing more about the structure of the problem we can choose other vectors $|a\rangle$ that will allow us to find a target state faster. (Techniques for doing this are given in [18].)

In order to understand what action the operator Q performs, let us define $|t\rangle$ as the normalized projection of $|a\rangle$ onto the space of target states, that is

$$|t\rangle=\frac{1}{v}\sum_{i=1}^{r}\langle t_i|a\rangle|t_i\rangle \quad \text{with } v^2=|\langle t|a\rangle|^2=\sum_{i=1}^{r}|\langle t_i|a\rangle|^2. \tag{8}$$

We can then rewrite Q as

$$Q=\sum_{i=1}^{r}|t_i\rangle\langle t_i|-\sum_{j=1}^{N-r}|l_j\rangle\langle l_j|+2|a\rangle\langle a|-4v|a\rangle\langle t|, \tag{9}$$

implying that Q only acts nontrivially on the space spanned by $|a\rangle$ and $|t\rangle$. Let us reexpress this by using the states $|t\rangle$ and $|l\rangle$ as an *orthonormal* basis for this subspace, where we have defined

$$|l\rangle=\frac{1}{\sqrt{1-v^2}}(|a\rangle-v|t\rangle). \tag{10}$$

The vector $|l\rangle$ is just the normalized projection of $|a\rangle$ onto the space of nontarget states ($\langle t|l\rangle=0$). The rest of the space (i.e., the space orthogonal to $|t\rangle$ and $|l\rangle$) can be broken up into the space of target states ($\mathcal{S}_{\mathcal{T}}$) and nontarget states ($\mathcal{S}_{\mathcal{L}}$). Using $|a\rangle=\sqrt{1-v^2}|l\rangle+v|t\rangle$, we can now write Q as

$$Q=\cos\phi(|t\rangle\langle t|+|l\rangle\langle l|)+\sin\phi(|t\rangle\langle l|-|l\rangle\langle t|)+I_T-I_L, \tag{11}$$

where $I_T=\Sigma_{i=1}^{r}|t_i\rangle\langle t_i|-|t\rangle\langle t|$ and $I_L=\Sigma_{j=1}^{N-r}|l_j\rangle\langle l_j|-|l\rangle\langle l|$ are the identity operators on $\mathcal{S}_\mathcal{T}$ and $\mathcal{S}_\mathcal{L}$, respectively, and $\phi\equiv\arccos[1-2v^2]$. From this we can see that Q is just a simple rotation matrix on the two-dimensional space spanned by $|l\rangle$ and $|t\rangle$, and acts trivially on the rest of the space. Note that the operator Q has been independently shown by Jozsa [16] to be an exact rotation in the special case of one solution and with $|a\rangle$ given by Eq. (7).

An arbitrary starting superposition $|s\rangle$ for the algorithm can be written as

$$|s\rangle=\alpha|t\rangle+\beta e^{ib}|l\rangle+|s_t\rangle+|s_l\rangle, \tag{12}$$

where the states $|s_t\rangle$ and $|s_l\rangle$ (which must have a norm less than one if the state $|s\rangle$ is to be properly normalized overall) are the components of $|s\rangle$ in $(\mathcal{S}_\mathcal{T})$ and $(\mathcal{S}_\mathcal{L})$, respectively. Also, α,β, and b are positive real numbers. After n applications of Q on an arbitrary starting superposition $|s\rangle$, we have

$$\begin{aligned}Q^n|s\rangle=&[\alpha\cos(n\phi)+\beta e^{ib}\sin(n\phi)]|t\rangle+[\beta e^{ib}\cos(n\phi)\\&-\alpha\sin(n\phi)]|l\rangle+|s_t\rangle+(-1)^n|s_l\rangle.\end{aligned} \tag{13}$$

If we measure this state our probability of success (i.e., measuring a target state) will be given by two terms. The first term is the squared magnitude of $Q^n|s\rangle$ projected into the space $\mathcal{S}_\mathcal{T}$. It is equal to $\langle s_t|s_t\rangle$ and is unchanged by Q. The second term is the squared magnitude of the component of $|t\rangle$ which is given by

$$\begin{aligned}g(n)&\equiv|\langle t|Q^n|s\rangle|^2\\&=|\alpha\cos(n\phi)+\beta e^{ib}\sin(n\phi)|^2\\&=\frac{\alpha^2+\beta^2}{2}+\frac{\alpha^2-\beta^2}{2}\cos(2n\phi)+\alpha\beta\cos(b)\sin(2n\phi)\\&=\frac{\alpha^2+\beta^2}{2}-\frac{1}{2}|\alpha^2+\beta^2e^{2ib}|\cos(2n\phi+\psi),\end{aligned} \tag{14}$$

where $\psi\equiv\arccos[(\beta^2-\alpha^2/|\alpha^2+\beta^2e^{2ib}|]$. This is the term that is affected by Q, and is the term we wish to maximize. The probability of success after n iterations of Q acting on $|s\rangle$ is thus

$$p(n,r,N)=\langle s_t|s_t\rangle+g(n). \tag{15}$$

Assuming that n is continuous (an assumption that we will justify shortly) the maxima of $g(n)$, and hence the maxima of the probability of success of Grover's algorithm, are given by the following:

$$n_j=\frac{-\psi+(1+2j)\pi}{2\phi},\quad j=0,1,2,\ldots. \tag{16}$$

The value of $g(n)$ at these maxima is given by

$$g(n_j)=\frac{\alpha^2+\beta^2}{2}+\frac{1}{2}|\alpha^2+\beta^2e^{2ib}|. \tag{17}$$

In practice, the optimal n must be an integer and typically the n_j's are not integers. However, since $g(n)$ can be written as

$$g(n_j\pm\delta)=g(n_j)-\phi^2|\alpha^2+\beta^2e^{2ib}|\delta^2+O(\delta^4) \tag{18}$$

around n_j and most interesting problems will have $v\ll1$ and hence $\phi\simeq2v\ll1$, simply rounding n_j to the nearest integer will not significantly change the final probability of success. So, we have

$$p(n_{max},r,N)=\frac{\alpha^2+\beta^2}{2}+\frac{1}{2}|\alpha^2+\beta^2e^{2ib}|+\langle s_t|s_t\rangle-O(v^2) \tag{19}$$

as the probability of measuring a target state after n_{max} applications of Q.

III. RECOVERING THE SPECIAL CASES

As a check on our fully generalized formula for the probability of success after n iterations, we attempt to recover the corresponding formulas obtained in the analyses of Biham *et al.* (for a fixed unitary operator and an arbitrary starting superposition) [20] and Grover (for an arbitrary unitary operator and a fixed starting superposition) [19].

In the case of Biham *et al.*, the starting state is arbitrary but the averaging state $|a\rangle$ is given by

$$|a\rangle=\frac{1}{\sqrt{N}}\sum_{x\in\mathcal{X}}|x\rangle. \tag{20}$$

In this case

$$\begin{aligned}v&=\sqrt{\frac{r}{N}},\\|t\rangle&=\frac{1}{\sqrt{r}}\sum_{i=1}^{r}|t_i\rangle,\\|l\rangle&=\frac{1}{\sqrt{N-r}}\sum_{i=1}^{N-r}|l_i\rangle.\end{aligned} \tag{21}$$

In the analysis of [20] they use $\overline{k}(0)$ and $\overline{l}(0)$ to represent the average amplitudes, in $|s\rangle$, of the target and nontarget states, respectively, and σ_k and σ_l to represent the standard deviations of those amplitudes. With some algebra one can see that the following relationships connect our notation to theirs:

$$\begin{aligned}&\alpha\to\overline{k}(0)\sqrt{r},\quad \beta e^{ib}\to\overline{l}(0)\sqrt{N-r},\\&\langle s_t|s_t\rangle\to r\sigma_k^2,\quad \langle s_l|s_l\rangle\to(N-r)\sigma_l^2,\\&\phi\to\omega,\quad \psi\to2\,\mathrm{Re}[\phi],\\&n\to t,\quad n_0\to T.\end{aligned} \tag{22}$$

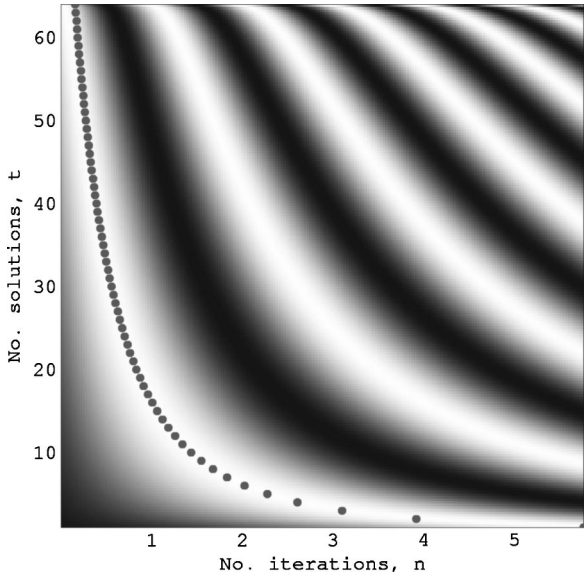

FIG. 1. Plot of the probability of success of Grover's algorithm after n iterations of amplitude amplification when there are r solutions amongst $N=64$ possibilities. White regions correspond to probability 1, black regions correspond to probability 0. Note that the success probability is periodic in the number of amplitude amplification iterations for a fixed number of solutions.

By substituting these relationships into Eqs. (14), (16), and (19), one reproduces the results of [20].

The second special case, in which $|a\rangle$ (with respect to which inversion is done) is an unknown normalized vector, while $|s\rangle$ is given by

$$|s\rangle=|a\rangle=\sqrt{1-v^2}|l\rangle+v|t\rangle, \tag{23}$$

was considered by Grover. Hence, $\alpha=v,\beta=\sqrt{1-v^2}$ and $b=0$. Also, $|s_t\rangle=|s_l\rangle=0$. These substitutions lead to $\psi=\phi$. Plugging this into Eqs. (16) and (19), we get

$$n_{max}=\frac{\pi}{2\phi}-\frac{1}{2}=\frac{\pi}{4v}-\frac{1}{2}-\frac{\pi v}{24}+O(v^2) \tag{24}$$

and

$$p(n_{max})=1-O(v^2) \tag{25}$$

which agree with the results of [19]. If we examine Eq. (15) in this case, we get

$$p(n)=\frac{1-\cos[(1+2n)\phi)]}{2}=\sin^2[(1+2n)\phi/2] \tag{26}$$

as the probability of measuring a target state after n iterations of Q [see Fig. 1 for a plot of $p(n)$ in the case of the Hadamard transformation].

IV. APPLICATION OF THE FORMULA FOR $P(N)$

Next, we show how to apply our analytic formula for the probability of success after n iterations, $P(n)$, to slightly speed up the quantum unstructured search algorithm. Although the speedup we obtain is not dramatic, it is worth making the point that it is possible at all as Zalka has proved, correctly, that Grover's algorithm is exactly optimal [22]. Many people have assumed, therefore, that it is impossible to beat Grover's algorithm. However, by combining techniques of quantum computing with those of classical computing, we show that it is possible to do a little bit better than Grover's algorithm on average. The result we report was apparently discovered previously by Boyer *et al.* [21] and later by Zalka [23] in the case where $|a\rangle$ is a uniform superposition [as in Eq. (7)]. It is shown here to persist for the more general case of an arbitrary $|a\rangle=|s\rangle$, which is the case treated in [19].

We consider a punctuated quantum search algorithm that works as follows:

Algorithm: Punctuated Quantum Search
(1) Run the quantum search algorithm for n iterations.
(2) Read the memory register.
(3) If the result is a target state halt; else, reset the register to the starting superposition and return to step (1).

The average time $T_{avg}(n)$ it will take to find a target state if we stop the generalized quantum search algorithm after n iterations of Q is

$$T_{avg}(n)=\sum_{i=1}^{\infty}[1-p(n)]^{i-1}p(n)i \quad n=\frac{n}{p(n)} \tag{27}$$

$$=\frac{2n}{1-\cos[(1+2n)\phi]}.$$

We can find the optimal strategy, i.e., the best number of iterations to use before we attempt to measure the register, by minimizing the expected running time T_{avg}. To do this, we set the derivative of T_{avg} to zero and solve for $n=n_{opt}$,

$$\frac{\partial T_{avg}}{\partial n}=\frac{2-2\cos[(1+2n)\phi]-4n\phi\sin[(1+2n)\phi]}{\{(1-\cos[(1+2n)\phi]\}^2}=0. \tag{28}$$

Typically, n will be much larger than one, so we can make the approximation $(1+2n)(\phi/2)\simeq n\phi\equiv x$, so that we obtain

$$1-\cos 2x=2x\sin 2x,$$

$$2\sin^2 x=4x\sin x\cos x, \tag{29}$$

$$2x=\tan x,$$

which gives $x_{opt}=1.1656$ as the lowest positive solution. This solution corresponds to the minimum of the function. Hence the optimal value of n is

$$n_{opt}\simeq\frac{x_{opt}}{\phi}=\frac{1.1656}{\phi}. \tag{30}$$

This value of n gives a probability of success of

$$p(n_{opt})=\sin^2 x_{opt}=0.8446 \tag{31}$$

at each measurement, and corresponds to an average number of iterations of

$$T_{avg}(n_{opt}) \simeq \frac{1}{\phi}\frac{x_{opt}}{\sin^2 x_{opt}} = \frac{1.3801}{\phi}. \tag{32}$$

This must be compared to $\pi/2\phi = 1.5708/\phi$ iterations if we run Grover's algorithm until the probability is maximal. Thus, we get a 12% reduction of the average computation time by making use of a punctuated algorithm.

It is interesting to note that, if we restrict the analysis some more to the case where $|a\rangle$ is a uniform superposition and where there is only one target state, then we have $\phi = 2/\sqrt{N}$, so that $T_{avg}(n_{opt}) \simeq 0.6900\sqrt{N}$. This is faster than the lower bounds in [21,24,25], and [22], but we are using a somewhat different model. They are looking at the minimum time it would take without measuring to find a solution with certainty up to errors from rounding n_{max} to the nearest integer. Instead, the model we use here allows for punctuated measurements and resets of the quantum search algorithm. Nevertheless, the punctuated quantum search algorithm is faster on average. Note that we have assumed that the time it takes to measure, check if a solution was reached, and reset the algorithm is negligible. This is reasonable as checking a solution only requires one function call.

The punctuated quantum search algorithm has another advantage in that it is less sensitive to decoherence. If we wait until we have the maximal probability of measuring a target state, then we must maintain coherence for $1.5708/\phi$ steps as compared to only $1.1655/\phi$ steps for the fastest measure and restart method. This is because we do not need to maintain coherence through the measurement stage of this method. In fact, the punctuated search that takes the same number of steps on average as the standard or maximal probability method (i.e., $\pi/2\phi = 1.5708/\phi$ steps) need only maintain coherence for $\pi/4\phi = 0.7854/\phi$ steps at a time. This represents only 50% of the coherence time required in the standard Grover method, and corresponds to waiting for a 50% probability of success and then measuring.

V. *K*-PARALLEL QUANTUM SEARCH

A way to speed up Grover's algorithm still further is to have a society of k computational agents all running Grover's algorithm independently at the same time. This is promising because the standard deviation

$$\sigma_T = \frac{n}{p(n)}\sqrt{[1-p(n)]} \tag{33}$$

in the computation time of punctuated quantum search is fairly large, and hence having multiple searches running may give a considerable speed up.

Suppose that we know that there are exactly r solutions amongst N candidates. Given $p(n,r,N)$, the probability of success for a single agent after n iterations, we can boost the success probability by using k agents acting in parallel. In particular, the probability that at least one agent, in a society of k independent agents, succeeds after each agent has undergone n iterations is given by

$$p_k(n) = 1 - [1-p(n)]^k. \tag{34}$$

Thus the expected cost, $T_{avg}^{(k)}$, of performing a k-parallel quantum search is given by

$$T_{avg}^{(k)}(n) = \sum_{j=1}^{\infty} p_k(n)[1-p_k(n)]^{j-1} jn = \frac{n}{p_k(n)}$$

$$= \frac{n}{1-\cos^{2k}\left((1+2n)\frac{\phi}{2}\right)}. \tag{35}$$

As in Eq. (27) we can find the value of n that minimizes the expected cost. To find the mimimum, we find where $\partial T_{avg}^{(k)}(n)/\partial n$ is equal to zero. This derivative is given by

$$\frac{\partial T_{avg}^{(k)}(n)}{\partial n} = \frac{1-\cos^{2k}\left((1+2n)\frac{\phi}{2}\right)\left[1+2k\,n\phi\tan\left((1+2n)\frac{\phi}{2}\right)\right]}{\left[1-\cos^{2k}\left((1+2n)\frac{\phi}{2}\right)\right]^2}. \tag{36}$$

For $r/N \ll 1$, i.e., when there are very few solutions amongst the items searched over, we have $\phi = \arccos(1-2r/N) \approx 2\sqrt{r/N}$. As before, substituting $x \equiv n\phi \simeq (1+2n)\phi/2$ and realizing that $n \gg 1$, we obtain

$$\frac{\partial T_{avg}^{(k)}(n)}{\partial n} \approx \frac{1-\cos^{2k}(x)[1+2kx\tan(x)]}{[1-\cos^{2k}(x)]^2}. \tag{37}$$

In order to find the minimum, we thus have to solve the transcendental equation

$$1-\cos^{2k}(x) = 2k\,x\cos^{2k}(x)\tan(x). \tag{38}$$

The variable $x < 1$ provided $n < \frac{1}{2}(\sqrt{N/r}-1)$. We know that we can solve the problem with near certainty if we iterate Grover's algorithm to the maximum probability state in $O[(\pi/4)\sqrt{N/r}]$ iterations. Hence, for a large enough number of parallel search agents, k, there is a reasonable chance that the optimum number of iterations, $n_{opt}(r,N,k)$ at which the expected search cost is minimized, satisfies the criterion that $x < 1$. We therefore expand Eq. (37) as a series approximation in x about $x = 0$. Actually, it appears that x scales as $O(1/\sqrt{k})$, so it tends to 0 as k tends to infinity. If we make such an expansion up to order x^2, we get

$$\frac{\partial T_{avg}^{(k)}(n)}{\partial n} \simeq \frac{1}{kx^2}\left(-1+\frac{3k-1}{6}x^2+\frac{5k^2-1}{20}x^4+O(x^6)\right). \tag{39}$$

As $\partial T_{avg}^{(k)}(n)/\partial n = 0$ is a second-order equation in x^2, it can be solved analytically. Three of the roots are nonphysical but

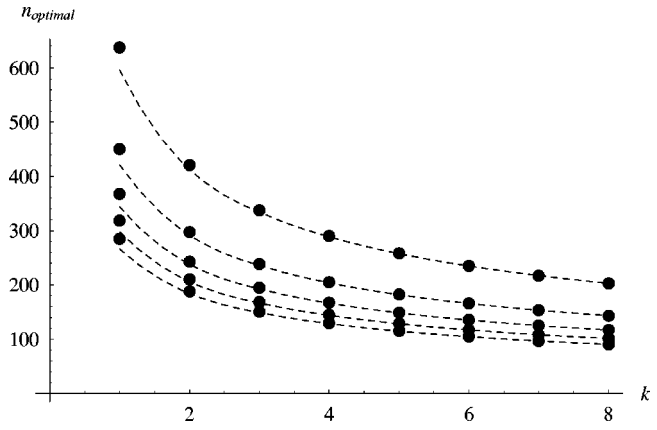

FIG. 2. Plot of the optimal number of iterations to use in k-parallel quantum search as a function of the degree of parallelism k for $r=1$ to $r=5$ solutions (top to bottom in the figure) for the case of a database of size $N=2^{20}$. The dashed curves correspond to the optima as predicted by our approximate formula for $n_{opt}(r,N,k)$. The points correspond to the exact optima obtained by numerical methods.

one corresponds to an approximation to the true minimum of $T_{avg}^{(k)}(n)$. Specifically, we find that $T_{avg}^{(k)}(n)$ is minimized when x is given by

$$x_{opt} \simeq \sqrt{\frac{5-15k+\sqrt{5}\sqrt{-31-30k+225k^2}}{-3+15k^2}}. \tag{40}$$

We note that $x<1$ for all $k\geqslant 2$, so that the derivation of the optimum formula is self-consistent. This expression for x_{opt} can be expanded in $1/k^{1/2}$, giving

$$x_{opt} \simeq 1.1118\frac{1}{k^{1/2}} + 0.0829\frac{1}{k^{3/2}} + O\left(\frac{1}{k^{5/2}}\right). \tag{41}$$

Using $\phi \simeq 2v = 2\sqrt{r/N}$ and Eq. (41), one gets the corresponding expression for $n_{opt}=x_{opt}/\phi$, i.e., the predicted optimal number of iterations for each of k quantum searches acting independently in parallel. In Fig. 2, this formula is shown to be in very good agreement with the exact result, obtained by numerical optimization.

Scaling of the parallel quantum search for large N and k

Now, if we are only interested in the scaling in N and k of the optimal number of iterations and expected computation time, it is enough to consider the expansion of $\partial T_{avg}^{(k)}(n)/\partial n$ [Eq. (39)] up to order $O(1)$. This simply yields

$$x_{opt} \simeq \sqrt{\frac{6}{3k-1}} \simeq O\left(\frac{1}{\sqrt{k}}\right). \tag{42}$$

This formula is only valid at the limit of large k, when x_{opt} tends to zero. The corresponding expression for the optimal number of amplitude amplification iterations is

$$n_{opt} = \frac{x_{opt}}{\phi} \simeq O\left(\frac{1}{\phi\sqrt{k}}\right). \tag{43}$$

We can then estimate the expected cost for an optimal k-parallel quantum search

$$T_{avg}^{(k)}(n,r,N) = \frac{1}{\phi}\frac{x_{opt}}{1-\cos^{2k}(x_{opt})}. \tag{44}$$

Again, using the series expansion around $x=0$, that is, $x/[1-\cos^{2k}(x)]=1/(kx)+O(x)$, we get

$$T_{avg}^{(k)}(n,r,N) \simeq \frac{1}{\phi}\frac{1}{kx_{opt}} \simeq O\left(\frac{1}{\phi\sqrt{k}}\right). \tag{45}$$

Remembering that $\phi \approx \sqrt{r/N}$, we conclude that T_{avg} scales as $O(\sqrt{N/rk})$. Thus, using k agents in parallel simply amounts to having each of them performing a search in a restricted space of size N/k, so that the gain in computation time is of order $O(\sqrt{k})$. Interestingly, this gain is not as good as when parallelizing a classical algorithm.[1] Accordingly, the cumulative time $T_{cumul}=kT_{avg}$, i.e., the sum of the time that all agents spend on quantum search, is *increased* by a factor $O(\sqrt{k})$ with respect to the case of a single agent ($k=1$).

Our results have implications for the design of prototypical quantum computers. If it is possible to maintain coherence indefinitely, for example, by building fault tolerance into the computer and by using quantum error correction schemes, our analysis suggests that it is better to use a single agent quantum search. This strategy minimizes the net computational resources expended in solving the problem. However, if coherence time *is* limited, as it most likely will be in prototypical quantum computers, then a parallel punctuated quantum search strategy becomes necessary, with the degree of parallelism set by the desired computation time and desired probability of success. The computational time can be made small by making the degree of parallelism sufficiently large but, of course, at the expense of greater net computational resources being expended on solving the problem.

Let us now consider the situation where the coherence time τ is fixed by some practical considerations, regardless the value of N and r. The number of agents k must then of the order of $O(N/r\tau^2)$ for the parallel time not to exceed the coherence time. This is an interesting result as it implies that the number of agents decreases *quadratically* for an increas-

[1]In the latter case, a computation time of order $O(N/r)$ is ideally reduced to $O(N/rk)$ by using k agents in parallel, so that one has a speedup of order $O(k)$.

ing τ. The classical counterpart would be a linear law only.[2] On the other hand, the bad result comes if we reexpress the cumulative computation time for k agents with this value of k:

$$T_{cumul} = kT_{avg} = k\tau = O\left(\frac{N}{r\tau}\right). \tag{46}$$

This means that we loose the square root speedup of Grover's algorithm (i.e., T_{cumul} does not scale as $\sqrt{N}$) whenever the coherence time is fixed. In order to exploit Grover's quantum speedup, the coherence time τ must necessarily increase as $\sqrt{N}$, i.e., as the square root of the size of the search space.

VI. CONCLUSIONS

In this paper we have shown how to generalize the analysis of an unstructured quantum search to incorporate the effects of an arbitrary starting superposition and an arbitrary unitary operator (or, equivalently, arbitrary state $|a\rangle$) *simultaneously*. We have also shown that, rather than iterating the amplitude amplification operator until the maximum probability of success is attained, i.e., after $O(0.785398\sqrt{N})$ iterations, it is better to measure after only $O(0.6900\sqrt{N})$ iterations (when the probability of success is only 84%). This punctuated strategy is approximately 12% faster than Grover's algorithm on average, and requires a shorter coherence time.

Moreover, an even better quantum search algorithm can be obtained by running k independent quantum searches in parallel, stopping as soon as any of the quantum searches finds a solution. We find that the optimal k-parallel punctuated quantum search strategy is different from that of single agent punctuated quantum search strategy. In general, the higher the degree of classical parallelism the less (parallel) time is needed to perform the quantum computation. This intuition is captured in Eq. (41), which gives the explicit connection between the optimal number of amplitude amplification iterations $n_{opt} = x_{opt}/\phi$ and the degree of parallelism k. This result is of practical utility to experimental realization of a quantum search algorithm. In particular, in any physical embodiment of a quantum search, there will be some natural coherence time beyond which the computation becomes unreliable. Of course, quantum error correction and fault-tolerant computation allow this time to be extended greatly, arguably indefinitively, if the individual error probability per gate operation can be made sufficiently small. Nevertheless, in practice, this might be extraordinarily difficult to achieve. Instead, if we can predict the degree of parallelism needed so that the quantum search has a good chance of completing within the natural coherence time of the physical system being used as the quantum computer, then the strategy of massive parallelism might provide a realistic alternative to relying solely on quantum error correction. Thus, we see the classical parallelism as an adjunct to quantum error correction rather than a replacement for it. Our results in Sec. V expose precisely the space/time tradeoff between quantum coherent computing and classical parallelism, at least in the context of unstructured quantum search.

Note. Some of the results obtained in this paper have been derived independently by Zalka in a revised (and unpublished [23]) version of Ref. [22].

ACKNOWLEDGMENTS

This work was supported by the President's Fund of the California Institute of Technology, the NASA/JPL Center for Integrated Space Microsystems, NASA Advanced Concepts Office, and the NASA Information and Computing Research Technologies Program.

[2]Classically, if the parallel computation time for each agent is restricted to τ, then the number of agents k should scale as $O(N/r\tau)$.

[1] N. Gershenfeld and I.L. Chuang, Science **275**, 350 (1997).
[2] D. G. Cory, M. D. Price, and T. F. Havel, Physica D **120**, 82 (1998).
[3] D.G. Cory *et al.*, e-print quant-ph/9802018.
[4] I.L. Chuang, N. Gershenfeld, M.G. Kubinec, and D.W. Leung, Proc. R. Soc. London, Ser. A **454**, 447 (1998).
[5] I.L. Chuang, L.M.K. Vandersypen, X. Zhou, D.W. Leung, and S. Lloyd, Nature (London) **393**, 143 (1998).
[6] J.A. Jones, M. Mosca, R.H. Hansen, Nature (London) **393**, 344 (1998).
[7] I.L. Chuang and Y. Yamamoto, Phys. Rev. A **52**, 3489 (1995); Phys. Rev. Lett. **76**, 4281 (1996); also available at http://xxx.lanl.gov/archive/quant-ph/9505011.
[8] J. D. Franson and T. B. Pittman, *An Optical Approach to Quantum Computing*, in Ref. [14].
[9] J.P. Dowling, Phys. Rev. A **57**, 4736 (1998).
[10] D. Deutsch and R. Jozsa, Proc. R. Soc. London, Ser. A **439**, 553 (1992).
[11] P. W. Shor, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science*, edited by S. Goldwasser (IEEE Computer Society Press, New York, 1994), pp. 124–134.
[12] L. K. Grover, *A Fast Quantum Mechanical Algorithm for Database Search*, Proceedings of the 28th Annual ACM Symposium on the Theory of Computing (ACM Press, New York, 1996), p. 212.
[13] N. J. Cerf, L. K. Grover, and C. P. Williams, available at http://xxx.lanl.gov/archive/quant-ph/9806078.
[14] A. Fijany and C. P. Williams, *Quantum Wavelet Transforms: Fast Algorithms and Complete Circuits*, in Proceedings of the First NASA Conference on Quantum Computing and Quantum Communications, Palm Springs, CA, 1998, Lecture Notes in Sci. Vol. 1509 (Springer-Verlag, New York, 1998); also available at http://xxx.lanl.gov/archive/quant-ph/9809004.
[15] G. Brassard, P. Hoyer, and A. Tapp, available at http://xxx.lanl.gov/archive/quant-ph/9805082.

[16] Richard Jozsa, Los Alamos e-print http://xxx.lanl.gov/archive/quant-ph/9901021.
[17] P. G. Kwiat, J. R. Mitchell, P. D. D. Schwindt, and A. G. White, Los Alamos e-print http://xxx.lanl.gov/archive/quant-ph/9905086.
[18] L. K. Grover, available as http://xxx.lanl.gov/archive/quant-ph/9711043.
[19] L.K. Grover, Phys. Rev. Lett. **80**, 4329 (1998).
[20] E. Biham, O. Biham, D. Biron, M. Grassl, and D. A. Lidar, available at http://xxx.lanl.gov/archive/quant-ph/9807027.
[21] M. Boyer, G. Brassard, P. Hoyer, and A. Tapp, Fortschr. Phys. **46**, 493 (1998); also available at http://xxx.lanl.gov/archive/quant-ph/9605034.
[22] C. Zalka, Los Alamos e-print archive http://xxx.lanl.gov/archive/quant-ph/9711070.
[23] C. Zalka (private communication).
[24] Lov K. Grover, Los Alamos e-print http://xxx.lanl.gov/archive/quant-ph/9809029.
[25] C.H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, in *Strengths and Weaknesses of Quantum Computing*, in special issue of SIAM J. Comput. **26**, 1510 (1997); available at http://xxx.lanl.gov/archive/quant-ph/9701001.