Implementation of multipartite unitary operations with limited resources

Dominic W. Berry
Department of Physics, The University of Queensland, Brisbane, Queensland 4072, Australia
and Centre for Quantum Computer Technology, Macquarie University, Sydney, NSW 2109, Australia
(Received 29 August 2006; revised manuscript received 10 November 2006; published 30 March 2007)

A general method for implementing multipartite unitary operations that are close to the identity using a small amount of entanglement and classical communication is presented. For cases where methods were previously known, our approach requires less entanglement. In addition, it is shown that, via compression, the average communication required from all parties except one may be reduced to the same as the entanglement. This improves upon previous methods which required a large amount of communication that did not scale down with the strength of the interaction.

DOI: 10.1103/PhysRevA.75.032349 PACS number(s): 03.67.Hk, 03.65.Ud

I. INTRODUCTION

Much of quantum information processing relies on the ability to perform operations between subsystems. In the case where these subsystems do not directly interact, it is necessary to achieve the operation by indirect means. For example, one could use a third subsystem which interacts with both, or use entanglement. In the case where the operation is only weakly entangling, it would be useful to be able to implement it using only a small amount of entanglement.

A SWAP operation may be applied using two teleportations [1], requiring two ebits of entanglement and two bits of communication in each direction. Similarly, a controlled NOT (CNOT) operation may be applied using one ebit and one bit of communication in each direction [2]. More general two-qubit unitaries may be implemented using a small amount of entanglement when they are weakly entangling [3,4]. Reference [3] gives a scheme to implement operations of the form $e^{-i\alpha\sigma_x\otimes\sigma_x}$ using average entanglement for small α of 5.9793α. This is considerably more than the entanglement of 1.9123α that these operations can be used to create [5], suggesting that this scheme can be improved upon.

References [3,4] showed how to use the scheme for $e^{-i\alpha\sigma_x\otimes\sigma_x}$ to implement general two-qubit unitaries. A simple generalization to multipartite operations of the form $e^{-i\alpha\sigma_x^{\otimes N}}$ was given in Ref. [4]. Reference [6] presented a range of schemes to implement restricted classes of operations. Alternative schemes for operations of the form $e^{-i\alpha\sigma_{n_A}\otimes\sigma_{n_B}}$, but allowing a significant probability of failure, are presented in Refs. [7,8]. A scheme for Hamiltonians that are a tensor product of self-inverse operators was given in Ref. [9].

Here we present a general method of implementing weakly entangling multipartite unitary operations using a small amount of entanglement or communication. For the $e^{-i\alpha\sigma_x\otimes\sigma_x}$ operation as considered by Refs. [3,4], the entanglement consumption is reduced. The entanglement consumption is dramatically reduced for some values of α, and is also reduced in the limit of small α. This method also reduces the average communication required from all parties except one to be the same as the entanglement. In contrast, the method of Refs. [3,4] requires a large amount of communication from all parties, even in the limit of a weak interaction.

This paper proceeds as follows. The general problem of implementing operations using entanglement and communication is discussed in Sec. II. The improved method for implementing evolution under Hamiltonians of the form $\sigma_z^{\otimes N}$ is presented in Sec. III. A method of using compression over multiple operations to reduce the communication required is given in Sec. IV. A method of applying these results to general Hamiltonians is given in Sec. V, and an alternative approach is given in Sec. VI. Conclusions are given in Sec. VII.

II. GENERAL PROBLEM

The situation considered is that there are N parties, each with a number of subsystems, and operations performed by the parties are assumed to have negligible cost. The unitary we wish to implement acts upon one subsystem in possession of each of the parties, and we quantify the entanglement and classical communication between the parties.

The most general method of implementing a multipartite unitary operation, U, using entanglement and classical communication may be described as follows. The N parties share an entangled resource state and perform many rounds of the following process. Each party performs a local unitary followed by a partial measurement. The parties then communicate classical information about the measurement results and repeat the process with local unitaries and measurements possibly depending on the transmitted information.

There are two scenarios that one may consider. One may consider a single implementation of the unitary, and consider the expectation value of the resources required. That is, for a probabilistic protocol, one may require different resources for different measurement results, and one weights them according to their probability of being required. The scheme of Cirac, Dür, Kraus, and Lewenstein (CDKL) [3,4] quantifies resources in this way.

Alternatively, one may consider many implementations, M, of the multipartite unitary, but with a fixed set of resources. To take account of the possibility that the resources are insufficient, one may allow implementation with nonunit fidelity. More formally, let us denote the entangled resource state used for implementation of $U^{\otimes M}$ by $|\psi\rangle_M$, and the map on the target and resource state by $\mathcal{E}_M$. This map corresponds to the average over all possible measurement results in a probabilistic protocol.

Given a measure of multipartite entanglement $\vec{E}$ and mea-

sure of classical communication $\vec{C}$, we introduce the *sufficient rate vector* $\vec{R}$. This is a vector giving the average entanglement and communication that is sufficient to implement the unitary U. The sufficient rate vector is closely related to the concept of achievable rates for operations [10], and it is clear that the achievable rates cannot be larger than the sufficient rates.

We define the sufficient rate vector in the following way. For all $\epsilon, \delta > 0$, there exists a scheme for implementation of $U^{\otimes M}$ using communication $\vec{C}$ such that $[\vec{C}, \vec{E}(|\psi\rangle_M)] \leq M\vec{R}(1+\delta)$ (where the brackets denote the concatenation of the vectors), and for all target states, ρ,

$$F\{\mathrm{Tr}_{\mathrm{res}}[\mathcal{E}_M(\rho \otimes |\psi\rangle_M)], U^{\otimes M}\rho(U^{\dagger})^{\otimes M}\} \geq 1-\epsilon. \quad (1)$$

Here F denotes the fidelity, and the partial trace removes the component of the system corresponding to the resource state. Note that the use of δ in the definition here differs from the capacity definitions in [10]. We use δ to make the sufficient rate vector more consistent with the single-copy case.

It is easily seen that the expected resource consumption for the single-copy scheme corresponds to a sufficient rate vector. If the expected resource consumption for the single-shot scheme is $\vec{R}$, then for all $\delta > 0$, the probability of the resources required for implementation of $U^{\otimes M}$ exceeding $M\vec{R}(1+\delta)$ may be made arbitrarily small for large M. We denote the probability of the resources required exceeding those provided by p_{fail}, and the corresponding output state by ρ_{fail}. The output state after discarding the measurement results is then

$$\rho_{\mathrm{out}} = p_{\mathrm{fail}}\rho_{\mathrm{fail}} + (1-p_{\mathrm{fail}})U^{\otimes M}\rho(U^{\dagger})^{\otimes M}. \quad (2)$$

The fidelity between ρ_{out} and the desired output state is no less than $1-p_{\mathrm{fail}}$. As p_{fail} may be made arbitrarily small for large M, $\vec{R}$ must be a sufficient rate vector. It is clear that analogous results hold for the expected resource consumption for multiple implementations of U. That is, if the expected resource consumption for implementation of $U^{\otimes M}$ is $\vec{R}$, then $\vec{R}/M$ is a sufficient rate vector.

In this work we focus on the asymptotic case, as this is required by the method used for reduction of the average communication. As the expected resource consumption for the single-copy case corresponds to sufficient rate vectors for the asymptotic case, our results are directly comparable to those of Cirac, Dür, Kraus, and Lewenstein (CDKL) [3,4]. The scheme presented here goes beyond the CDKL scheme in three ways.

(1) The entanglement consumption is reduced.

(2) The communication in most directions is close to the entanglement consumption.

(3) It applies to general multipartite Hamiltonians on multilevel systems.

III. ENTANGLEMENT CONSUMPTION

We initially consider the Hamiltonian $\sigma_z^{\otimes N}$ for N parties, so the unitary we wish to apply is $U(\alpha)=\exp(i\alpha\sigma_z^{\otimes N})$. This Hamiltonian is equivalent, under local unitaries, to the Hamiltonian considered in Sec. IX of Ref. [4]. Here we give an improved scheme which reduces the entanglement consumed. In this section we give the scheme in terms of a single implementation; from Sec. II, the results also apply to the asymptotic case.

The N parties use the resource state

$$|\psi(\beta)\rangle = \cos(\beta)|0\rangle^{\otimes N} + i\sin(\beta)|1\rangle^{\otimes N}. \quad (3)$$

This resource state is more general than that used in Refs. [3,4], in that arbitrary β is allowed. It is this generality that allows us to improve upon the entanglement consumption. Also note that this state requires half as many qubits.

In the following we use the "stator" notation, introduced in Ref. [6]. A stator represents the action of applying an operator and appending a state. A general stator is of the form $\Sigma_k|\phi_k\rangle\otimes U_k$, and acts on state $|\psi\rangle$ as

$$\left(\sum_k |\phi_k\rangle \otimes U_k\right)|\psi\rangle = \sum_k |\phi_k\rangle \otimes (U_k|\psi\rangle). \quad (4)$$

The parties now apply a four step process.

Step 1. Each party applies a controlled-Z operation between the entangled resource and their component of the target system. This yields the stator

$$\cos(\beta)|0\rangle^{\otimes N}\otimes \mathbb{1}^{\otimes N} + i\sin(\beta)|1\rangle^{\otimes N}\otimes\sigma_z^{\otimes N}. \quad (5)$$

Step 2. Parties 1 to $N-1$ apply Hadamard operations followed by computational basis measurements on their component of the resource state. After removing the measured qubits, the resulting stator is

$$\cos(\beta)|0\rangle\otimes\mathbb{1}^{\otimes N} \pm i\sin(\beta)|1\rangle\otimes\sigma_z^{\otimes N}, \quad (6)$$

where the sign is $+$ if the number of measurement results equal to 1 is even, and $-$ otherwise.

Step 3. Party N performs a correction on their component of the resource state based on the measurement results. If the number of measurement results equal to 1 is odd, then they perform a σ_z operation on the remaining portion of the resource state. We then obtain the stator with the plus in Eq. (6).

Step 4. Party N applies a projection measurement on their component of the resource state. Projection onto the state $|\phi\rangle=\cos(\gamma)|0\rangle+\sin(\gamma)|1\rangle$ yields the unitary operation proportional to

$$\cos(\beta)\cos(\gamma)\mathbb{1}^{\otimes N} + i\sin(\beta)\sin(\gamma)\sigma_z^{\otimes N}. \quad (7)$$

For failure, projection onto the state $|\phi^{\perp}\rangle=\sin(\gamma)|0\rangle - \cos(\gamma)|1\rangle$ is obtained, giving an operation proportional to

$$\cos(\beta)\sin(\gamma)\mathbb{1}^{\otimes N} - i\sin(\beta)\cos(\gamma)\sigma_z^{\otimes N}. \quad (8)$$

The measurement in Step 4 is chosen such that, for the correct measurement result, the unitary $U(\alpha)$ is applied. This will be obtained provided

$$\tan(\beta)\tan(\gamma) = \tan(\alpha). \quad (9)$$

For the incorrect measurement result, the unitary is of the form $U(\alpha')$, but with $\alpha' \neq \alpha$. To obtain the correct result, we now wish to implement the unitary $U(\alpha-\alpha')$. Combined

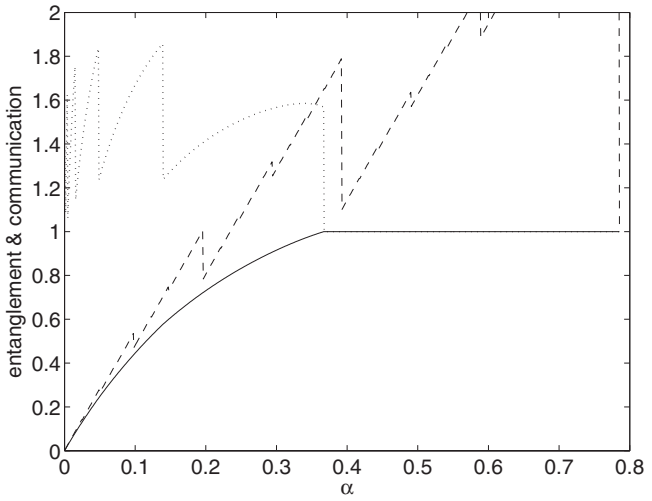

FIG. 1. The average entanglement consumed to implement the operation $U(\alpha)$ as a function of α. The case where the resource states have been numerically optimized is shown as the solid line, and the case for the scheme of Ref. [3] is shown as the dashed line. The communication from party N for the numerically optimized scheme is shown as the dotted line. The numerically optimized scheme improves on the entanglement consumption for the scheme of Ref. [3] for all nontrivial values of α.

with the previous incorrect unitary $U(\alpha')$, this will result in the desired $U(\alpha)$. To achieve this, one requires a new resource state $|\psi(\beta')\rangle$ (where β' is not in general equal to β) and repeats the process.

If the incorrect measurement result is obtained repeatedly, one can obtain the correct unitary by using the resource state $(|0\rangle^{\otimes N}+|1\rangle^{\otimes N})/\sqrt{2}$. One can deterministically perform any unitary of the form $U(\alpha)$ with this resource state. To see this, consider steps 1–4 with $\beta=\pi/2$ and $\gamma=\alpha$. Then, in the case of failure one simply applies the local unitary $U(\pi/2)=i\sigma_z^{\otimes N}$.

In the following the terminology "stage" is used to mean steps 1–4. The notation β_l is used for the parameter for the entangled resource state used in stage l (i.e., after $l-1$ failures). In addition, α_l is used for the parameter for the unitary we wish to implement in stage l. That is, $\alpha_1=\alpha$, $U(\alpha_2)$ is the unitary we require for the correction after one failure, and so forth. We also use L for the maximum total number of stages.

If $\alpha=\pi/2^n$ and $\beta_l=\alpha_l$, this method is similar to that of CDKL. At each stage the probability of success is equal to 1/2, and for repeated failure we need to implement $U(\pi/2)$, which may be performed locally. The resource state is different from that used by CDKL, but the average entanglement consumed is identical.

To improve upon the entanglement consumption we adjust the entanglement of the resource states. By using resource states with larger entanglement, the probability for success is increased, and the average entanglement consumed is reduced. It is straightforward to numerically optimize for the resource states that minimize the average entanglement consumed for a given α. The resulting average entanglement is plotted in Fig. 1 as a function of α. For the results presented here the inbuilt MATLAB minimization routine "fminsearch" was used to minimize the average entanglement consumption as a function of the β_l.

The entanglement consumed for the CDKL scheme is also shown in Fig. 1. For values of α not of the form $\alpha=\pi/2^n$, an expansion as in Eq. (11) of Ref. [4] was used. For all values of α (except the trivial $\alpha=0$) the entanglement consumed for the numerically optimized scheme is less than that for the CDKL scheme. In some cases this improvement is dramatic.

In the limit of small α the entanglement consumed is approximately 5.6418α, as compared to 5.9793α for the CDKL scheme [13]. This improvement is not as dramatic, but it is surprising as our initial investigations suggested that the asymptotic entanglement consumption could not be reduced below that for the CDKL scheme.

To show that the entanglement consumed is approximately 5.6418α in the limit of small α, we may use the following approach. First we determine the entanglement required to implement $U(A)$ for values of A over a range of a factor of 2. Here we take $A\in[\pi/2^{20},\pi/2^{19})$. Now to implement $U(\alpha)$ for $\alpha<\pi/2^{20}$, we select an $A\in[\pi/2^{20},\pi/2^{19})$ such that $\alpha=A/2^n$. Then, rather than numerically optimizing for the best intermediate entangled states, we select the intermediate entangled states with $\beta_l=\alpha 2^{l-1}$ for $1\leqslant l\leqslant n$, so $\beta_l=\alpha_l$. That is, up to the point where the correction required is A, we use the scheme with success probability 1/2 for each measurement. If the correction $U(A)$ is needed, it may be achieved using the numerically optimized scheme.

Via this approach, the average entanglement consumed to implement $U(\alpha)$ is

$$\mathcal{E}(\alpha)=\mathcal{E}(A)2^{-n}+\sum_{l=1}^{n}2^{1-l}E(\alpha_l), \tag{10}$$

where $\mathcal{E}(\alpha)$ is the entanglement consumed to implement $U(\alpha)$, and $E(\alpha)$ is the entanglement of the resource state $|\psi(\alpha)\rangle$. We have $E(\alpha)=h[\sin^2(\alpha)]$ where h is the binary entropy function $h(p)=-p\log_2(p)-(1-p)\log_2(1-p)$.

This entanglement measure is the entropy of the reduced density operator for one subsystem. This is a consistent entanglement measure for Schmidt decomposable multipartite pure states. It is possible to apply similar methods to the bipartite case to perform entanglement concentration and dilution between these states and the standard states $(|0\rangle^{\otimes N}+|1\rangle^{\otimes N})/\sqrt{2}$ [11].

The first term in Eq. (10) is the probability for n failures (2^{-n}) multiplied by the entanglement required to implement $U(A)$. The sum is the entanglement of the state with $\beta_l=\alpha_l$ for stages 1 to n multiplied by the probability. Equation (10) gives the upper bound

$$\frac{\mathcal{E}(\alpha)}{\alpha}<\frac{\mathcal{E}(A)+\sum_{k=1}^{\infty}2^kE(A2^{-k})}{A}. \tag{11}$$

The expression on the right-hand side (RHS) of Eq. (11) is independent of n, and only depends on A. This expression is plotted as a function of A in Fig. 2. It can be seen that this

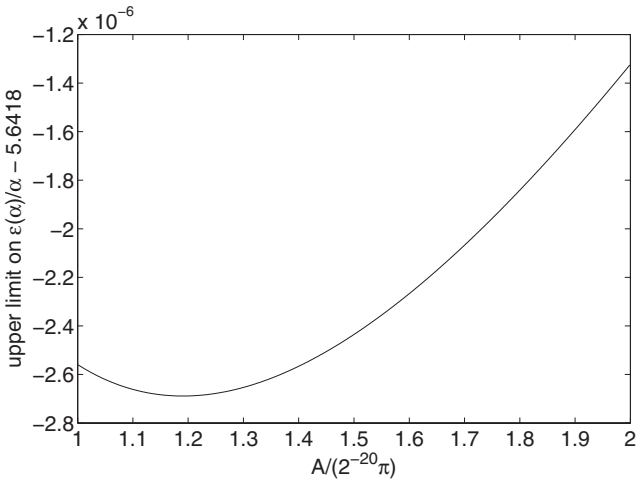

FIG. 2. The upper limit on $\mathcal{E}(\alpha)/\alpha$ [the RHS of Eq. (11)] minus 5.6418 as a function of A.

expression does not exceed 5.6418 for this range of A. Thus we find that, in the limit of small α, the entanglement consumed need not exceed 5.6418α.

As an alternative proof which does not depend on Fig. 2, we may expand α as $\alpha=\Sigma_n a_n \pi/2^n$ with $a_n \in \{0,1\}$. For $\alpha < \pi/2^{20}$, each $U(\pi/2^n)$ may be implemented using the approach given above with average entanglement no larger than $5.6418\pi/2^n$ [using the value of the upper limit (11) for $A=\pi/2^{20}$]. Hence the average total entanglement is no more than 5.6418α.

IV. CLASSICAL COMMUNICATION

The next issue to consider is that of the classical communication required. We consider implementation of $U^{\otimes M}$, so at stage l the set of entangled resource states is of the form

$$|\Psi_l\rangle = [\cos(\beta_l)|0\rangle^{\otimes N} + i\sin(\beta_l)|1\rangle^{\otimes N}]^{\otimes M_l}, \tag{12}$$

where M_l is the number of entangled states used in stage l. Let $p(l)$ be the probability of requiring the entangled resource state in stage l of the single-shot protocol (i.e., the probability of $l-1$ failures). For small δ we take $M_l=\lfloor Mp(l)(1+\delta)\rfloor$. The qubits used for this state are depicted in Fig. 3. We may alternatively express this state as

$$|\Psi_l\rangle = \sum_{\mathbf{i}} \mu_{\mathbf{i}} |\mathbf{i}\rangle^{\otimes N}, \tag{13}$$

where $\mathbf{i}=(i_1,\ldots,i_{M_l})$, $\mu_{\mathbf{i}}=\Pi_m \sin^{i_m}(\beta_l)\cos^{1-i_m}(\beta_l)$ and $|\mathbf{i}\rangle = |i_1\rangle\otimes\cdots\otimes|i_{M_l}\rangle$.

Now we make an approximation to the state $|\Psi_l\rangle$ by retaining only typical sequences of the i. Denoting the set of typical $\mathbf{i}$ by S_l, the approximate state is

$$|\tilde{\Psi}_l\rangle \propto \sum_{\mathbf{i}\in S_l} \mu_{\mathbf{i}} |\mathbf{i}\rangle^{\otimes N}. \tag{14}$$

It is a standard result for typical sequences that, for all $\epsilon,\delta>0$, there exists an M_l such that the fidelity is at least $1-\epsilon/2L$ with the number of elements in S_l no more than $2^{M_l E(\beta_l)(1+\delta)}$. The entanglement of the state $|\tilde{\Psi}_l\rangle$ does not exceed $M_l E(\beta_l)(1+\delta)$, whereas it is $M_l E(\beta_l)$ for the exact state. Therefore, in the limit of small δ, we do not increase the entanglement consumption by using this approximate state.

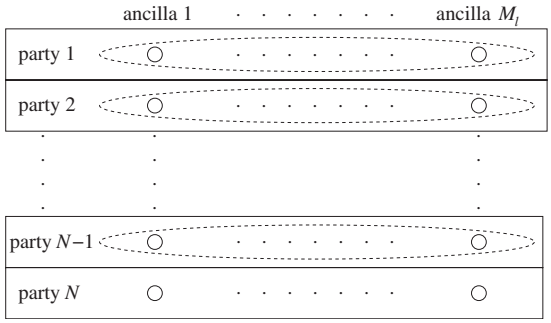

FIG. 3. A diagram of the ancilla qubits held by the N parties. The qubits are represented by circles, and the horizontal rectangles show the qubits held by each party. The columns correspond to each resource state $\cos(\beta_l)|0\rangle^{\otimes N}+i\sin(\beta_l)|1\rangle^{\otimes N}$ used for independent implementations of $U(\alpha)$. The dashed ellipses show the qubits upon which each party performs their joint measurement in the basis (15).

Because fidelity does not decrease under completely positive trace-preserving (CPTP) maps, the fidelity between the final state and the output for the exact state must be at least $1-\epsilon/2L$. In the case where we use approximate resource states with typical sequences at each stage, the final output must have fidelity at least $1-\epsilon/2$ with the state obtained with the exact resource states, ρ_{out}. Provided M is sufficiently large that the probability of the resources required exceeding those provided, p_{fail}, does not exceed $\epsilon/2$, then the fidelity with the desired state, ρ, is at least $1-\epsilon$.

Now, to reduce the average communication required, we replace Step 2 with a joint measurement in the Fourier transform basis on the states $|\mathbf{i}\rangle$ for typical sequences by each party from 1 to $N-1$. That is, each party performs a measurement in the basis

$$|j\rangle = e^{i2\pi B(\mathbf{i})j/\|S_l\|}|\mathbf{i}\rangle, \tag{15}$$

where B is a function which assigns unique integers from 0 to $\|S_l\|-1$ for each $\mathbf{i}\in S_l$. The qubits upon which each party performs this joint measurement are depicted in Fig. 3 [14].

Denoting the measurement result obtained by party n by j_n, it is necessary for party N to perform the phase correction

$$\sum_{\mathbf{i}\in S_l} e^{i2\pi B(\mathbf{i})\Sigma_n j_n/\|S_l\|}|\mathbf{i}\rangle\langle\mathbf{i}|. \tag{16}$$

In order for party N to be able to perform this correction, it is necessary for party N to be able to determine the sum $\Sigma_n j_n$ (modulo $\|S_l\|$). There are many different ways that this information may be communicated to party N. For example, party n may communicate j_n directly to party N. Alternatively, party 1 may communicate j_1 to party 2, then party 2 may communicate $j_1+j_2 \bmod \|S_l\|$ to party 3, and so forth. The communication from each party $1,\ldots,N-1$ does not exceed $M_l E(\beta_l)(1+\delta)$ bits. Using the expression for M_l, this does not exceed $Mp(l)E(\beta_l)(1+\delta)^2$. The average total communi-

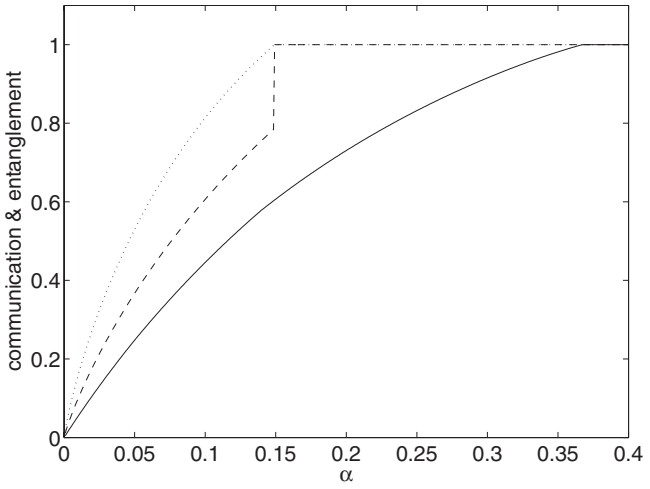

FIG. 4. Results for the scheme optimized for minimum classical communication from the last party. This communication is shown as the dotted line, and the corresponding entanglement required is shown as the dashed line. The entanglement for the entanglement-optimized scheme is shown as the solid line for comparison.

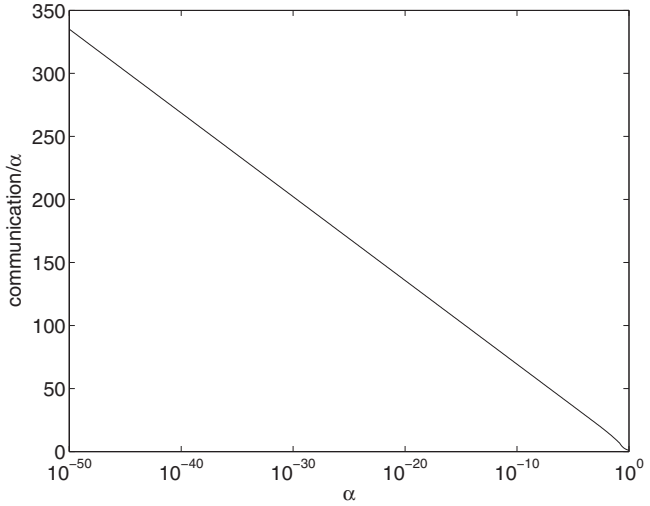

FIG. 5. The ratio of the communication from the final party to the interaction strength α.

cation from each party $1,\ldots,N-1$ does not exceed $\Sigma_l p(l)E(\beta_l)(1+\delta)^2=\mathcal{E}(\alpha)(1+\delta)^2$. Thus in the limit of small δ, this communication is the same as the entanglement consumed.

After the phase correction has been performed, we simply perform Step 4 as above. Now it is necessary to communicate the positions of the measurement successes from party N to the other parties. The probability of failure at this step is $p(l+1)/p(l)$, and we can achieve some reduction of the communication required using compression. The communication required is $M_l h[p(l+1)/p(l)](1+\delta)$, which does not exceed $Mp(l)h[p(l+1)/p(l)](1+\delta)^2$.

In general, the probability of failure, $p(l+1)/p(l)$, is given by

$$\cos(2\beta_l)\sin^2\gamma_l+\sin^2\beta_l, \tag{17}$$

where γ_l is the parameter used for the projection measurement in stage l. Except for deterministic implementation with $\beta_l=\pi/4$, we have $\beta_l\leqslant\pi/4$, so $\cos(2\beta_l)>0$, and $p(l+1)/p(l)$ is greater than $\sin^2\beta_l$. This implies that $h[p(l+1)/p(l)]$ is greater than $E(\beta_l)$, so the communication required from the last party is larger than the entanglement. When we determine the total communication from party N for the numerically optimized scheme, we find that it is much larger than the entanglement, and does not scale down with α (see Fig. 1).

It is also possible to select the entangled resources to minimize the average communication required from the last party. In this case numerical results indicate that the optimal scheme uses an entangled resource with $\tan(\beta)=\tan^2(\alpha)$. For small α, the probability of failure is small. In the case of failure we simply implement the scheme with 1 ebit of entanglement and 1 bit of communication in each direction. The corresponding classical communication is shown in Fig. 4.

The classical communication is still larger than the entanglement for the scheme optimized for minimum entanglement consumption. However, it does go to zero in the limit of small α. Unfortunately the ratio of the classical communication to α is not bounded for small α (see Fig. 5). The ratio increases approximately as $|\log\alpha|$. This means that a scheme based on breaking the evolution up into a large number of short time intervals will require a large amount of communication from the last party.

V. GENERAL HAMILTONIANS

In the previous sections we have shown how to implement evolution under a Hamiltonian of the form $\sigma_z^{\otimes N}$. Now we show how to implement a general tensor product Hamiltonian of the form $H=H_1\otimes\cdots\otimes H_N$. To do this, we apply a method similar to that in Ref. [12]. First note that it is possible to diagonalize the H_j via local unitaries. Taking $\Delta=\|H\|$, the diagonalized form of H is

$$H_{\mathrm{diag}}=\Delta\bigotimes_j \mathrm{diag}(a_{1,j},a_{2,j},a_{3,j},\ldots,a_{d_j-1,j},a_{d_j,j}), \tag{18}$$

where $a_{l,j}\in[-1,1]$, and d_j is the dimension of the subsystem which H_j acts upon. In the following we consider simulation of this diagonalized form.

This Hamiltonian may be simulated by the Hamiltonian $\sigma_z^{\otimes N}$ in a similar way as in Ref. [12]. The chain of simulations used is

$$\begin{aligned}\sigma_z^{\otimes N}&\mapsto H''_{A_1B_1}\otimes\sigma_z^{\otimes N-1}\\&\mapsto H''_{A_1B_1}\otimes H''_{A_2B_2}\otimes\sigma_z^{\otimes N-2}\\&\ \ \vdots\\&\mapsto\bigotimes_j H''_{A_jB_j}\mapsto H_{\mathrm{diag}}/\Delta,\end{aligned} \tag{19}$$

where $H''_{A_jB_j}=\mathrm{diag}(a_{1,j},-a_{1,j},\ldots,a_{d_j,j},-a_{d_j,j})$. We use the A_j subscripts to indicate the subsystems upon which we wish to

implement H, and B_j to indicate ancilla subsystems.

In order to perform this chain of simulations, we need to, in general, perform the simulation

$$\bigotimes_{n=1}^{j-1} H''_{A_nB_n} \otimes \sigma_z^{\otimes(N-j+1)} \mapsto \bigotimes_{n=1}^{j} H''_{A_nB_n} \otimes \sigma_z^{\otimes(N-j)}. \quad (20)$$

To do this, we first append the d_j dimensional ancilla B_j so the jth term in the tensor product is diag$(1,-1,1,-1,\ldots,1,-1)$. Now we take $p_l=(a_{l,j}+1)/2 \in [0,1]$, and define the local unitaries U_l which exchange the $(2l-1)$th and $(2l)$th basis vectors of A_jB_j. To simulate $\bigotimes_{n=1}^{j} H''_{A_nB_n} \otimes \sigma_z^{\otimes(N-j)}$ for small time δt, we simply apply U_l at time $p_l\delta t$ and again at time δt. In this way, we may apply each successive simulation in the chain (19). The final simulation in the chain may be achieved by simply restricting to the appropriate subspace.

This method is a generalization of the method for bipartite tensor product Hamiltonians given in Ref. [12]. In Ref. [12] the simulation is reversible. This is because the diagonal Hamiltonian may be expanded as a sum of local Hamiltonians and a Hamiltonian for which the maximum and minimum eigenvalues have the same magnitude. The local Hamiltonians may be ignored because they can be implemented locally. One can then use H_{diag}/Δ to simulate $\sigma_z^{\otimes N}$ by restricting to the subspace for the maximum and minimum eigenvalues. In the multipartite case we cannot use a similar method because the additional terms will be multipartite Hamiltonians on $N-1$ subsystems, rather than local Hamiltonians.

The case of self-inverse Hamiltonian evolution, as in Ref. [9], is particularly simple. All the $a_{l,j}$ are ± 1, so it is possible to achieve the entire simulation chain without intermediate unitaries simply by restricting to the appropriate subspaces. Thus self-inverse Hamiltonian evolution is equivalent to evolution under $\sigma_z^{\otimes N}$.

In order to implement evolution under H for a time interval t, we divide the time interval into m intervals of length t/m. For m sufficiently large, the chain of simulations above will be accurate. The expectation value for the consumed entanglement is then no more than $5.6418t\|H\|$. In fact, the expectation value is approximately equal to this, because each simulation of $U(\alpha)$ is for $\alpha \ll 1$. As above, the asymptotic average communication from each of the parties $1,\ldots,N-1$ may be made equal to this value.

Implementing the Hamiltonian for a large number of short time intervals does not cause problems in taking the appropriate limits to obtain the average entanglement and communication. We may choose the number of intervals, m, to be large enough to obtain fidelity (ignoring sources of error other than the finite time interval) of $1-\epsilon/2$. Then we choose the number of copies, M, to be large enough that the evolution over each individual interval is implemented with fidelity at least $1-\epsilon/2m$. The overall fidelity is then at least $1-\epsilon$.

However, the large number of time intervals does cause the communication required from party N to diverge. As discussed in Sec. IV, the communication from this party does not scale down with the length of the time interval. As the number of intervals m must approach infinity in the limit of high fidelity, the average communication from party N diverges. Even if one optimizes for minimal communication at each step (as in Fig. 4), the communication required for an interval of length δt scales as $|\log(\delta t)|\,\delta t$, so the communication from party N scales as $\log m$. As we require $m \to \infty$, the communication from party N still diverges in the limit of high fidelity.

We can also use the above methods to implement evolution under general multipartite Hamiltonians on multilevel systems. Any Hamiltonian may be expressed in the form $H=\Sigma_k H_k$, where each H_k is a tensor product Hamiltonian. We then simply implement $(\Pi_k e^{-itH_k/m})^m$. For large m, this simulation may be made arbitrarily accurate. The average entanglement consumed and communication required (except for communication from one of the parties) is approximately $5.6418t\Sigma_k\|H_k\|$. General Hamiltonians will require a large number of terms in the sum, so $\Sigma_k\|H_k\|$ may be much larger than $\|H\|$. However, regardless of the number of terms in the sum, the Hamiltonian may be simulated using resources scaling linearly in t (except for the communication from one of the parties).

In general, any weakly entangling unitary is equivalent to evolution under a Hamiltonian for small t. It is straightforward to show that a weakly entangling unitary can be simplified by local unitaries to be close to the identity. Any unitary is equivalent to evolution under a Hamiltonian; if the unitary is close to the identity, then the evolution is for short time. Therefore the above results show that a weakly entangling unitary can be implemented with a small amount of entanglement and communication (except from one party).

VI. ALTERNATIVE METHOD

Now we present an alternative method which, although not as efficient in terms of entanglement consumption, addresses the problem of the large amount of communication from the final party being required. We may express a general multipartite unitary in the form

$$U = \sum_k \lambda_k V_k^{(1)} \otimes V_k^{(2)} \otimes \cdots \otimes V_k^{(N)}, \quad (21)$$

where the $V_k^{(j)}$ are local unitaries. To implement this operation, we use an entangled resource state of the form

$$|\Psi_U\rangle = \sum_k \mu_k |k\rangle_{B_1} \otimes \cdots \otimes |k\rangle_{B_N}. \quad (22)$$

We use B_j for the subsystems for the resource state, and A_j for the subsystems for the system state (which we wish to apply U to).

Each party j performs the controlled operation $\Sigma_k |k\rangle_{B_j}\langle k| \otimes V_k^{(j)}$, where $V_k^{(j)}$ acts upon A_j. As a result, we have the stator

$$\sum_k \mu_k |k\rangle_{B_1} \otimes V_k^{(1)} \otimes \cdots \otimes |k\rangle_{B_N} \otimes V_k^{(N)} \quad (23)$$

acting upon the system.

Parties 1 to $N-1$ perform local measurements in the Fourier transform basis on their components of the resource state. After communicating the results to party N, and a phase correction by party N, the stator is

$$\sum_k \mu_k V_k^{(1)} \otimes \cdots \otimes |k\rangle_{B_N} \otimes V_k^{(N)}. \tag{24}$$

Party N then performs a projective measurement on their component of the resource state, with success corresponding to projection onto the state $\Sigma_k \nu_k |k\rangle$. Choosing $\mu_k \nu_k^* \propto \lambda_k$, U is implemented in the case of success.

In the case of failure a different unitary operation is performed. This case may be corrected for by teleporting the portions of the state in subsystems A_1 to A_{N-1} to party N, performing the appropriate unitary, then teleporting these portions of the state back to their respective subsystems.

This approach may be used to address the issue of the large amount of communication being required from one of the parties. As discussed in Sec. V, the communication required from the last party is not bounded for the first approach. On the other hand, using the approach given in this section, we can ensure that all resources required approach zero in the limit as U approaches the identity. Let us choose the operator basis such that $V_0^{(j)}=\mathbb{1}$. Then, in the limit as U approaches the identity, λ_0 approaches 1 and the other λ_k approach 0. Now we select $\mu_k \propto \nu_k \propto \sqrt{\lambda_k}$, and use a two stage process where we use the resource state $|\Psi_U\rangle$ in the first stage, and if there is failure then we simply implement the required operation deterministically using teleportation.

In the limit $U \to \mathbb{1}$, μ_0 approaches 1, so the entanglement of the resource state $|\Psi_U\rangle$ approaches 0. As in Sec. IV we may use compression such that the communication required from parties 1 to $N-1$ is close to this entanglement in the first stage. Using this choice for μ_k we find that the probability of failure approaches zero in the limit $U \to \mathbb{1}$. The contribution to the resource requirements from the second stage therefore approaches zero. In addition, it is possible to use compression to reduce the average communication required for the positions of the failures. In the limit as the probability of failure approaches zero, this communication requirement approaches zero.

Thus we see that via this approach, all the resource requirements: entanglement and communication from each of the parties, can be made to approach zero in the limit as the unitary approaches the identity. This approach is less efficient than the first approach in terms of most of the resources. The resource consumption is not linear in the effective evolution time, and instead scales in a similar way as in Fig. 4. The advantage of this approach is that all the resource requirements may be made to approach zero in the limit $U \to \mathbb{1}$, without the exception of the communication from one party, as in the case of the first approach.

VII. CONCLUSION

We have presented a scheme for implementing evolution under general multipartite Hamiltonians on multilevel systems. This scheme is both more efficient and more general than previously published schemes [3,4,6–9]. The entanglement required for this scheme is less than that for Refs. [3,4], in some cases dramatically so. In the case of tensor product Hamiltonians, the entanglement required is approximately $5.6418t\|H\|$. In contrast, the scheme of CDKL [3,4] (which applies to more restricted cases) requires entanglement of $5.9793t\|H\|$ for small t. This improvement is surprising, as our initial studies suggested that the CDKL scheme could not be improved upon.

For multiple implementations of the operation, the average communication from all parties except one may be reduced to be equal to the entanglement consumption. This is a dramatic improvement over previous methods, which required a large amount of communication from each party. For the scheme of Refs. [3,4] the communication required approaches infinity in the limit of weak interactions. We have also presented an alternative approach such that the communication required from all parties approaches zero as the unitary approaches the identity.

Our scheme may also be applied to completely arbitrary multipartite Hamiltonians for multilevel systems. Previously published methods only considered restricted classes of operations, such as two-qubit unitaries [3,4]. For our scheme the resource requirements, except for communication from one party, scale linearly with the evolution time. If the Hamiltonian is the sum of a large number of tensor product Hamiltonians, then the scaling constant will be large, though the resource requirements still approach zero for small evolution times.

This research raises a number of interesting issues. Although this scheme improves on the resource consumption, the entanglement and communication requirements are still significantly above the capacities for the operations. This means that, either it is possible to improve the resource consumption further, or there is a fundamental irreversibility in that operations require more resources to implement than they can produce. In the case of the entanglement, the scheme was optimized over this class of schemes, so it would require a dramatically different approach to obtain further improvements.

ACKNOWLEDGMENTS

This research has been supported by the Australian Research Council. The author is grateful to Jason Twamley, Andrew Doherty, and Barry Sanders for fruitful discussions.

[1] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, Phys. Rev. Lett. **70**, 1895 (1993).
[2] J. Eisert, K. Jacobs, P. Papadopoulos, and M. B. Plenio, Phys. Rev. A **62**, 052317 (2000).
[3] J. I. Cirac, W. Dür, B. Kraus, and M. Lewenstein, Phys. Rev. Lett. **86**, 544 (2001).
[4] W. Dür and J. I. Cirac, Phys. Rev. A **64**, 012317 (2001).
[5] W. Dür, G. Vidal, J. I. Cirac, N. Linden, and S. Popescu, Phys. Rev. Lett. **87**, 137901 (2001).
[6] B. Reznik, Y. Aharonov, and B. Groisman, Phys. Rev. A **65**, 032312 (2002).
[7] B. Groisman and B. Reznik, Phys. Rev. A **71**, 032322 (2005).
[8] L. Chen and Y.-X. Chen, Phys. Rev. A **71**, 054302 (2005).
[9] H.-S. Zeng, Y.-G. Shan, J.-J. Nie, and L.-M. Kuang, e-print quant-ph/0508054.
[10] C. H. Bennett, A. W. Harrow, D. W. Leung, and J. A. Smolin, IEEE Trans. Inf. Theory **49**, 1895 (2003).
[11] C. H. Bennett, S. Popescu, D. Rohrlich, J. A. Smolin, and A. V. Thapliyal, Phys. Rev. A **63**, 012307 (2001).
[12] A. M. Childs, D. W. Leung, and G. Vidal, IEEE Trans. Inf. Theory **50**, 1189 (2004).
[13] Note that this is the *asymptotic* entanglement consumption for the scheme of CDKL, not just an upper bound.
[14] Each party performs the joint measurement on its own qubits; at no stage is a joint measurement on qubits possessed by different parties required.