

Entanglement consumption of instantaneous nonlocal quantum measurements

S R Clark^{1,2,4}, A J Connor², D Jaksch^{1,2} and S Popescu³

¹ Centre for Quantum Technologies, National University of Singapore,
3 Science Drive 2, Singapore 117543, Singapore

² Clarendon Laboratory, University of Oxford, Parks Road,
Oxford OX1 3PU, UK

³ H H Wills Physics Laboratory, University of Bristol, Tyndall Avenue,
Bristol BS8 1TL, UK

E-mail: s.clark@physics.ox.ac.uk

New Journal of Physics **12** (2010) 083034 (35pp)

Received 6 April 2010

Published 13 August 2010

Online at <http://www.njp.org/>

doi:10.1088/1367-2630/12/8/083034

Abstract. Relativistic causality has dramatic consequences on the measurability of nonlocal variables and poses the fundamental question of whether it is physically meaningful to speak about the value of nonlocal variables at a particular time. Recent work has shown that by weakening the role of the measurement in preparing eigenstates of the variable, it is in fact possible to measure all nonlocal observables instantaneously by exploiting entanglement. However, for these measurement schemes to succeed with certainty, an infinite amount of entanglement must be distributed initially and all this entanglement is necessarily consumed. In this work, we sharpen the characterization of instantaneous nonlocal measurements by explicitly devising schemes in which only a finite amount of the initially distributed entanglement is ever utilized. This enables us to determine an upper bound to the average consumption for the most general cases of nonlocal measurements. This includes the tasks of state verification, where the measurement verifies if the system is in a given state, and verification measurements of a general set of eigenstates of an observable. Despite its finiteness, the growth of entanglement consumption is found to display an extremely unfavourable exponential of an exponential scaling with either the number of qubits needed to contain the Schmidt rank of the target state or the total number of qubits in the system for an operator measurement. This scaling is seen to be a consequence of the combination of the generic exponential

⁴ Author to whom any correspondence should be addressed.

scaling of unitary decompositions combined with the highly recursive structure of our scheme required to overcome the no-signalling constraint of relativistic causality.

Contents

1. Introduction	2
2. Framework	5
3. Teleportation and instantaneous nonlocal unitaries	8
3.1. Teleportation	8
3.2. The Vaidman scheme	10
4. Finite consumption scheme	13
4.1. Pauli rotation chain	13
4.2. Concatenation of rotation chains	15
4.3. Average entanglement consumption	17
5. State verification measurements	19
5.1. Two-qubit states	19
5.2. Bipartite multi-qubit states	21
6. Instantaneous measurements of nonlocal operators	23
6.1. Two-qubit observables	24
6.2. Bipartite multi-qubit observables	26
7. Conclusions	27
Acknowledgments	28
Appendix A. Computing the average entanglement consumption	28
Appendix B. Uniformly controlled rotations	31
Appendix C. Constructing a Schmidt superposition state	32
References	34

1. Introduction

The formal compatibility of quantum mechanics with special relativity is highly nontrivial [1] and is in many ways quite miraculous [2]. Perhaps the most well-known difficulty in combining these formalisms arises from the so-called ‘collapse’ of a quantum state associated with the measurement process, and in particular the instantaneity of this change. This problem is highlighted in its most simple form by considering two observers, Alice and Bob, who are spacelike separated. Conventional wisdom holds that any self-adjoint operator that can be defined for Alice and Bob’s joint system is measurable in principle [3]. But in fact, most such operators represent nonlocal variables, meaning that they cannot be written in the form $A \otimes B$, where A and B are self-adjoint operators acting on Alice and Bob’s local Hilbert spaces, respectively. Earlier on, it was recognized that if such nonlocal variables were instantaneously measurable, in the standard sense in some Lorentz frame, then violations of relativistic causality arise (see figure 1 for a description of this effect for ideal measurements on two separated spin-1/2 particles). In 1931, Landau and Peierls [4] claimed that this observation implied, quite generally, the impossibility of measuring any nonlocal variable at a well-defined time, and even




Figure 1. (a) In a scenario envisaged by Landau and Peierls a particle is initially localized at a point A and is subject to an ideal momentum measurement at some later time $t = t_0$. The effect of this measurement is to instantaneously collapse the particle's wavefunction into a momentum eigenstate which subsequently redistributes the particle's probability amplitude throughout all space. There is then a nonzero probability of finding the particle at a location B which is spacelike separated from A . (b) A simpler formulation of this causality violation can be framed using spin- $\frac{1}{2}$ particles (see section 2). Here we consider a device that can perform an instantaneous ideal measurement of the magnitude of the total spin squared $J^2 = \sum_{k=x,y,z} (\sigma_k^A + \sigma_k^B)^2$ of two spacelike separated spins. Depending on the local outcomes a and b via some function f , the state after the measurement will be projected into one with a well-defined value $J^2 = f(a, b)$. (c) If such a measuring device exists it would violate relativistic causality. Suppose Alice and Bob prepare a state $|\uparrow_A\rangle|\uparrow_B\rangle$ in the distant past and arrange to measure J^2 at time t_0 . If just prior to the J^2 measurement at time $t_0 - \epsilon$ Alice flips her spin, then a measurement of σ_z^B by Bob just after the J^2 measurement at time $t_0 + \epsilon$ will yield \uparrow and \downarrow with equal probability. Since the time interval 2ϵ can be made arbitrarily small, Alice can use the J^2 measurement to send superluminal signals.

went so far as to postulate a new uncertainty principle to this effect. Thus, a common consensus arose that it only made sense to speak of local variables as observables in relativistic quantum mechanics.

It was only in 1980 that this conjecture was finally refuted by Aharonov and Albert [5, 6] who explicitly constructed a scheme for measuring certain nonlocal variables (e.g. the Bell operator; see section 2) instantaneously without contradicting causality. In contrast to previous studies, their measurement scheme explicitly introduced entangled probes whose quantum correlations enable nonlocal properties of the system to become correlated to local properties of the measuring device. This means that by combining the correlated local outcomes of the two observers, at some point in the future when their light cones have intersected, the final nonlocal measurement result can be revealed. Their discovery had serious implications for the notion of states and observables in relativistic quantum mechanics. It immediately disproved the previously held covariant state reduction postulate [7] whose validity was dependent on only local variables being measurable. It also showed that no covariant succession of states at a given time can be associated with the system, since observers in different Lorentz frames will have conflicting accounts of the reduction process, which cannot be reconciled within any single covariant state history. This far-reaching conclusion culminated in their proposing that to take

account of changes to a state vector, induced by local or nonlocal measurement processes, it is required that the wavefunction ceases being a function of spacetime and instead becomes a functional on the set of spacelike hypersurfaces [8, 9].

Further work [10] then detailed explicit methods for measuring nonlocal variables such as $A + B$ and modular sums like $(A + B) \bmod a$, where a is a desired eigenvalue. It was later proven in generality by Popescu and Vaidman [11] that any conceivable measurement requires the erasure of local information (within the relevant degrees of freedom) in order to be compatible with causality. For a standard nondemolition measurement, to satisfy this additional requirement there is a dramatic restriction on what is measurable. For the case of two spin- $\frac{1}{2}$ particles, causality limits the measurability of operators to those with either trivial direct product eigenstates or maximally entangled Bell states. The measurability of the latter is permitted, because the reduced density matrix of either spin is always proportional to the unit matrix. A surprising consequence of this result is that even nonlocal variables with product eigenstates (see equation (8) in section 6.1) are not measurable [12], showing further that standard quantum measurements can be nonseparable in a way not entirely captured by the notion of entanglement⁵. The information erasure theorem [11] indicates that causal measurement schemes for almost all nonlocal variables cannot be a standard von Neumann measurement. Such measurements leave the system undisturbed if it was in an eigenstate of the observable before the measurement and play a dual role of both observing a quantity and preparing the system in an eigenstate of the corresponding observable [14, 15]. It is now recognized that this framework, which was the basis of the Landau and Peierls conjecture, is too restrictive to decide whether a nonlocal variable attains the status of a physical observable. Instead an operator's measurability should be determined in a broader paradigm of verification measurements [16]–[18]. A verification measurement can confirm with certainty whether the system is in an eigenstate of an observable at a given time, but does not necessarily leave the system in an eigenstate after it is completed. These measurements are therefore destructive and nonrepeatable.

Recent work in the context of gauge theories has further highlighted the fundamental implications of how measurability is defined [19]. In particular for gauge theories, which are used to describe all elementary particles, it is common to characterize gauge field configurations by Wilson loop operators. These manifestly nonlocal quantities are taken to be basic observables in gauge theory. Yet it was shown that the nondemolition measurement of spacelike Wilson loops in a relativistic non-Abelian gauge theory violates causality and that instead only verification measurements are possible [19]. From a different perspective, it has also been shown recently that the use of additional ancillary resources can dramatically alter the properties of nonlocal measurements, for example by revealing Bell-inequality violations in delocalized single-particle mode entanglement that would otherwise be prohibited by super-selection rules [20]–[22]. Indeed, by moving both to verification measurements and exploiting ancilla it has been found that there are no causal restrictions on what variables can be measured. Firstly, using methods devised for remote probabilistic rotations [23], it was shown that all observables of two spin- $\frac{1}{2}$ particles can be measured instantaneously [16]. Secondly, a method based on teleportation [24] was devised which demonstrates the instantaneous measurability of all observables for multipartite systems of arbitrary dimension [17, 25]. These studies have

⁵ This situation is in contrast to the usual *nonlocality without entanglement* scenario where the constraint is on quantum resources and unlimited classical communication is assumed [13].

therefore answered affirmatively that the instantaneous measurement of all nonlocal variables⁶ can be achieved without contradicting quantum mechanics and causality. Thus, in principle, all nonlocal variables are valid physical observables and so within this framework the conventional wisdom is reestablished.

A critical ingredient in these measurement schemes is entanglement. However, since the main aim of those schemes [16, 17] was disproving causal restrictions, they were not concerned with limiting the amount of entanglement consumed. As a result, to guarantee success in the most general cases, these schemes require an unlimited supply of entanglement to be initially distributed between Alice and Bob, and all of this entanglement is necessarily consumed. Here we go beyond this by systematically addressing the latter issue, namely the entanglement consumption. Firstly, we explicitly devise a scheme which significantly optimizes that devised by Vaidman in [17], where only a finite amount of the initial entanglement is ever consumed on average. This enables us to sharpen the characterization of instantaneous nonlocal measurements by quantifying the cost of nonlocal measurement tasks. Specifically, we determine an upper bound to the average consumption for state verification, where the measurement verifies if the system is in a given state, and for the verification measurement of a general set of eigenstates of an observable. Secondly, it is straightforward to show from our scheme that by only allowing a finite amount of the initial entanglement, in addition to a finite average consumption, the measurement can still proceed with certainty but will suffer a bounded error on its statistics.

The structure of this paper is as follows. In section 2, we lay out the framework we shall use in this study, and describe the approach to nonlocal measurements taken with specific attention paid to the Bell measurement example. This is followed in section 3 by a brief review of teleportation as an ingredient in instantaneous protocols and an outline of the pioneering work by Vaidman [17]. The main component of this work, what we call *rotation chains*, is introduced in section 4. In this section, the protocol for a single chain is described in detail and is shown to have a finite average entanglement consumption. In addition, it is explained how these chains can be concatenated to implement arbitrarily complex nonlocal unitaries and the scaling of the average entanglement consumption with the number of chains is also found. The remainder of the paper then utilizes these tools for several nonlocal measurement problems. Firstly, in section 5 it is applied to state verification measurements starting with an arbitrary two-qubit state before generalizing to an arbitrary finite-sized bipartite multi-qubit system where the scaling of entanglement consumption with the Schmidt rank of target state is obtained. Secondly, the state verification scheme is expanded in section 6 to enable the simultaneous verification of any set of orthogonal eigenstates constituting a full operator measurement. Again two-qubit observables are considered in detail, followed by a bipartite multi-qubit system where the scaling in entanglement consumption with the system size is determined. Finally, in section 7, we conclude and comment on open problems for future work.

2. Framework

Let us now describe in more detail the framework used within this study. We shall exclusively consider both the principal system and measuring probes as being composed of two distinguishable parts built up from spin- $\frac{1}{2}$ particles (qubits) and each localized in different

⁶ One caveat to this, which applies to this work as well, is variables related to fermionic degrees of freedom that are spatially delocalized [17, 26].

regions of space occupied by Alice and Bob which are spacelike separated. While this is not the most general scenario, it has proven to be particularly well suited for investigating quantum measurements and nonlocality [11, 27, 28]. The local regions themselves are assumed to be small enough to neglect causality restrictions within them but large enough compared to the Compton wavelength to neglect relativistic effects such as pair creation. Relativistic causality then enters due to the scale of the distances between the two parts of the system and otherwise the formalism of nonrelativistic quantum mechanics can be used. Within this setting, we shall consider measurement schemes that are localizable⁷ quantum operations. This means that they can be composed of arbitrary local operations between the local parts of the system (we assume that all local operations are equally easy to apply) and entangled resources that were shared prior to the measurement but do not utilize any classical communication. Localizable operations are manifestly causal, although curiously not all causal operations are themselves localizable [29].

In general, a nonlocal measurement requires previously arranged cooperative actions of Alice and Bob which can be broken into three steps. In the first step, which will be seen to be essential, suitably entangled ancilla systems must be prepared and distributed to the parties. Secondly, each party performs a local operation, such as unitaries and ideal (irreversible) projective measurements, on their part of the principal system and entangled ancillae. Thirdly, the classical information extracted by both parties in the second step is transmitted to a central location C where the readout of the result is completed⁸. These steps are summarized in figure 2(a). Since the local operations that act on parts of the system and measuring device in step two can proceed without waiting or knowing the outcomes of actions performed by the other party they can in principle be performed in an arbitrarily small time. Thus, when we speak of an ‘instantaneous measurement’, we are referring to the particular Lorentz frame where both observers performed their actions at time t_0 . Since we are interested in examining questions of causality, as opposed to covariance, we shall continue to use the terminology of quantum states and confine our description to this Lorentz frame. At the end of step two, both Alice and Bob are in possession of a set of indelible local classical bits. In accordance with the information erasure theorem [11], these local outcomes can only specify which eigenvalue of the nonlocal variable the system had at time t_0 once they are combined later at a point C in the future light cones of both observers. As a consequence, although the measurement was instantaneous and completed in step two at time t_0 , the result is not necessarily known instantaneously by either party and can only be reconstructed much later at step three. Despite these features nonlocal verification measurements retain the usual requirements that (i) when the system is in an eigenstate of the observable the outcome corresponding to that eigenstate is produced with certainty, and the linearity of quantum mechanics then ensures that (ii) for a general superposition of eigenstates the corresponding eigenvalues are observed with the appropriate quantum probabilities.

The features of a nonlocal measurement just discussed are best outlined by a concrete example. In figure 2(b), a nonlocal demolition measuring scheme for the Bell operator of

⁷ Following earlier work [29] the relevance of our results for quantum field theory should be understood as applying to the idealization that the external probe variables are ‘heavy’ with rapidly decaying correlations, while the field variables are ‘light’. In this situation, the notion of localizability, which requires a strict separation between field and probe, is credible.

⁸ A more general scenario can permit quantum information to be transmitted. This would enable so-called *exchange measurements* [10] to occur where the principal system is swapped into the measuring device, essentially freezing its state, and is then later measured at C . In this case the measurement has not really occurred until the last step and its outcome did not exist at time t_0 .

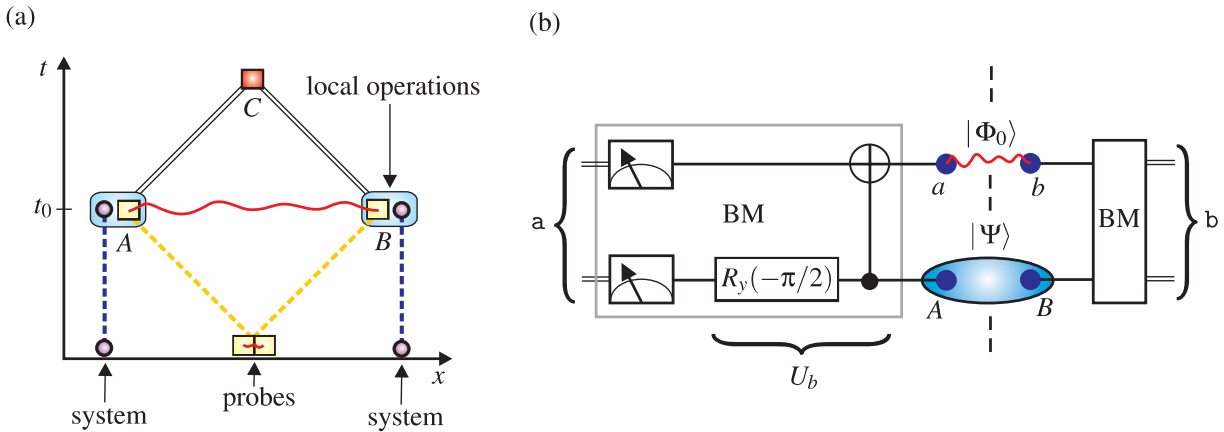


Figure 2. (a) A depiction of the type of measurement scheme considered in this work. The scheme intends to measure a property, at a given time t_0 , of a system composed of two spacelike separated parts at regions A and B. To do so, measuring probes are prepared some time earlier, possibly in an entangled state (signified by the wiggly line), and transported to the two locations A and B. Once the probes arrive at these locations at time t_0 , local operations are performed between the parts of the system and probe at each region, resulting in local classical information. This information is then transmitted to a location C where the future light cones of A and B intersect. The overall outcome of the instantaneous nonlocal measurement is then deduced at C, but pertains to the system at time t_0 . (b) As an example of such a measurement scheme a circuit diagram is shown for the demolition nonlocal measurement of the Bell operator on two qubits utilizing one maximally entangled ancilla $|\Phi_0\rangle$. The vertical dashed line delineates the two regions and highlights that all operations in this circuit are local. In this example, both local operations correspond to a Bell measurement. To aid the explanation of this measurement given in section 3, we show on the left-hand side a Bell measurement composed of a unitary U_b and single-qubit measurements [15] whose outcomes together give a binary encoding of the overall $\{0, 1, 2, 3\}$ result. For subsequent diagrams we shall simply denote local Bell measurements by a ‘BM’ box with four outcomes and not be concerned with its internals, be they a joint measurement projecting on to local Bell states or the single-qubit form given above. The *global* outcome for the instantaneous nonlocal Bell measurement is then designated by the addition modulo 4 of the local results $c = a \oplus b$.

two qubits is shown. The Bell operator possesses the nondegenerate maximally entangled eigenstates

$$|\Phi_0\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B + |1\rangle_A |1\rangle_B),$$

$$|\Phi_1\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A |1\rangle_B + |1\rangle_A |0\rangle_B),$$

$$|\Phi_2\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B - |1\rangle_A |1\rangle_B),$$

$$|\Phi_3\rangle = \frac{i}{\sqrt{2}} (|0\rangle_A |1\rangle_B - |1\rangle_A |0\rangle_B).$$

For the measurement scheme shown in figure 2(b), a pair of ancilla qubits in the state $|\Phi_0\rangle$ have been previously distributed. Such maximally entangled pairs will form the resource for all of the schemes studied in this work. The measurement of the Bell operator then proceeds by each party performing a local Bell measurement between their half of the system and ancilla pair. Since it is a demolition verification measurement, once it is completed the local parts of the system and ancilla are left in direct product states with equal unbiased probabilities. Thus, in accordance with causality, local information in the relevant degrees of freedom is erased and the local reduced density matrix is maximally mixed at all times. As a result, the local outcomes $a, b \in \{0, 1, 2, 3\}$ reveal no information about the global outcome in isolation. Instead they are correlated nonlocally, with the final outcome being $c = a \oplus b$, where \oplus is modulo 4 addition. We shall explain how this measurement scheme works shortly in section 3. While this demolition Bell-measurement scheme shares many features with the more general schemes about to be introduced, we mention for completeness that, with the use of an additional entangled pair and a suitable modification of the circuit in figure 2(b), a non-demolition scheme can be devised [5, 6, 14]. The information erasure theorem [11] proves that this is the only nonlocal variable of two qubits that possesses a non-demolition measurement scheme, because the reduced density matrix of any of its eigenstates for either party is maximally mixed.

3. Teleportation and instantaneous nonlocal unitaries

To generalize the Bell measurement just described to a more general nonlocal measurement scheme, it turns out to be very convenient to describe the local operations performed by both parties in terms of the instantaneous part of the teleportation protocol [24]. In this section, we shall describe teleportation within the framework outlined above and also detail earlier work by Vaidman [17] that demonstrated how, through a prearranged recursive structure, it enables the instantaneous measurement of any nonlocal variable.

3.1. Teleportation

As is well known, the teleportation [24] of an arbitrary state of d qubits $|\Psi\rangle$ can be accomplished by local operations and classical communication if Alice and Bob share one half of d maximally entangled two-qubit states $|\Phi_0\rangle$. This follows from the identity

$$|\Psi\rangle_{A_1 A_2 \dots A_d} \otimes |\Phi_0\rangle_{a_1 b_1} \otimes |\Phi_0\rangle_{a_2 b_2} \otimes \dots \otimes |\Phi_0\rangle_{a_d b_d}$$

$$= \frac{1}{2^d} \sum_{\mathbf{m}} |\Phi_{\mathbf{a}_1}\rangle_{A_1 a_1} \otimes |\Phi_{\mathbf{a}_2}\rangle_{A_2 a_2} \otimes \dots \otimes |\Phi_{\mathbf{a}_d}\rangle_{A_d a_d} \sigma_{\mathbf{a}} |\Psi\rangle_{b_1 b_2 \dots b_d}.$$

Here we have designated a tensor product of Pauli operators⁹ over the system of d qubits as $\sigma_{\mathbf{a}} = \sigma_{\mathbf{a}_1} \otimes \sigma_{\mathbf{a}_2} \otimes \dots \otimes \sigma_{\mathbf{a}_d}$, where $\mathbf{a} = (\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_d)$ is a d -dimensional vector of outcomes

⁹ We will refer to tensor products of Pauli operators as a Pauli string operator or a simply as a Pauli distortion depending on the context.

$a_j \in \{0, 1, 2, 3\}$ and we numerically index σ_k with $\sigma_0 = \mathbb{1}$ while $1 \mapsto x, 2 \mapsto z$ and $3 \mapsto y$. Teleportation is then achieved by Alice measuring each pair of qubits A_j, a_j in the Bell basis $|\Phi_k\rangle$ with the outcome fixing the element a_j in \mathbf{a} . Overall, this collapses Bob's d qubits to the state $|\Psi\rangle$, modulo a Pauli distortion $\sigma_{\mathbf{a}}$ determined by Alice's equiprobable measurement outcomes \mathbf{a} . The full teleportation protocol is finished by Alice transmitting $2d$ classical bits to Bob specifying the vector of outcomes \mathbf{a} so he can remove the Pauli distortion $\sigma_{\mathbf{a}}$ and recover $|\Psi\rangle$ with certainty. Since we will be exclusively concerned with instantaneous operations this last step, which necessarily takes a finite amount of time to implement, will never be performed and we shall from now on use the term *teleportation* to describe the Bell measurement part only. The cost of instantaneity is the unavoidable presence of the equiprobable Pauli distortion $\sigma_{\mathbf{a}}$, which due to the identity

$$\frac{1}{4^d} \sum_{\mathbf{a}} \sigma_{\mathbf{a}} |\Psi\rangle \langle \Psi| \sigma_{\mathbf{a}} = \frac{1}{2^d} \mathbb{1},$$

preserves relativistic causality by completely scrabbling the reduced density matrix of the local system. As we shall see, despite the distortion, teleportation can nonetheless be exploited to achieve instantaneous nonlocal operations.

The simple instantaneous Bell measurement in fact already highlights some essential properties of nonlocal measurements and readily allows us to identify a bipartite multi-qubit generalization. Specifically, in figure 2(b), we can interpret Bob's local Bell measurement (on the right) as a teleportation of his half of the system to Alice yielding an outcome \mathbf{b} . In a step that will be shared by all schemes in this work he has localized their initially distributed system. Alice can then attempt to apply a unitary that maps the locally unmeasurable set of eigenstates into the trivially measurable direct product set $|0\rangle|0\rangle, |0\rangle|1\rangle, |1\rangle|0\rangle$ and $|1\rangle|1\rangle$. Were Alice to apply a general unitary U , its effect would be confounded by the Pauli distortion $\sigma_{\mathbf{b}}$ on her receiving qubit unknown to her. For a Bell measurement the required unitary $U_{\mathbf{b}}$, depicted in figure 2(b), is a member of a special class of unitaries in this regard. Specifically, for a multi-qubit system with a distortion $\sigma_{\mathbf{b}}$, there are a set of unitaries $U \in \mathcal{S}$ which satisfy $U\sigma_{\mathbf{b}} = \sigma_{\mathbf{b}'}U$, in which a Pauli distortion $\sigma_{\mathbf{b}}$ can be propagated through them at the expense of possibly changing to a different Pauli distortion $\sigma_{\mathbf{b}'}$. This set of unitaries \mathcal{S} is called stabilizers and can be constructed, up to a global phase, from quantum circuits containing only CNOT, Hadamard and phase gates [15].

Since both the CNOT gate and the rotation $R_y(-\pi/2) = \exp(i\pi\sigma_y/4)$ are stabilizers, the effect of $U_{\mathbf{b}}(\mathbb{1} \otimes \sigma_{\mathbf{b}})$ in the Bell measurement is summarized as $U_{\mathbf{b}}, (\mathbb{1} \otimes \sigma_x)U_{\mathbf{b}}, (\sigma_x \otimes \sigma_z)U_{\mathbf{b}}$ and $(\sigma_x \otimes \sigma_y)U_{\mathbf{b}}$ for $\mathbf{b} = (0, 1, 2, 3)$, respectively. The scheme terminates, once $U_{\mathbf{b}}$ is applied, with a measurement of both qubits in the computational basis (i.e. z -axis). Since Pauli distortions simply map direct product states between themselves once $\sigma_{\mathbf{b}}$ has been propagated through $U_{\mathbf{b}}$, it induces a benign, but causality-preserving, nondeterministic mapping between the Bell and direct product bases. This final measurement in the fixed z -axis is therefore certain to complete the scheme.

With this observation we can immediately construct a nonlocal instantaneous verification measurement for any operators on any number of qubits whose eigenstates are all stabilizer states (or states that are locally equivalent to them). In addition to Bell states, this class includes some of the most well-studied multi-qubit entangled states such as the Greenberger–Horne–Zeilinger (GHZ) state [30], cluster states [31, 32] and more generally graph states [33]. The entanglement consumption of stabilizer measurements, analogous to the Bell measurement,




Figure 3. A schematic diagram of a nonlocal verification measurement scheme for any observable on $2d$ qubits whose complete set of eigenstates are locally equivalent to stabilizer states. The unitaries V_A^\dagger and V_B^\dagger are arbitrary and account for the local equivalence. The unitary $U \in \mathcal{S}$ is a stabilizer and is applied to the system once it has been localized by teleportation. The entire set of qubits at Alice's location can be measured in the z -axis to complete the scheme. Generalization to unequal distributions of qubits and multi-party scenarios is straightforward.

is then simply the minimum number of ebits needed to localize the system. A schematic diagram of a stabilizer measurement protocol is given in figure 3. For the most general nonlocal measurements the unitary U required will not be a stabilizer. Our goal is therefore to devise a scheme that, after the system has been localized by teleportation, enables U to be applied while still propagating any Pauli distortions to the end. As we shall see, for arbitrary unitaries U , this is a highly nontrivial and expensive task.

3.2. The Vaidman scheme

The general nonlocal measurement scheme devised by Vaidman [17] starts in the same way as Bell measurement in figure 2(b) by Bob teleporting his half of the system to Alice. Without any loss of generality, we focus on a system of two qubits. Since Alice and Bob's aim is to measure some nonlocal variable O with eigenstates $|o_1\rangle$, $|o_2\rangle$, $|o_3\rangle$ and $|o_4\rangle$, they devise a unitary U transformation that maps these eigenstates to the measurable direct product basis as

$$\begin{aligned} U |o_1\rangle &= |0\rangle|0\rangle, & U |o_2\rangle &= |0\rangle|1\rangle, \\ U |o_3\rangle &= |1\rangle|0\rangle, & U |o_4\rangle &= |1\rangle|1\rangle. \end{aligned}$$

Given the system was initially in that state $|\Psi\rangle_{AB}$, the state of Alice's qubit A and the ancilla qubit a_1 , representing the receiving qubit of the teleportation from Bob, is now in a state $\sigma_{b_1}|\Psi\rangle_{Aa_1}$. The first step of the scheme is for Alice to simply apply U to qubits A and a_1 . With a probability of $1/4$, Bob's teleportation will be non-distorting with $b_1 = 0$ and Alice would have successfully mapped the eigenstates of O to the measurable direct product basis. For the other three distortions, the resulting unitaries $U\sigma_x$, $U\sigma_y$ and $U\sigma_z$ will not in general map the eigenstates to the direct product basis (unless of course U happens to be a stabilizer). Since Alice

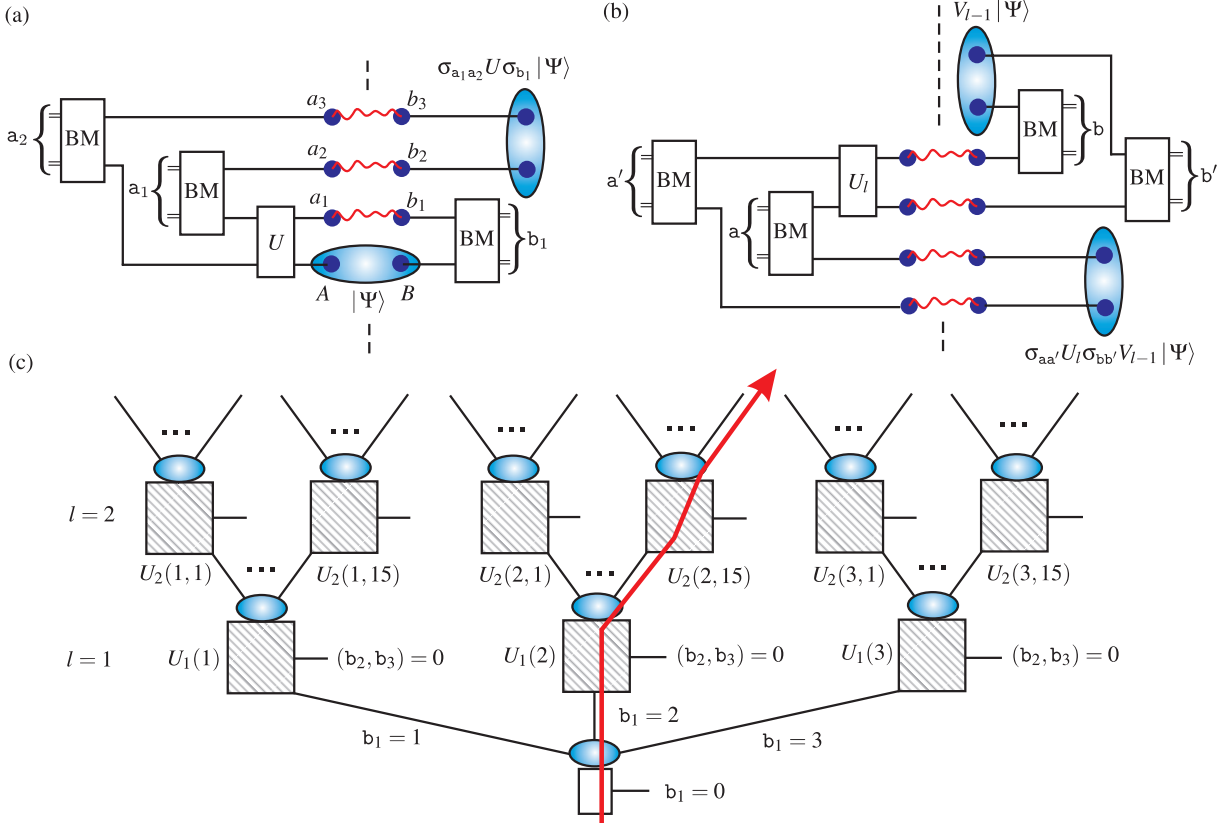


Figure 4. (a) At the start of the Vaidman scheme, Bob teleports his half of the system to Alice who then applies the unitary U between her half and the teleported qubit. Since Alice does not know whether this action was successful she teleports the entire system back to Bob. (b) A cluster mentioned in the main text. Given Bob has a state $V_{l-1} |\Psi\rangle$ with some accumulative unitary V_{l-1} applied to it, he can teleport the entire system to Alice who can apply a correction U_l and teleport it back. If Bob's teleportation is nondistorting so $b = b' = 0$, then Alice has successfully corrected the returned state. (c) A tree diagram of the Vaidman scheme. The root of the tree is the initial step of the scheme depicted in (a). The hatched boxes spawning from the root represent clusters identical to that depicted in (b). Depending on his teleportation outcomes, Bob traverses the tree structure utilizing only a specific path of clusters corresponding to his history of outcomes (e.g. illustrated by the arrow). See the main text for a more detailed description of the scheme.

has no knowledge of b_1 , she has no choice but to teleport the entire system of two qubits back to Bob. For brevity we shall from now on call a complete teleportation of the system, regardless of the number qubits, a *channel*. This initial step of the scheme is shown in figure 4(a).

At Bob's side he expects the return of the system and on the fortuitous occasion that his first teleportation gave $b_1 = 0$ he can be assured that Alice successfully applied U to $|\Psi\rangle$ leaving his two ancilla qubits b_2 and b_3 in the state $\sigma_{a_1 a_2} U |\Psi\rangle_{b_2 b_3}$. Just as with the Bell measurement scheme the final mapping is to the trivial direct product basis modulo a subsequent Pauli distortion. Thus, despite the fact that Bob has no knowledge of the outcomes a_1 and a_2 , he can go ahead

and immediately measure the qubits in the z -axis completing the verification measurement of O . For the cases where $b_1 \neq 0$, Bob knows that Alice did not apply U in isolation. To allow Alice the opportunity to correct this mistake the scheme from here on adopts a tree-like structure. The root of this tree is the first teleportation just described. Above this there are now three branches each labelled by the possible distorting values b_1 might take. Each branch leads to a *cluster* whose structure is illustrated in figure 4(b). A cluster simply contains a teleportation channel for Bob to send back the system to Alice and a corresponding channel for Alice to return it. Depending on the actual value of b_1 , Bob traverses the corresponding branch of the tree and sends the system back to Alice via the channel in that branch's cluster. He will never act on clusters in any other branch of the scheme.

At the receiving end of the incoming channel in each of the three clusters, Alice will have, if the cluster was used, a state $\sigma_{b_2 b_3} \sigma_{a_1 a_2} U \sigma_{b_1} |\Psi\rangle$. She can now infer the value of b_1 from the cluster's label. Under the assumption that Bob's teleportation in that cluster was nondistorting, so $b_2 = b_3 = 0$, she can devise a correction unitary $U_1(b_1)$, dependent on b_1 , obeying

$$U_1(b_1) \sigma_{a_1 a_2} U \sigma_{b_1} = U.$$

Thus, with a probability of $1/16$, Alice will undo the previous unitary and distortions and map the eigenstates of O to the direct product basis. To complete the cluster, she teleports the resulting two-qubit system back to Bob via a corresponding return channel as shown in figure 4(b). Since Alice does not know which, if any, of the clusters were used she must perform this b_1 -dependent correction on all three clusters.

The situation for Bob is now identical to the first round but with a smaller probability of success. If $b_2 = b_3 = 0$, then as before he can immediately measure the incoming qubits on the cluster he used and complete the measurement. For the other 15 possible distorting outcomes, Bob knows Alice's correction will have failed. To overcome this failure, the same strategy is applied. Each of the clusters in the first level of the tree spawns 15 new branches, one for each possible distortion by Bob's previous teleportation, again leading to a new cluster. From his current position in the tree Bob now traverses the appropriate branch dependent on b_2, b_3 and teleports the system back to Alice through the channel in that branch's cluster. For Alice the situation is now that she has 45 incoming channels to operate on since she has no knowledge of Bob's actual path through the tree. For each cluster in this second level, she can continue to guarantee a $1/16$ chance of success by again devising a unitary $U_2(b_1, b_2, b_3)$ obeying

$$U_2(b_1, b_2, b_3) \sigma_{a_4 a_5} U_1(b_1) \sigma_{b_2 b_3} \sigma_{a_1 a_2} U \sigma_{b_1} = U.$$

The labels on the tree structure provide Alice with a complete history of distortions that Bob would have induced had he traversed those branches and this is essential for her to be able to construct a correction. The only knowledge Alice lacks is the nature of Bob's last teleportation and her correction only works on the assumption that it is nondistorting. The scheme therefore continues in the same way following an exponentially growing tree structure, depicted in figure 4(c). The measurement is completed once Bob has performed a nondistorting teleportation.

So long as this scheme is repeated to infinite depth it can, quite remarkably, ensure that at some point along the path traversed by Bob the unitary U is applied with certainty, modulo some proceeding Pauli distortions. At this termination point Bob can then complete the measurement. The cost of achieving this task, however, is unbounded. In particular, the division of labour is highly skewed since Alice must operate on all branches, whose number grows exponentially

with the level, and unlike Bob has no termination condition. This means that the infinite amount of entanglement that was initially distributed to form the scheme's tree structure is necessarily consumed without exception. In return for this effort, however, Alice has complete control of the unitary U eventually implemented and need only decide what it is immediately before she starts her actions. The scheme generalizes straightforwardly for d qubits but with a probability of success 4^{-d} at each level and $4^d - 1$ new branches spawning to the next level. Additionally, since Alice does all the correction work, the scheme can be readily adapted to work with any number of parties [17]. This is a property shared with stabilizer measurements since there only one party is needed to perform the stabilizer circuit. The Vaidman scheme provides a constructive proof that an instantaneous measurement scheme, which is guaranteed to succeed, can be devised for any nonlocal variable and nonetheless be compatible with both quantum mechanics and causality. For the remainder of this work we describe a scheme, based on a simple but significant modification of Vaidman's, that can similarly be used to measure any nonlocal variable and is guaranteed to succeed, but consumes only a finite amount of entanglement on average.

4. Finite consumption scheme

The essential adjustment we make to Vaidman's scheme is that rather than attempting to apply the desired unitary U directly at each step, we instead decompose U into a sequence of simpler unitaries and attempt to apply these individually in separate rounds of the scheme. These simpler unitaries are Pauli rotations $R_{\mathbf{j}}(\theta) = \exp(-i\theta\sigma_{\mathbf{j}}/2) = \cos(\frac{1}{2}\theta)\mathbb{1} - i\sin(\frac{1}{2}\theta)\sigma_{\mathbf{j}}$ involving the exponential of a Pauli string operator $\sigma_{\mathbf{j}}$, introduced earlier in section 3.1, by an angle θ . In this section, we will concentrate on implementing Pauli rotations and give explicit examples of decomposing general unitaries U in terms of them later when we discuss specific applications in section 5 and section 6. Despite Pauli rotations not being stabilizers (aside from when $\theta = \pi/2$), they do have extremely advantageous properties with respect to Pauli distortions. As we shall now describe, this can be exploited to yield a scheme where both parties have local termination conditions and only a finite amount of the initial entanglement is ever consumed.

4.1. Pauli rotation chain

The basic component of all our measurement schemes is a *rotation chain*, which applies a single Pauli rotation $R_{\mathbf{j}}(\theta)$ designated by an angle θ and a nontrivial vector \mathbf{j} specifying the Pauli string known to both parties. A rotation chain is composed of a sequence of teleportation channels in which the entire system of d qubits is teleported together back and forth in an alternating direction between Alice and Bob, as depicted in figure 5. The starting point of a rotation chain is the familiar situation where one party, say Alice, possesses the entire system. Initially, when the system was distributed, it was in some state $|\Psi\rangle$; however, the actual state of the system at Alice's location contains a Pauli distortion $\sigma_{\mathbf{b}_1}|\psi\rangle$ defined by a vector \mathbf{b}_1 known only to Bob. This distortion is taken to have arisen from earlier teleportations (such as previous rotation chains as described shortly in section 4.2) and so we take all 4^d possible vectors \mathbf{b}_1 as equiprobable¹⁰. As we have seen, the presence of this distortion generally results in $4^d - 1$ different errors if Alice tried to apply the complete unitary U directly. In a rotation chain, Alice

¹⁰ For Bob's initial teleportation which localizes the system, $\sigma_{\mathbf{b}_1}$ has zero elements for all of Alice's qubits and so is only equiprobable over a subset of $4^{d/2}$ strings. However, since Alice will attempt to apply a Pauli rotation on the entire system, the effect of this type of $\sigma_{\mathbf{b}_1}$ is identical.




Figure 5. A schematic diagram of a Pauli rotation chain used to implement a unitary $R_j(\theta)$. As a result of previous teleportations Alice possesses the entire system, but in a state $\sigma_{b_1} |\Psi\rangle$, where \mathbf{b}_1 is known only to Bob. As a result, she cannot be certain that she has applied $R_j(\theta)$ directly to the state $|\Psi\rangle$. The scheme depicted shows that by exploiting a sequence of directed teleportation channels the entire system can be bounced back and forth between Alice and Bob such that there is a probability of $\frac{1}{2}$ at each step that either party possesses a state $R_j(\theta) |\Psi\rangle$ modulo a preceding Pauli distortion. This strategy is a bipartite multi-qubit generalization of a similar single-qubit scheme presented in [17]. See the main text for a more detailed description of the scheme.

instead applies $R_j(\theta)$ and only one of two possibilities occurs:

$$R_j(\theta) \sigma_{b_1} |\Psi\rangle = \begin{cases} \sigma_{b_1} R_j(\theta) |\Psi\rangle, & \text{for } \mathbf{b}_1 \in \mathbf{c}(\mathbf{j}), \\ \sigma_{b_1} R_j(-\theta) |\Psi\rangle, & \text{for } \mathbf{b}_1 \in \bar{\mathbf{c}}(\mathbf{j}). \end{cases} \quad (1)$$

Since $\mathbf{j} \neq (0, 0, \dots, 0)$, and so never designates a string of identity operators, we denote here $\mathbf{c}(\mathbf{j})$ as the set of $4^d/2$ vectors specifying Pauli strings that commute with σ_j , while $\bar{\mathbf{c}}(\mathbf{j})$ is the other half of the total set of vectors that anti-commute with σ_j . In the latter case, propagation of the rotation through the distortion σ_{b_1} results in a sign change. Thus, Alice has a probability of $\frac{1}{2}$, independent of d , to have implemented the correct rotation on the initial state. Moreover, the only error she can make is to rotate in the wrong direction. Since she has no knowledge of her success, she must teleport the entire system back to Bob via the first channel shared between them.

At Bob's side he immediately applies the unitary σ_{b_1} , corresponding to the initial distortion, to the incoming qubits. If $\mathbf{b}_1 \in \mathbf{c}(\mathbf{j})$, then his initial distortion was commuting and the incoming

qubits will be in a state $\sigma_{\mathbf{b}_1} \sigma_{\mathbf{a}_1} \sigma_{\mathbf{b}_1} R_j(\theta) |\Psi\rangle$, where $\sigma_{\mathbf{a}_1}$ is a new distortion induced by Alice's teleportation. Bob therefore knows that the incoming qubits have had, modulo a subsequent distortion, the correct rotation applied to their initial state. He then keeps these qubits ready for further operations (see section 4.2) or a measurement. His actions for this chain are then terminated. If $\mathbf{b}_1 \in \bar{\mathbf{c}}(\mathbf{j})$ then his initial distortion was anti-commuting and he knows that Alice performed $R_j(-\theta)$ instead. Following a strategy outlined in [34], Bob can attempt to correct this, under a previously agreed assumption that \mathbf{a}_1 is commuting, by applying a new *double angle* rotation $R_j(2\theta)$ to the qubits. This gives $R_j(2\theta) \sigma_{\mathbf{b}_1} \sigma_{\mathbf{a}_1} \sigma_{\mathbf{b}_1} R_j(-\theta) |\Psi\rangle$, which is the desired state only if $\mathbf{a}_1 \in \mathbf{c}(\mathbf{j})$ is a commuting distortion. To overcome his lack of knowledge regarding \mathbf{a}_1 , Bob teleports all d qubit back to Alice via the next channel in the chain.

The situation for Alice is now identical to Bob's just described. She immediately applies $\sigma_{\mathbf{a}_1}$ to the incoming qubits. If $\mathbf{a}_1 \in \mathbf{c}(\mathbf{j})$ is commuting, she can be certain that, if it was necessary, Bob succeeded in correcting her rotation. In this case, the state of her system is $\sigma_{\mathbf{a}_1} \sigma_{\mathbf{b}_2} \sigma_{\mathbf{b}_1} \sigma_{\mathbf{a}_1} \sigma_{\mathbf{b}_1} R_j(\theta) |\Psi\rangle$, where $\sigma_{\mathbf{b}_2}$ is a new distortion induced by Bob's teleportation back. This final state is of the required form so she keeps the qubits and terminates her actions in this chain. If $\mathbf{a}_1 \in \bar{\mathbf{c}}(\mathbf{j})$, Bob's rotation causes an accumulative error of $R_j(-3\theta)$. Alice attempts to correct this, again under the assumption that his last distortion is $\sigma_{\mathbf{b}_2}$, by applying another double angle rotation $R_j(4\theta)$. Note that she does not need to assume or know anything about earlier distortions by Bob, such as $\sigma_{\mathbf{b}_1}$, since it appears twice in the accumulative distortion. She then teleports the qubits back via the next channel and the scheme continues. Schematic diagrams of these steps in the rotation chain scheme are given in figure 5.

Note that both Alice and Bob have a probability of $\frac{1}{2}$ of implementing the jointly agreed rotation at each step and can both determine their success by local outcomes. Since the actions of both parties terminate there is zero probability that the chain continues indefinitely and so only a finite amount of the initial entanglement is ever consumed. A disadvantage of joint termination is that as a rotation chain proceeds both parties lose knowledge of where the appropriately transformed qubits finally reside. Instead the actual pathway taken by the system is only reconstructed by the combination of Alice and Bob's local classical records. The manner in which the rotation chain deals with the Pauli distortions caused by teleportation is very reminiscent of one-way quantum computing [31, 32]. There the indeterminism of single-qubit measurements used to drive the computation produces Pauli distortions at intermediate stages, which, via minor adjustments in the subsequent operations, are propagated to the end of the computation. Their effect is then to simply alter the interpretation of the final output measurements. If further rotations are required, then, as we shall show in the next section, distortions can continue to be propagated to the end.

4.2. Concatenation of rotation chains

Let us now suppose that Alice and Bob wish to apply a further rotation $R_k(\xi)$ to the d -qubit state $R_j(\theta) |\Psi\rangle$. To do this, they can use a second rotation chain, which applies $R_k(\xi)$ to the output from the first $R_j(\theta)$. However, since the first rotation chain has multiple opportunities for terminating successfully on both Alice and Bob's side, a second chain must be available separately for each of these exit points to cover all eventualities. This gives a tree structure of concatenated chains like that shown in figure 6. Following this figure, suppose that Alice exits the first chain first on her q th opportunity. The d -qubits she then possesses will be in a state carrying a large accumulative distortion dependent on its history up to that point through the




Figure 6. A schematic diagram of the concatenation of two Pauli rotation chains used to implement a unitary $R_k(\xi)R_j(\theta)$ modulo a preceding Pauli distortion. For each possible exit from the $R_j(\theta)$ rotation chain there is a second $R_k(\xi)$ chain. The boxes containing arrows in this figure represent the entire rotation chain protocol depicted in figure 5. Arrows that end with a \times indicate that the originating party has not participated in the protocol for this specific chain. While both parties participate in the first chain only one of the secondary chains has overlapping actions of Alice and Bob. In this figure, Alice exits the first chain on her q th opportunity, while Bob exits on his p th where $p > q$. The dashed ‘L’-shaped line indicates the actual path taken by the principal system in this case. Actions performed by either party not intersecting this line do not contribute to the final outcome, but the no-signalling restriction requires that they are performed so that all eventualities are covered and the desired unitary is implemented with certainty.

first chain as

$$\sigma_{a_q} \sigma_{b_{q+1}} \sigma_{b_q} \sigma_{a_q} \sigma_{a_{q-1}} \cdots \sigma_{b_2} \sigma_{b_1} \sigma_{a_1} \sigma_{b_1} R_j(\theta) |\Psi\rangle. \quad (2)$$

She can go ahead and engage these qubits with the designated second rotation chain for this exit point which, in an identical way to the first, will apply $R_k(\xi)$. Since the accumulative distortion in equation (2) contains two of every previous distortion, except for $\sigma_{b_{q+1}}$, the criterion for

Alice's success in applying the second rotation $R_k(\xi)$ is based only on Bob's last teleportation, via $\mathbf{b}_{q+1} \in \mathbf{c}(\mathbf{k})$, and not on the complete history. This is in stark contrast to the Vaidman scheme.

As depicted in figure 6, both Alice and Bob must perform all the necessary steps for each of the second chains covering all possible exit points of the other party up to the point where they themselves exit from the first chain. This ensures that if the other party was successful before them the overall scheme still succeeds with certainty. Since all the first chain and all those spawning from it have zero probability of continuing indefinitely the overall scheme also has a finite average consumption. It is also clear that this concatenation can continue, albeit at increasing expense, for any finite sequence of rotations to be applied to the d -qubit initial state, and still retain a finite average entanglement consumption. We now examine more precisely what this consumption is.

4.3. Average entanglement consumption

To measure the consumption, we count the number of channels that are required on average. A detailed description of this calculation is given in appendix A. In summary, we find that the average channel consumption for a single rotation chain is $\langle c_1 \rangle = 5$, while concatenation of further rotation chains results in a rapid growth as $\langle c_2 \rangle = 20$, $\langle c_3 \rangle = 59$, $\langle c_4 \rangle = 156$ and so on. These channel averages $\langle c_n \rangle$ give the average consumption of entanglement, measured in ebits, once they are multiplied by d . By utilizing the recursive structure of the protocol the average channel consumption $\langle c_n \rangle$ can be approximated, in the limit of a large number of concatenated rotations n , by the exponential growth

$$\langle c_n \rangle \approx C \phi^n, \quad (3)$$

where $C = (10 + 7\sqrt{2})/4$ and $\phi = 1 + \sqrt{2}$. As shown in figure 7(b), the fit of this approximation to the exact consumption for $n > 4$ demonstrates that it is very good for all but the smallest n .

We saw earlier that a rotation by an angle $\pi/2$ has the special property that $R_j(\pi/2)$ is a stabilizer. In this case, no rotation chain steps are required. More generally, a chain involving a binary angle $\theta = \pi/2^D$ not only implements the desired rotation when a commuting distortion occurs, but also when a sequence of $D - 1$ erroneous rotations are made since the required double angle correction reduces to $R_j(\pi/2)$. Thus, for rotations with a binary angle, the chain terminates with certainty in a finite number of teleportations [16, 17]. An example of a single rotation chain applicable to an angle that is any odd multiple of $\pi/32$ is given in figure 7(a). The total number of initial channels that must be available for n concatenated rotation chains, each of length D , is finite but grows exponentially with n as

$$C_{\text{init}} = \frac{(D - 1)[(D - 1)^n - 1]}{D - 2}. \quad (4)$$

In an identical way to the $D \rightarrow \infty$ case the average consumption $\langle c_n \rangle$ of this initial resource can be computed. In figure 7(b), both C_{init} and $\langle c_n \rangle$ are shown for $D = 3$ and 7. For $D \leq 3$ both the initial resource and the average consumption remain below the average consumption for $D \rightarrow \infty$. For $D > 3$, the average consumption $\langle c_n \rangle$ rapidly converges to the $D \rightarrow \infty$ limit and the initial resources grow far beyond it.

While we have shown a finite average consumption in general, a practically relevant question arises as to what effect the restriction to finite initial resources has for general rotations. One strategy for doing this is to simply truncate continuous angle rotation chains to some maximum number of iterations. Indeed, if this strategy is applied to the Vaidman scheme, by

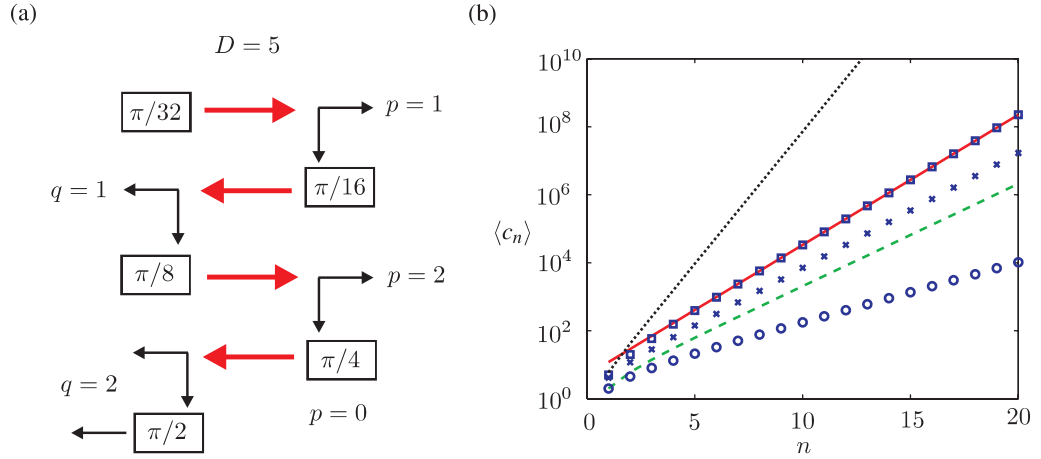


Figure 7. (a) An example of a rotation chain for a binary angle $\theta = \pi/2^D$ with $D = 5$. The scheme is guaranteed to terminate on Alice's second step $q = 2$, since the correction is a rotation by $\pi/2$ that always succeeds modulo a proceeding Pauli distortion. For D odd there is an outcome $p = 0$ corresponding to when Bob never succeeds. (b) The average channel consumption $\langle c_n \rangle$ for the scheme performing n successive rotations of the form $R_{j_n}(\theta_n) \dots R_{j_1}(\theta_1)$. The exact calculation of $\langle c_n \rangle$ for non-binary θ_j angles (i.e. infinite length rotation chains) is shown (\square) as well as the pure exponential approximation (solid line) given in equation (3). The exact $\langle c_n \rangle$ is also shown for binary angles with $D = 3$ (\circ) and $D = 7$ (\times), along with the total number of initial channels C_{init} that must be available in both cases as the dashed line and dotted line, respectively.

limiting its tree depth, it results in its having a finite consumption equal to its finite initial resources. This approach, however, introduces a possibility that the measurement will fail completely and yield no result. Binary angle rotation chains present a more elegant means of exploring the implications of finite initial resources for our scheme. Rather than truncating continuous angle rotation chains, we instead consider a more interesting and relevant scenario where the desired rotation angle is discretized to a multiple of a binary angle that matches the maximum allowed number of iterations. Given the decomposition of the desired final unitary U into a sequence of Pauli rotations, the new scheme performs rotations about the nearest binary angle $\tilde{\theta} = \lceil 2^D \theta / \pi \rceil \pi / 2^D$ to the exact angle θ . In this way, we obtain a measurement scheme constructed from finite initial resources, consuming only a fraction of those on average, which is guaranteed to succeed at the expense of only implementing an approximation of U . Since the finite scheme no longer maps the eigenstates of our desired observable O to the direct product basis, a crucial question is then how much the measurement statistics of our approximation differ from the exact case. For a single rotation $R_j(\theta)$, the error can be defined as

$$\begin{aligned} E(\theta, \tilde{\theta}) &= \max_{|\Psi\rangle} \left\| (R_j(\theta) - R_j(\tilde{\theta})) |\Psi\rangle \right\|, \\ &= \max_{|\Psi\rangle} \left\| (\mathbb{1} - e^{-(i/2)\Delta\theta\sigma_j}) |\Psi\rangle \right\|, \end{aligned}$$

where the maximum is taken over all normalized states $|\Psi\rangle$ and $\Delta\theta = \tilde{\theta} - \theta$. The error $E(\theta, \tilde{\theta})$ can be shown [15] to bound the absolute difference between the probabilities P and \tilde{P} for the

outcome of any positive operator valued measurement on $R_j(\theta)|\Psi\rangle$ and $R_j(\tilde{\theta})|\Psi\rangle$, respectively, as $|P - \tilde{P}| \leq 2E(\theta, \tilde{\theta})$. An upper bound to $E(\theta, \tilde{\theta})$ can be obtained by assuming the maximum deviation for $\Delta\theta = \pi/2^{D+1}$, which gives

$$E(\theta, \tilde{\theta}) \leq \sqrt{2} \sqrt{1 - \cos\left(\frac{\pi}{2^{D+2}}\right)} \approx \frac{\pi}{4} 2^{-D}$$

and shows that the error decreases exponentially with D . For a sequence of n rotation chains implementing the binary approximation to U , an important result from quantum computation [35] shows that the overall error is at most the sum of the errors of the individual rotations and so the exponential suppression of the measurement error is retained. We will now finish this work by applying rotation chains to a variety of basic measurement problems. Our results will mostly concentrate on the average entanglement consumption of continuous angle rotation chains but can be equally viewed as an upper bound to the average consumption of any finite binary angle scheme.

5. State verification measurements

Our first application of the tools developed in section 4 is to state verification measurements. A verification of a given state $|\Psi\rangle$ means that the measurement always yields a ‘yes’ result if the system is in the state $|\Psi\rangle$ and a ‘no’ result if the system is in any orthogonal state $|\Psi_\perp\rangle$. If the initial state is a superposition, then the appropriate probabilities for ‘yes’ and ‘no’ results follow from the linearity of quantum mechanics. No assumptions are made about the final state of the system, so there is no requirement that $|\Psi\rangle$ itself is undisturbed by the verification measurement.

5.1. Two-qubit states

To begin, we present a simple scheme that performs a demolition verification of any two-qubit state $|\Psi\rangle \in \mathbb{C}^2 \otimes \mathbb{C}^2$ split between two parties A and B. The construction of a verification scheme for $|\Psi\rangle$ follows from its corresponding Schmidt decomposition

$$|\Psi\rangle = \cos\left(\frac{1}{2}\theta\right) |\phi_0\rangle_A |\phi_0\rangle_B + \sin\left(\frac{1}{2}\theta\right) |\phi_1\rangle_A |\phi_1\rangle_B, \quad (5)$$

where $|\phi_k\rangle$ are Alice’s (Bob’s) local Schmidt states and we have parameterized the corresponding Schmidt coefficients according to an angle $0 \leq \theta \leq \pi/2$. To verify $|\Psi\rangle$, our scheme implements the inverse of the quantum circuit, shown in figure 8(a), that prepares $|\Psi\rangle$ locally. Starting from the initial state $|0\rangle_A |0\rangle_B$ this circuit performs a rotation $R_y(\theta)$ into the state $|\Lambda_1\rangle = \cos(\frac{1}{2}\theta)|0\rangle + \sin(\frac{1}{2}\theta)|1\rangle$ for qubit A, applies a CNOT gate U_{cn} between the pair of qubits controlled by A and is then followed by the product of single-qubit unitaries $V_A \otimes V_B$ that map the computational basis of each qubit into the respective local Schmidt basis as $|k\rangle \mapsto |\phi_k\rangle$, with $k \in \{0, 1\}$.

To invert this, the verification scheme therefore starts with Alice and Bob performing the local unitary transformations V_A^\dagger and V_B^\dagger . Bob then teleports his half of the system B to Alice leaving qubit A and her ancilla qubit a in her possession in the distorted state $\sigma_{0b} V_A^\dagger \otimes V_B^\dagger |\Psi\rangle_{Aa}$

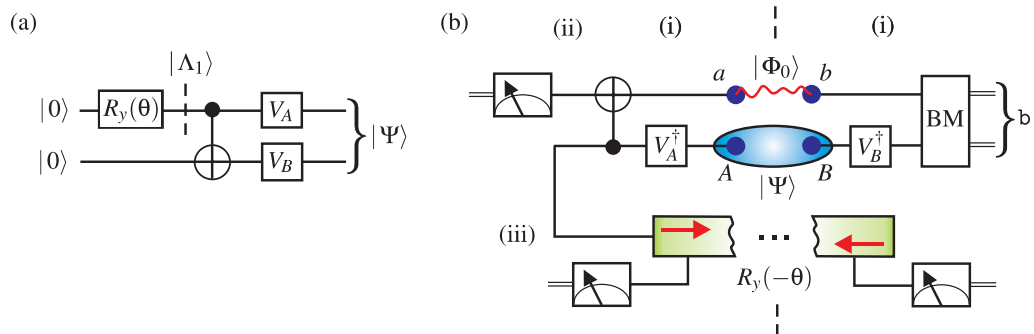


Figure 8. (a) A local circuit that constructs an arbitrary two-qubit state $|\Psi\rangle$ from a standard initial state $|0\rangle|0\rangle$. Firstly, a rotation $R_y(\theta)$ is applied to qubit A forming a single-qubit state $|\Lambda_1\rangle$ composed of a superposition with real amplitudes corresponding to the Schmidt coefficients of $|\Psi\rangle$. This is then followed by a CNOT gate controlled by qubit A and then two arbitrary single-qubit unitaries V_A and V_B are applied, which rotate the computational basis into the required local Schmidt basis of $|\Psi\rangle$. (b) A nonlocal instantaneous verification of the state $|\Psi\rangle$ essentially reverses the circuit shown in (a). In step (i), the inverses of the local unitaries V_A and V_B are applied and the qubit B is teleported to Alice. In step (ii), Alice then applies the CNOT and measures out the received qubit. The most complicated step is (iii) where the rotation $R_y(-\theta)$ is applied to qubit A. This is implemented via a single-qubit rotation chain followed by a measurement of the successful output.

described by his Bell measurement outcome b . This is shown in figure 8(b) and is labelled as step (i). Alice now applies a CNOT gate between qubits A and a . Since the CNOT gate is a stabilizer any distortion can be propagated past it at the expense of spreading the distortion over the control qubit. Regardless of this, Alice can be certain that she has implemented, up to a distortion, $U_{\text{cn}} V_A^\dagger \otimes V_B^\dagger |\Psi_0\rangle = |\Lambda_1\rangle_A |0\rangle_a$ and disentangled qubit A from qubit a . She can then measure qubit a completing step (ii) in figure 8(b). The distortion σ_{0b} ensures that the outcome reveals no information to Alice.

Alice must now map the remaining qubit A, with certainty, into the z -axis so it too can be measured. To achieve this, she needs to apply a rotation $R_y(-\theta)$. Her situation is identical to the scenario considered in section 4.1 and can be readily dealt with using one single-qubit rotation chain as shown in step (iii) of figure 8(b). The average entanglement $\langle e \rangle$ consumed by this nonlocal two-qubit state verification scheme has no dependence on the value of θ except when it is a binary angle. In particular, for a maximally entangled state with $\theta = \pi/2$, precisely 1 ebit is required, while the partially entangled states with $\theta = \pi/4$ or $\theta = \pi/8$ need precisely 2 and 3 ebits to be verified, respectively. Binary angles $\theta = (2m-1)\pi/2^D$, with m integer, have a consumption

$$\langle e_D \rangle = 6 + 2^{2-D} + 2^{-D/2} \left\{ \frac{7}{\sqrt{2}} [-1 + (-1)^D] - 5[1 + (-1)^D] \right\}. \quad (6)$$

For any angle θ that is not binary $\langle e \rangle = 6$ ebits on average and is independent of the entropy of entanglement of the state $|\Psi\rangle$. As expected, this consumption is the asymptotic limit $D \rightarrow \infty$ of equation (6).

Although this measurement scheme was devised to verify a single state $|\Psi\rangle$, the ‘no’ results do in fact verify a special set of states in the orthogonal complement,

$$\begin{aligned} |\Psi_1\rangle &= \cos\left(\frac{1}{2}\theta\right) |\phi_0\rangle_A |\phi_1\rangle_B - \sin\left(\frac{1}{2}\theta\right) |\phi_1\rangle_A |\phi_0\rangle_B, \\ |\Psi_2\rangle &= \sin\left(\frac{1}{2}\theta\right) |\phi_0\rangle_A |\phi_0\rangle_B + \cos\left(\frac{1}{2}\theta\right) |\phi_1\rangle_A |\phi_1\rangle_B, \\ |\Psi_3\rangle &= \sin\left(\frac{1}{2}\theta\right) |\phi_0\rangle_A |\phi_1\rangle_B - \cos\left(\frac{1}{2}\theta\right) |\phi_1\rangle_A |\phi_0\rangle_B, \end{aligned} \quad (7)$$

which are related to $|\Psi\rangle$ in the same way that the Bell states are related to $|\Phi_0\rangle$. The scheme is therefore a verification measurement of an operator possessing these states, along with $|\Psi\rangle$, as eigenstates. When $\theta = \pi/2$ the scheme is the demolition verification measurement of the Bell operator already presented in figure 2(b). We shall consider shortly in section 6 the more complicated task of simultaneously verifying an arbitrary set of eigenstates.

5.2. Bipartite multi-qubit states

The verification scheme for two-qubit states can be generalized for any state $|\Psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ split between two parties A and B, where $\mathcal{H}_A = (\mathbb{C}^2)^{\otimes v}$ and $\mathcal{H}_B = (\mathbb{C}^2)^{\otimes w}$ are tensor products of qubits. Again the scheme operates by performing a nonlocal unitary U , which maps $|\Psi\rangle$ to a locally measurable state as $U|\Psi\rangle \mapsto |0, 0, \dots, 0\rangle$, modulo Pauli distortions. As with two qubits the scheme focuses on the Schmidt decomposition of $|\Psi\rangle$, which now takes the form

$$|\Psi\rangle = \sum_{\alpha=1}^{\chi} \lambda_{\alpha} |\phi_{\alpha}\rangle_A |\phi_{\alpha}\rangle_B,$$

where $\chi \leq \min(2^v, 2^w)$ is the Schmidt rank designating the number of nonzero λ_{α} Schmidt coefficients satisfying $\sum_{\alpha} \lambda_{\alpha}^2 = 1$, and $|\phi_{\alpha}\rangle$ are Alice’s (A) or Bob’s (B) local Schmidt states. Before starting the verification scheme, Alice and Bob use this canonical form for the target state to determine local unitaries V_A^{\dagger} and V_B^{\dagger} , which can be applied to their v - and w -qubit subsystems, respectively, to map their local Schmidt states into the computational basis. For either party, this takes the form

$$V^{\dagger} |\phi_{\alpha}\rangle = |\vec{\alpha}, 0, \dots, 0\rangle,$$

where $\vec{\alpha}$ is a $d = \lceil \log_2(\chi) \rceil$ dimensional binary vector representing the integer index α and $|\vec{\alpha}\rangle$ is a d -fold tensor product of the σ_z eigenstates $|0\rangle$ and $|1\rangle$. The action of the V ’s on the orthogonal complement to the subspace spanned by the local Schmidt states $|\phi_{\alpha}\rangle$ can be defined arbitrarily. The resulting state $|\psi\rangle = V_A^{\dagger} \otimes V_B^{\dagger} |\Psi\rangle$ is then entirely contained in the smallest possible subspace of the original $(v+w)$ -qubit system composed of two equal-sized d -qubit subsystems at A and B. Once this initial compression is performed the remaining $v-d$ and $w-d$ qubits at Alice and Bob’s location containing (some of) the orthogonal complement to $|\psi\rangle$ can be immediately measured in the computational basis. Any outcome other than $|0\rangle$ for each qubit indicates an immediate ‘no’ result.

Having mapped the target state $|\Psi\rangle$ to $|\psi\rangle$, the scheme then continues by implementing the inverse of the circuit which locally constructs $|\psi\rangle$ from the $2d$ -qubit initial state $|0, 0, \dots, 0\rangle$. Specifically this construction circuit begins by creating a superposition state $|\Lambda_d\rangle$ on the first d qubits (generalizing $|\Lambda_1\rangle$ from earlier) of the form

$$|\Lambda_d\rangle = \sum_{\vec{x}} \lambda_{\vec{x}} |\vec{x}\rangle,$$

where $\lambda_{\vec{x}}$ are the real Schmidt coefficients of $|\Psi\rangle$ indexed by the d -dimensional binary vector \vec{x} and appropriately padded with zeros if necessary. This type of superposition state can be formed by a cascade $F_1^0 F_2^1 F_3^2 \dots F_d^{d-1}$ of so-called uniformly controlled rotations (see appendix B and [36] for details) about the y -axis acting on the first set of d qubits. Once the state $|\Lambda_d\rangle \otimes |0, \dots, 0\rangle$ has been generated, a *staircase* sequence of CNOT gates is applied between pairs of qubits from the first set of d and the second set of d (see figure C.1). This then constructs the canonical Schmidt form for the state $|\psi\rangle$ as

$$|\psi\rangle = \sum_{\vec{x}} \lambda_{\vec{x}} |\vec{x}\rangle \otimes |\vec{x}\rangle.$$

A more detailed description of this circuit is given in appendix C where it is shown explicitly for $d = 4$ qubits in figure C.1.

Given this construction circuit the verification scheme proceeds with Bob teleporting his d qubits to Alice. She then implements the sequence of CNOT gates locally on the $2d$ qubits in her possession. Since this part of the circuit is a stabilizer it is guaranteed to succeed but will propagate Pauli distortions originally confined to the ancilla qubits receiving Bob's half of the system to Alice's half. This leaves a state of the form $\sigma_j |\Lambda_d\rangle \otimes |0, \dots, 0\rangle$ in Alice's possession, but with only Bob knowing j . Since the CNOTs have successfully disentangled the two halves the qubits originating from Bob are now in a product state in the computational basis and can be measured immediately. Alice is now left with her d qubits, which require the final sequence of uniformly controlled rotations to be applied. The decomposition of the cascading sequence of uniformly controlled rotation into Pauli rotations requires $2^d - 1$ distinct gates (figure B.1(b) shows this decomposition for F_3^2), which, as expected, is identical to the number of independent rotation angles defining $|\Lambda_d\rangle$. The scheme then implements these rotations by concatenating rotation chains. The complete nonlocal verification scheme for $|\Psi\rangle$ is shown in figure 9.

Combining the scaling in the number of rotations with that of the average consumption $\langle c_n \rangle$ for concatenated rotations in equation (3) yields an exponential of an exponential scaling

$$\langle e \rangle = C d \phi^{2^d - 1} \text{ ebits},$$

with the minimum number of qubits d required to contain the Schmidt rank of the target state $|\Psi\rangle$. While this consumption lacks any dependence on the values of the angular parameters (excluding binary angles), it does depend on the entanglement in $|\Psi\rangle$ as measured by the Schmidt rank. Note that the choice of unitary U that can implement a verification of a single state $|\Psi\rangle$ is not unique. However, the choice made in this scheme is unique in the sense that it is defined only by the nonlocal parameters of the target state itself and is therefore the most economical. Also, similar to the two-qubit case, this state verification scheme is also a verification measurement of a special operator whose complete set of eigenstates spanning the orthogonal complement also happens to be mapped to locally measurable states.

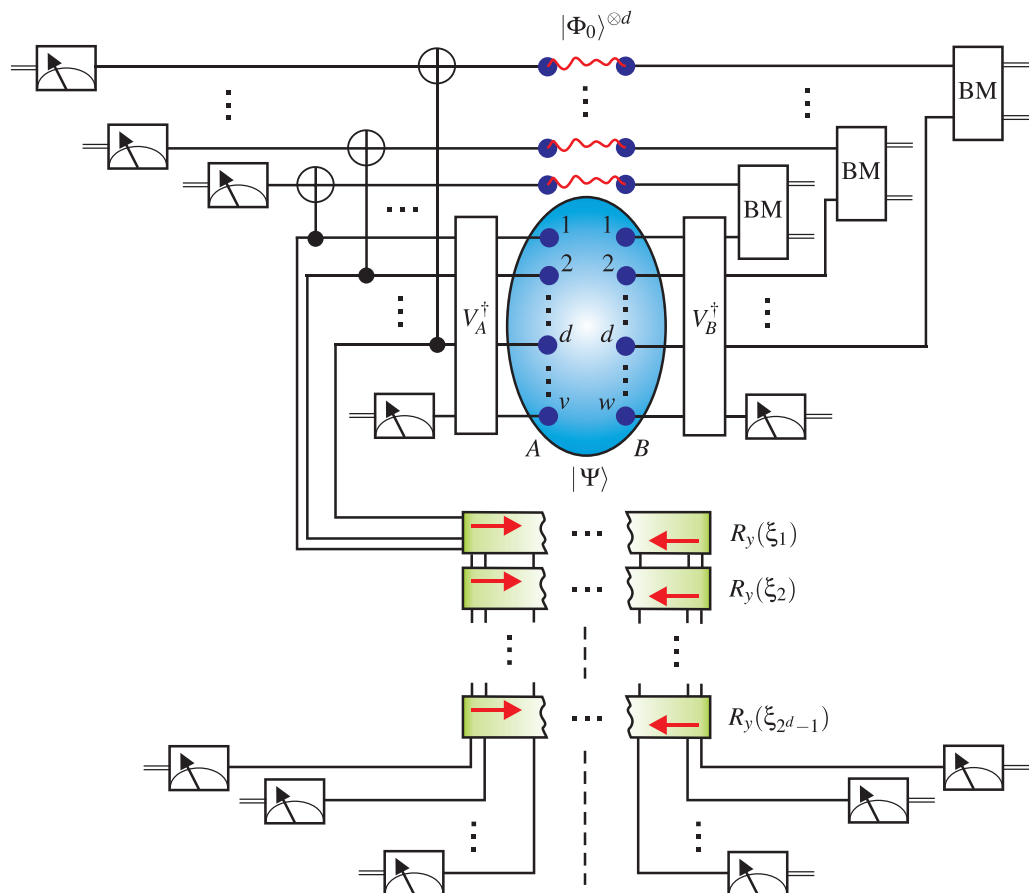


Figure 9. The nonlocal verification scheme for a state $|\Psi\rangle \in (\mathbb{C}^2)^{\otimes v} \otimes (\mathbb{C}^2)^{\otimes w}$. After performing local unitaries V_A^\dagger and V_B^\dagger that map the target state $|\Psi\rangle$ to a $2d$ -qubit state $|\psi\rangle$ and measuring out the orthogonal complement, Bob teleports his d qubits to Alice. Following the inverse of the circuit in figure C.1 Alice performs a sequence of CNOT gates between her d qubits and those received from Bob, with the latter being measured immediately afterwards. Alice then inverts the sequence of uniformly controlled rotations in the y -axis that produce $|\Lambda_d\rangle$ via $2^d - 1$ concatenated d -qubit rotation chains. Note that although the concatenated rotation chains are drawn sequentially they should be understood as forming a massively recursive structure. The output from the final chain is then measured. We have also ignored here the optimization that successive sets of rotation chains act on smaller number of qubits due to the structure of the circuit in figure C.1.

6. Instantaneous measurements of nonlocal operators

We now generalize the measurement schemes introduced so far to perform a simultaneous demolition verification of each of the nondegenerate eigenstates of an arbitrary nonlocal observable O . Our strategy is again to implement a nonlocal unitary U that maps each eigenstate of O into a different computational basis state which is then locally measurable. Unlike the state verification scheme, which is already a special class of operator measurement, here we are interested in complete generality.

6.1. Two-qubit observables

Before outlining a scheme for the most general case, we first describe some schemes for special classes of eigenstates for two qubits. A particularly interesting class of observables is those with a *twisted* eigenbasis,

$$\begin{aligned}
 |\Psi_0\rangle &= |0\rangle_A |0\rangle_B, \\
 |\Psi_1\rangle &= |0\rangle_A |1\rangle_B, \\
 |\Psi_2\rangle &= |1\rangle_A \left[\sin\left(\frac{1}{2}\theta\right) |0\rangle_B + e^{i\varphi} \cos\left(\frac{1}{2}\theta\right) |1\rangle_B \right], \\
 |\Psi_3\rangle &= |1\rangle_A \left[\cos\left(\frac{1}{2}\theta\right) |0\rangle_B - e^{i\varphi} \sin\left(\frac{1}{2}\theta\right) |1\rangle_B \right].
 \end{aligned} \tag{8}$$

Despite these eigenstates being product states, it has been shown that if ideal measurements of this basis were possible it would allow violations of causality [12, 17]. Unlike the direct (or untwisted) product basis, a verification measurement of the twisted product basis requires entanglement [17]. As seen in figure 10(a), the circuit that generates this basis locally can straightforwardly yield the nonlocal measurement scheme in figure 10(b) which utilizes just one single-qubit rotation chain. The average entanglement consumption for this basis is dependent on the eigenstate requiring $\langle e \rangle = 4$ ebits for $|\Psi_0\rangle$ and $|\Psi_1\rangle$ (where no rotation is needed), or $\langle e \rangle = 6$ ebits for $|\Psi_2\rangle$ and $|\Psi_3\rangle$. In this way, the measurement of the twisted basis is very similar to the eigenbasis in equation (7) encountered for state verification. There the eigenbasis was composed of equally but partially entangled eigenstates and needed $\langle e \rangle = 6$ ebits for all eigenstates. The consumption for entangled eigenstates, however, grows quickly even with a slight generalization. For instance, adding an identical relative phase $e^{i\varphi}$ to all of the basis states in equation (7) necessitates the concatenation of two single-qubit rotation chains (first for the z -axis and second for the y -axis) and elevates the consumption to $\langle e \rangle = 21$ ebits. Generalizing further gives an eigenbasis composed of partially but unequally entangled eigenstates with differing relative phases, as

$$\begin{aligned}
 |\Psi_0\rangle &= \sin\left(\frac{1}{2}\theta_1\right) |0\rangle_A |0\rangle_B + \cos\left(\frac{1}{2}\theta_1\right) e^{i\varphi_1} |1\rangle_A |1\rangle_B, \\
 |\Psi_1\rangle &= \cos\left(\frac{1}{2}\theta_1\right) |0\rangle_A |0\rangle_B - \sin\left(\frac{1}{2}\theta_1\right) e^{i\varphi_1} |1\rangle_A |1\rangle_B, \\
 |\Psi_2\rangle &= \sin\left(\frac{1}{2}\theta_2\right) |0\rangle_A |1\rangle_B + \cos\left(\frac{1}{2}\theta_2\right) e^{i\varphi_2} |1\rangle_A |0\rangle_B, \\
 |\Psi_3\rangle &= \cos\left(\frac{1}{2}\theta_2\right) |0\rangle_A |1\rangle_B - \sin\left(\frac{1}{2}\theta_2\right) e^{i\varphi_2} |1\rangle_A |0\rangle_B,
 \end{aligned}$$

described by four real parameters. From the local preparation circuit shown in figure 10(c), which contains two uniformly controlled rotations, a total of four Pauli rotations are required (one for each parameter). The corresponding nonlocal measurement scheme is shown in figure 10(d). The first three rotations require two-qubit chains, whereas the last acts only on the second qubit and so can be reduced to a single-qubit chain. Following the calculation in

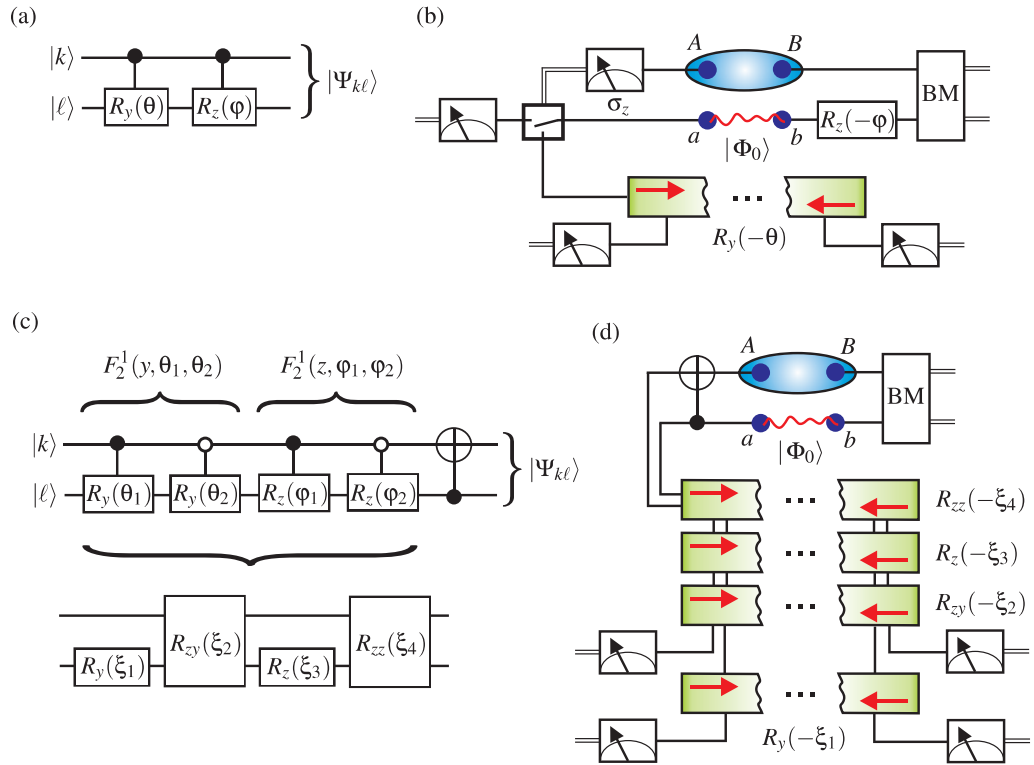


Figure 10. (a) A local circuit that constructs the twisted basis set $|\Psi_{k\ell}\rangle$ from the computational basis $|k\rangle|\ell\rangle$. Here the indices k and ℓ are bits that together as $k\ell$ are a binary representation of the $\{0, 1, 2, 3\}$ index used in equation (8). Firstly, a controlled rotation $R_y(\theta)$ is applied to form the twist and then another rotation $R_z(\varphi)$ is performed to introduce the phase. (b) The nonlocal measurement scheme that performs the inverse of (a). The phase can be removed locally by Bob, while the final controlled rotation $R_y(-\theta)$ can be replaced by a classical control. If a rotation is required, it is implemented by a rotation chain with the output being measured in the z -axis. (c) The local circuit that constructs a general partially entangled basis set from $|k\rangle|\ell\rangle$. Dependent on the state $|k\rangle$ the second qubit is rotated about the y - and z -axis by different angles according to two uniformly controlled rotations, followed by a CNOT gate that entangles them. The uniformly controlled rotations can be decomposed into the sequence of Pauli rotations shown. (d) The nonlocal measurement scheme that performs the inverse of (c). A concatenated sequence of two-qubit rotation chains is applied implementing the inverse of the Pauli rotation decomposition in (c). Note that the final rotation is applied to one qubit only.

appendix A, the average entanglement consumption for the measurement of this eigenbasis is $\langle e \rangle = 224$ ebits.

To devise a scheme to deal with the most general eigenbasis, we require a circuit composed only of Pauli rotations, each of which can be handled with rotation chains, that can build a general $SU(4)$ unitary U . For two-qubits this can be accomplished by using the so-called Cartan

decomposition [34, 37] of an SU(4) unitary as

$$U = (V_A \otimes V_B) e^{(i/2)\xi_1\sigma_1\otimes\sigma_1} e^{(i/2)\xi_2\sigma_2\otimes\sigma_2} e^{(i/2)\xi_3\sigma_3\otimes\sigma_3} (W_A \otimes W_B),$$

where V_A , V_B , W_A and W_B are single-qubit SU(2) gates, and $\pi/2 \geq \xi_1 \geq \xi_2 \geq |\xi_3| \geq 0$. The Cartan decomposition has been extremely popular in recent work [38]–[40] on quantum circuits, since it beautifully exposes the nonlocal content of any two-qubit unitary. Rather than needing to consider all 15 real parameters the classification of two-qubit unitaries reduces to the three coordinates (ξ_1, ξ_2, ξ_3) and allows the set of locally inequivalent gates to be characterized geometrically as points within a tetrahedron [39]. In the context of nonlocal measurements, the first pair of unitaries W_A and W_B can be trivially applied by each party locally before the start of the scheme. If the last pair of single-qubit unitaries are then decomposed as a sequence of rotations $V = R_z(\alpha)R_y(\beta)R_z(\gamma)$, we see that the U is expressed entirely in terms of Pauli rotations. Furthermore, since our final measurement after U will be in the z -axis, the latter R_z rotation for either of the local V unitaries is not necessary. This leaves seven real parameters relevant for the nonlocal measurement.

To compute the average entanglement consumption in this most general case, we perform one optimization. Rather than simply concatenating seven two-qubit rotation chains (which would consume $2\langle c_7 \rangle + 1 = 4719$ ebits on average), we instead split up the qubits after the three nonlocal gates and perform the final two single-qubit rotations on them separately and simultaneously¹¹. A simple modification of the calculation in appendix A shows that this splitting gives a consumption equivalent to five two-qubit rotation chains and so the average entanglement consumption for the most general two-qubit observable is $2\langle c_5 \rangle + 1 = 787$ ebits. Finally, recall from section 4.3 that the average consumptions quoted above are upper-bounds to those that would be attained if the angles involved were binary. For example the twisted basis measurement instead consumes at most an average of $\langle e \rangle = 3$ ebits if $\theta = (2m - 1)\pi/8$, where m is an integer.

6.2. Bipartite multi-qubit observables

The situation for mapping the eigenstates of a nonlocal d -qubit observable to the computational basis is less clear due to the lack of an optimal quantum circuit construction for arbitrary SU(2^d) unitaries. On general grounds, an exponential number of one-parameter Pauli rotations is expected to be required, since an SU(2^d) unitary is defined by $4^d - 1$ reals. However, as the two-qubit case illustrates, not all these parameters are relevant for nonlocal measurements. Recent work [36] on quantum circuits allows us to identify (although not optimally) some of these redundant local parameters and, moreover, provides an explicit construction of such a circuit decomposition in terms of Pauli rotations. By exploiting a cosine–sine decomposition recursively, a circuit composed only of uniformly controlled rotations was devised in [36]. So far for $d \geq 4$ qubits this construction represents the most efficient circuit decomposition in terms of the number of CNOT and elementary single-qubit gates needed. For our purposes, the important aspects of this construction are that an SU(2^d) unitary can be formed from a circuit of $2^{d+1} - 2$ uniformly controlled rotations F_{d-1}^d , alternating between the y - and z -axes, followed by a cascade of d uniformly controlled rotations $F_1^0 F_2^1 \dots F_{d-1}^d$ involving sequentially decreasing numbers of qubits and all in the z -axis. To illustrate this a complete decomposition [36] of

¹¹ Splitting the qubits up can only be done once they never need to interact again. Once separated the qubits progress along different pathways through the scheme and no party knows precisely where the actual pair are located.




Figure 11. The quantum circuit for an arbitrary $d = 3$ qubit $SU(2^d)$ unitary U in terms of 17 uniformly controlled rotation gates (see [36] and appendix B for more details of these gates). The first $2^{d+1} - 2 = 14$ gates alternate between rotations in the z - and y -axis. The last d gates, which are shaded, are all in the z -axis and for nonlocal measurements where U is to map our desired eigenstates to the direct product basis, this final cascade of gates can be ignored.

a $d = 3$ qubit gate is shown in figure 11. If this type of decomposition is used in a nonlocal measurement scheme, then the final cascade (shaded in the example in figure 11) can be ignored since all qubit measurements terminating the circuit are performed in the computational basis. As shown in appendix C, each F_{d-1} gate requires 2^{d-1} Pauli rotations, equal to the number of reals defining it. Thus, using this circuit construction, $4^d + 2^d$ concatenated d -qubit rotation chains are needed to implement an arbitrary $SU(2^d)$ unitary. An exponential of an exponential scaling with the number of qubits d again arises for the average entanglement consumption $\langle e \rangle$ for a nonlocal measurement of a bipartite multi-qubit observable.

7. Conclusions

In this work, we have studied in detail the average entanglement consumption for both nonlocal state verification and operator measurements. The approach applied was similar to that of earlier work [17] where teleportation was employed to first localize the system and then used in a multi-round protocol to implement the mapping U from a general set of states into a locally measurable set. The central advancement here is that in contrast to previous schemes [16, 17] this can be done by consuming only a finite amount of entanglement on average, even in the most general cases, while continuing to succeed with certainty. The reason for this is that the application of U is broken up into a sequence of Pauli rotations $R_j(\theta)$. By expressing teleportation in terms of Pauli distortions, a decomposition of this type has the privileged feature that distortions at each step either leave the operation $R_j(\theta)$ intact or produce only one type of failure, namely $R_j(-\theta)$. This enabled us to construct a rotation chain scheme with a bipartite termination condition that applies $R_j(\theta)$ with certainty and consumes only a finite amount of the initial entanglement on average. Moreover, we showed how rotation chains can be concatenated forming a recursive structure that permits arbitrarily complex sequences of them to be applied while retaining a finite average consumption overall. As an aside this result also shows that bipartite distributed quantum computation can be performed instantaneously with only a finite average entanglement consumption. Interesting comparisons with distributed cluster state one-way quantum computation could be made [31, 32].

Despite the finiteness of the entanglement consumption, its growth is found to display an extremely unfavourable exponential of an exponential scaling with either the Schmidt rank of the state to be verified, or size of the system on which the nonlocal observable acts. While this scaling is scheme dependent, there is good reason to believe that it is fundamental to the

underlying problem. Indeed, both the complexity of constructing circuit decompositions of a general unitary and the recursive protocol required to overcome the no-signalling constraints individually display exponential scaling. Whether causality forces any conceivable nonlocal measurement scheme to have this combination of scalings is an open problem.

Our aim here has been to prove that general nonlocal measurements can be accomplished with certainty while consuming on average only a finite amount of entanglement. While achieving this, the resulting scheme has not been proven to be optimal. Specifically, the consumption in our schemes has no dependence on the actual value of the various rotation angles that appear, beyond the special case of binary angles. Instead, the consumption is always averaged over integer units of ebits and the resulting measure of complexity of the required unitary is coarse-grained to simply counting the number of nontrivial rotation angles specifying it. It is possible that a more efficient scheme can be devised where the entangled resources are qubit pairs that are partially entangled, in a way that is linked to the rotation angles, thereby providing a tailored resource and an angle-dependent entanglement consumption even for continuous angles.

Another important deficiency of the schemes presented is that they do not yet represent a practical deterministic measurement procedure due to the infinite amount of entanglement that needs to be initially distributed. Here a finite average consumption arises, because we have introduced termination conditions for both parties. The requirement for an infinite amount of initial distributed resources appears to be of a different origin, namely the continuous real parameters that appear in the problem. An important exception to this was shown to occur for angles that are binary fractions of π where only a finite amount of initial entanglement is needed [16]. The measurement of the Bell operator is an extreme example of this. Using this result we considered the experimentally relevant case where arbitrary rotation angles are discretized to binary angles. We showed that this results in nonlocal measurement, which is still certain to succeed, but requires only finite initial resources. The resulting measurement performed is an approximation to the exact one and we bounded the error of this procedure. Although not proven it appears unlikely that an exact protocol exists for the most general measurement which succeeds with certainty and requires only a finite amount of initial entanglement. Finally, unlike Vaidman's scheme [17] and stabilizer measurements, our rotation chain methods do not easily generalize to more than two parties and so an interesting open problem is whether all multi-party nonlocal measurements can be done with a finite average entanglement consumption.

Acknowledgments

SRC and DJ thank the National Research Foundation and the Ministry of Education of Singapore for support. DJ acknowledges support from the ESF program EuroQUAM (EPSRC grant EP/E041612/1), the EPSRC (UK) through QIP IRC (GR/S82176/01) and the European Commission under the Marie Curie programme through QIPEST. AJC thanks Keble College, Oxford for providing funding.

Appendix A. Computing the average entanglement consumption

In this section, the calculation of the average entanglement consumption for the rotation chains used in our scheme is described. We first calculate the average number of channels (complete

teleportations of the system) required for the implementation of a single rotation. Using this result and the recursive structure of the scheme, we then calculate the average consumption for two rotations and finally generalize this to an arbitrary number of rotations concatenated together.

A.1. A single rotation chain

Following the discussion in section 4.1, we consider a chain where Alice possesses the entire system initially and begins the protocol as in figure 5. The probability that Alice terminates at her q th step while Bob terminates at his p th step is given by $(\frac{1}{2})^{p+q}$. The consumption of channels is governed by the last party to terminate and is denoted by c_1 . If Alice terminates last and at her q th opportunity, then the total number of channels used will be $c_1 = 2q$ and, hence, even. Likewise Bob terminating last at his p th opportunity gives a consumption $c_1 = 2p - 1$ and is odd. Note that a rotation chain has a minimum consumption of two channels and consequently for Bob to terminate last we require $p \geq 2$. The average number of teleportations $\langle c_1 \rangle$ is then easily calculated as a sum of the case when Alice exits the chain first and the case when Bob exits the chain first, as

$$\langle c_1 \rangle = \sum_{q=1}^{\infty} \sum_{p=1}^q \left(\frac{1}{2} \right)^{q+p} 2q + \sum_{p=2}^{\infty} \sum_{q=1}^{p-1} \left(\frac{1}{2} \right)^{q+p} (2p - 1) = 5.$$

In order to calculate the average consumption when two rotation chains are concatenated, we need to introduce two more average consumptions of a single chain. Specifically, we define $\langle a_1 \rangle$ as the average consumption when only the initiating party is actually performing any actions on the chain and likewise $\langle b_1 \rangle$ for the case where only the receiving party is performing any actions on the chain. These are readily computed as

$$\langle a_1 \rangle = \sum_{q=1}^{\infty} \left(\frac{1}{2} \right)^q 2q = 4, \quad \text{and} \quad \langle b_1 \rangle = \sum_{p=1}^{\infty} \left(\frac{1}{2} \right)^p (2p - 1) = 3.$$

In the next two subsections we shall generalize these quantities to $\langle a_n \rangle$, $\langle b_n \rangle$ and $\langle c_n \rangle$ to designate the corresponding average consumptions for n chains concatenated where only the initiating party, only the receiving party, or both parties are performing actions from the start, respectively.

A.2. Two rotation chains concatenated

As with a single chain, the consumption for two concatenated rotation chains breaks into two cases depending on whether Alice or Bob exits the first chain last. In figures 6 and A.1, the latter situation is illustrated with Alice exiting the first chain on the q th opportunity, while Bob exits on his p th, with $p > q$. Up to her exit point Alice must play the role of the receiving party on all the second chains Bob has available to him at his exit points. Since, for this case, Bob has not used any of these chains Alice consumes $q\langle b_1 \rangle$ channels on average through these redundant actions. Similarly, Bob must be a receiving party for all $p - 1$ of Alice's second chains up to his exit point p consuming $(p - 2)\langle b_1 \rangle + \langle c_1 \rangle$ channels on average, with the $\langle c_1 \rangle$ accounting for the fact that one second channel (the q th) was used by both parties. At his exit point Bob will

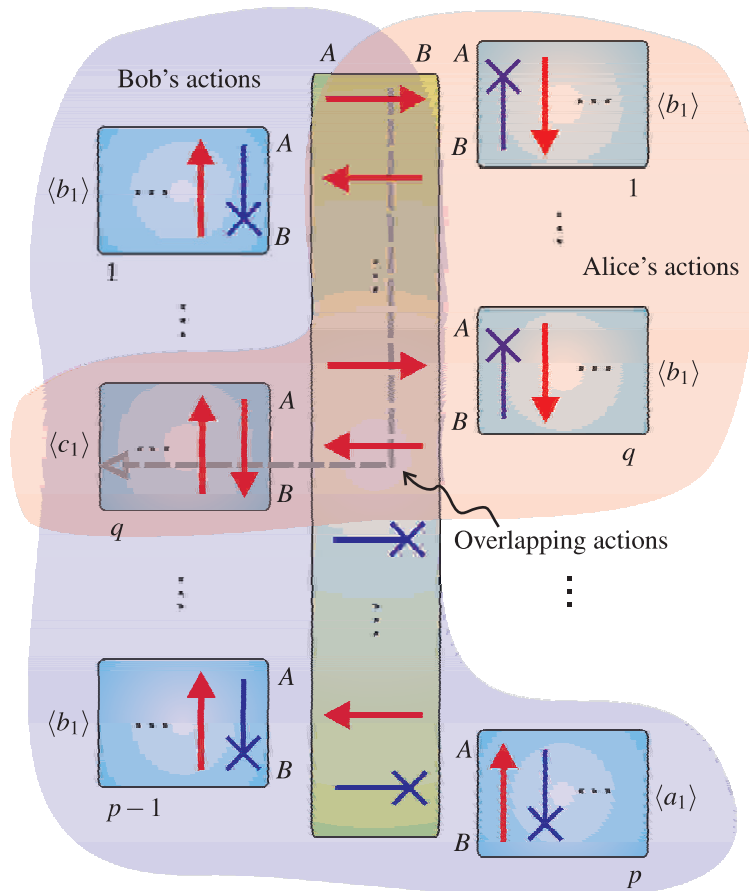


Figure A.1. Another version of figure 6 but with the consumption quantities $\langle a_1 \rangle$, $\langle b_1 \rangle$ and $\langle c_1 \rangle$ labelled for the appropriate second chains. The independent actions of Alice and Bob are shaded for the case where Bob exits the first chain last. The dashed 'L'-shaped line indicates the actual path taken by the principal system and lies where their actions overlap.

consume a further $\langle a_1 \rangle$ channels for the second chain which only he acts on. Finally, since Bob exits last (so $p \geq 2$) the consumption of channels in the first chain will be $2p - 1$. Performing the analogous counting of channels for the opposite case where Alice exits the first chain last and averaging over all the exit points p, q of the first chain with probabilities $(\frac{1}{2})^{q+p}$ gives

$$\langle c_2 \rangle = 5 + \langle c_1 \rangle + \langle a_1 \rangle + 2\langle b_1 \rangle = 20.$$

It is clear from this that the quantity $\langle r_2 \rangle = \langle a_1 \rangle + 2\langle b_1 \rangle$ represents the cost of recursion within the protocol, which in this case doubles the consumption from that expected for two independent rotation chains. We can similarly compute the one-party consumptions for two rotations as $\langle a_2 \rangle = 4 + \langle a_1 \rangle + 2\langle b_1 \rangle$ and $\langle b_2 \rangle = 3 + \langle a_1 \rangle + \langle b_1 \rangle$.

A.3. Concatenating n rotation chains

The generalization to n concatenated rotation chains can be computed straightforwardly by using the recursive structure of the protocol. The calculation proceeds in an identical way to

two chains except that each second chain itself is now regarded as a sequence of $n - 1$ chains. This gives the linked recurrence relations for the component consumptions

$$\langle c_n \rangle = 5 + \langle c_{n-1} \rangle + \langle a_{n-1} \rangle + 2\langle b_{n-1} \rangle,$$

$$\langle a_n \rangle = 4 + \langle a_{n-1} \rangle + 2\langle b_{n-1} \rangle,$$

$$\langle b_n \rangle = 3 + \langle a_{n-1} \rangle + \langle b_{n-1} \rangle.$$

After denoting the recursive consumption as $\langle r_n \rangle = \langle a_{n-1} \rangle + 2\langle b_{n-1} \rangle$, we see that it obeys a closed recurrence relation $\langle r_n \rangle = 3\langle r_{n-1} \rangle - \langle r_{n-2} \rangle - \langle r_{n-3} \rangle$. Given the recursive consumptions $\langle r_1 \rangle = 0$, $\langle r_2 \rangle = 10$ and $\langle r_3 \rangle = 34$, a closed solution for $\langle r_n \rangle$ can be found as

$$\langle r_n \rangle = A\phi^n + B\left(\frac{-1}{\phi}\right)^n - 7,$$

where $A = (3 + 2\sqrt{2})/2$, $B = (3 - 2\sqrt{2})/2$ and $\phi = 1 + \sqrt{2}$. For all but the smallest n , the recursive consumption $\langle r_n \rangle$ is well approximated by only the first term and so, as might be anticipated, displays a pure exponential growth with n . Since the total average consumption is $\langle c_n \rangle = 5n + \sum_{k=1}^n \langle rk \rangle$, it also displays a pure exponential scaling asymptotically as

$$\langle c_n \rangle \sim A\left(\sum_{k=1}^n \frac{1}{\phi^k}\right)\phi^n \approx C\phi^n,$$

where $C = (10 + 7\sqrt{2})/4$. As shown in figure 7 this approximation to the exact consumption is already very good once $n > 4$.

Appendix B. Uniformly controlled rotations

We make repeated use of a special sequence of multi-qubit controlled rotation gates, which, following the nomenclature of [36], is called a uniformly controlled rotation. This gate is denoted as $F_n^k(\mathbf{a}, \vec{\theta})$ and signifies a k -fold controlled rotation of some qubit n about the three-dimensional axis \mathbf{a} by one of the 2^k different rotation angles contained in $\vec{\theta} = (\theta_1, \theta_2, \dots, \theta_{2^k})$. The uniformly controlled rotation where qubits $1, \dots, n-1$ are the controls and qubit n is the target has a matrix representation

$$F_n^{n-1}(\mathbf{a}, \vec{\theta}) = \begin{pmatrix} R_{\mathbf{a}}(\theta_1) & & \\ & \ddots & \\ & & R_{\mathbf{a}}(\theta_{2^{n-1}}) \end{pmatrix}.$$

This gate is motivated by its easily interpreted action, namely it can be seen to implement a different rotation angle on qubit n dependent on each of the 2^{n-1} basis configurations of the control qubits. In figure B.1(a), the circuit defining F_4^3 is shown. For our applications we shall exclusively consider rotations in either the y -axis $F_n^k(y)$ or z -axis $F_n^k(z)$ and we will frequently use a construction that decomposes such F_n^k 's into 2^k single-parameter Pauli




Figure B.1. (a) The circuit of multi qubit controlled rotations that constructs the uniformly controlled rotation $F_4^3(\mathbf{a}, \vec{\theta})$ about an axis \mathbf{a} defined by an eight-component vector of angles $\vec{\theta}$. The gate symbol we use for a uniformly controlled rotation is on the left with grey circles. (b) A decomposition in terms of Pauli rotations is shown for a uniformly controlled rotation $F_3^2(y, \vec{\theta})$ about the y -axis and defined by a four-component vector $\vec{\theta}$. The corresponding Pauli rotation angles given are related to the four angles in $\vec{\theta}$ as $\xi_1 = -(\theta_1 + \theta_2 + \theta_3 + \theta_4)/8$, $\xi_2 = (\theta_2 - \theta_1 - \theta_3 + \theta_4)/8$, $\xi_3 = (\theta_3 + \theta_4 - \theta_1 - \theta_2)/8$ and $\xi_4 = (\theta_2 + \theta_3 - \theta_1 - \theta_4)/8$. This decomposition readily generalizes to uniformly controlled rotations involving larger numbers of qubits.

rotations. Specifically for a uniformly controlled rotation $F_n^k(y, \vec{\theta})$ this construction involves performing a single-qubit rotation R_y on qubit n , followed by two-qubit rotations R_{zy} between each of the k control qubits and qubit n , followed by three qubit rotations R_{zzy} between every pair of the k control qubits and qubit n , and so on until a final rotation $R_{zz\cdots zy}$ is performed involving all the k control qubits and qubit n . For this example, the Pauli strings for the rotations always specify a σ_y on qubit n and σ_z on any of the k control qubits. Each of the 2^k rotations involves a different rotation angle that itself is a linear combination of the angles in $\vec{\theta}$. A detailed example of this decomposition for a single $F_3^2(y)$ gate is given in figure B.1(b), while in figure 10(b) a decomposition for the pair of gates $F_2^1(y)F_2^1(z)$ is depicted.

Appendix C. Constructing a Schmidt superposition state

For bipartite multi-qubit state verification in section 5.2, a circuit is required that generates a normalized state in a superposition of all 2^d computational basis states with arbitrary real amplitudes. Such a superposition state can be parameterized in terms of $2^d - 1$ angles $0 \leq \theta_j \leq \pi$ with amplitudes given by

$$\lambda_{\vec{x}} = \cos\left(\frac{1}{2}\Theta_{x_1}^{[1]}\right) \cos\left(\frac{1}{2}\Theta_{x_1x_2}^{[2]}\right) \cdots \cos\left(\frac{1}{2}\Theta_{x_1x_2\cdots x_d}^{[d]}\right),$$




Figure C.1. The quantum circuit for constructing the state $|\psi\rangle$ composed of eight qubits. The first part of the circuit constructs the four-qubit state $|\Lambda_4\rangle$ using a cascade of uniformly controlled rotations. The resulting state $|\Lambda_4\rangle$ is a normalized superposition of each of the 2^4 computational basis states with real amplitudes parameterized by the 15 angles $\theta_0, \dots, \theta_{14}$. These angles are chosen to be the Schmidt coefficients of the target state $|\psi\rangle$. The final part of the circuit performs a sequence of CNOT gates between the first four qubits and an additional four creating the canonical Schmidt form for $|\psi\rangle$. This circuit can be readily generalized to larger numbers of qubits. In particular, the structure of the first part of the circuit is based on taking the previous $k-1$ qubits in the state $|\Lambda_{k-1}\rangle$, adding qubit k in the state $|0\rangle$ and performing a further uniformly controlled rotation F_k^{k-1} . The resulting k qubits are then expanded to the state $|\Lambda_k\rangle$ and an additional 2^{k-1} angles are introduced into its parametrization. See appendix C for more details.

where the angles Θ are defined from θ via

$$\begin{aligned}
 \Theta_0^{[1]} &= \theta_0 & \Theta_{00}^{[2]} &= \theta_1 & \Theta_{000}^{[3]} &= \theta_3 & \dots \\
 \Theta_1^{[1]} &= \theta_0 - \pi & \Theta_{01}^{[2]} &= \theta_1 - \pi & \Theta_{001}^{[3]} &= \theta_3 - \pi \\
 & & \Theta_{10}^{[2]} &= \theta_2 & \Theta_{010}^{[3]} &= \theta_4 \\
 & & \Theta_{11}^{[2]} &= \theta_2 - \pi & \Theta_{011}^{[3]} &= \theta_4 - \pi \\
 & & & & \Theta_{100}^{[3]} &= \theta_5 \\
 & & & & \Theta_{101}^{[3]} &= \theta_5 - \pi \\
 & & & & \Theta_{110}^{[3]} &= \theta_6 \\
 & & & & \Theta_{111}^{[3]} &= \theta_6 - \pi.
 \end{aligned}$$

For example, when $d = 1$ this reduces to $\lambda_0 = \cos(\frac{1}{2}\theta_0)$ and $\lambda_1 = \sin(\frac{1}{2}\theta_0)$, while for $d = 2$ we have $\lambda_0 = \cos(\frac{1}{2}\theta_0)\cos(\frac{1}{2}\theta_1)$, $\lambda_1 = \cos(\frac{1}{2}\theta_0)\sin(\frac{1}{2}\theta_1)$, $\lambda_2 = \sin(\frac{1}{2}\theta_0)\cos(\frac{1}{2}\theta_2)$ and $\lambda_3 = \sin(\frac{1}{2}\theta_0)\sin(\frac{1}{2}\theta_2)$. This parameterization of coefficients naturally arises from a sequence of uniformly controlled rotations defined in appendix B. The construction of a state $|\Lambda_d\rangle$ is then achieved by a cascade of uniformly controlled rotations, all around the y -axis, involving an incrementally increasing subset $1, \dots, k$ of the d qubits as F_k^{k-1} giving a circuit $F_1^0 F_2^1 F_3^2 \dots F_d^{d-1}$. In figure C.1, the circuit building $|\Lambda_4\rangle$ is shown. This figure also shows that as each successive qubit k is added it becomes entangled with the subset of $k - 1$ qubits in the state $|\Lambda_{k-1}\rangle$ previously rotated, leaving an enlarged total state $|\Lambda_k\rangle$ that is completely defined by the $2^k - 1$ independent angles θ_j .

References

- [1] Peres A and Terno D R 2004 *Rev. Mod. Phys.* **76** 93
- [2] Peskin M and Schröder D 1995 *An Introduction to Quantum Field Theory* (New York: Perseus Books)
- [3] Dirac P A M 1958 *Principles of Quantum Mechanics* (Oxford: Oxford University Press)
- [4] Landau L and Peierls R 1931 *Z. Phys.* **69** 56
- [5] Aharonov Y and Albert D Z 1980 *Phys. Rev. D* **21** 3316
- [6] Aharonov Y and Albert D Z 1981 *Phys. Rev. D* **24** 359
- [7] Hellwig K E and Kraus K 1970 *Phys. Rev. D* **1** 566
- [8] Aharonov Y and Albert D Z 1984 *Phys. Rev. D* **29** 223
- [9] Aharonov Y and Albert D Z 1984 *Phys. Rev. D* **29** 228
- [10] Aharonov Y, Albert D Z and Vaidman L 1986 *Phys. Rev. D* **34** 1805
- [11] Popescu S and Vaidman L 1994 *Phys. Rev. A* **49** 4331
- [12] Groisman B and Vaidman L 2001 *J. Phys. A: Math. Gen.* **34** 6881
- [13] Bennett C H, DiVincenzo D P, Fuchs C A, Mor T, Rains E, Shor P W, Smolin J A and Wootters W K 1993 *Phys. Rev. A* **59** 1070
- [14] Breuer H-P and Petruccione F 2002 *The Theory of Open Quantum Systems* (Oxford: Oxford University Press)
- [15] Nielsen M A and Chuang I 2000 *Quantum Computation and Quantum Information* (Cambridge: Cambridge University Press)
- [16] Groisman B and Reznik B 2002 *Phys. Rev. A* **66** 022110
- [17] Vaidman L 2003 *Phys. Rev. Lett.* **90** 010402
- [18] Sorkin R D 1993 *Directions in General Relativity* vol 2, ed B L Hu and T A Jacobson (Cambridge: Cambridge University Press) (arXiv:gr-qc/9302018)
- [19] Beckman D, Gottesman D, Kitaev A and Preskill J 2002 *Phys. Rev. D* **65** 065022
- [20] Ashhab S, Maruyama K, Brukner C and Nori F 2009 *Phys. Rev. A* **80** 062106
- [21] Heaney L and Anders J 2009 *Phys. Rev. A* **80** 032104
- [22] Paterek T, Kurzynski P, Oi D and Kaszlikowski D 2010 (arXiv:1004.5184)
- [23] Reznik B, Aharonov Y and Groisman B 2002 *Phys. Rev. A* **65** 032312
- [24] Bennett C H, Brassard G, Crépeau C, Jozsa R, Peres A and Wootters W K 1993 *Phys. Rev. Lett.* **70** 1895
- [25] Groisman B, Reznik B and Vaidman L 2003 *J. Mod. Opt.* **50** 943
- [26] Aharonov Y and Vaidman L 2000 *Phys. Rev. A* **61** 052108
- [27] Bohm D and Aharonov Y 1957 *Phys. Rev.* **108** 1070
- [28] Bell J S 1964 *Physics* **1** 195
- [29] Beckman D, Gottesman D, Nielsen M A and Preskill J 2001 *Phys. Rev. A* **64** 052309
- [30] Greenberger D M, Horne M A and Zeilinger A 1989 *Bell's Theorem, Quantum Theory, and Conceptions of the Universe* (Dordrecht: Kluwer Academic) (arXiv:0712.0921)
- [31] Raussendorf R and Briegel H J 2001 *Phys. Rev. Lett.* **86** 5188

- [32] Raussendorf R, Browne D E and Briegel H J 2003 *Phys. Rev. A* **68** 022312
- [33] Hein M, Eisert J and Briegel H J 2004 *Phys. Rev. A* **69** 062311
- [34] Cirac J I, Dür W, Kraus B and Lewenstein M 2001 *Phys. Rev. Lett.* **86** 544
- [35] Bernstein E and Vazirani U 1997 *SIAM J. Comput.* **26** 1411 (arXiv:quant-ph/9701001)
- [36] Möttönen M, Vartiainen J J, Bergholm V and Solomaa M M 2004 *Phys. Rev. Lett.* **93** 130502
- [37] Dür W, Vidal G, Cirac J I, Linden N and Popescu S 2001 *Phys. Rev. Lett.* **87** 137901
- [38] Kraus B and Cirac J I 2001 *Phys. Rev. A* **63** 032308
- [39] Zhang J, Vala J, Sasty S and Whaley K B 2003 *Phys. Rev. A* **67** 042313
- [40] Makhlin Y 2002 *Quantum Inf. Process.* **1** 243