

Theory of quantum secret sharing

Daniel Gottesman*

Microsoft Corporation, One Microsoft Way, Redmond, Washington 98052

(Received 18 October 1999; published 16 March 2000)

I present a variety of results on the theory of quantum secret sharing. I show that any mixed state quantum secret sharing scheme can be derived by discarding a share from a pure state scheme, and that the size of each share in a quantum secret sharing scheme must be at least as large as the size of the secret. I show that the only constraints on the existence of quantum secret sharing schemes with general access structures are monotonicity (if a set is authorized, so are larger sets) and the no-cloning theorem. I also discuss some aspects of sharing classical secrets using quantum states. In this situation, the size of each share can sometimes be half the size of the classical secret.

PACS number(s): 03.67.Dd

I. INTRODUCTION

In a classical secret sharing scheme, some sensitive classical data are distributed among a number of people such that certain sufficiently large sets of people can access the data, but smaller sets can gain no information about the shared secret. For instance, a possible application is to share the key for a joint checking account shared by many people. No individual is able to withdraw money, but sufficiently large groups can use the account.

One particularly symmetric variety of secret sharing scheme is called a *threshold scheme*. A (k, n) classical threshold scheme has n shares, of which *any* k are sufficient to reconstruct the secret, while any set of $k-1$ or fewer shares has no information about the secret. Blakely [1] and Shamir [2] showed that threshold schemes exist for all values of k and n with $n \geq k$.

It is also possible to consider more general secret sharing schemes which have an asymmetry between the power of the different shares. For instance, one might consider a scheme with four shares A , B , C , and D . Any set containing A , B , and C or A and D can reconstruct the secret, but any other set of shares has no information. In this example, the presence of A is essential to reconstructing the secret, but not sufficient — A needs the help of either D or both B and C . This particular scheme can be constructed by taking a $(5, 7)$ threshold scheme, and assigning three shares to A , two shares to D , and one share to each of B and C , but other schemes exist which cannot be constructed by bundling together shares of a threshold scheme. The list of which sets are able to reconstruct the secret is called an *access structure* for the secret sharing scheme. It turns out that a secret sharing scheme exists for any access structure, provided it is monotone [3] — i.e., if a set S can reconstruct the secret, so can all sets containing S .

With the advent of quantum computation, it is possible that quantum information may someday be as commonplace as classical information, and we may wish to protect it the same ways as we protect classical information. Using quan-

tum secret sharing [4], we could perhaps create joint checking accounts containing quantum money [5], or share hard-to-create ancilla states [6], or perform a secure distributed quantum computation. The authors of Ref. [4] showed some basic results about quantum secret sharing schemes, including the existence of quantum threshold schemes. A quantum $((k, n))$ threshold scheme (the use of double parentheses distinguishes it from a classical scheme) exists provided the no-cloning theorem is satisfied — i.e., $n/2 < k \leq n$. In this paper, I will prove some further results about quantum secret sharing schemes with general access structures, including the fact that the no-cloning theorem and monotonicity provide the only restriction on the existence of quantum secret sharing schemes.

Another possible application of quantum states to secret sharing is to create secret sharing schemes sharing classical data using quantum states [7,8]. This could allow, for instance, for more secure distribution of the shares of the scheme. I will show below that it can also produce more efficient schemes: in any purely classical scheme, the size of each important share must be at least as large as the size of the secret, whereas, using quantum states to share a classical secret, we can sometimes make each share *half* the size of the secret.

In the theory of classical secret sharing, one sometimes considers schemes which do not completely hide the secret from unauthorized groups of people, or from which the secret cannot be perfectly reconstructed even by authorized sets. I will not consider the quantum generalizations of such schemes. I will only consider the theory of perfect secret sharing schemes, in which the data are either completely revealed or completely hidden, with no middle ground.

II. QUANTUM SECRET SHARING

I will begin by reviewing some results from Ref. [4] which will form the basis of much of the later discussion. In a perfect quantum secret sharing scheme, any set of shares is either an *authorized* set, in which case someone holding all of those shares can exactly reconstruct the original secret, or an *unauthorized* set, in which case someone holding just those shares can acquire no information at all about the se-

*Electronic address: gottesma@microsoft.com

cret quantum state (that is, the density matrix of an unauthorized set is the same for all encoded states). For a generic state split up into a number of shares, most sets will be neither authorized nor unauthorized — quantum secret sharing schemes form a special set of states.

One constraint on quantum secret sharing schemes is an obvious one inherited from classical schemes. Any secret sharing scheme must be *monotone*. That is, if we increase the size of a set, it cannot switch from authorized to unauthorized (the indicator function which is 0 for unauthorized sets and 1 for authorized sets is monotonic).

As we shall see in Sec. III, the only other constraint on quantum secret sharing schemes is the no-cloning theorem [9,10]. We cannot make two copies of an unknown quantum state. Therefore, we cannot distribute the shares of quantum secret sharing scheme into two disjoint authorized sets (each of which could produce a copy of the original state). Since every set is either authorized or unauthorized, this implies the complement of an authorized set is always an unauthorized set.

A *pure state* quantum secret sharing scheme encodes pure state secrets as pure states (when all of the shares are available). A *mixed state* quantum secret sharing scheme may encode some or all pure states of the secret as mixed states. Pure state schemes have some special properties, as a consequence of the following theorem, but the general quantum secret sharing scheme is a mixed state scheme. Theorem 1 and corollary 2 first appeared in Ref. [4].

Theorem 1: Let C be a subspace of a Hilbert space \mathcal{H} which can be written as tensor product of the Hilbert spaces of various coordinates. Then C corrects erasure errors¹ on a set K of coordinates iff

$$\langle \phi | E | \phi \rangle = c(E) \quad (1)$$

(independent of $|\phi\rangle \in C$) for all operators E acting on K . A pure state encoding of a quantum secret is a quantum secret sharing scheme iff the encoded space corrects erasure errors on unauthorized sets and it corrects erasure errors on the complements of authorized sets.

Proof: The first equivalence follows from the theory of quantum error-correcting codes. To recover the original secret on an authorized set, we must be able to compensate for the absence of the remaining shares, which is to say from an erasure error on the complement of the authorized set. Condition (1) implies that measuring any Hermitian operator on the coordinates K gives us no information about which state in C we have. This means the density matrix on K does not depend on the state, which is precisely the condition we need an unauthorized set to satisfy. ■

As a corollary, we find that pure state schemes are only possible for a highly restricted class of access structures.

TABLE I. Adding authorized sets to obtain a maximal access structure. After adding BD and CD , every subset of $ABCD$ either contains an authorized set or is contained in an unauthorized set.

	Authorized	Complement (unauthorized)
Start	ABC AD	D BC
Add	BD	AC
Add	CD	AB

Corollary 2: In a pure state quantum secret sharing scheme, the authorized sets are precisely the complements of the unauthorized sets.

Proof: By the no-cloning theorem, the complement of an authorized set is always an unauthorized set. By theorem 1, for a pure state scheme, we can correct erasure errors on any unauthorized set. This means we can reconstruct the secret in the absence of those shares; that is, the complement is an authorized set. ■

In other words, in a pure state scheme, any division of the shares into two sets will produce one authorized set and one unauthorized set. In a mixed state scheme, we could sometimes obtain two unauthorized sets. (The no-cloning theorem ensures we never obtain two authorized sets.) If some subdivision gives us two unauthorized sets, we can define a new access structure by making one of them authorized. Of course, the new access structure should be monotone, which may require us to create additional authorized sets at the same time.

For instance (see Table I), if we started with the access structure ABC or AD from Sec. I (any set containing A , B , and C is authorized, as is any set containing both A and D), we could add the set BD (so any set containing B and D is also now authorized — this means that BCD becomes authorized as well). We wish to continue to satisfy the no-cloning theorem, so we never add a new authorized set contained in the complement of an existing authorized set. For instance, in the example, we could not have added BC as an authorized set, since its complement AD is already authorized.

In short, we start out with a list of authorized sets. We make a list of their complements; nothing on this list can ever become authorized without violating the no-cloning theorem. There will be a number of sets left over. We choose one, put it on the authorized list, and put its complement on the list of complements. We can keep doing this for a while, but eventually, every set will be on one of the two lists, and we will be forced to stop. Thus an access structure where the authorized and unauthorized sets are complements of each other is a *maximal* quantum access structure.

In the example, after adding BD , we could add CD as an authorized set. The authorized sets are now all sets containing ABC , AD , BD , or CD . Now we must stop — every set contains one of these four sets or is contained in the complement of one of them. Pure state schemes and maximal access structures may seem like a very special situation, but in fact

¹An erasure error is a general error on a known coordinate. For instance, it replaces the coordinate with a state $|e\rangle$ orthogonal to the regular Hilbert space. Recall that a quantum error-correcting code of distance d can correct $d-1$ erasure errors or $\lfloor (d-1)/2 \rfloor$ general errors.

they play a central role in the theory of quantum secret sharing because of the following theorem.

Theorem 3: Every mixed state quantum secret sharing scheme can be described as a pure state quantum secret sharing scheme with one share discarded. The access structure of the pure state scheme is unique.

Proof: Given a superoperator that maps the Hilbert space \mathcal{S} of the secret to density operators on \mathcal{H} (which is a tensor product of the Hilbert spaces of the various shares), we can extend the superoperator to a unitary map from \mathcal{S} to $\mathcal{H} \otimes \mathcal{E}$ for some space \mathcal{E} . We assign this additional Hilbert space to the extra share. In other words, we can “purify” the mixed state encoding by adding an extra share. The original mixed state scheme is produced by discarding the extra share. I claim that the new pure state encoding is a quantum secret sharing scheme.

Sets on the original shares remain authorized or unauthorized, as they were before adding \mathcal{E} . Given a set T including the extra share, look at the complement of T , which is a set not including \mathcal{E} and is thus either authorized or unauthorized (in the new scheme as well as the old). For instance, if we purify the scheme (ABC or AD) by adding a fifth share E , the complement of CDE is unauthorized, while the complement of DE is authorized. If the complement is authorized, then we can correct for erasures on T , and condition (1) holds for T — we can obtain no information about the secret from T , and T is unauthorized. If the complement of T is unauthorized, we can correct erasures on the complement. Therefore, we can reconstruct the state with just T , and T is authorized. Thus the new scheme is secret sharing. It is clear from the argument that any other purification of the mixed state scheme would produce the same access structure. ■

In Ref. [4], we presented a class of quantum secret sharing schemes where every share had the same size as the secret. One might wonder if it is possible to do better. For instance, can we make one share much smaller than the secret, possibly at the cost of enlarging another share? The answer is no, provided we only consider important shares (unimportant shares never make a difference as to whether a set is authorized or unauthorized).

Theorem 4: The dimension of each important share of a quantum secret sharing scheme must be at least as large as the dimension of the secret.

Proof: We need only prove the result for pure state schemes. By theorem 3, the result for mixed state schemes will follow.

Let S be an important share in a pure state quantum secret sharing scheme. Then there is an unauthorized set T such that $T \cup \{S\}$ is authorized. Share the state $|0\rangle$ and give the shares of T to Bob and the remaining shares (including S) to Alice. By corollary 2, Alice’s shares form an authorized set; she can correct for erasures on T . By theorem 6 below, this means Alice can perform any operation she likes on the secret without disturbing Bob’s shares. She can equally well perform quantum interactions between the secret and other quantum states held by her. In particular, if Alice has state $|\psi\rangle$ from a Hilbert space of dimension s (the size of the secret), she can coherently swap it into her shares of the secret sharing scheme, which now encodes the state $|\psi\rangle$.

Then Alice sends just the share S to Bob. Bob now has an authorized set, so he can reconstruct $|\psi\rangle$. Therefore, by theorem 5 below, share S must have had dimension at least s as well. ■

The above proof depends on two theorems of interest outside the theory of quantum secret sharing. The first is obvious, and it is also true; it has not, to my knowledge, appeared before in the literature.

Theorem 5: Even in the presence of preexisting entanglement, sending an arbitrary state from a Hilbert space of dimension s requires a channel of dimension s .

Proof: (This proof is due to Michael Nielsen [11].) Assume that in addition to whatever entanglement is given, Alice and Bob share a cat state $\sum |i\rangle_A |i\rangle_B$ of dimension s . Using a straightforward variant of superdense coding [12], Alice can encode one of s^2 classical states in this cat state. Now Alice transmits her half of the cat state to Bob, using the preexisting entanglement if it helps. Bob can now reconstruct the classical state, so by the bounds on superdense coding [13], Alice must have used a channel of dimension s . ■

The second theorem is more interesting. It says that if Alice can read a piece of quantum data, she can also change it any way she likes, without disturbing any entanglement of the encoding with the outside. There will be no way to tell that the data has been changed.

Theorem 6: Suppose a superoperator \mathcal{S} maps a Hilbert space H to density operators on $A \otimes B$, and \mathcal{S} restricted to A (that is, traced over B) is invertible (by quantum operation). Then for any unitary $U: H \rightarrow H$, there exists a unitary operation $V: A \rightarrow A$ such that $V \circ \mathcal{S} = \mathcal{S} \circ U$.

Proof: We can extend the superoperator \mathcal{S} to a unitary operator W and enlarge B with the necessary extra dimensions. If V works for W , it will also work for \mathcal{S} . Since W is invertible on A , the image subspace corrects erasure errors on B , and

$$\langle \psi | E | \psi \rangle = c(E) \quad (2)$$

for any operator E acting on B , where $c(E)$ is independent of $|\psi\rangle \in W(H)$. Choose a basis $|j\rangle_B$ for B . Given any state $|\psi\rangle$ in the image of W , we can write it as

$$|\psi\rangle = \sum |\psi_j\rangle_A |j\rangle_B. \quad (3)$$

(The states $|\psi_j\rangle$ are not necessarily orthogonal, although we could have made them orthogonal for any single $|\psi\rangle$.) If we let E be a projection on the basis states of B , or a projection on the basis states followed by a permutation of those basis states, Eq. (2) implies that the inner products $\langle \psi_i | \psi_j \rangle$ are independent of $|\psi\rangle$. Therefore, there is a unitary operation V acting on A that takes any set of states $|\psi_j\rangle_A$ for $|\psi\rangle \in W(H)$ to the set of states $|\phi_j\rangle_A$ for any state $|\phi\rangle \in W(H)$. In fact, V will map $|\psi\rangle$ to $|\phi\rangle$.

More generally, and by the same logic, given any two bases of $W(H)$, there will be a unitary V on A that takes one to the other. Given $U: H \rightarrow H$, we can define U as mapping a basis $|v_i\rangle$ to basis $|w_i\rangle$. Then define $V: A \rightarrow A$ as an operator

that maps $W|v_i\rangle$ to $W|w_i\rangle$, and the theorem follows. ■

The theorem says, in a sense, that there is no such thing as a secure quantum ROM (“read-only memory”). The obvious classical analog of this theorem is false: If Alice and Bob each has one copy of a piece of classical information, then Alice can read the information, but cannot change it without alerting Bob. The fact that she can in the quantum case can be seen as an aspect of no-cloning. If Alice has a complete copy of the data, there cannot be any extra information about it elsewhere; thus Alice can safely alter the data. I conclude this section with an easy theorem that will be needed in the construction of a general access structure.

Theorem 7: If S_1 and S_2 are quantum secret sharing schemes, then the scheme formed by concatenating them (expanding each share of S_1 as the secret of S_2) is also secret sharing.

The reason this requires proof is that, due to some nonlocal quantum effect, it might have been possible to get more information from sets in two copies of S_2 than can be accessed from just one of the sets.

Proof: By theorem 3, we need only consider pure state schemes. Then the concatenated scheme S is a pure state scheme too. Suppose we have some set of shares T . We can write it as the union $\cup T_i$, where T_i is a set on the i th copy of S_2 . Consider the set U of copies on which T_i is authorized. U is either an authorized or an unauthorized set of S_1 . If it is authorized, then our large set T is certainly authorized — we reconstruct the copies of S_2 in U , and use U to reconstruct the original secret.

If U is unauthorized, we look at the complement of T . It can be written as a union $\cup T'_i$, where T'_i is the complement of T_i in its copy of S_2 . T'_i is authorized whenever T_i is unauthorized. Therefore, the set of copies on which T'_i is authorized is the complement of U , which is authorized. Thus the complement of T is authorized, so T is unauthorized. ■

Clearly the proof works equally well for more complicated concatenation schemes, with multiple levels or with a different scheme S_2 for each share of S_1 . Also note that if we bundle shares together (assigning two or more shares to the same person), the result is still a secret sharing scheme.

III. CONSTRUCTION OF A GENERAL ACCESS STRUCTURE

This section will be devoted to proving that monotonicity and the no-cloning theorem provide the only restrictions on the existence of quantum secret sharing schemes. The same result was shown by Smith [14] by adapting a classical construction using monotone span programs. The construction given here is far from optimal in terms of the share sizes of the resulting schemes.

Theorem 8: A quantum secret sharing scheme exists for an access structure S iff S is monotone and satisfies the no-cloning theorem (i.e., the complement of an authorized set is an unauthorized set). For any maximal quantum access structure S , a pure state scheme exists.

It will be helpful to first understand an analogous classical construction [3]. Any access structure can be written in a

disjunctive normal form, which is the OR of a list of authorized sets. For our standard example, with authorized sets ABC and AD , the normal form is $(A \text{ AND } B \text{ AND } C) \text{ OR } (A \text{ AND } D)$. This normal form provides a construction in terms of threshold schemes — the AND gate corresponds to a $(2,2)$ threshold scheme (which has one authorized set $A \text{ AND } B$), while the OR gate corresponds to a $(1,2)$ threshold scheme (for which $A \text{ OR } B$ is authorized). Then by concatenating the appropriate set of threshold schemes, we obtain a construction for the original access structure.

In the quantum case, this construction fails, because by the no-cloning theorem, there is no $((1,2))$ quantum threshold scheme. A single authorized set (such as $A \text{ AND } B \text{ AND } C$) still corresponds to a quantum threshold scheme [a $((3,3))$ scheme in this case], but to take the OR of these authorized sets, we will have to do something different. We will replace the $((1,2))$ scheme with $((r, 2r-1))$ schemes (which correspond to majority functions instead of OR). r of the shares will be the individual authorized sets of the desired access structure, and the other $r-1$ shares will be from another access structure that is easier to construct.

The full construction is recursive. Given constructions of access structures for $n-1$ shares, we will construct all maximal access structures for n shares. From maximal access structures on n shares we will be able to construct all access structures on n shares. We can start from the base case of 1 share, which just has the trivial $((1,1))$ access structure. The construction will assume we know how to create threshold schemes, for instance using the construction in Ref. [4].

Given any maximal access structure S on n shares, consider the access structure S' obtained by discarding one share. Certainly S' is still monotone and still satisfies the no-cloning theorem. Therefore, by the inductive hypothesis, we have a construction for the access structure S' . Now, following the proof of theorem 3, add an additional share to S' , putting it in an overall pure state. By the proof of theorem 3, we know the resulting scheme is in fact a quantum secret sharing scheme. It is not hard to see that S is the unique access structure produced this way. For instance, the maximal access structure $ABC \text{ OR } AD \text{ OR } BD \text{ OR } CD$ can be formed by purifying the (mixed state) scheme with access structure ABC [just a $((3,3))$ threshold scheme].

Now suppose we are given a general quantum access structure S on n shares. We describe this access structure by a list of its minimal authorized sets A_1, A_2, \dots, A_r . As mentioned above, A_i by itself defines a quantum access structure — a $((k,k))$ threshold scheme, in fact, if A_i contains k shares.

S has a total of r minimal authorized sets. Let us take a $((r, 2r-1))$ quantum threshold scheme, and expand each of its shares using another secret sharing scheme. Share i , for $i=1, \dots, r$, is expanded using the threshold scheme associated with the set A_i . Shares $r+1$ through $2r-1$ will all be expanded using another secret sharing scheme S' .

S' will be a pure state scheme, with a maximal access structure which can be achieved by adding authorized sets to S . That means when A is an authorized set of S (so it contains

some A_i), it is also an authorized set of S' . Therefore, we can reconstruct the last $r-1$ shares of the $((r, 2r-1))$ scheme, as well as at least one of the first r shares, so A is an authorized set for the concatenated scheme.

Conversely, if we have a set B which does not include any of the sets A_i , we do not have an authorized set for any of the schemes A_i . B might be an authorized set for the scheme S' , but that only gives us authorized sets for at most $r-1$ shares of the $((r, 2r-1))$ scheme. Therefore, B is an unauthorized set. This shows that the access structure of the concatenated scheme is exactly S , completing the construction.

As an example, consider this construction applied to the access structure ABC OR AD . The three rows represent shares of a $((2,3))$ scheme, so authorized sets on any two rows suffice to reconstruct the secret. Repeated letters imply bundling, so A gets a share from each of the first two rows, as well as one from the third row.

$$((2,3)) \text{ scheme } \begin{cases} ((3,3)):A,B,C \\ ((2,2)):A,D \\ S'. \end{cases} \quad (4)$$

The first two rows are threshold schemes. S' is a maximal access structure containing $\{A,B,C\}$ and $\{A,D\}$. For instance, in this case, S' could be the scheme ABC OR AD OR BD OR CD which we constructed earlier; or we could just use the trivial scheme with authorized set $\{A\}$ (give A the secret).

I noted in Sec. I that this particular scheme can be easily constructed directly from a $((5,7))$ threshold scheme. However, not all access structures can be made by bundling together shares of a threshold scheme (for instance, $ABCD$ OR ADE OR BCD cannot be so constructed² — E would have to get more shares of the threshold scheme than B since ADE is authorized while ABD is not, but BCD is authorized while CDE is not), while the recursive construction always works.

IV. SHARING CLASSICAL SECRETS

We can also use quantum states to share classical secrets, a process previously considered in Refs. [7] and [8]. Many of the theorems proved above will fail in this situation. For instance, superdense coding [12] provides an example of a $(2,2)$ threshold scheme where each share is a single qubit, but the secret is two classical bits: the four Bell states $|00\rangle \pm |11\rangle$, $|01\rangle \pm |10\rangle$ encode the four possible two-bit numbers, and for all four states, each qubit is completely random. This $(2,2)$ scheme is a pure state scheme, yet does not satisfy corollary 2, and the share size is smaller than the size of the secret. Neither is possible for a purely classical scheme or for a purely quantum scheme. Another difference is that there is

²For quantum access structures, threshold schemes suffice for fewer than five shares, whereas for classical access structures, there are examples where they fail for four shares. This is because the four-share classical examples would violate the no-cloning theorem.

no rule against copying classical data, so, for instance, (k,n) threshold schemes are allowed with $k < n/2$. We can write down conditions for a pure state scheme of this sort to be secret sharing, along the lines of theorem 1.

Theorem 9: Suppose we have a set of orthonormal states $|\psi_i\rangle$ encoding a classical secret. Then a set T is an unauthorized set iff

$$\langle \psi_i | F | \psi_i \rangle = c(F) \quad (5)$$

(independent of i) for all operators F on T . T is authorized iff

$$\langle \psi_i | E | \psi_j \rangle = 0 \quad (i \neq j) \quad (6)$$

for all operators E on the complement of T .

Note that only the basis states $|\psi_i\rangle$ appear in theorem 9, whereas in theorem 1, the condition had to hold for all $|\psi\rangle$ in a Hilbert space. This is the source of the difference between classical and quantum secrets — the former hides just a set of orthogonal states, while the latter hides all superpositions of those states.

Proof: On an unauthorized set, we should be able to acquire no information about which state $|\psi_i\rangle$ we have. This is precisely condition (5). On an authorized set, we need to be able to correct for the erasure of the qubits on the complement. This is equivalent to being able to distinguish the state $|\psi_i\rangle$ from the state $|\psi_j\rangle$ with an arbitrary operator applied to the complement of T . That is, it is equivalent to condition (6). ■

Note that purely classical secret sharing schemes can be considered as a particular special case of sharing classical data with quantum states — every encoding in a purely classical scheme is just a mixture of tensor products of basis states. Purely classical secret sharing schemes are always mixed state schemes, since classically, there is no way to hide information without randomness.

Superdense coding provided an example where using quantum data allowed a factor of 2 improvement in space over any classical scheme. It turns out that this is the best we can do.

Theorem 10: The dimension of each important share of a classical secret sharing scheme must be at least as large as the square root of the dimension of the secret. The total size of each authorized set must be at least as large as the secret.

This means that a $2n$ -bit secret requires shares of at least n qubits.

Proof: The proof is quite similar to the proof of theorem 4, which gives the corresponding result for quantum secret sharing schemes. We create the quantum state corresponding to the shared secret 0. If it is a mixed state scheme, we include any extra qubits needed to purify it (the result may not be a secret sharing scheme, however — theorem 3 need not hold). If S is the share under consideration, and T is an unauthorized set such that $T \cup \{S\}$ is authorized, give T to Bob, and all the other shares (including S and the extra purifying qubits) to Alice.

Bob has no information about the secret; $\langle \psi_i | E | \psi_i \rangle$ is independent of i . Therefore, as in the proof of theorem 6, Alice can perform, without access to Bob's qubits, a trans-

formation between $|\psi_0\rangle$ (the current state) and $|\psi_i\rangle$ for any i . Then she sends the share S to Bob, who now has an authorized set, and can reconstruct i . We have sent a secret of dimension s using prior entanglement and the share S , which by the bounds on superdense coding [13] must therefore have dimension at least \sqrt{s} . Those bounds also show the size of the channel plus preexisting entanglement must be s , so the size of the full authorized coalition is at least s . ■

Note that we used an analogue of theorem 6 in the proof. As I noted before, the general case of theorem 6 is not true here: Since the secret is classical, we could make two copies of it. Then one copy is sufficient to read it, but both are needed to change it without leaving a trace. In fact, the version of the theorem we have used is just the proof that perfect quantum bit commitment is impossible [15,16] — Bob has no information about the state, so Alice can change the state to whatever she likes.

Besides being an interesting result about secret sharing schemes, this theorem is useful in analyzing other cryptographic concepts. For instance, it shows that there is no useful unconditionally secure cryptographic memory protocol, which can only be unlocked with a key, which we would want to be much smaller than the stored data. Such a protocol would be a (2,2) secret sharing scheme, so the theorem requires that the key be at least half the size of the data.

Theorem 10 can be easily modified to show that in any purely classical scheme, each important share must be at least dimension s , not \sqrt{s} . This follows because if Alice and Bob are just sending classical states back and forth, they need a channel of dimension s to send the secret rather than dimension \sqrt{s} . We have already seen one example where this improvement is achievable using quantum states.

When else can we get this factor of 2 improvement in the number of qubits per share? I do not have a full answer to this question. Certainly for a (1, n) threshold scheme, no improvement is possible, since each authorized coalition (each single share) must be as large as the secret. For many other threshold schemes, however, an improvement is possible.

Theorem 11: A (k,n) threshold scheme exists sharing a classical secret of size $s = p^2$ with one qubit (a p -dimensional quantum state) per share whenever $n \leq 2k - 2$, $p \geq n$, and p is prime.

Before giving the proof, I will review some basic facts about quantum and classical error-correcting codes which will be needed in the construction. A classical linear $[n,k,d]$ code encodes k bits in n bits and corrects $d - 1$ erasure errors. Classical codes must satisfy the Singleton bound $d \leq n - k + 1$. A code C where the bound is met exactly is called a MDS code (for “maximum distance separable”), and has some interesting properties. The dual C^\perp of C (composed of those words which have vanishing inner product with all words of C) is also a MDS code. When C is an $[n,k,n-k+1]$ code, C^\perp is an $[n,n-k,k+1]$ code. The codewords of the dual code form the rows of the parity check matrix. By measuring the parities specified by the parity check matrix, we can detect errors — any parity which is nonzero signals an error. In addition, in a MDS code, there is a codeword

with support exactly on the set T for any set T of size d . See, for instance, Chap. 11 of Ref. [17] for a discussion of MDS codes.

Quantum codes can frequently be described in terms of a stabilizer [18,19]. The stabilizer of a code is an Abelian group consisting of those tensor products of Pauli matrices which fix every quantum codeword. That is, the codewords live in an eigenspace of all elements of the stabilizer. If the stabilizer contains 2^a elements, it is generated by just a elements, and if we have n qubits, the code encodes $n - a$ qubits. We usually consider the $+1$ eigenspace of the stabilizer generators, but we could instead associate an arbitrary sign to each generator. Tensor products of Pauli matrices have eigenvalues ± 1 , so each set of signs will specify a different coding subspace of the same size.

Stabilizer codes can be easily generalized to work over higher-dimensional spaces [20]. We replace the regular Pauli matrices with their analogs for p -dimensional states $X:|j\rangle \mapsto |j+1\rangle$, $Z:|j\rangle \mapsto \omega^j|j\rangle$, and powers of X and Z [arithmetic is now modulo p , and $\omega = \exp(2\pi i/p)$]. The eigenvalues of X , Z , and their products and tensor products are powers of ω , so instead of associating a sign with each generator of the stabilizer, we should instead associate a power of ω .

There is a standard construction, known as the Calderbank-Shor-Steane (CSS) construction [21,22], which takes two binary classical error-correcting codes and produces a quantum code. This construction generalizes easily to qubits. Take the parity check matrix of the first code C_1 and replace j with X^j , interpreting the rows as generators of the stabilizer. Take the parity check matrix of the second code C_2 and replace j with Z^j , again interpreting rows as generators of the stabilizer. The stabilizer must be Abelian — this produces a constraint on the two classical codes, namely that $C_2^\perp \subseteq C_1$. If C_1 is an $[n,k_1,d_1]$ code and C_2 is an $[n,k_2,d_2]$ code, the corresponding CSS code will be an $[[n,k_1+k_2-n,\min\{d_1,d_2\}]]$ quantum code.

Now consider the classical polynomial code D_r whose coordinates are $(f(\alpha_1), \dots, f(\alpha_n))$. $\alpha_1, \dots, \alpha_n$ are n distinct elements of \mathbb{Z}_p (recall that $p \geq n$), and f runs over polynomials of degree up to r .³ There are $r+1$ coefficients to specify f , so D_r encodes $r+1$ bits. Given the function evaluated at $r+1$ locations, we can use polynomial interpolation to reconstruct the polynomial. In other words, even if $n - (r+1)$ coordinates of the code are missing, we can reconstruct the $r+1$ coefficients specifying the polynomial. Thus this is an $[n,r+1,n-r]$ classical code — a MDS code. Also note that $D_r \subseteq D_{r+1}$.

The codes D_r provide good examples of purely classical secret sharing schemes [2]. If we choose the first r coefficients of the polynomial at random, any set of just r coordinates will contain no information about the remaining coefficient, so we get an $(r+1,n)$ threshold scheme. Applying the CSS construction to the codes D_r and D_{r-1}^\perp [23,24] simi-

³For an appropriate choice of the α_i s, D_r is a Reed-Solomon code or an extended Reed-Solomon code.

larly produces good examples of quantum secret sharing schemes [4]. With this background, we are now ready to tackle the construction.

Proof of theorem 11: We will produce a class of secret sharing schemes which use one qubit for each share and encode two classical pits, whereas any purely classical scheme could only encode one pit. We will use the classical codes D_r to create p^2 related CSS quantum codes with certain useful properties. The secret sharing scheme will encode each of the p^2 classical states as the mixture of all states in the corresponding code from this family.

Lemma: The parity check matrix for the code D_{r-1} includes a row R such that for any set of $r+1$ coordinates, there is a linear combination of rows of D_{r-1} with support exactly on that set of coordinates. R appears in the linear combination with coefficient 1. Similarly, the dual code D_q^\perp has, in its parity check matrix, a row S which appears with coefficient 1 in a linear combination with support on any given set of $n-q$ coordinates.

For instance, we can take $n=4$, $r=2$, $q=1$, $p=5$. D_1 has generator matrix

$$G = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 0 & 1 & 2 & 3 \end{pmatrix} \quad (7)$$

(generated by polynomials 1 and x), and D_1^\perp has generator matrix

$$G' = \begin{pmatrix} 2 & 4 & 1 & 3 \\ 3 & 0 & 1 & 1 \end{pmatrix}. \quad (8)$$

(The parity check matrix of D_1 is the generator matrix of D_1^\perp and vice versa.) By subtracting j times the first row of G from the second row of G , we obtain a vector with support on the three-element set excluding coordinate j . Similarly, by adding some multiple of the first row of G' to the second row of G' , we can obtain a vector with support on any three coordinates.

Proof of lemma: The codes D_r and D_q^\perp are linear, so we only need prove the coefficients of rows R and S are nonzero — then some rescaling will always give the result with coefficient 1.

Since D_r is a MDS code of distance $n-r$, its dual is a MDS code of distance $r+2$. Thus, the parity check matrix of D_r (which is also the generator matrix of D_r^\perp) has a linear combination of rows with support on any set of $r+2$ coordinates, but no linear combination of rows has weight $r+1$ or less. Since D_{r-1} is included in D_r , but encodes one fewer pit, the parity check matrix of D_{r-1} is just the parity check matrix of D_r with one row R added. That parity check matrix has a linear combination of rows with support on any set of $r+1$ coordinates. Since no linear combination of rows of D_r^\perp has weight $r+1$, each of the weight $r+1$ linear combinations must include a component of row R . A similar argument gives the result for D_q^\perp . ■

Now suppose we create the CSS code corresponding to the two classical codes D_{r-1} and D_q^\perp . We require that $q = n-r-1$, $2r \geq n$. Then $q < r$, so $D_q \subseteq D_{r-1}$, and we have

a quantum code. We are given two classical pits a and b to share among n parties. Assign a phase ω^a to the generator R corresponding to row R of D_{r-1} and a phase ω^b to the generator S corresponding to row S of D_q^\perp . All the other generators have phase $+1$. Create the density matrix formed by a uniform mixture over states in the subspace specified by this stabilizer. There are p^2 of these mixed states.

Claim: The set of mixed states described above define a (k, n) threshold scheme encoding two classical pits, with $k = r+1 = n-q$.

For instance, in the case $n=4$, $r=2$, $q=1$, $p=5$, we obtain the stabilizers

$$\begin{aligned} & X^2 X^4 X X^3, \\ \omega^a & X^3 I X X, \\ & Z Z Z Z, \\ \omega^b & I Z Z^2 Z^3, \end{aligned} \quad (9)$$

with $\omega = \exp(2\pi i/5)$. The claim is that this gives a $(3, 4)$ secret sharing scheme. I now proceed to establish the claim, which will prove theorem 11.

For any set T of k coordinates, there will be an element MR of the stabilizer with support on that set of coordinates, where M contains no factors of R or S . This follows from the lemma: There is a linear combination $M+R$ of rows of the parity check matrix of D_{r-1} with support on T . This linear combination translates to an element of the stabilizer — the rows of the parity check matrix become generators of the stabilizer, addition of two rows becomes multiplication of the corresponding generators, and scalar multiplication of a row becomes taking the corresponding generator to the appropriate power.

Since MR has support on T , we can measure its eigenvalue with access only to T . M is a product of generators which are not R or S , so the state has eigenvalue $+1$ for M , and it has eigenvalue ω^a for MR . Thus the eigenvalue of MR tells us a . Similarly, there is an element NS of the stabilizer with support on T , with N having no factors of R or S . We can measure the eigenvalue of NS , and it tells us b . Thus, any set of at least k coordinates is an authorized set.

A particular value of the secret is encoded as a uniform distribution over states in the stabilizer code described above. Thus, the density matrix corresponding to the secret is the projection on the subspace which is left fixed by the stabilizer. That is,

$$\rho(ab) = \prod_i (I + M_i + M_i^2 + \cdots + M_i^{p-1}) \quad (10)$$

$$= \sum_{M \in S} M \quad (11)$$

(normalized appropriately). The M_i are the generators of the stabilizer S . Assume the appropriate phase is included in M in this sum (this means that if we wish M to have eigenvalue ω , we include it as $\omega^{-1}M$, which has eigenvalue $+1$).

Suppose T is a set of $k-1$ or fewer coordinates. The

density matrix of T is the trace of $\rho(ab)$ over the complement of T . Now, X and Z , and all nontrivial products of X and Z , have trace 0. Thus, the only terms in the expression for $\rho(ab)$ which contribute to the trace are those coming from M with weight $\leq k-1$. But the parity check matrices for D_{r-1} and D_q^\perp contain no rows or linear combination of rows of weight less than k . Thus the density matrix of T is

just the identity, regardless of the value of ab . Thus T is unauthorized, proving the theorem. ■

ACKNOWLEDGMENTS

I would like to thank Richard Cleve, Hoi-Kwong Lo, Michael Nielsen, and Adam Smith for helpful discussions.

-
- [1] G. Blakely, Proc. AFIPS **48**, 313 (1979).
 - [2] A. Shamir, Commun. ACM **22**, 612 (1979).
 - [3] J. Benaloh and J. Leichter, in *Proceedings Crypto '88*, edited by S. Goldwasser (Springer-Verlag, New York, 1990), p. 27.
 - [4] R. Cleve, D. Gottesman, and H.-K. Lo, Phys. Rev. Lett. **82**, 648 (1999).
 - [5] S. Wiesner, SIGACT News **15**, 78 (1983).
 - [6] D. Gottesman and I. Chuang, Nature (London) **402**, 390 (1999).
 - [7] M. Hillery, V. Bužek, and A. Berthiaume, Phys. Rev. A **59**, 1829 (1999).
 - [8] A. Karlsson, M. Koashi, and N. Imoto, Phys. Rev. A **59**, 162 (1999).
 - [9] W. K. Wootters and W. H. Zurek, Nature (London) **299**, 802 (1982).
 - [10] D. Dieks, Phys. Lett. A **92**, 271 (1982).
 - [11] M. Nielsen (private communication).
 - [12] C. H. Bennett and S. J. Wiesner, Phys. Rev. Lett. **69**, 2881 (1992).
 - [13] R. Cleve, W. van Dam, M. Nielsen, and A. Tapp, e-print quant-ph/9708019.
 - [14] A. Smith, e-print quant-ph/0001087.
 - [15] D. Mayers, Phys. Rev. Lett. **78**, 3414 (1997).
 - [16] H.-K. Lo and H.-F. Chau, Phys. Rev. Lett. **78**, 3410 (1997).
 - [17] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes* (North-Holland, New York, 1977).
 - [18] D. Gottesman, Phys. Rev. A **54**, 1862 (1996).
 - [19] A. R. Calderbank, E. M. Rains, P. W. Shor, and N. J. A. Sloane, Phys. Rev. Lett. **78**, 405 (1997).
 - [20] E. Knill, e-print quant-ph/9608048.
 - [21] A. R. Calderbank and P. W. Shor, Phys. Rev. A **54**, 1098 (1996).
 - [22] A. Steane, Proc. R. Soc. London, Ser. A **452**, 2551 (1996).
 - [23] D. Aharonov and M. Ben-Or, in *Proceedings of the 29th Annual ACM Symposium on Theory of Computation* (ACM, New York, 1998), p. 176.
 - [24] D. Aharonov and M. Ben-Or, e-print quant-ph/9906129.