Approximate quantum error-correcting codes

Claude Crépeau *, Daniel Gottesman †, Adam Smith ‡

Abstract

It is a standard result in the theory of quantum error-correcting codes that no code of length n can fix more than $n/4$ arbitrary errors, regardless of the dimension of the coding and encoded Hilbert spaces. However, this bound only applies to codes which exactly correct errors. Naively, one might expect that correcting errors to very high fidelity would only allow small violations of this bound. However, this intuition is incorrect: we construct in this paper quantum error-correcting codes capable of correcting up to $n/2 - 1$ arbitrary errors with fidelity exponentially close to 1. This demonstrates a severe distinction between exact quantum error correction and approximate quantum error correction.

* School of Computer Science, McGill University, Montréal (Qc), Canada H3A 2A7. e-mail: `crepeau@cs.mcgill.ca`. Supported in part by Québec's MRST and Canada's NSERC. Some of this research was done while the author was visiting MSRI, Berkeley CA.

† Perimeter Institute, Waterloo, Ontario, Canada N2J 2W9. e-mail: `dgottesman@perimeterinstitute.ca`. Some of this research was done while the author was supported by the Clay Mathematics Institute, and some while the author was visiting MSRI, Berkeley CA.

‡ M.I.T., Lab. for Computer Science, Cambridge MA 02139, USA. e-mail: `asmith@theory.lcs.mit.edu`. Supported in part by U.S. Army Research Office Grant DAAD19-00-1-0177. Some of this research was done while the author was visiting McGill University.

1 Introduction

Quantum computers are likely to be highly susceptible to errors from a variety of sources, much more so than classical computers. Therefore, the study of quantum error correction is vital not only to the task of quantum communications but also to building functional quantum computers. In addition, quantum error correction has many applications to quantum cryptography. For instance, the methods of quantum error correction and fault-tolerant quantum computation can be used to perform multiparty secure quantum computations [8]. For all of these reasons, it is interesting to study bounds on the performance of quantum error-correcting codes (QECCs) in various scenarios.

It is an immediate result of the no-cloning theorem [15] that no quantum error-correcting code of length n can fix $n/2$ erasures because that would imply that we could reconstruct two copies of an encoded quantum state from two halves of the full codeword. This statement is valid regardless of the dimension of the coding Hilbert space.

Another well known result from the theory of quantum error correction is that a length n code can fix t arbitrary single position errors if and only if it can fix $2t$ erasure errors [10]. This follows immediately from the quantum error-correction conditions [10]

$$\langle\psi_i|E_a^\dagger E_b|\psi_j\rangle = C_{ab}\delta_{ij} \tag{1}$$

(for basis encoded states $\{|\psi_i\rangle\}$ and correctable errors $\{E_a\}$) and implies that no QECC of length n can fix more than $n/4$ arbitrary errors, regardless of the dimension of the coding and encoded Hilbert spaces.

In this paper, we show the existence of QECCs of length n that can fix $n/2-1$ arbitrary single position errors with fidelity exponentially close to 1. That is, *approximate* quantum error-correcting codes have the capability of correcting errors in a regime where no *exact* QECC will function. This is important for a few reasons:

- It suggests it may be possible to build approximate QECCs which are highly efficient and yet useful in common error correction scenarios, improving on exact QECCs for the same scenarios. In most cases, exact reconstruction of the quantum state is not necessary, so a more efficient approximate QECC would be welcome.
- The approximate QECCs we construct *always* have efficient decoding algorithms since this procedure depends only on the erasure correction capabilities of a related stabilizer code. Correcting erasures on a stabilizer QECC is always easy as it only requires solving a linear system of equations. This is a highly attractive property, since finding efficient codes with good decoding algorithms is generally a difficult problem.
- It demonstrates that the connection between correcting general errors and erasure errors breaks down for approximate QECCs. This calls into question some of the basic foundations of the theory of quantum error correction, as it suggests there is no sensible notion of distance for an approximate quantum error-correcting code.
- While the codes we present in this paper are not particularly useful for traditional quantum error correction, they may allow verifiable quantum secret sharing (VQSS) and secure multiparty quantum computation (MPQC) beyond previously known bounds (although we do not claim here to construct such protocols). Indeed, the codes we present here show that the purported proof for an upper bound for VQSS which we claimed in a previous paper [8] only holds for exact VQSS.
- It demonstrates that there can be a dramatic difference in behavior between the exact performance of some quantum-mechanical task and approximate performance of the task, even when the approximation is exponentially good. A similar divergence between exact and approximate bounds has recently been seen in the context of private quantum channels [11]. These examples serve as a caution to the entire quantum information community when dealing with approximate performance of quantum protocols.

The idea of using a randomized encoding algorithm is not new in QECC. In particular [5] have devised codes that can correct more (malicious) errors on average than any deterministic QECC. However, their model significantly differs from ours in one of two ways: they assume either that the errors occur at random or that the code is randomly agreed by the coder and the decoder but is kept secret from the adversarial noise source. This model does not seem suitable in applications such as VQSS and MPQC [8]. In our model no secret is shared by the coder and decoder. However, part of our code can be viewed as providing a way for the coder to information-theoretically encrypt the necessary secret. (This is possible since the adversary only has access to part of the transmitted state, though it could be any part.)

The codes we build are somewhat peculiar since they will be linear but of *fractional* dimension. We now introduce some notation for such objects. A classical $[n, k, d]_Q$ code is a linear subspace of dimension k of length n vectors of $\mathbb{F}_Q$-components such that for any two distinct codewords c_1, c_2 we have that c_1 differs in at least d (out of n) positions from c_2. Let $\mathbb{F}_q$ be a subfield of $\mathbb{F}_Q$ such that $q = p^m$ and $Q = p^M$ for some prime p and integers $m < M$. An $[n, k, d]_Q$ code contains p^{kM} codewords (it is a $(n, Q^k, d)_Q$ linear code). If it contains only q codewords associated to the information words of $\mathbb{F}_q$ then the resulting code will be an $[n, m/M, d]_Q$ code (it is a $(n, q, d)_Q$ linear code). Since $m < M$ this code has fractional dimension, and it is linear only over the small field $\mathbb{F}_q$, not over $\mathbb{F}_Q$.

Stabilizer QECCs encoding k Q-dimensional registers in n Q-dimensional registers with quantum distance d (i.e., they are capable of correcting $d-1$ erasure errors or $\lfloor (d-1)/2 \rfloor$ general errors) are conventionally denoted with the notation $[[n, k, d]]_Q$, or $((n, Q^k, d))_Q$ for nonstabilizer codes. As in the classical case, we may define stabilizer QECCs with fractional dimension $[[n, m/M, d]]_Q$ to be $((n, q, d))_Q$ stabilizer codes. The codes we construct will be approximate QECCs and will not have a well-defined notion of distance, but we will use this notation to indicate codes that can correct $d-1$ erasure errors exactly, but may correct more than $\lfloor (d-1)/2 \rfloor$ general errors approximately.

Note that the no-cloning argument above also trivially applies to QECCs of fractional dimension. Similarly, the relation between error correction and erasure correction mentioned above also applies to such fractional dimension codes: exactly the same proof can be used [10]. This implies that no (fractional dimension) QECC of length n can fix more than $n/4$ arbitrary errors, regardless of the dimension of the coding Hilbert space.

In clear contrast, the main result of this paper is the construction of $[[n, \Theta(\frac{1}{n+s}), n/2]]_{\Theta(n(n+s))}$ QECCs that can correct $n/2 - 1$ arbitrary errors with fidelity at least $1 - 2^{-\Omega(s)}$.

2 Preliminaries

Classical Authentication In the classical setting, an authentication scheme is defined by a pair of functions $A : \mathcal{K} \times M \to C$ and $V : \mathcal{K} \times C \to M \times \{\text{valid}, \text{invalid}\}$ such that for any message $\mu \in M$ and key $k \in \mathcal{K}$ we have *completeness*

$$V_k(A_k(\mu)) = \langle \mu, \text{valid} \rangle$$

and that for any opponent O, we have *soundness*

$$\text{Prob}\left[V_k(O) \in \{\langle \hat{\mu}, \text{invalid} \rangle | \hat{\mu} \in M\}\right] \geq 1 - 2^{-\Omega(t)}$$

$$\text{for any message } \mu \in M, \text{Prob}\left[V_k(O(A_k(\mu))) \in \{\langle \mu, \text{valid} \rangle\} \cup \{\langle \hat{\mu}, \text{invalid} \rangle | \hat{\mu} \in M\}\right] \geq 1 - 2^{-\Omega(t)}$$

where $t = \lg \#C - \lg \#M$ is the security parameter creating the tradeoff between the expansion of the messages and the security level. Note that we only consider information-theoretically secure schemes, not schemes that are based on computational assumptions.

Wegman and Carter [6] introduced several constructions for such schemes; more recently, Gemmel and Naor [9] introduced a nearly optimal construction using $5t + \lg m$ bits of key (where m is the number of bits of the message). This compares quite well to the known lower bound of $t + \lg m$ for such a result [9]. The same work [6] also introduced a technique to re-use an authentication function several times by using one-time-pad encryption on the tag, so that an opponent cannot learn *anything* about the particular key being used by $\mathcal{A}$ and $\mathcal{B}$. Thus, at a marginal cost of only t secret key bits per authentication, the confidentiality of the authentication function h is guaranteed and thus may be re-used. (an aribtrary number of times).

The canonical instances that satisfy this definition come through the notion of a Family of Universal Hash Functions (FUHF) which is a class H of hash functions from M to T ($C = M \times T$) such that

- for any message $\mu \in M$ and tag $\tau \in T$, when h is chosen at random from H,

$$\text{Prob}\,[h(\mu) = \tau] = 1/\#T.$$

- for any two distinct messages $\mu_1, \mu_2 \in M$, when h is chosen at random from H,

$$\text{Prob}\,[h(\mu_1) = h(\mu_2)] = 1/\#T.$$

Here the encoding function $E(\mu) = \mu|h(\mu)$ and the decoding function $D(\mu|\tau) = \begin{cases} \langle\mu, \text{valid}\rangle & \text{if } \tau = h(\mu) \\ \langle\mu, \text{invalid}\rangle & \text{otherwise} \end{cases}$.

Wegman and Carter [6] have suggested several FUHFs for the special case $M = \{0,1\}^m$, $T = \{0,1\}^t$. For instance the class $H = \{h|h(m) = am\rfloor_t \oplus b, a \in GF(2^m), b \in GF(2^t)\}$ where m is considered as an element of $GF(2^m)$ and where $\rfloor_t$ means "truncated to the t least significant bits" is a FUHF with key size $m + t$.

Wegman and Carter also introduced a technique to re-use an authentication function several times by using one-time-pad encryption on the output of the function, so that an opponent cannot learn *anything* about the particular function which is used by $\mathcal{A}$ and $\mathcal{B}$. Thus, at a marginal cost of only t secret key bits per authentication, the confidentiality of the authentication function h is guaranteed and thus may be re-used, on and on.

For the remainder of this paper, we assume the reader is familiar with the basic notions and notation of quantum computing (see textbooks such as [12]).

Quantum Authentication At an intuitive level, a quantum authentication scheme is a keyed system which allows $\mathcal{A}$ to send a state ρ to $\mathcal{B}$ with a guarantee: if $\mathcal{B}$ accepts the received state as "good", the fidelity of that state to ρ is almost 1. Moreover, if the adversary makes no changes, $\mathcal{B}$ should always accept, and the fidelity should be exactly 1.

However, a reasonable definition for quantum authentication requires a tradeoff between $\mathcal{B}$'s chances of accepting, and the expected fidelity of the received system to $\mathcal{A}$'s initial state given his acceptance: as $\mathcal{B}$'s chance of accepting increases, so should the expected fidelity.

There is no reason to use the languages of both probability and fidelity here: for classical tests, fidelity and probability of acceptance coincide. With this in mind we first define what constitutes a quantum authentication scheme, and then give a definition of security:

Definition 1 *A* quantum authentication scheme *(*QAS*) is a pair of polynomial time quantum algorithms A and V together with a set of* classical *keys $\mathcal{K}$ such that:*

- *A takes as input an m-qubit message system M and a key $k \in \mathcal{K}$ and outputs a transmitted system C of $m + t$ qubits.*
- *V takes as input the (possibly altered) transmitted system $\hat{C}$ and a classical key $k \in \mathcal{K}$ and outputs two systems: a m-qubit message state $\hat{M}$, and a single (verdict) qubit V which indicates acceptance or rejection. The classical basis states of V are called $|\text{ACC}\rangle$, $|\text{REJ}\rangle$ by convention.*

For any fixed key k, we denote the corresponding super-operators by A_k and V_k.

Note that $\mathcal{B}$ may well have measured the qubit V to see whether or not the transmission was accepted or rejected. Nonetheless, we think of V as a qubit rather than a classical bit since it will allow us to describe the joint state of the two systems $\hat{M}, V$ with a density matrix.

There are two conditions which should be met by a quantum authentication protocol. On the one hand, in the absence of intervention, the received state should be the same as the initial state and $\mathcal{B}$ should accept.

On the other hand, we want that when the adversary does intervene, $\mathcal{B}$'s output systems have high fidelity to the statement "either $\mathcal{B}$ rejects or his received state is the same as that sent by $\mathcal{A}$". One difficulty with this is that it is not clear what is meant by "the same state" when $\mathcal{A}$'s input is a mixed state. It turns out that it is sufficient to define security in terms of pure states; one can deduce an appropriate statement about the fidelity of mixed or entangled states.

Given a pure state $|\psi\rangle \in \mathcal{H}_M$, consider the following test on the joint system $\hat{M}, V$: output a 1 if the first m qubits are in state $|\psi\rangle$ *or* if the last qubit is in state $|\text{REJ}\rangle$ (otherwise, output a 0). The projectors corresponding to this measurement are

$$\begin{aligned} P_1^{|\psi\rangle} &= |\psi\rangle\langle\psi| \otimes I_V + I_{\hat{M}} \otimes |\text{REJ}\rangle\langle\text{REJ}| \\ &\quad - |\psi\rangle\langle\psi| \otimes |\text{REJ}\rangle\langle\text{REJ}| \\ P_0^{|\psi\rangle} &= (I_{\hat{M}} - |\psi\rangle\langle\psi|) \otimes (|\text{ACC}\rangle\langle\text{ACC}|) \end{aligned}$$

We want that for all possible input states $|\psi\rangle$ and for all possible interventions by the adversary, the expected fidelity of V's output to the space defined by $P_1^{|\psi\rangle}$ is high. This is captured in the following definition of security.

Definition 2 *A* QAS *is secure with error ϵ for a state $|\psi\rangle$ if it satisfies:*
Completeness: *For all keys $k \in \mathcal{K}$: $V_k(A_k(|\psi\rangle\langle\psi|)) = |\psi\rangle\langle\psi| \otimes |\text{ACC}\rangle\langle\text{ACC}|$*
Soundness: *For all super-operators $\mathcal{O}$, let $\rho_{\mathcal{B}}$ be the state output by $\mathcal{B}$ when the adversary's intervention*[1] *is characterized by $\mathcal{O}$, that is:*

$$\rho_{\mathcal{B}} = \mathbb{E}_k\Big[V_k(\mathcal{O}(A_k(|\psi\rangle\langle\psi|)))\Big] = \frac{1}{|\mathcal{K}|}\sum_k V_k(\mathcal{O}(A_k(|\psi\rangle\langle\psi|)))$$

where " $\mathbb{E}_k$" means the expectation when k is chosen uniformly at random from $\mathcal{K}$. The QAS *has soundness error ϵ for $|\psi\rangle$ if:*

$$\text{Tr}\left(P_1^{|\psi\rangle}\rho_{\mathcal{B}}\right) \geq 1 - \epsilon$$

A QAS *is secure with error ϵ if it is secure with error ϵ for all states $|\psi\rangle$.*

We will actually want authentication protocols that have an additional composibility property: If (A_k, V_k) is a QAS for key k, then the concatenated protocol

$$\left(\bigotimes_{i=1}^{n} A_{k_i}, \bigotimes_{i=1}^{n} V_{k_i}\right) \tag{2}$$

should be a QAS for the key $(k_1, \ldots, k_n)$, with the understanding that the concatenated verification protocol rejects if any of the tensor components rejects (i.e., the concatenated verdict qubit is the AND of the individual verdict qubits, identifying $|\text{ACC}\rangle$ with TRUE and $|\text{REJ}\rangle$ with FALSE).

Quantum authentication protocols satisfying definition 2 were constructed in [3]. We do not know if the above compositibility property follows in general from definition 2, but the protocols constructed in [3] certainly do. This follows because they are constructed from stabilizer purity testing codes (PTCs), which clearly statisfy a corresponding property (if Q_k is a stabilizer PTC with error ϵ, then $\bigotimes_{i=1}^{n} Q_{k_i}$ is a stabilizer PTC with error $n\epsilon$).

3 Definition of approximate QECC (AQECC)

At an intuitive level, an approximate quantum error-correcting code allows $\mathcal{A}$ to send a state ρ to $\mathcal{B}$ with a guarantee: the fidelity of the state received by $\mathcal{B}$ to ρ is almost 1.

Let $q = p^m$ and $Q = p^N$ for some prime p and integers $m < N$. We first define what constitutes an AQECC over $\mathbb{F}_Q$, and then give a definition of correctness:

Definition 3 *An* approximate quantum error correcting code *(AQECC) is a pair of polynomial time quantum algorithms E (encoder) and D (decoder) such that:*

- *E takes as input a m-quqit message system M and outputs a (mixture of) codeword(s) C of n quQits.*

[1] We make no assumptions on the running time of the adversary.

- *D takes as input the (possibly altered) transmitted system $\hat{C}$ and outputs a m-quqit message state $\hat{M}$.*

We will define the correctness of an AQECC on pure states, but it follows from a result of [4] that the output of the AQECC also has high fidelity to an input which is mixed or part of an entangled state.

Given a pure state $|\psi\rangle \in \mathcal{H}_M$, consider the following test on the system $\hat{M}$: output a 1 if the first k qu$q$its are in state $|\psi\rangle$ (otherwise, output a 0). The projectors corresponding to this measurement are

$$\begin{aligned} P_\psi &= |\psi\rangle\langle\psi| \\ P_\psi^\perp &= (I_{\hat{M}} - |\psi\rangle\langle\psi|) \end{aligned}$$

We want that for all possible input states $|\psi\rangle$ and for all possible interventions by the adversary, the expected fidelity of $\mathcal{B}$'s output to the space defined by P_ψ is high. This is captured in the following definition of correctness.

Definition 4 *An AQECC is t-correct with error ϵ for a state $|\psi\rangle$ if for all super-operators $\mathcal{O}$ acting on at most t quQits,*

$$\mathrm{Tr}\left(P_\psi \rho_{\mathcal{B}}\right) \geq 1 - \epsilon,$$

where $\rho_{\mathcal{B}}$ is the state output by $\mathcal{B}$ when the adversary's intervention[2] is characterized by $\mathcal{O}$, that is:

$$\rho_{\mathcal{B}} = D(\mathcal{O}(E(|\psi\rangle\langle\psi|))).$$

A AQECC is t-correct with error ϵ if it is t-correct with error ϵ for all states $|\psi\rangle$.

4 A length 3 quantum code approximately correcting one arbitrary error

We start with a small example, from a well known code. The code c corrects one erasure error:

$$\begin{aligned} |0\rangle &\rightarrow |000\rangle + |111\rangle + |222\rangle \\ |1\rangle &\rightarrow |012\rangle + |120\rangle + |201\rangle \\ |2\rangle &\rightarrow |021\rangle + |102\rangle + |210\rangle \end{aligned} \tag{3}$$

Let $H_1 \otimes H_2 \otimes H_3$ be the coding space of the original code

$$c|\psi\rangle \in H_1 \otimes H_2 \otimes H_3,$$

and let (A_k, V_k) be a quantum authentication scheme as constructed in [3].

We construct a three-component code c' as follows:

$$\begin{aligned} c'|\psi\rangle = &\left\langle A_{k_1}(H_1), k_2, k_3 \right\rangle, \\ &\left\langle A_{k_2}(H_2), k_1, k_3 \right\rangle, \\ &\left\langle A_{k_3}(H_3), k_2, k_1 \right\rangle. \end{aligned} \tag{4}$$

Let $H_1' \otimes H_2' \otimes H_3'$ be the coding space of the new code

$$c'|\psi\rangle \in H_1' \otimes H_2' \otimes H_3'$$

Note that k_1, k_2, and k_3 are random *classical* strings which we use as keys for the quantum authentication protocol A_k. Thus, the H_i's contain both quantum and classical information. Intuitively, we use the QAS to ensure that an adversary cannot change the quantum state of a single register without being detected; thus, we can transform general errors into erasure errors, allowing us to correct one faulty register out of three (no exact QECC can do this). Then we distribute the authentication keys among the three registers so that $\mathcal{B}$ can recover them. We must, however, do so in a way that prevents an adversary with access to a single register from either learning the key applying to her own register (which would allow her to change the quantum state) or from preventing reconstruction of the classical keys.

[2] We make no assumptions on the running time of the adversary.

Theorem 1 *If A_k is a* QAS *secure with error ϵ then c' is a 1-correct AQECC with error prob.* $\text{poly}(\epsilon)$*, correcting one arbitrary error.*

We postpone the proof of this theorem to Section 5 where we will prove more general result.

4.1 Reconstruction

In all cases, the reconstruction has two phases. First we reconstruct the classical keys and use them to verify and decode the quantum authentications. This may result in discarding one register, but at least two remain, which is enough for the erasure-correcting code to recover the original encoded state.

- All k_i's agree in H_1', H_2', H_3':
  Recover k_i from either $H_j', j \neq i$, check that $A_{k_i}(H_i)$ properly authenticates H_i. If one authentication fails, ignore the improperly authenticated H_i and reconstruct the valid codeword as $c|\psi\rangle \in H_1 \otimes H_2 \otimes H_3$ using the erasure recovery algorithm from both $H_j, j \neq i$.
- Some H_i' disagrees with H_j', H_h' on both keys k_h and k_j:
  Discard register i, which must be corrupted. Recover k_j from H_h' and k_h from H_j', and decode the authentications $A_{k_j}(H_j)$ and $A_{k_h}(H_h)$ (which should both pass, since only one register can fail). Reconstruct the valid codeword as $c|\psi\rangle \in H_1 \otimes H_2 \otimes H_3$ using the erasure recovery algorithm from H_j and H_h.
- H_i' and H_j' disagree on key k_h, while H_h' agrees with everyone:
  Either register i or j is corrupt. Get k_i and k_j from H_h' and check that $A_{k_i}(H_i)$ properly authenticates H_i, and that $A_{k_j}(H_j)$ properly authenticates H_j. If neither fail, reconstruct the valid codeword as $c|\psi\rangle \in H_1 \otimes H_2 \otimes H_3$ using the erasure recovery algorithm from H_i and H_j. If one fails, say $A_{k_i}(H_i)$, then conclude register i is corrupt and recover k_h from H_j', decode $A_{k_h}(H_h)$, and reconstruct the valid codeword as $c|\psi\rangle \in H_1 \otimes H_2 \otimes H_3$ using the erasure recovery algorithm from H_h and H_j.

Other cases cannot arise, since only one register can have been changed from the original encoding.

5 A general n-component approximate QECC family correcting up to $d - 1 < n/2$ arbitrary errors

In order to generalize the above construction to cases with n registers, we need to systemize the distribution of the classical keys. Recall that we needed two conditions: First, the adversary should not be able to learn the classical key for her register, but the receiver $\mathcal{B}$ should be able to reconstruct the keys. Second, the adversary should not be able to interfere with $\mathcal{B}$'s reconstruction of the keys.

We ensure the first condition by encoding the keys in a classical secret sharing scheme [14]. Then to achieve the second condition, we further authenticate the shares of the classical secret sharing scheme using a classical authentication scheme. Of course, this requires further classical keys; in particular, we introduce one for each ordered pair of distinct registers, and use all the keys ℓ_{ij} to authenticate share j. Note that this construction is essentially a simplification of those in [13, 7]. They essentially produced approximate error-correcting codes for classical data, on the way to building multi-party computing protocols.

Let $\mathcal{Q}$ be a QECC that can correct $d-1 < n/2$ arbitrary erasure errors: $\mathcal{Q} = [[n, k, d]]$. Such a code can be constructed over sufficiently large dimension Q; for instance, use a polynomial quantum code [1]. The coding space of $\mathcal{Q}$ is defined as

$$\mathcal{Q}|\psi\rangle \in H_1 \otimes H_2 \otimes H_3 \otimes ... \otimes H_n.$$

We assume $\dim(H_1) = \dim(H_2) = ... = \dim(H_n)$.

We construct a new code $\mathcal{Q}'$ over larger Hilbert spaces that can correct $d-1 < n/2$ arbitrary errors except with small probability. Register i of the n-component code $\mathcal{Q}'$ contains the following:

$$\langle A_{k_i}(H_i), s_i, [\ell_{ij}(\forall j \neq i)], [h_{\ell_{ji}}(s_i)(\forall j \neq i)]\rangle, \tag{5}$$

where we have used the classical authentication scheme (in systematic form):

$$m, \ell \to (m, h_\ell(m)), \tag{6}$$

which has error ϵ, and $(s_1, \ldots, s_n) \in_R SS_{n,d}(k_1, \ldots, k_n)$, a secret sharing scheme such that any $d-1$ s_i's contains no information about $(k_1, \ldots, k_n)$ whereas any d of those s_i's completely define $(k_1, \ldots, k_n)$.

For instance, the $n = 3$ case of this construction is as follows:

$$\begin{aligned} c'|\psi\rangle \quad = \quad & \langle A_{k_1}(H_1), s_1, [\ell_{12}, \ell_{13}], [h_{\ell_{21}}(s_1), h_{\ell_{31}}(s_1)]\rangle\,, \\ & \langle A_{k_2}(H_2), s_2, [\ell_{21}, \ell_{23}], [h_{\ell_{12}}(s_2), h_{\ell_{32}}(s_2)]\rangle\,, \\ & \langle A_{k_3}(H_3), s_3, [\ell_{31}, \ell_{32}], [h_{\ell_{13}}(s_3), h_{\ell_{23}}(s_3)]\rangle\,. \end{aligned} \tag{7}$$

Let $H_1' \otimes H_2' \otimes ... \otimes H_n'$ be the coding space of the new code

$$\mathcal{Q}'|\psi\rangle \in H_1' \otimes H_2' \otimes ... \otimes H_n'$$

Theorem 2 *If A_k is a* QAS *secure with error ϵ then $\mathcal{Q}'$ is a QECC correcting $d-1$ arbitrary errors except with prob. $O(2^{d/2}\sqrt{n\epsilon})$.*

5.1 Reconstruction

The reconstruction procedure is similar to that for the previous protocol, but slightly more involved, since we must verify the classical authentications as well. Rather than breaking the procedure into different cases, in this version of the protocol, we can systematically go through four steps: First, verify the classical authentications and discard any invalid classical share. Second, reconstruct the keys k_i. Third, verify and decode the quantum authentications. Fourth, discard any invalid quantum register and reconstruct the encoded quantum state.

1. Verify classical authentications:
   For each s_i, consider it valid if at least half its authentications are correct according to $\ell_{ji}, j \neq i$. Discard any share s_i which is not valid.
2. Reconstruct the keys k_i:
   Up to $d-1$ shares s_i can have been discarded in the first stage, so at least $n-d+1 \geq n/2+1 > d$ shares remain. Use these to reconstruct $(k_1, \ldots, k_n)$. If the remaining shares are not all consistent with a single value of the secret, $\mathcal{B}$ aborts and outputs the quantum state $|0\rangle$.
3. Verify and decode the quantum authentications:
   Use the key k_i to verify and decode the quantum authentication $A_{k_i}(H_i)$.
4. Reconstruct the encoded quantum state:
   Discard any registers which failed the quantum authentication, and use the remaining registers to reconstruct the valid codeword as $c|\psi\rangle \in H_1 \otimes \ldots \otimes H_n$ using the erasure recovery algorithm. (At most $d-1$ have been discarded.) If the remaining registers are not consistent with a single quantum codeword, $\mathcal{B}$ aborts and outputs the quantum state $|0\rangle$.

Proof:

Clearly, if no errors occurred, the above procedure will exactly reconstruct the original encoded state. We need to show, however, that it still approximately reconstructs the state when there are up to $d-1$ arbitrary errors in unknown locations. Let S be the set of registers attacked by the adversary, and let T be the remaining (i.e., correct) registers.

Note that the first two steps are purely classical. The adversary must output classical bit strings for the registers in S. However, if she alters any of the shares s_i, $\mathcal{B}$ will reject it in step 1 unless she successfully forges at least one authentication $h_{\ell_{ji}}(s_i)$ (for $j \in T$). This again is a purely classical task, and by the information-theoretic security of the classical authentication protocol, her probability of successfully doing so is at most ϵ per attempt, thus yielding a total probability bounded by $(n-d+1)\epsilon \leq n\epsilon$ for the $n+1-d$ parallel attempts.

Conversely, any share s_i from T will always be accepted by $\mathcal{B}$, since it will pass at least those authentications $h_{\ell_{ji}}(s_i)$ for $j \in T$, which comprise at least half of all of its authentications. Therefore, in stage 1, $\mathcal{B}$ always keeps at least $n-d+1$ shares. With probability at least $1-dn\epsilon$, the values s_i of the kept shares are the same as when the state was encoded, and therefore are all consistent with a single value of the secret. That is, for any strategy by the adversary, $\mathcal{B}$'s probability of aborting in stage 2 is at most $dn\epsilon$. Otherwise, $\mathcal{B}$ reconstructs the correct set of keys $(k_1, \ldots, k_n)$.

Now, the adversary is left with the task of mounting a quantum attack against the QAS protecting the quantum part of each register. In doing so, she is limited by the security condition for a QAS: While the authenticated quantum state itself provides some information about the keys k_i, the adversary would need additional information in order to successfully attack the scheme, and the classical secret sharing scheme ensures that the adversary, with access to only $d-1$ shares, cannot get *any* additional information.

Therefore, we know that the quantum authentications $A_{k_i}(H_i)$ will, when decoded by $\mathcal{B}$, produce states with fidelity at least $1-\epsilon$ to the subspace formed by the input state and $|\text{REJ}\rangle$. Furthermore, by the composibility property of our QAS, a similar condition (with fidelity $1-m\epsilon$) holds for any set of m registers.

Now, in stage 3, $\mathcal{B}$ measures $|\text{ACC}\rangle$ or $|\text{REJ}\rangle$ and keeps only those registers with $|\text{ACC}\rangle$. Let this set of registers be V. We know $T \subseteq V$, but V might be strictly larger than T, depending on the adversary's attack. That allows the possiblity that the reconstructed state might not be exactly the original encoded state, or that reconstruction might be impossible due to an inconsistency in the accepted registers.

We know from proposition 12 of [3] that the density matrix held by $\mathcal{B}$, conditioned on acceptance, has fidelity at least $1-\epsilon/p_{\text{ACC}}$ to the input density matrix, where p_{ACC} is the probability of $\mathcal{B}$ accepting the state. Therefore, let p_W be the probability that, if $\mathcal{B}$ were to test only the set W, he would accept that set. Some sets W will have large p_W, whereas for others, p_W will be quite small. However, p_W is at least as large as the probability that $\mathcal{B}$'s complete accepted set $V = W$.

Therefore, let us consider the probability that V, the set of accepted registers, has $p_V < 2^{-(d-1)/2}\sqrt{n\epsilon}$. Since $\mathcal{B}$ always accepts the registers in T, the only question is whether $\mathcal{B}$ accepts a given set of registers in S. There are therefore at most 2^{d-1} possible subsets that $\mathcal{B}$ could accept that have probability $p_V < 2^{-(d-1)/2}\sqrt{n\epsilon}$. Thus, the probability that $\mathcal{B}$ actually accepts one of these sets is at most $2^{(d-1)/2}\sqrt{n\epsilon}$. Otherwise, the fidelity of the reconstructed state on V to the actual input on those registers is at least $1-2^{(d-1)/2}\sqrt{n\epsilon}$. That means that, with probability at least $1-2^{(d-1)/2}\sqrt{n\epsilon}$, $\mathcal{B}$ reconstructs in stage 4 a state that has fidelity at least $1-2^{(d-1)/2}\sqrt{n\epsilon}$ to the original input state. Since there was also a probability ϵ that the adversary could force the protocol to abort at stage 2 of the reconstruction, the overall fidelity of the final reconstructed state to the initial input state is at least

$$(1-dn\epsilon) \cdot \left[1-2^{(d-1)/2}\sqrt{n\epsilon}\right]^2. \tag{8}$$

□

5.2 Specific examples

Let be a power of two $n = 2^m$. Let C_1 be a $[2^m, 2^{m-1}, 2^{m-1}+1]_{2^m}$ Extended Reed-Solomon code and C_2 be a $[2^m, 2^{m-1}+1, 2^{m-1}]_{2^m}$ Extended Reed-Solomon code.

The CSS code obtained from C_1, C_2 can correct $2^{m-1}-1 = n/2-1$ erasure errors:

$$\mathcal{Q} = [[2^m, 1, 2^{m-1}]]_{2^m} = [[n, 1, n/2]]_n.$$

The related code obtained from our construction would need $2m + 20s + 2\log m \in \Theta(m+s)$ bits [9, 3] of quantum authentication key per component, to obtain $\epsilon < 2^{-\Omega(m+s)}$ error probability. Each component also contains a secret share of n such keys, thus needs $w \in \Theta(2^m(m+s))$ bits per s_i [14]. The classical authentications will each require $\Theta(s+\log w)$ bits [9] per key for a total of $\Theta(2^m(s+\log w))$ per component to obtain $\epsilon < 2^{-\Omega(m+s)}$ error probability. To summarize, a total of $\Theta(2^m(m+s))$ bits per component are necessary to obtain error probabilities $\epsilon < 2^{-\Omega(m+s)}$.

The resulting fidelity

$$(1-dn\epsilon) \cdot \left[1-2^{(d-1)/2}\sqrt{n\epsilon}\right]^2 \in (1-2^{-\Omega(s)}) \cdot \left[1-2^{(d-1)/2}2^{-\Omega(s)}\right]^2$$

is exponentially (in s) close to 1 as long as $s \in \Omega(n) \subseteq \Omega(d)$.

We conclude that the resulting code is an $[[n, \Theta(\frac{1}{n+s}), n/2]]_{\Theta(n(n+s))}$ QECC correcting $n/2-1$ arbitrary errors with fidelity at least $1-2^{-\Omega(s)}$.

6 Discussion and open questions

We have constructed quantum error correcting codes that are capable of correcting general errors when up to half the registers are affected. This contrasts considerably with known upper bounds that limit a QECC to correcting errors on less than one-fourth of all registers. The price for being able to violate this bound is that we only correct the state approximately; however, we do so with exponentially good fidelity.

In general, extrapolating from exact performance of a quantum task to approximate performance is dangerous, but possible. Factors of the dimension may arise, and since the dimension is exponential in the number of qubits, dramatically different behavior becomes possible. This phenomenon is likely behind the performance of our codes, and suggests that high-fidelity AQECCs are only possible when working in high dimension. It remains an interesting open question, however, if it is actually possible to construct AQECCs with both high efficiency and high fidelity in the usual model where the encoded subspace has the same dimension as each register of the code.

Our codes instead consist of a small logical subspace and large registers containing both quantum and classical information. As such, they are not so useful for practical problems in quantum error correction, but do serve as an interesting in-principle demonstration of the potential power of approximate error correction. In addition, they may be directly useful for VQSS and MPQC. Any such construction must be more complex, however, to take account of dishonest senders and receivers, and to allow the participants in the protocol to alter a state in the correct way without altering it in any unapproved manner. Indeed, it remains possible that the prior bound of $n/4$ cheaters does in fact restrict VQSS and MPQC; however, we have shown here that the existing proof of that bound does not apply to VQSS and MPQC protocols which only guarantee approximate reconstruction of the quantum state.

Acknowledgements

Thanks to Umesh Vazirani for helpful discussions.

References

[1] D. Aharonov and M. Ben-Or. Fault tolerant quantum computation with constant error. In *Proc. of 29th STOC*, pages 176–188, El Paso, Texas, 4–6 May 1997. This is a preliminary version of [2].

[2] D. Aharonov and M. Ben-Or. Fault tolerant quantum computation with constant error rate. quant-ph/9906129. Preliminary version in *STOC '97*. Submitted to SIAM J. Comp., June 1999.

[3] H. Barnum, C. Crépeau, D. Gottesman, A. Tapp, and A. Smith. Authentication of quantum messages. In proceedings of The 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS'02), November 16 - 19, 2002 Vancouver, BC, Canada, pages 449–458. Complete version (including Proposition 12): Quantum Physics, abstract quant-ph/0205128, 22 pages, May 2002.

[4] H. Barnum, E. Knill and M. A. Nielsen, "On Quantum Fidelities and Channel Capacities," quant-ph/980901, IEEE Trans.Info.Theor. 46 (2000) 1317-1329.

[5] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, "Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels.", Phys. Rev. Lett. 76 (1996) 722-725, Quantum Physics, abstract quant-ph/9511027.

[6] J. L. Carter and M. N. Wegman, "New hash functions and their use in authentication and set equality", *Journal of Computer and System Sciences*, Vol. 22, 1981, pp. 265 – 279.

[7] R. Cramer, I. Damgård, S. Dziembowski, M. Hirt and T. Rabin. Efficient Multiparty Computations with Dishonest Minority In *Proc. of EUROCRYPT 1999*.

[8] C. Crépeau, D. Gottesman, and A. Smith. Secure multi-party quantum computation. In Proceedings of 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada. ACM, 2002, pages 643–652.

[9] P. Gemmell and M. Naor. Codes for interactive authentication. In D. R. Stinson, editor, *CRYPTO*, volume 773 of *Lecture Notes in Computer Science*, pages 355–367. Springer, 1994.

[10] D. Gottesman, "An Introduction to Quantum Error Correction.", Quantum Physics, abstract quant-ph/0004072, 15 pages, talk given at AMS Short Course on Quantum Computation.

[11] P. Hayden, D. Leung and A. Winter, Personal communications, 2002-2003.

[12] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[13] T. Rabin and M. Ben-Or. Verifiable Secret Sharing and Multiparty Protocols with Honest Majority. In *Proc. of STOC 1989*, p. 73–85.

[14] A. Shamir. How to share a secret. *Communications of the ACM*, 22:612–613, 1979.

[15] W. K. Wootters and W. H. Zurek, "A single quantum cannot be cloned", *Nature* 299, 802, 1982 .