Entanglement wedge reconstruction using the Petz map

Chi-Fang Chen,[a,c] **Geoffrey Penington**[b] **and Grant Salton**[a,b,c]

[a]*Department of Physics, Stanford University,*
Stanford CA 94305, U.S.A.

[b]*Stanford Institute for Theoretical Physics, Stanford University,*
Stanford CA 94305, U.S.A.

[c]*Institute for Quantum Information and Matter, Caltech,*
Pasadena, CA 91125, U.S.A.

E-mail: chifang@caltech.edu, geoffp@stanford.edu, gsalton@caltech.edu

ABSTRACT: At the heart of recent progress in AdS/CFT is the question of subregion duality, or entanglement wedge reconstruction: which part(s) of the boundary CFT are dual to a given subregion of the bulk? This question can be answered by appealing to the quantum error correcting properties of holography, and it was recently shown that robust bulk (entanglement wedge) reconstruction can be achieved using a universal recovery channel known as the *twirled Petz map*. In short, one can use the twirled Petz map to recover bulk data from a subset of the boundary. However, this map involves an averaging procedure over bulk and boundary modular time, and hence it can be somewhat intractable to evaluate in practice. We show that a much simpler channel, the Petz map, is sufficient for entanglement wedge reconstruction for any code space of fixed finite dimension — no twirling is required. Moreover, the error in the reconstruction will always be non-perturbatively small. From a quantum information perspective, we prove a general theorem extending the use of the Petz map as a general-purpose recovery channel to subsystem and operator algebra quantum error correction.

Contents

1 Introduction

The AdS/CFT correspondence — a duality between a gravitational theory in asymptotically AdS space, and a conformal field theory in one fewer spatial dimension — has enjoyed extraordinary success in the twenty years since it was first proposed [1]. An important problem in AdS/CFT is that of subregion duality: which subregion of the CFT (if any) is dual to a given subregion of the bulk spacetime? Recently, it was discovered that the bulk-to-boundary map in AdS/CFT defines a quantum error correcting code [2, 3]. In light of this profound revelation, the problem of subregion duality can be rephrased: which subregion of the bulk spacetime can be 'reconstructed' from a given subregion of the boundary?

Over the course of the last five years, this question has been answered [4–6]: the bulk region encoded in an arbitrary boundary region A is the so-called *entanglement wedge* of A, denoted a. Within a single, static timeslice, the entanglement wedge[1] of A is the bulk region bounded by A itself and its Ryu-Takayanagi surface χ_A [10], which is the minimal area bulk surface anchored to the boundary of A; see figure 1. Given any bulk operator ϕ_a lying within the entanglement wedge a, there exists a boundary operator Φ_A acting only on the boundary region A, which approximately reproduces the action of the bulk operator ϕ_a. The task of finding such an operator Φ_A is known as entanglement wedge reconstruction.

The conjecture of entanglement wedge reconstruction was developed in [11–13] and established with increasing levels of rigour in [4–6, 14]. It was shown in [9, 15] that the error in the reconstruction can be made non-perturbatively small at large boundary gauge group rank N, or equivalently small gravitational coupling G_N.

---

[1]This definition is valid only within a single, static timeslice of a static bulk spacetime, or at a moment of time reversal symmetry. More generally, and more formally, the covariant Ryu-Takayanagi surface χ_A is defined to be the smallest surface of *extremal* area homologous to A [7]. The entanglement wedge is then the bulk domain of dependence of *any* achronal bulk surface bounded by A and χ_A. At higher orders in perturbation theory, one should use the quantum extremal surface, which extremizes the Ryu-Takayanagi formula $\mathcal{A}/4G_N + S_{\text{bulk}}$, where S_{bulk} is the bulk entanglement entropy, rather than simply the classical area $\mathcal{A}$ [8, 9].

Figure 1. An operator ϕ_a, acting on the entanglement wedge a of $A = A_1 \cup A_2$, can be reconstructed on the boundary region A by the map $\mathcal{D}_A^\dagger : M_a \to M_A$. The solid interior curves represent the RT surface of A and the entire shaded region forms the entanglement wedge a (restricted to a single timeslice). The darker gray areas are the entanglement wedges of A_1 and A_2 individually, and also together form the causal wedge of A. Since the operator ϕ_a is not in the causal wedge of A, we cannot reconstruct it simply by using the bulk and boundary equations of motion; the more sophisticated machinery of quantum error correction is required. Moreover, ϕ_a can only be reconstructed on $A = A_1 \cup A_2$; neither A_1 nor A_2 alone contains any information about ϕ_a.

The realization that bulk reconstruction can be rephrased in the language of quantum error correction [2] paved the way for most of this success. Bulk operators in AdS/CFT are only well defined within the "code subspace" $\mathcal{H}_{\text{code}}$ of states with the correct smooth bulk geometry. If we let $J : \mathcal{H}_{\text{code}} \to \mathcal{H}_{CFT}$ be the isometry embedding this code subspace into the larger CFT Hilbert space, entanglement wedge reconstruction can be rephrased as the task of finding a decoding channel $\mathcal{D}$ that can recover from the noisy channel $\mathcal{N} = [J(\cdot)J^\dagger]_A$, where ρ_A is the restriction of the boundary state ρ to region A. More specifically, entanglement wedge reconstruction is equivalent to the existence of a decoding channel $\mathcal{D}$ such that, for all states ρ in the bulk code space,

$$\mathcal{D} \circ \mathcal{N}(\rho) \approx \rho_a, \tag{1.1}$$

where ρ_a is the restriction of the bulk state ρ to the entanglement wedge a.[2] If such a decoding channel exists, then we can use the adjoint channel $\mathcal{D}^\dagger$, defined by

$$\text{Tr}[\mathcal{D}^\dagger(\phi)\,\sigma] = \text{Tr}[\phi\,\mathcal{D}(\sigma)], \tag{1.2}$$

for all operators ϕ and states σ, to map bulk operators ϕ_a to boundary reconstructions $\Phi_A = \mathcal{D}^\dagger(\phi_a)$ with support only in region A. Since

$$\text{Tr}(\Phi_A\,\rho) = \text{Tr}[\phi_a\,\mathcal{D} \circ \mathcal{N}(\rho)] \approx \text{Tr}(\phi_a\,\rho), \tag{1.3}$$

the expectation values of ϕ_a and Φ_A approximately agree for all states $\rho \in S(\mathcal{H}_{\text{code}})$. It can be shown that this is also true for higher point correlators [5].

---

[2]Here, restriction can be thought of as a partial trace, but in an operator algebra quantum error correction picture, this is really a restriction of a state to a subalgebra.

Interestingly, the entanglement wedge a may contain regions outside of the 'causal wedge' of A (the intersection of the past and future of the boundary domain of dependence of A). Given a bulk operator ϕ in the causal wedge of a region A, it is well-understood how to reconstruct the operator ϕ within the boundary region A, given only the bulk and boundary equations of motion, using the so-called HKLL procedure [16, 17]. However, it was only by introducing the tool of quantum error correction that we have begun to understand that the entire entanglement wedge (rather than just the causal wedge) can indeed be reconstructed from region A.

The first clue that a boundary region A encodes more than just the causal wedge actually comes from the Ryu-Takayanagi (RT) formula [10, 18]. Including the leading quantum correction [19], the RT formula states that the entanglement entropy S_A of any boundary region A is given by

$$S_A = \frac{\mathcal{A}(\chi_A)}{4G_N} + S_{\text{bulk}}, \tag{1.4}$$

where $\mathcal{A}(\chi_A)$ is the area of the RT surface χ_A and S_{bulk} is the bulk entanglement entropy associated to the entanglement wedge of A. The entanglement entropy, although not an actual observable, is therefore a quantity that depends only on the reduced density matrix of the state on region A, but depends on the entire entanglement wedge in the bulk. Somewhat remarkably, (1.4) alone is, in fact, sufficient to imply the existence of decoding channels $\mathcal{D}$ that can be used for entanglement wedge reconstruction [4, 5]. The key intermediate step, which was shown in [14], is that (1.4) implies an approximate equality between bulk and boundary relative entropies.

Unfortunately, even though it is, at this point, very well established that entanglement wedge reconstruction is possible in principle (and hence that decoding channels $\mathcal{D}$ must exist), it has proved somewhat challenging to find constructions that work for bulk operators lying outside the causal wedge (and hence for which we cannot use the HKLL prescription) and that are both explicit and practical. An explicit, if somewhat impractical, general construction was given in [2, 4]. However, this construction relies on the unphysical assumption of exact quantum error correction, which does not exist at finite N.

It was shown in [6] that the evolution of bulk operators in bulk modular time is related via the extrapolate dictionary to the evolution of boundary operators in boundary modular time. Since bulk modular evolution should be linear in the free field approximation $N \to \infty$, one might hope to expand a bulk operator at any point in the entanglement wedge in terms of the modular evolution of operators at the boundary of the wedge, and thus derive an explicit formula for the boundary representation of the bulk operator. However, as yet, the details of this expansion remain unknown, and it is not even clear how to show rigorously that one should exist at all.

Finally, it was demonstrated in [5], using the tools of approximate operator algebra quantum error correction, that robust entanglement wedge reconstruction can be achieved using the so-called *twirled Petz map* [20], even at finite N. The twirled Petz map is an example of a "universal recovery channel" — a general purpose decoding map that lets one approximately recover from the action of a quantum channel. That is, given a quantum

channel $\mathcal{N}$ and a fixed state σ, the goal is to find a recovery channel $\mathcal{R}_{\sigma,\mathcal{N}}$ such that $\mathcal{R}_{\sigma,\mathcal{N}} \circ \mathcal{N}[\rho] \approx \rho$ for all ρ. The twirled Petz map is one such recovery channel $\mathcal{R}_{\sigma,\mathcal{N}}$, defined to be

$$\mathcal{R}_{\sigma,\mathcal{N}}(\rho) = \int dt\, \frac{\pi}{2}\, (\cosh(\pi t) + 1)^{-1}\, \sigma^{\frac{1-it}{2}} \mathcal{N}^{\dagger} \left([\mathcal{N}(\sigma)]^{-\frac{1-it}{2}} \rho\, [\mathcal{N}(\sigma)]^{-\frac{1+it}{2}} \right) \sigma^{\frac{1+it}{2}}. \quad (1.5)$$

If we replace σ by the maximally mixed state τ, the expression simplifies significantly. We can use the twirled Petz map for bulk reconstruction by setting the channel $\mathcal{N}$ to be $\mathcal{N} = [J(\cdot)J^{\dagger}]_A$. With the simple choice $\sigma = \tau$, this leads to the boundary reconstruction Φ_A of a bulk operator ϕ_a as

$$\Phi_A = \mathcal{R}_{\tau,\mathcal{N}}^{\dagger}(\phi_a) = \frac{1}{d_{\text{code}}} \int dt\, \frac{\pi}{2}\, (\cosh(\pi t) + 1)^{-1}\, \tau_A^{-\frac{1-it}{2}} [J\phi_a J^{\dagger}]_A \tau_A^{-\frac{1+it}{2}}, \quad (1.6)$$

where $\tau_A = \mathcal{N}(\tau)$. Even though choosing the reference state to be maximally mixed has simplifed the expression, it still involves a twirling or averaging over boundary modular time with the precisely chosen weighting $\pi/2\,(\cosh(\pi t) + 1)^{-1}$.

In this paper, we will show that such averaging is unnecessary for code spaces of any fixed finite dimension in the semiclassical limit $N \to \infty$ and $G_N \to 0$. Instead it is sufficient to use the much simpler *Petz map* [21] reconstruction

$$\Phi_A = \frac{1}{d_{\text{code}}} \tau_A^{-1/2} [J\phi_a J^{\dagger}]_A\, \tau_A^{-1/2}. \quad (1.7)$$

We are hopeful that this more tractable recovery map should prove significantly easier to evaluate explicitly; we discuss the challenges and prospects of doing so in section 4. For other examples of situations where twirling is unnecessary and the Petz map itself is sufficient, see [22–25].

Our strategy for proving the efficacy of the Petz map for entanglement wedge reconstruction builds on work by Barnum and Knill [26], who showed that, for ordinary subspace quantum error correction, the Petz map will always have an *average* decoding error that is almost as small as the average error of the optimal decoding channel. Roughly speaking, the Petz map is always 'pretty good'. We extend these results to subsystem and operator algebra quantum error correcting codes and then show that the average error can always be used to bound the worst-case error, so long as the dimension of the code space does not grow too quickly in the limit of large N. (We discuss very large code spaces such as those of black hole microstates briefly in section 4.)

In section 2, we formalize the problem of entanglement wedge reconstruction in the language of quantum error correction and show how to adapt the results of Barnum and Knill to prove that reconstruction is possible using the Petz map. Our main technical result is a general theorem in quantum error correction, the proof of which is given in section 3, and an application of which simplifies the problem of entanglement wedge reconstruction. Section 4 consists of a brief discussion of potential applications and extensions of our work.

2 Entanglement wedge reconstruction and the Petz map

In order to apply information-theoretic techniques to the problem of entanglement wedge reconstruction, we first need to rephrase our task in the language of quantum information. We employ the same framework used in [5] — finite-dimensional approximate operator algebra quantum error correction.

The AdS/CFT correspondence is a duality between a boundary conformal field theory with Hilbert space $\mathcal{H}_{CFT}$, and a bulk quantum gravity theory. In principle, if AdS/CFT is supposed to be a true duality between theories, the 'bulk' Hilbert space should be isomorphic to the boundary Hilbert space $\mathcal{H}_{CFT}$. However, a complete, non-perturbative, microscopic description of the entire Hilbert space from a purely bulk perspective, if one exists, remains unknown. Moreover, any such Hilbert space would be dominated by large black holes. Instead, we are normally only interested in a small 'code subspace' $\mathcal{H}_{\text{code}}$ of states with a smooth semiclassical bulk geometry; for example, we might consider small bulk perturbations about the vacuum state. We therefore have an isometry $J : \mathcal{H}_{\text{code}} \to \mathcal{H}_{CFT}$. Equivalently, we can consider the quantum channel $\mathcal{J}(\cdot) = J(\cdot)J^\dagger$ which maps bulk density matrices to boundary density matrices. As it turns out, none of our results rely on $\mathcal{J}$ being an isometry as opposed to a more general quantum channel.

For simplicity, we assume that both $\mathcal{H}_{\text{code}}$ and $\mathcal{H}_{CFT}$ are finite-dimensional. In the case of $\mathcal{H}_{\text{code}}$, this is justified by the fact that we cannot include arbitrarily high energy excitations in the bulk without causing significant backreaction and eventually creating a black hole. In the case of $\mathcal{H}_{CFT}$, we should be able to regularize the boundary theory in the UV, while only affecting bulk physics close to the boundary. Of course, the real value of these assumptions for our purposes is that they allow us to apply known results from the large literature on finite-dimensional quantum error correction.

We denote the algebra of observables on the Hilbert space $\mathcal{H}_{\text{code}}$ by $\mathcal{B}(\mathcal{H}_{\text{code}})$ and the algebra of observables on $\mathcal{H}_{CFT}$ by $\mathcal{B}(\mathcal{H}_{CFT})$. The entanglement wedge a has an associated von Neumann subalgebra $\mathcal{M}_a \overset{i}{\hookrightarrow} \mathcal{B}(\mathcal{H}_{\text{code}})$, consisting of bulk observables that act only on a; similarly, the boundary region A is associated with a von Neumann subalgebra $\mathcal{M}_A \overset{i}{\hookrightarrow} \mathcal{B}(\mathcal{H}_{CFT})$. We use the notation from [5], where the space of density matrices on a von Neumann subalgebra $\mathcal{M}$ acting on a Hilbert space $\mathcal{H}$ is denoted by $S(\mathcal{M}) \cong S(\mathcal{H}) \cap \mathcal{M}$. This space is isomorphic to the space of positive normalized linear functionals on the algebra. See the appendix of [5] for more details.

The question of entanglement wedge reconstruction can then be rephrased as the question of whether the channel $\mathcal{N} = [\mathcal{J}(\cdot)]_A$ forms an approximate error-correcting code for the algebra $\mathcal{M}_a$. Here, the restriction channel $[\cdot]_A$ simply projects the density matrix onto the algebra $\mathcal{M}_A$. In other words, entanglement wedge reconstruction is possible if (and only if) there exists a decoding channel $\mathcal{D} : S(\mathcal{M}_A) \to S(\mathcal{M}_a)$ such that

$$\mathcal{D} \circ \mathcal{N}(\rho) \approx \rho_a, \tag{2.1}$$

for all states $\rho \in S(\mathcal{H}_{\text{code}})$; the restriction ρ_a is the projection of ρ onto $\mathcal{M}_a$. For subsystem algebras, this corresponds to taking a partial trace over the complementary subsystem and

Figure 2. In the Heisenberg picture, $\mathcal{M}_a \overset{i}{\hookrightarrow} \mathcal{B}(\mathcal{H}_{\text{code}})$ and $\mathcal{M}_A \overset{i}{\hookrightarrow} \mathcal{B}(\mathcal{H}_{CFT})$ are von Neumann subalgebras acting on the code space $\mathcal{H}_{\text{code}}$ and CFT Hilbert space $\mathcal{H}_{CFT}$ respectively. The Heisenberg channel $\mathcal{J}^\dagger = J^\dagger(\cdot)J$ maps boundary observables to their projection in the code space. The task of entanglement wedge reconstruction is to find a Heisenberg decoding channel $\mathcal{D}^\dagger : \mathcal{M}_a \to \mathcal{M}_A$ that maps bulk observables ϕ_a in $\mathcal{M}_a$ to boundary observables Φ_A in $\mathcal{M}_A$. When projected into the code space using $\mathcal{J}^\dagger$, the boundary observable Φ_A should reproduce the original bulk observable ϕ_a. In the Schrödinger picture, the directions of all channels are reversed. The channel $\mathcal{J}$ now maps bulk states to the corresponding boundary states. The Heisenberg channels i_a and i_A, which embed the von Neumann subalgebras $\mathcal{M}_a$ and $\mathcal{M}_A$ into the larger algebras of observables $\mathcal{B}(\mathcal{H}_{\text{code}})$ and $\mathcal{B}(\mathcal{H}_{CFT})$, are the adjoints of the restriction maps onto $S(\mathcal{M}_a)$ and $S(\mathcal{M}_A)$ respectively. Finally, the decoding channel $\mathcal{D} : S(\mathcal{M}_A) \to S(\mathcal{M}_a)$ satisfies $\mathcal{D}[\mathcal{J}(\cdot)_A] \approx (\cdot)_a$.

hence agrees with the usual notion of a reduced density matrix; operator algebra quantum error correction therefore generalizes subsystem quantum error correction.

In the Heisenberg (adjoint) picture, this condition becomes

$$\mathcal{N}^\dagger \circ \mathcal{D}^\dagger(\phi_a) = \mathcal{J}^\dagger \circ \mathcal{D}^\dagger(\phi_a) \approx \phi_a. \tag{2.2}$$

Note that, since the adjoint of the restriction channel is simply the embedding of the subalgebra in the larger algebra of observables, $\mathcal{N}^\dagger(O_A) = \mathcal{J}^\dagger(O_A)$ for all operators $O_A \in \mathcal{M}_A$. In other words, $\Phi_A = \mathcal{D}^\dagger(\phi_a)$ acts in approximately the same way as ϕ_a:

$$\text{Tr}(\Phi_A \mathcal{J}(\rho)) \approx \text{Tr}(\phi_a \rho) \tag{2.3}$$

The complete setup, in both the Schrödinger and Heisenberg pictures, is shown in figure 2.

It was argued in [5] that the twirled Petz map provides an example of a decoding map with an error that is perturbatively suppressed in $1/N$. It was then shown in [15] that there must exist some decoding channel $\mathcal{D}$ with a non-perturbatively small error; however, this argument was non-constructive. Both results relied heavily on the approximate equality between bulk and boundary relative entropies found in [14]. A refined statement of this approximate relative entropy equality was derived in [9], which allows one to show the existence of a decoding channel that is accurate to all orders in perturbation theory. Here, we shall simply take as our starting assumption the existence of *some* good decoding channel $\mathcal{D}'$; we will not need to know any details of this channel. We can therefore use the result of [15] to assume that the decoding error when using this channel is non-perturbatively small. The following theorem, which we prove in section 3, then implies that the Petz map is also a good decoding channel:

Theorem 1. *Let $\mathcal{M}_a \overset{i}{\hookrightarrow} \mathcal{B}(\mathcal{H}_{\mathrm{code}})$ be a von Neumann subalgebra acting on the code space $\mathcal{H}_{\mathrm{code}}$ with dimension d_{code}, let $\mathcal{N}$ be a quantum channel, and suppose that there exists a channel $\mathcal{D}'$ such that*

$$\|\mathcal{D}' \circ \mathcal{N}(\rho) - \rho_a\|_1 < \delta \,.$$

Let

$$\mathcal{P}_{\tau, \mathcal{N}} := \frac{1}{d_{\mathrm{code}}} \mathcal{N}^\dagger \left[\mathcal{N}(\tau)^{-1/2} (\cdot) \mathcal{N}(\tau)^{-1/2} \right]$$

be the Petz map with maximally mixed reference state τ. Then

$$\|\mathcal{P}_{\tau, \mathcal{N}} \circ \mathcal{N}(\rho)|_a - \rho_a\|_1 \le d_{\mathrm{code}} \sqrt{8\delta}. \tag{2.4}$$

Note that our bound on the error when reconstructing the reduced state using the Petz map $\mathcal{P}_{\tau, \mathcal{N}}$ is significantly higher than the original error δ. Not only is the error proportional to $\sqrt{\delta}$, but it is also proportional to the dimension of the code space d_{code}. As we shall see in section 3, the square root appears because of inefficiencies in converting between trace distances and fidelities using the Fuchs-van de Graaf inequalities [27], while the factor of d_{code} appears in order to convert a bound on the average-case error into a bound on the worst-case error. Nevertheless, so long as the error using the original decoding channel $\mathcal{D}'$ is non-perturbatively small, the Petz map error will also be non-perturbatively small, provided the dimension of the code space does not grow superpolynomially in N. For most code spaces of interest, such as perturbations about the vacuum, the code space dimension will be $O(1)$, and so this factor of code space dimension is not problematic. We discuss very large code spaces, such as those containing large numbers of black hole microstates, briefly in section 4. However, so long as we confine ourselves to perturbative excitations of quantum fields in a fixed gravitational background, the Petz map can always be trusted — no twirling is required.

3 Proof of Theorem 1

The spirit of Theorem 1 follows that of Barnum and Knil [26], who proved the following theorem for ordinary subspace quantum error correction:

Theorem 2 (Barnum-Knill [26]). *Given any pair of quantum channels $\mathcal{D}'$, $\mathcal{N}$, and ensemble of commuting density matrices (p_k, ρ_k) whose sum $\sum_k p_k \rho_k = \rho$, the Petz map*

$$\mathcal{P}_{\rho, \mathcal{N}}[\cdot] := \rho^{1/2} \mathcal{N}^\dagger \left[\mathcal{N}(\rho)^{-1/2} (\cdot) \mathcal{N}(\rho)^{-1/2} \right] \rho^{1/2}$$

with reference state ρ, satisfies

$$\sum_k p_k F(\rho_k, \mathcal{P}_{\mathcal{N}, \rho} \circ \mathcal{N}) \ge \left(\sum_k p_k F(\rho_k, \mathcal{D}' \circ \mathcal{N}) \right)^2. \tag{3.1}$$

Here, the entanglement fidelity $F(\sigma, \mathcal{Z})$ is defined by

$$F(\sigma, \mathcal{Z}) := \langle \sigma | \, V_{\mathcal{Z}}^\dagger \left(|\sigma\rangle \langle \sigma| \otimes \mathbb{1}_E \right) V_{\mathcal{Z}} \, |\sigma\rangle \,,$$

where $|\sigma\rangle \in \mathcal{H}_{\mathrm{code}} \otimes \mathcal{H}_R$ is a purification of $\sigma \in S(\mathcal{H}_{\mathrm{code}})$ and $V_{\mathcal{Z}} : \mathcal{H}_{\mathrm{code}} \to \mathcal{H}_{\mathrm{code}} \otimes \mathcal{H}_E$ is a Stinespring extension of the channel $\mathcal{Z} : S(\mathcal{H}_{\mathrm{code}}) \to S(\mathcal{H}_{\mathrm{code}})$.

If we now assume that $\mathcal{D}'$ is a recovery channel for $\mathcal{N}$ that works with high fidelity, then Theorem 2 states that $\mathcal{P}_{\mathcal{N},\rho} \circ \mathcal{N}$ is close to the identity when measured using the average entanglement fidelity of an ensemble $\{\rho_k\}$ with average state ρ. Note that, unlike our Theorem 1, there is no factor of d_{code} in the size of the error for the Petz map $\mathcal{P}_{\mathcal{N},\rho}$ as compared to the original decoding channel $\mathcal{D}'$. Instead, (3.1) implies that the error, measured using the *average entanglement fidelity*, increases by at most a factor of two.[3] The factor of d_{code} will appear when we convert this average error into a worst-case error.

For concreteness, let us write down an explicit basis for the von Neumann subalgebra $\mathcal{M}_a$. The exact description of $\mathcal{J}$ and $\mathcal{M}_A$ (and hence $\mathcal{N}$) are unimportant for our purposes. It is a fact about finite-dimensional von Neumann algebras (see, for example, [28]) that we can always find a set of Hilbert spaces $\mathcal{H}_\alpha$ and $\mathcal{H}_{\bar\alpha}$, parameterized by α, such that

$$\mathcal{M}_a = \bigoplus_\alpha \mathcal{B}(\mathcal{H}_\alpha) \otimes \mathbb{1}_{\bar\alpha},$$
$$\mathcal{H}_{\text{code}} = \bigoplus_\alpha \mathcal{H}_\alpha \otimes \mathcal{H}_{\bar\alpha}. \tag{3.2}$$

Note that

$$\sum_\alpha d_\alpha d_{\bar\alpha} = d_{\text{code}}, \tag{3.3}$$

where d_α, $d_{\bar\alpha}$ and d_{code} are the dimensions of $\mathcal{H}_\alpha$, $\mathcal{H}_{\bar\alpha}$ and $\mathcal{H}_{\text{code}}$ respectively. In this basis, any state $\rho_a \in S(\mathcal{M}_a)$ can be parameterized as

$$\rho_a = \sum_\alpha p_\alpha \rho_\alpha \otimes \tau_{\bar\alpha} = \sum_{\alpha, i_\alpha} p_\alpha p_{i_\alpha}^{(\alpha)} |i_\alpha\rangle \langle i_\alpha| \otimes \tau_{\bar\alpha}, \tag{3.4}$$

where the states $\tau_{\bar\alpha} \in S(\mathcal{H}_{\bar\alpha})$ are maximally mixed, $\rho_\alpha \in S(\mathcal{H}_\alpha)$ are normalized density matrices, p_α and $p_i^{(\alpha)}$ are normalized probability distributions, and $|i_\alpha\rangle$ forms an orthonormal basis for $\mathcal{H}_\alpha$.

We now have all the ingredients we need to begin a proof of Theorem 1. Let $\mathcal{Z} = \mathcal{P}_{\tau,\mathcal{N}} \circ \mathcal{N}$. We first note that $\mathcal{Z}$ is a self-adjoint superoperator. For any operator ϕ,

$$\text{Tr}[\phi\,\mathcal{Z}(\rho)] = \frac{1}{d_{\text{code}}} \text{Tr}\left[\phi\,\mathcal{N}^\dagger\left(\mathcal{N}(\tau)^{-1/2}\,\mathcal{N}(\rho)\,\mathcal{N}(\tau)^{-1/2}\right)\right] \tag{3.5}$$

$$= \text{Tr}[\mathcal{Z}(\phi)\rho] = \text{Tr}[\phi\,\mathcal{Z}^\dagger(\rho)]. \tag{3.6}$$

Hence we have that $\mathcal{Z} = \mathcal{Z}^\dagger$. Note that this argument relied crucially on our choice for the reference state in the Petz map $\mathcal{P}_{\tau,\mathcal{N}}$ as the maximally mixed state.

Now, let $\phi_a \in \mathcal{M}_a$ be a Hermitian operator, which we can assume to have eigendecomposition

$$\phi_a = \lambda_{i_\alpha} |i_\alpha\rangle \langle i_\alpha|. \tag{3.7}$$

---

[3]An entanglement fidelity $F(\rho, \mathcal{D} \circ \mathcal{N}) = 1$ implies perfect recovery of a purification of ρ. Hence, we can naturally quantify the recovery error when decoding using the channel $\mathcal{D}'$ by

$$\delta = 1 - \sum_k p_k F(\rho_k, \mathcal{D}' \circ \mathcal{N}).$$

The equivalent error measure, when decoding using the Petz map $\mathcal{P}_{\rho,\mathcal{N}}$, will then be bounded by $2\,\delta$.

We can bound the operator norm

$$\|\mathcal{Z}^\dagger(\phi_a) - \phi_a\|_\infty \leq \|\mathcal{Z}^\dagger(\phi_a) - \phi_a\|_1 \tag{3.8}$$

$$\leq \sum_{\alpha, i_\alpha} |\lambda_{i_\alpha}| \, \|(\mathcal{Z}^\dagger - \mathbb{1})[|i_\alpha\rangle \langle i_\alpha| \otimes \mathbb{1}_{\bar\alpha}]\|_1 \tag{3.9}$$

$$= \sum_{\alpha, i_\alpha} |\lambda_{i_\alpha}| \, d_{\bar\alpha} \|\mathcal{Z}[\rho_\alpha^{i_\alpha}] - \rho_\alpha^{i_\alpha}\|_1, \tag{3.10}$$

where the first inequality uses the monotonicity of the Schatten p-norms, the second inequality used the triangle inequality, and in the final equality we factored out $d_{\bar\alpha}$ so that $\rho_\alpha^{i_\alpha} = |i_\alpha\rangle \langle i_\alpha| \otimes \tau_{\bar\alpha}$ are normalized states, and more importantly we used the fact that the channel $\mathcal{Z}$ is self-adjoint. We now simply need to bound the average trace norm error of the channel $\mathcal{Z}$ on states $\rho_a \in S(\mathcal{M}_a)$. This is quadratically controlled by Theorem 2:

Proposition 2.1.

$$\sum_{i_\alpha, \alpha} \frac{d_{\bar\alpha}}{d_{\text{code}}} \|\mathcal{Z}[\rho_\alpha^{i_\alpha}] - \rho_\alpha^{i_\alpha}\|_1^2 \leq 8\delta \tag{3.11}$$

Proof. We first note that

$$\sum_{i_\alpha, \alpha} \frac{d_{\bar\alpha}}{d_{\text{code}}} \rho_\alpha^{i_\alpha} = \tau. \tag{3.12}$$

Hence

$$\sum_{i_\alpha, \alpha} \frac{d_{\bar\alpha}}{d_{\text{code}}} \|\mathcal{Z}[\rho_\alpha^{i_\alpha}] - \rho_\alpha^{i_\alpha}\|_1^2 \leq 4 \sum_{i_\alpha, \alpha} \frac{d_{\bar\alpha}}{d_{\text{code}}} (1 - F(\rho_\alpha^{i_\alpha}, \mathcal{Z}[\rho_\alpha^{i_\alpha}])) \tag{3.13}$$

$$\leq 4 - 4(\sum_\alpha \frac{d_{\bar\alpha}}{d_{\text{code}}} F(\rho_\alpha^{i_\alpha}, \mathcal{D}' \circ \mathcal{N}[\rho_\alpha^{i_\alpha}]))^2 \tag{3.14}$$

$$\leq 4 - 4 \left(\sum_\alpha \frac{d_{\bar\alpha}}{d_{\text{code}}} \left(1 - \frac{1}{2} \|\mathcal{D}' \circ \mathcal{N}[\rho_\alpha^{i_\alpha}] - \rho_\alpha^{i_\alpha}\|_1 \right) \right)^4 \tag{3.15}$$

$$\leq 8\delta, \tag{3.16}$$

where the first inequality uses one of the Fuchs-van de Graaf inequalities [27], the second uses (3.12) and Theorem 2, the fourth again uses the Fuchs-van de Graaf inequalities, and the fifth uses our assumption $\|\mathcal{D}'\mathcal{N}(\rho) - \rho_a\|_1 < \delta$ and (3.3). □

Applying Proposition 2.1 to (3.8), we find

$$\sum_{\alpha, i_\alpha} |\lambda_{i_\alpha}| \, d_{\bar\alpha} \|\mathcal{Z}[\rho_\alpha^{i_\alpha}] - \rho_\alpha^{i_\alpha}\|_1 \leq \|\phi_a\|_\infty \sum_{\alpha, i_\alpha} \sqrt{d_{\bar\alpha} d_{\text{code}}} \cdot \sqrt{\frac{d_{\bar\alpha}}{d_{\text{code}}}} \|\mathcal{Z}[\rho_\alpha^{i_\alpha}] - \rho_\alpha^{i_\alpha}\|_1 \tag{3.17}$$

$$\leq \|\phi_a\|_\infty \sqrt{\sum_{\alpha, i_\alpha} d_{\bar\alpha} d_{\text{code}}} \cdot \sqrt{8\delta} \tag{3.18}$$

$$= \|\phi_a\|_\infty d_{\text{code}} \cdot \sqrt{8\delta}, \tag{3.19}$$

where, in the first inequality, we used the fact that $\|\phi_a\|_\infty \geq |\lambda_{i_\alpha}|$ for all λ_{i_α} and, in the second inequality, we used the Cauchy-Schwarz inequality. We therefore find that

$$\|\mathcal{Z}^\dagger(\phi_a) - \phi_a\|_\infty \leq \|\phi_a\|_\infty d_{\text{code}} \sqrt{8\delta}.$$

Now, since

$$\|\mathcal{Z}(\rho)_a - \rho_a\|_1 = \sup_{\phi_a \in \mathcal{B}(\mathcal{H}_a)} \frac{1}{\|\phi_a\|_\infty} \text{Tr}([\mathcal{Z}^\dagger(\phi_a) - \phi_a]\rho) \leq \sup_{\phi_a \in \mathcal{B}(\mathcal{H}_a)} \frac{1}{\|\phi_a\|_\infty} \|\mathcal{Z}^\dagger(\phi_a) - \phi_a\|_\infty, \tag{3.20}$$

we immediately arrive at our desired result

$$\|(\mathcal{P}_{\tau,\mathcal{N}} \circ \mathcal{N}[\rho])_a - \rho_a\|_1 = \|(\mathcal{Z}[\rho])_a - \rho_a\|_1 \leq d_{\text{code}} \sqrt{8\delta}, \tag{3.21}$$

for any state $\rho \in S(\mathcal{H}_{\text{code}})$.

Note that we could have directly seen from Proposition 2.1 (using the triangle inequality) that for any state $\rho_a \in S(\mathcal{M}_a)$ we have

$$\|\mathcal{Z}(\rho_a)_a - \rho_a\|_1 \leq \sqrt{8\delta d_{\text{code}}}. \tag{3.22}$$

However, although this is a tighter bound than (3.21), it only applies to states in the code space that are of the form given in (3.4). In the Heisenberg picture, we want our reconstructed operator to work for all states in the code space — not just states in $S(\mathcal{M}_a)$.

The same problem of extending reconstruction from states $\rho_a \in S(\mathcal{M}_a)$ to all states $\rho \in S(\mathcal{H}_{\text{code}})$ was previously encountered for the twirled Petz map in [5]. It was shown that the approximate equality between bulk and boundary relative entropies [14] implies that any state $\rho \in S(\mathcal{H}_{\text{code}})$ satisfies

$$\mathcal{N}(\rho) \approx \mathcal{N}(\rho_a). \tag{3.23}$$

Hence (3.22) implies that, for all states $\rho \in S(\mathcal{H}_{\text{code}})$, we have

$$\|\mathcal{Z}(\rho)_a - \rho\|_1 \leq \sqrt{8\delta d_{\text{code}}} + \varepsilon, \tag{3.24}$$

where ε is independent of d_{code} and $\varepsilon \to 0$ in the limit $N \to \infty$. However, (3.23) really only holds because of the complementary recovery property of AdS/CFT. Not only does region A learn everything about the entanglement wedge a, it also learns nothing about the complementary bulk region $\bar{a}$, which is the entanglement wedge of region $\bar{A}$. In general, operator algebra quantum error correcting codes will not even approximately satisfy (3.23) — consider, for example, the case where $\mathcal{N}$ is the identity channel and $\mathcal{M}_a$ is any proper subalgebra of the algebra of observables $\mathcal{B}(\mathcal{H}_{\text{code}})$. It follows that (3.24) is specific to holographic codes. In contrast, Theorem 1 is a completely general fact about operator algebra quantum error correction. *Theorem 1 is therefore a true extension of the range of validity of the Petz map as a general-purpose, approximate recovery channel to operator algebra and subsystem codes.*

4 Discussion

In this work, we proved a theorem in quantum error correction about the quality of decoding using the Petz map as a general recovery channel. Our theorem generalizes the work of Barnum and Knill [26] to the case of operator algebra quantum error correction, and subsystem quantum error correction.

By applying our theorem to AdS/CFT, we showed that entanglement wedge reconstruction can be achieved using the Petz map, so long as the dimension of the code space we expect to be able to reconstruct is not too large. In particular, the Petz map constitutes a good recovery map provided the code space dimension does not grow superpolynomially in the limit of large N. In practice, this is almost always the case for code spaces of interest.

It is worth commenting briefly on the major exception to this rule: code spaces containing large numbers of black hole microstates.[4] The entropy of such code spaces may, in general, be $O(1/G_N)$. Hence the dimension of the code space may be exponential in N. However, as yet, the only black hole microstates that we understand are generic, equilibrium microstates. For code spaces made out of such microstates, we would expect the worst-case and average reconstruction errors to approximately agree, even though the in-principle large code space dimension means that very large differences between these two fidelities are possible. It is therefore reasonable to hope that the Petz map will still be valid for entanglement wedge reconstruction. On the other hand we should not trust any semi-classical description of non-generic, finely-tuned black hole microstates, and thus entanglement wedge reconstruction might not be possible for such states. As such, there are no known situations in which entanglement wedge reconstruction is possible, yet we cannot use the Petz map to achieve it.

While we emphasized the utility of the Petz map over other reconstruction techniques, we have not made any serious attempt to actually evaluate the Petz map in particular cases. Though the Petz map is much simpler to write down and, in principle, evaluate than the twirled Petz map, there still remain significant obstacles to doing so. Let us briefly discuss the challenges involved. We wish to explicitly evaluate

$$\Phi_A = \frac{1}{d_{\text{code}}} \tau_A^{-1/2} [J \phi_a J^\dagger]_A \, \tau_A^{-1/2}. \tag{4.1}$$

The operator $J \phi_a J^\dagger$ can be found by taking the global HKLL boundary reconstruction Φ^{HKLL} and projecting it into the code space [5]

$$J \phi_a J^\dagger = P_{\text{code}} \Phi^{\text{HKLL}} P_{\text{code}}. \tag{4.2}$$

Therefore, the main challenge lies in finding the restriction of this operator to region A. For simplicity, we assume, in accordance with common practice (though not with reality) that the CFT Hilbert space factorizes as $\mathcal{H}_{CFT} \cong \mathcal{H}_A \otimes \mathcal{H}_{\bar{A}}$ with $\mathcal{M}_A \cong \mathcal{B}(\mathcal{H}_A)$; the restriction map is then simply a partial trace over $\mathcal{H}_{\bar{A}}$. One difficulty arises because the HKLL procedure gives an operator Φ that is not localized in time. To take the partial

---

[4]For a detailed discussion of this topic see, for example, [15].

trace over region $\bar{A}$, we need to use the Heisenberg equations of motion to rewrite Φ_A in terms of operators at time zero.[5] Such operators will in general be very complicated and hard to evaluate. Essentially, the obstruction is simply the usual obstruction to evaluating quantities that are not protected by symmetry on the boundary side of AdS/CFT. Strongly coupled quantum field theories are simply hard to deal with; thankfully, there also exists a weakly coupled bulk.

Acknowledgments

References

[1] J.M. Maldacena, *The large N limit of superconformal field theories and supergravity*, *Int. J. Theor. Phys.* **38** (1999) 1113 [*Adv. Theor. Math. Phys.* **2** (1998) 231] [`hep-th/9711200`] [INSPIRE].

[2] A. Almheiri, X. Dong and D. Harlow, *Bulk locality and quantum error correction in AdS/CFT*, *JHEP* **04** (2015) 163 [`arXiv:1411.7041`] [INSPIRE].

[3] F. Pastawski, B. Yoshida, D. Harlow and J. Preskill, *Holographic quantum error-correcting codes: toy models for the bulk/boundary correspondence*, *JHEP* **06** (2015) 149 [`arXiv:1503.06237`] [INSPIRE].

[4] X. Dong, D. Harlow and A.C. Wall, *Reconstruction of bulk operators within the entanglement wedge in gauge-gravity duality*, *Phys. Rev. Lett.* **117** (2016) 021601 [`arXiv:1601.05416`] [INSPIRE].

[5] J. Cotler, P. Hayden, G. Penington, G. Salton, B. Swingle and M. Walter, *Entanglement wedge reconstruction via universal recovery channels*, *Phys. Rev.* **X 9** (2019) 031011 [`arXiv:1704.05839`] [INSPIRE].

[6] T. Faulkner and A. Lewkowycz, *Bulk locality from modular flow*, *JHEP* **07** (2017) 151 [`arXiv:1704.05464`] [INSPIRE].

[7] V.E. Hubeny, M. Rangamani and T. Takayanagi, *A covariant holographic entanglement entropy proposal*, *JHEP* **07** (2007) 062 [`arXiv:0705.0016`] [INSPIRE].

---

[5]For a more detailed discussion of similar issues, see [5].

[8] N. Engelhardt and A.C. Wall, *Quantum extremal surfaces: holographic entanglement entropy beyond the classical regime*, *JHEP* **01** (2015) 073 [arXiv:1408.3203] [ɪɴSPIRE].

[9] X. Dong and A. Lewkowycz, *Entropy, extremality, euclidean variations and the equations of motion*, *JHEP* **01** (2018) 081 [arXiv:1705.08453] [ɪɴSPIRE].

[10] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [hep-th/0603001] [ɪɴSPIRE].

[11] B. Czech, J.L. Karczmarek, F. Nogueira and M. Van Raamsdonk, *The gravity dual of a density matrix*, *Class. Quant. Grav.* **29** (2012) 155009 [arXiv:1204.1330] [ɪɴSPIRE].

[12] M. Headrick, V.E. Hubeny, A. Lawrence and M. Rangamani, *Causality & holographic entanglement entropy*, *JHEP* **12** (2014) 162 [arXiv:1408.6300] [ɪɴSPIRE].

[13] A.C. Wall, *Maximin surfaces and the strong subadditivity of the covariant holographic entanglement entropy*, *Class. Quant. Grav.* **31** (2014) 225007 [arXiv:1211.3494] [ɪɴSPIRE].

[14] D.L. Jafferis, A. Lewkowycz, J. Maldacena and S.J. Suh, *Relative entropy equals bulk relative entropy*, *JHEP* **06** (2016) 004 [arXiv:1512.06431] [ɪɴSPIRE].

[15] P. Hayden and G. Penington, *Learning the alpha-bits of black holes*, *JHEP* **12** (2019) 007 [arXiv:1807.06041] [ɪɴSPIRE].

[16] A. Hamilton, D.N. Kabat, G. Lifschytz and D.A. Lowe, *Local bulk operators in AdS/CFT and the fate of the BTZ singularity*, *AMS/IP Stud. Adv. Math.* **44** (2008) 85 [arXiv:0710.4334] [ɪɴSPIRE].

[17] I.A. Morrison, *Boundary-to-bulk maps for AdS causal wedges and the Reeh-Schlieder property in holography*, *JHEP* **05** (2014) 053 [arXiv:1403.3426] [ɪɴSPIRE].

[18] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, *JHEP* **08** (2013) 090 [arXiv:1304.4926] [ɪɴSPIRE].

[19] T. Faulkner, A. Lewkowycz and J. Maldacena, *Quantum corrections to holographic entanglement entropy*, *JHEP* **11** (2013) 074 [arXiv:1307.2892] [ɪɴSPIRE].

[20] M. Junge, R. Renner, D. Sutter, M.M. Wilde and A. Winter, *Universal recovery maps and approximate sufficiency of quantum relative entropy*, *Ann. Henri Poincaré* **19** (2018) 2955.

[21] M. Ohya and D. Petz, *Quantum entropy and its use*, Springer Science & Business Media, Berlin Germany (2004).

[22] K. Li and A. Winter, *Squashed entanglement*, **k**-*extendibility, quantum markov chains and recovery maps*, *Found. Phys.* **48** (2018) 910 [ɪɴSPIRE].

[23] A.M. Alhambra, S. Wehner, M.M. Wilde and M.P. Woods, *Work and reversibility in quantum thermodynamics*, *Phys. Rev.* **A 97** (2018) 062114 [arXiv:1506.08145].

[24] A.M. Alhambra and M.P. Woods, *Dynamical maps, quantum detailed balance and the Petz recovery map*, *Phys. Rev.* **A 96** (2017) 022118 [arXiv:1609.07496] [ɪɴSPIRE].

[25] M. Lemm and M.M. Wilde, *Information-theoretic limitations on approximate quantum cloning and broadcasting*, *Phys. Rev.* **A 96** (2017) 012304 [arXiv:1608.07569] [ɪɴSPIRE].

[26] H. Barnum and E. Knill, *Reversing quantum dynamics with near-optimal quantum and classical fidelity*, *J. Math. Phys.* **43** (2002) 2097 [quant-ph/0004088].

[27] C.A. Fuchs and J. Van De Graaf, *Cryptographic distinguishability measures for quantum-mechanical states*, *IEEE Trans. Inf. Theory* **45** (1999) 1216 [quant-ph/9712042].

[28] D. Harlow, *The Ryu-Takayanagi Formula from Quantum Error Correction*, *Commun. Math. Phys.* **354** (2017) 865 [arXiv:1607.03901] [ɪɴSPIRE].