

Learning the Alpha-bits of black holes

Patrick Hayden and Geoffrey Penington

*Stanford Institute for Theoretical Physics, Stanford University,
450 Jane Stanford Way, Stanford, CA, U.S.A.*

E-mail: phayden@stanford.edu, geoffp@stanford.edu

ABSTRACT: When the bulk geometry in AdS/CFT contains a black hole, boundary subregions may be sufficient to reconstruct certain bulk operators if and only if the black hole microstate is known, an example of state dependence. Reconstructions exist for any microstate, but no reconstruction works for all microstates. We refine this dichotomy, demonstrating that the same boundary operator can often be used for large subspaces of black hole microstates, corresponding to a constant fraction α of the black hole entropy. In the Schrödinger picture, the boundary subregion encodes the α -bits (a concept from quantum information) of a bulk region containing the black hole and bounded by extremal surfaces. These results have important consequences for the structure of AdS/CFT and for quantum information. Firstly, they imply that the bulk reconstruction is necessarily only approximate and allow us to place non-perturbative lower bounds on the error when doing so. Second, they provide a simple and tractable limit in which the entanglement wedge is state dependent, but in a highly controlled way. Although the state dependence of operators comes from ordinary quantum error correction, there are clear connections to the Papadodimas-Raju proposal for understanding operators behind black hole horizons. In tensor network toy models of AdS/CFT, we see how state dependence arises from the bulk operator being ‘pushed’ through the black hole itself. Finally, we show that black holes provide the first ‘explicit’ examples of capacity-achieving α -bit codes. Unintuitively, Hawking radiation always reveals the α -bits of a black hole as soon as possible. In an appendix, we apply a result from the quantum information literature to prove that entanglement wedge reconstruction can be made exact to all orders in $1/N$.

KEYWORDS: AdS-CFT Correspondence, Gauge-gravity correspondence

ARXIV EPRINT: [1807.06041](https://arxiv.org/abs/1807.06041)

Contents

1	Introduction	1
2	Alpha-bits	6
2.1	Technical definitions	7
3	Alpha-bits from the Hawking radiation	10
4	Alpha-bits in the entanglement wedge	15
4.1	Entanglement wedge reconstruction from approximate decoupling	18
4.2	Entanglement wedges for code spaces containing black holes	22
5	Alpha-bits of BTZ black holes	25
6	Alpha-bits in tensor networks	28
7	The space of black hole microstates	32
8	Discussion	36
8.1	Upper bounds on error correction accuracy	36
8.2	State dependence of the entanglement wedge	38
8.3	State dependence of operators and the Papadodimas-Raju proposal	39
8.4	Explicit reconstruction of state-dependent bulk operators	43
8.5	Black holes are α -bit sup?	45
A	Operator algebra quantum error correction and the information disturbance tradeoff	46

1 Introduction

Recent work [1, 2] has made clear that the semiclassical limit of the AdS/CFT duality is perhaps best understood in the language of quantum error correcting codes.¹ Specifically, the Hilbert space of certain conformal field theories contains subspaces of states (we shall refer to these as the bulk Hilbert spaces or code spaces), which have a dual interpretation as a quantum field theory on a semiclassical gravitational background that is asymptotically anti-de Sitter space. The error correcting codes are the isometries from the bulk subspaces to the larger boundary space.

¹This is just one of a series of insights that have been reached about quantum gravity through applying the tools of quantum information [3–7].

It is generally understood that the error correction should only become exact in the limit of vanishing Newton’s constant $G_N \rightarrow 0$ or, equivalently, diverging gauge group rank $N \rightarrow \infty$. Nonetheless, out of convenience, many of the toy models and mathematical results [2, 8, 9] that have been developed for understanding the AdS/CFT error-correction paradigm involve finite dimensional Hilbert spaces with exact quantum error correction. In this paper, we show that such a framework is insufficient to capture some of the important aspects of error correction in AdS/CFT, which are only possible for finite-dimensional code spaces if the error correction is merely approximate, rather than exact.

The magnitude ε of these uncorrectable errors is very small — in fact they are non-perturbatively suppressed in the semiclassical limit (their Taylor expansion in G_N or $1/N$ has no non-zero terms). However, we show that the existence of these tiny, seemingly insignificant approximations makes possible key features of the AdS/CFT correspondence that provably could not otherwise exist.

Such features arise when the code space $\mathcal{H}_{\text{code}}$ of ‘nice’ states with smooth bulk geometries that we wish to be able to error correct includes states containing black holes. Since the Bekenstein-Hawking entropy becomes infinite in the classical limit $G_N \rightarrow 0$, the maximum dimension of a code space consisting of a large number of black hole microstates diverges, if the horizon area is held fixed in AdS units. As a result, there is no limit in which the error $\varepsilon \rightarrow 0$, the code space contains ‘all’ the black hole microstates (or even contains a constant fraction of the black hole entropy), and the dimension of the code space remains bounded. This should immediately make us cautious about believing that results based on the twin assumptions of exact error correction and finite-dimensional Hilbert spaces will continue to be valid in this context.

Suppose we consider a black hole in AdS/CFT, together with a boundary region A that consists of slightly over half of the entire boundary. The size of the entanglement wedge of region A (the region of the bulk that can be reconstructed from boundary region A) depends on whether the black hole is in a specific known microstate or if the black hole is an unknown state (which we can model as the thermal ensemble). For the known microstate, the black hole is contained in the entanglement wedge of A ; for the thermal ensemble, it lies outside. Because the two entanglement wedges are not the same, there exist bulk operators that can be reconstructed in boundary region A if and, more importantly, *only* if the state of the black hole is known. Such operators exist for every black hole microstate, but there exists no single boundary operator that works for all microstates. In other words, the boundary operator is state-dependent.

We emphasize that this notion of state dependent reconstructions should be distinguished from the idea that interior reconstructions are state dependent, as advocated for most prominently by Papadodimas and Raju [10, 11] (although there appear to be close connections between the two, see section 8.3). In the Papadodimas-Raju constructions, even a global reconstruction that is allowed to act on the entire boundary is necessarily state dependent. In contrast, in our case, the state dependence is only necessary when the reconstruction is restricted to acting on a subregion of the boundary. All our results can therefore be completely understood in the framework of quantum error correction; there is no conflict with the linearity of quantum mechanics.




Figure 1. A black hole with horizon area \mathcal{A}_0 in AdS-space. The boundary is separated into two regions, A and \bar{A} with shared boundary ∂A . There are two important bulk minimal surfaces with boundary ∂A . The minimal surface homologous to \bar{A} has area \mathcal{A}_1 , while the minimal surface homologous to A has area \mathcal{A}_2 .

More specifically, the existence of state-dependent operators for all states (but the absence of *state-independent* operators) corresponds in the Schrödinger picture to a weakened version of the usual notion of a quantum error correcting code, called universal subspace quantum error correction. This was studied in detail in [12]. In terms of the language introduced in [12], we say that region A encodes the *zero-bits* of the bulk region.

As the size of the boundary region A increases, less state dependence is required. A single boundary operator can now reconstruct bulk operators for code subspaces containing many different black hole microstates, although the reconstruction will still necessarily depend on the code subspace chosen.² As shown in figure 1, in general there will be two minimal surfaces with boundary ∂A that might form the RT surface for states in this code space. One will be homologous to \bar{A} (and hence contain the black hole in the entanglement wedge of A), while the other will be homologous to A . We say that they have areas \mathcal{A}_1 and \mathcal{A}_2 respectively; if region A is greater than half the boundary, then $\mathcal{A}_2 > \mathcal{A}_1$.

We show that there exists a single boundary reconstruction in region A of a given bulk operator that works for every state in a subspace, so long as the bulk operator lies within the entanglement wedge of A for every state in the subspace, including states that are entangled with a reference system. Equivalently, the bulk operator must lie inside the entanglement wedge of A for all states *including mixed states* with support only in the subspace.

If there were no black hole, this subtle distinction of requiring even mixed states to contain the bulk operator in their entanglement wedge would be unimportant. The Ryu-Takayanagi surface (or more precisely the quantum extremal surface [13]), which bounds

²We shall continue to use the term “state dependence” even when a single reconstruction can simultaneously work for many, but not all, states, because the alternative expression “code subspace dependence”, while more precise, is something of a mouthful.

the entanglement wedge, is defined as the surface that minimises the sum of $\mathcal{A}/4G_N$ (where \mathcal{A} is the area of the surface) and the bulk entropy S_{bulk} . If the bulk entropy

$$S_{\text{bulk}} = O(1),$$

this will always be the surface with area \mathcal{A}_1 , at least in the limit $G_N \rightarrow 0$. However, if we consider a subspace of black hole microstates of sufficiently large dimension d such that

$$\log d + \frac{\mathcal{A}_1}{4G_N} = \max(S_{\text{bulk}}) + \frac{\mathcal{A}_1}{4G_N} > \frac{\mathcal{A}_2}{4G_N}, \quad (1.1)$$

then the RT surface can jump to the surface with area \mathcal{A}_2 for states sufficiently entangled with the reference system (or sufficiently mixed). For such states, the entanglement wedge of A will no longer contain either the black hole or the bulk region between the two minimal surfaces but outside the black hole. It is not possible to find a single boundary operator reconstruction for operators in this region that will work for that entire subspace. However, if the dimension d is not sufficiently large for this to happen, then the entanglement wedge will always lie on the original surface. Hence, operators between the minimal surfaces (as well as operators acting on the black hole itself) can be reconstructed from the boundary region A .

In other words, any bulk operator lying between the two minimal surfaces can be reconstructed as a single boundary operator that works for *any* subspace of black hole microstates of dimension $d < e^{\alpha S}$ where $S = \mathcal{A}_0/4G_N$ is the Bekenstein-Hawking entropy. The dimensionless parameter

$$\alpha = \frac{\mathcal{A}_2 - \mathcal{A}_1}{\mathcal{A}_0}$$

is independent of G_N , so remains fixed in the semiclassical limit. (In accordance with our previous claim, it should be clear that the parameter α increases as the size of region A increases, until eventually $\alpha = 1$ and a single operator can work for all black hole microstates.) As above, this is an example of universal subspace quantum error correction. The difference now is that the dimension of the subspace which can be error-corrected is allowed to grow with the dimension of the larger space of all black hole microstates. If we again make use of terminology from [12], the region A now encodes the α -bits of the bulk region.

Nontrivial realisations of universal subspace quantum error correction are only possible when the error correction is approximate. In the exact setting, being able to correct all small subspaces automatically implies being able to correct arbitrary subspaces. Even if the error correction is approximate, the error in correcting subspaces of larger dimension can be bounded in terms of the error for subspaces with smaller dimension. However, the quality of the approximation degrades as the dimension of the subspaces increases. This makes it possible to find a limit in which the errors tend to zero for small subspaces, but stay order one for large subspaces, so long as in this limit the dimension of the full code space tends to infinity. We will see that the classical limit $G_N \rightarrow 0$ ($N \rightarrow \infty$) in AdS/CFT is an example of precisely this kind of limit. The seemingly insignificant, non-perturbatively

small errors make possible order one effects that continue to exist, and in fact become more sharply defined, even in the semiclassical limit.

In section 2, we review the basic construction of α -bits and universal subspace error correction. Section 3 then shows that the evaporation of black holes into Hawking radiation provides a natural example of a capacity-achieving α -bit code. In contrast to our usual intuition, black holes rush to reveal their α -bits in the Hawking radiation as quickly as they possibly can; in this sense the Hawking radiation contains as much (rather than as little) information as possible about the state of the black hole. Earlier work by Hayden and Preskill in [14] on information retrieval from evaporating black holes can be interpreted as the $\alpha = 0$ case of this more general fact.

In section 4, we develop the main result of the paper: the appearance in AdS/CFT of α -bit encodings for code spaces containing black holes. We develop the ideas sketched out above in significantly greater depth and precision. Section 5 includes more specific calculations for the simple case of a uncharged, non-rotating BTZ black hole in AdS₃. They show that the region between the minimal surfaces but outside the black hole horizon is always approximately AdS scale, regardless of the size of the black hole — at least in this simple case, it requires a large central charge CFT with a weakly curved gravity dual to have locality at small scales compared to the size of this ‘ α -bit’ region.

Section 6 explores how α -bit codes can arise in a basic tensor network toy model of AdS/CFT. In this context, the intuition behind the state dependence of the boundary operators can be made very clear; bulk operators have to be pushed through the black hole itself in order to reach the boundary, in a way that manifestly depends on the subspace of black hole microstates being considered. Meanwhile, section 7 provides more detailed, technical justifications that back up our assumptions about the existence of a code space of black hole microstates with the correct entropy.

Section 8 consists of an extended discussion on various aspects and implications of the paper. This discussion makes use of many of the results derived in the main sections of the paper, but can be read relatively independently. We show how α -bit codes can be used to put lower bounds on the uncorrectable error ε . Specifically, even though it is possible to make the error equal to zero to all orders in perturbation theory, we show that there must sometimes exist errors

$$\varepsilon > e^{-\eta/G_N}$$

for any $\eta > 0$. We argue that α -bit codes provide the most controlled setting in which we can understand a state-dependent entanglement wedge, before discussing tantalising connections and similarities between the α -bit codes we discuss and the Papadodimas-Raju proposal [10, 11] for the state dependence of operators behind the black hole horizon. We then briefly discuss explicit recovery maps for the α -bit codes and, lastly, observe that black holes provide examples of ‘explicit’ (as opposed to randomly generated) capacity-achieving α -bit codes for noiseless quantum channels, which until now had not been known.

Finally appendix A describes a result from quantum information that has so far not appeared in the quantum gravity literature [15]. It provides a generalisation of the Dong-Harlow-Wall condition [9] to approximate reconstruction and is sufficient to prove that




Figure 2. Alice has a quantum state $|\psi\rangle \in \mathcal{H}_{\tilde{A}}$ consisting of n qubits, for some large n . She adds a few qubits in a fixed state $|0\rangle$ (embedding $\mathcal{H}_{\tilde{A}}$ as a subspace of a slightly larger Hilbert space \mathcal{H}_A), applies a Haar random unitary U and then sends slightly more than half of the qubits to Bob, throwing the rest away. We say that Alice has sent the zerobits of the state $|\psi\rangle \in \mathcal{H}_{\tilde{A}}$ to Bob.

entanglement wedge reconstruction can be made exact to all orders in $1/N$. A special case of this result is used in section 4, but the theorem and some background are included in full because of the relevance to the wider literature on quantum error correction and AdS/CFT.

2 Alpha-bits

We begin with a basic review of the concept of universal subspace error correction and α -bits. For more detail see [12]. Because the definitions involved are quite technical, we will begin by illustrating the basic phenomenon we are trying to capture with a relatively simple example.

Suppose we apply a Haar-random unitary U to some large number n of qubits then throw away some fraction that is slightly less than half, as shown in figure 2. Call the input Hilbert space \mathcal{H}_A , the qubits that are kept \mathcal{H}_B and the qubits that are discarded \mathcal{H}_E . Now consider the fate of a typical pair of orthogonal pure states on \mathcal{H}_A in the limit of large n . Both will get mapped to states almost maximally entangled between \mathcal{H}_B and \mathcal{H}_E . Moreover, because \mathcal{H}_E is much smaller than \mathcal{H}_B , the reduced states on \mathcal{H}_E will be nearly maximally mixed and therefore effectively indistinguishable. For the same reason, the states on \mathcal{H}_B will have small rank relative to the dimension of \mathcal{H}_B , which leads to their being nearly orthogonal.

In fact, due to strong measure concentration effects in high dimension, those properties will hold not just for one pair of orthogonal states on \mathcal{H}_A , or two pairs, or even a countable number of pairs. It will hold for *all* pairs of orthogonal states in a subspace $\mathcal{H}_{\tilde{A}} \subseteq \mathcal{H}_A$ that is almost as large as \mathcal{H}_A in qubit terms: $n - o(n)$ qubits. More generally, the map from $\mathcal{H}_{\tilde{A}}$ to \mathcal{H}_B approximately preserves the pairwise distinguishability of states as measured by the trace distance despite shrinking the number of qubits by a factor of two [16, 17]. Because the dimension of \mathcal{H}_B is roughly the square root of the dimension of $\mathcal{H}_{\tilde{A}}$, this seems paradoxical. The resolution is that the map encodes the geometry of the unit sphere of the Hilbert space $\mathcal{H}_{\tilde{A}}$ into the space of *density matrices* on \mathcal{H}_B ; the pure state geometry is partially encoded into noise. “Sending the zero-bits of $\mathcal{H}_{\tilde{A}}$ ” can be roughly defined as

approximately preserving the geometry of the unit sphere of pure states. We shall provide a precise technical definition in section 2.1.

Returning to the example, while the full space $\mathcal{H}_{\bar{A}}$ has been transmitted in some sense, it is clearly not possible to perform approximate quantum error correction and completely reverse the effect of the channel; doing so would lead to the quantum capacity of a qubit being greater than one, which by recursion would make it infinite. Geometry preservation does have an operational consequence, however. If we restrict the states to *any* two-dimensional subspace of $\mathcal{H}_{\bar{A}}$, there is a decoding operation that will perform quantum error correction. The only catch is that the decoding operation will in general depend on the two-dimensional subspace that we wish to decode. Note, however, that the encoding $\mathcal{H}_{\bar{A}} \hookrightarrow \mathcal{H}_A$ and the channel do not. If we think of Alice sending Bob a state, then Bob has to know which two-dimensional subspace the state is in, but Alice does not.

What if Bob wishes to be able to decode larger subspaces of $\mathcal{H}_{\bar{A}}$? What fraction of the qubits need to be kept now? As one would expect, to decode the entire space $\mathcal{H}_{\bar{A}}$ they have to keep almost all of them. On the other hand, so long as they keep a fraction greater than $\frac{1+\alpha}{2}$ of the qubits, Bob can decode any subspace of up to αn qubits (or in other words any subspace of dimension at most $2^{\alpha n}$).³ We call the task of decoding any such subspace *universal approximate subspace error correction* and say that \mathcal{H}_B encodes the α -bits of $\mathcal{H}_{\bar{A}}$.⁴ A zero-bit is then simply the special case of an α -bit with $\alpha = 0$.⁵

2.1 Technical definitions

We now turn to a more formal definition. Readers satisfied by the level of rigour given above should feel free to proceed to the next section. An exact quantum error correcting code consists at its simplest of an encoding and transmission channel $\mathcal{N} : S(\mathcal{H}_A) \rightarrow S(\mathcal{H}_B)$ (where $S(\mathcal{H}_A)$ is the space of density matrices on \mathcal{H}_A) together with a decoding or recovery channel $\mathcal{D} : S(\mathcal{H}_B) \rightarrow S(\mathcal{H}_A)$ such that

$$\mathcal{D} \circ \mathcal{N} = \text{Id}_A.$$

In other words, given any state $\rho \in S(\mathcal{H}_A)$, it is possible to exactly recover the state ρ from the state $\mathcal{N}(\rho)$.

Suppose, as above, that we allow the receiver Bob to have some additional information about the state ρ ; he again knows that the state lies within some particular subspace of A . Obviously this can make the task of finding a recovery channel $\tilde{\mathcal{D}}$ considerably easier. Indeed, the normal approach used to make error correcting codes out of a noisy transmission

³Technically, the construction in [12] requires the use of shared randomness to achieve this rate but this can be eliminated by block coding.

⁴Note that the *number* of α -bits sent is determined by the dimension of $\mathcal{H}_{\bar{A}}$ rather than the dimension of the subspaces Bob wishes to decode. This is because the whole space $\mathcal{H}_{\bar{A}}$ is available to him; he just needs to make a choice about which subspace he is interested in. Furthermore, in the case of zero-bits, the subspace dimension is always $d = 2$, even though the amount of information sent clearly depends on the size of $\mathcal{H}_{\bar{A}}$, and so this is the only sensible definition.

⁵Since decoding a one-dimensional state is trivial, we need the subspaces to have dimension $2^{\alpha n} + 1$ to correctly recover the definition of zerobits when $\alpha = 0$. This change has negligible effect on the definition of α -bits for $\alpha > 0$; at most it can slightly increase the size of the error in recovering the state.

channel is to first apply an encoding channel that consists of an isometry from a (smaller) code space into the (larger) input space of the transmission channel, in such a way that the code space is possible to decode, even though the large input space is not.

However, in this case we don't just want Bob to be able to decode some particular subspace that has especially nice properties. We want him to be able to decode the state provided he knows *any* sufficiently small subspace that it is contained in. In the framework of exact quantum error correction, this would mean the existence of an exact decoding channel \mathcal{D}_S for any sufficiently small subspace \mathcal{H}_S .⁶ However, it turns out that, even if we only require any two-dimensional subspace to be error-correctable, the existence of a decoding channel for every subspace implies that the complete space \mathcal{H}_A can also be error-corrected. In other words there is no advantage to being able to use a different decoding channel for each subspace, if you still have to be able to decode any possible subspace exactly.

How can this be reconciled with our analysis above? The answer, of course, lies in our assumption that the error correction had to be exact. If we instead only require approximate error correction to be possible, the situation becomes completely different. In this case we only require that the decoding channel get back something very close to (rather than exactly) the original state. In other words we only require that,

$$\|\mathcal{D}_S \circ \mathcal{N} - \text{Id}\|_{\diamond} \leq \varepsilon$$

for some $\varepsilon \ll 1$.⁷

The Stinespring dilation theorem says that for any channel $\mathcal{N} : S(\mathcal{H}_A) \rightarrow S(\mathcal{H}_B)$ there exists an isometry $V : \mathcal{H}_A \rightarrow \mathcal{H}_B \otimes \mathcal{H}_E$ that is unique up to isomorphisms of \mathcal{H}_E such that for all states $\rho \in S(\mathcal{H}_A)$,

$$\mathcal{N}(\rho) = \text{Tr}_E V \rho V^\dagger.$$

We can then define the complementary channel $\mathcal{N}^c : S(\mathcal{H}_A) \rightarrow S(\mathcal{H}_E)$ by

$$\mathcal{N}^c(\rho) = \text{Tr}_B V \rho V^\dagger.$$

The subspace decoupling duality, proved in [12], states that, if there exists a decoding channel \mathcal{D}_S , for any subspace of dimension less than or equal to k , with error at most ε as above, then there exists a state $\sigma \in S(\mathcal{H}_E)$ such that for *all* states $\rho \in S(\mathcal{H}_R \otimes \mathcal{H}_A)$

$$\|(\mathcal{N}^c \otimes \text{Id}_R)\rho_{AR} - \sigma_E \otimes \rho_R\|_1 \leq \delta \tag{2.1}$$

where \mathcal{H}_R is a reference system whose dimension is also equal to k and $\delta \leq 8\sqrt{\varepsilon}$.⁸ In other words, the environment \mathcal{H}_E encodes almost no information about the state ρ . We

⁶We emphasize again that the encoding is not allowed to depend on \mathcal{H}_S . Otherwise the encoding channel can simply map \mathcal{H}_S to any fixed subspace, and the task becomes identical to ordinary error correction for the smaller space \mathcal{H}_S .

⁷We have bounded the error in terms of the diamond norm (see appendix A) here, but we could have equally used the operator norm since these bound each other in a dimension-independent way in the neighbourhood of the identity [18].

⁸The subspace decoupling duality can be derived almost immediately by applying Kretschmann et al.'s information-disturbance theorem [19] to arbitrary subspaces of dimension k .

say that the complementary channel \mathcal{N}^c is approximately k -forgetful. Conversely if the complementary channel \mathcal{N}^c is approximately k -forgetful with uncertainty δ , then there exists a decoding channel for any subspace of dimension at most k with error $\varepsilon \leq 2\sqrt{2\delta}$.

How important is the inclusion of a reference system with dimension k in (2.1)? By writing out Schmidt decompositions and using triangle inequality (see for example lemma 23 of [17]), one can easily show that, for any k ,

$$\max_{\rho \in \mathcal{S}(\mathcal{H}_R \otimes \mathcal{H}_A)} \|(\mathcal{N}^c \otimes \text{Id}_R)\rho_{AR} - \sigma_E \otimes \rho_R\|_1 \leq k \max_{\rho \in \mathcal{S}(\mathcal{H}_A)} \|\mathcal{N}^c(\rho) - \sigma_E\|_1. \quad (2.2)$$

In other words, δ grows at most linearly with the reference dimension k . Hence, as claimed above, exact universal subspace error correction is indeed equivalent to exact quantum error correction; if $\varepsilon = 0$ for $k = 2$, then we also have $\delta = 0$ for $k = 2$, and hence $\delta = \varepsilon = 0$ for all k .

However there are examples of quantum channels that saturate the bound (2.2) [17]. This means that, when the dimension of the complete space \mathcal{H}_A is very large, the forgetfulness δ and reconstruction error ε can end up being much larger for large subspace dimensions k than for small subspace dimensions. In other words, it may be possible to reconstruct any sufficiently small subspace with very high fidelity, while still being completely impossible to reconstruct the entire space.

To make precise statements without reference to epsilons and deltas, it is generally necessary to consider a limit where the dimension d of the code space tends to infinity, for example the classical limit of the space of black hole microstates. In general, the question of whether the error $\varepsilon \rightarrow 0$ will then depend on how the subspace dimension k scales with the dimension d of the complete space.

To take a trivial example, universal subspace error correction for subspaces of dimension at most

$$k = \frac{d}{t}$$

for some fixed t is equivalent to conventional error correction (i.e. either both have errors that tend to zero in some limit or neither does). However, if the subspace dimension grows sublinearly with d , universal subspace error correction is inequivalent to ordinary error correction. The most natural possibility to consider is that the dimension of the subspaces grows proportionally to d^α for some $0 \leq \alpha \leq 1$. If universal subspace error correction is possible with vanishing error for such a dimension we say that Bob has the α -bits of the state sent by Alice.

The case $\alpha = 0$, which we call *zero-bits*, corresponds to the ability to do universal subspace error correction for any constant dimension k that is independent of d . Just as we saw above when d/k was held fixed, the exact value of k does not matter (we generally take $k = 2$ for convenience) since, if $\varepsilon \rightarrow 0$ for any $k \geq 2$, it will also tend to zero for all fixed k , even though, for finite errors, the size of k will affect the size of the error ε . Formally, we define the α -bit decoding error based on subspaces of dimension $k = 2^{\alpha n} + 1$, since this formula gives $k = 2$ for zero-bits ($k = 1$ would be trivial) and scales as d^α for $\alpha > 0$. However, as we have discussed, the exact dimension of the subspace is essentially unimportant; we only care about how this dimension scales with n .

3 Alpha-bits from the Hawking radiation

In this section we argue based on simple qubit toy models that an evaporating black hole reveals its α -bits through its Hawking radiation as quickly as possible, saturating the α -bit capacity of a noiseless quantum channel. This is in sharp contrast with the usual notion that the Hawking radiation tries to hide information about the black hole state for as long as possible, but we shall show that the two ideas are not merely reconcilable but in fact equivalent. We generalise the arguments made by Hayden and Preskill in [14], which can be interpreted as describing the special case where $\alpha = 0$. This section is relatively self-contained and is not necessary to understand the main claims of the paper, which are developed in section 4; however, it is both of interest in its own right and features strong similarities to the way the α -bits of black holes are encoded in AdS/CFT — suggesting that the lessons from AdS/CFT may well be important in a significantly broader context.

It is often incorrectly implied that the Hawking radiation contains no information until the Page time, after which it begins to reveal the qubits of the black hole one by one. In fact there is good reason to think that (at least in simple toy models) not a single qubit of the black hole will be revealed, to an observer knowing nothing about the original black hole state, until the black hole has almost entirely evaporated.⁹ Instead, after the Page time, the α -bits of the entire black hole will be revealed for increasing values of α , until eventually all the qubits are revealed, essentially simultaneously, at the very end of the evaporation process. No particular subsystem is revealed before any other subsystem; however, increasingly large subspaces of the entire system become decodable.

Consider a large semiclassical black hole A in a pure microstate that is *already known* by some observer Bob. Alice wants to hide her diary D , a small quantum state, from Bob by dropping it into this black hole. After she has done so, Bob knows that the black hole is in some particular small-dimensional subspace of the large Hilbert space of black hole microstates — specifically the subspace of states that could have been created by the diary falling in. In the semiclassical limit, the dimension of the space of black hole microstates tends to infinity, while the dimension of the small subspace remains fixed.

The black hole is then allowed to evaporate into Hawking radiation. We assume Bob has a perfect understanding of the microscopic dynamics of the black hole and the ability to collect all the Hawking radiation that is emitted by it, as well as infinite computational power. However, even with these awesome powers, he has no ability to measure the internal black hole degrees of freedom themselves. How long does Bob have to wait in order to determine the original state of the diary with a high degree of confidence?

This question was studied in detail in [14]. Since the dynamics of the black hole interior are expected to be highly chaotic, Bob cannot hope for the small subspace of possible black hole states that could have been created by the diary to be especially easy to decode from the Hawking radiation. The problem is essentially equivalent to the question of whether

⁹More precisely, there will no tensor product factorisation of the black hole $\mathcal{H}_{\text{BH}} \cong \mathcal{H}_{\text{qubit}} \otimes \mathcal{H}_{\text{rest}}$ such that the reduced state on $\mathcal{H}_{\text{qubit}}$ can be determined from the Hawking radiation before the black hole has almost evaporated, even if the dimension of $\mathcal{H}_{\text{qubit}}$ is only two.

Bob is able to decode *any* arbitrary small subspace of black hole states. In our language, Bob needs to have access to the *zero-bits* of the black hole.

To conclusively answer the question of how information is encoded in the Hawking radiation (we shall assume throughout this paper that the evaporation is unitary), we would have to understand the exact details of the dynamics of the evaporation of black holes; these details are, of course, as yet unknown. However, considerable progress can be made using some fairly basic assumptions and arguments.

One toy model of the evaporation of a black hole is to add a few ancilla qubits (in order to make the process slightly thermodynamically irreversible) and then to apply a random unitary. However, this is exactly the model that we claimed in section 2 had its zero-bits encoded in slight more than half of the output qubits. This suggests that the zero-bits of the black hole should be encoded in any fraction of the Hawking radiation greater than one half. Indeed this is essentially the model and conclusion reached in [14].

A slightly more sophisticated model of a black hole would be to use an element of a unitary 2-design rather than a fully Haar random unitary. (All the elements of a 2-design can be chosen to be much less computationally complex than a generic Haar random unitary and hence could reasonably be applied within approximately the scrambling time.) Conveniently, so long as we model the black hole as applying a randomly sampled element of the 2-design with the choice of element known by Bob (and not simply as applying a generic element of the 2-design), we obtain exactly the model that was shown in [12] to saturate the α -bit capacity for general values of α .

Despite their popularity in the literature as toy models of black holes, such random unitary models all suffer from a significant flaw as models of real-life black hole evaporation. Specifically, rather than being only slightly thermodynamically irreversible, black hole evaporation is in general *highly* thermodynamically irreversible, with numerical calculations suggesting that the thermodynamic entropy increases by a factor of approximately 1.48 over the course of the evaporation process [20]. This has a number of important qualitative effects: for example, it means that the Page time, when the entropy of the radiation equals that of the black hole, occurs significantly before the black hole has lost half its entropy.

However we can in principle prevent this $O(S)$ thermodynamic entropy increase. For example, we can extract only a small amount of energy and entropy from the Hawking radiation, slightly reducing its temperature, and reflect the rest back into the black hole. Notably, this can be achieved very easily and naturally in AdS/CFT by simply add a weak local coupling to the boundary theory. Alternatively, all but the highest energy Hawking modes (with frequency well above the Hawking temperature) may be reflected back into the black hole by a potential barrier; this happens, for example, in near-extremal Reissner-Nordström black holes, for example. In the interests of simplicity, we shall therefore assume throughout this section that the black hole evaporates by some close-to-thermodynamically-reversible process (unless stated otherwise).

An alternative argument to the simplified toy models discussed above, but which reaches the same conclusion goes as follows. Rather than make any assumptions about the dynamics of the black hole itself, we can simply assume that the semiclassical result

of thermal Hawking radiation is correct — whenever this assumption is consistent with unitarity. Assuming that our evaporation process is close to being thermodynamically reversible, this implies that the semiclassical calculation should be accurate (and the Hawking radiation should look thermal) so long as we look at less than half of the Hawking radiation. However, the subspace decoupling duality discussed in section 2 means that this is equivalent to the zero-bits of the black hole being encoded in any fraction of the Hawking radiation *greater* than half.

A natural generalisation of the problem of reconstructing a small diary thrown into a known black hole is to replace the diary by a second smaller black hole D . Now the dimension of the diary Hilbert space is no longer small and fixed; instead it is exponential in $1/G_N$. Unlike the black hole that it is thrown into, the state of the black hole D is unknown to Bob. Let the horizon area of the black hole D be $\alpha\mathcal{A}$ where \mathcal{A} is the horizon area of the final *combined* black hole.

The subspace of possible black hole states that can be created upon throwing in the diary is no longer small, but instead grows in the semiclassical limit as $e^{\alpha S}$ where $S = \mathcal{A}/4G_N$. To determine the original state of the diary, Bob now needs access to the α -bits of the larger combined black hole.

Using the slightly more sophisticated version of the random unitary model, as well as the α -bit capacity results from [12], we see that to recover the α -bits of the black hole, Bob needs to obtain at least an

$$\frac{\alpha + 1}{2}$$

fraction of the Hawking radiation [12].

What about if we again try to argue from the principle that the Hawking radiation should look thermal whenever this is consistent with information preservation? To make use of the subspace decoupling duality, we now need to allow the black hole states to be entangled with a reference system \mathcal{H}_R of dimension $e^{\alpha S}$. We want to know whether the reference system \mathcal{H}_R , together with the part of the Hawking radiation \mathcal{H}_E which is thrown away, contains any information about the state of the black hole. We assume that the reduced density matrix of the state on $\mathcal{H}_E \otimes \mathcal{H}_R$ will look like the thermal ensemble on \mathcal{H}_E tensored with the reduced density matrix of the original state on \mathcal{H}_R , so long as such a state can be purified by the remaining fraction p of the Hawking radiation \mathcal{H}_B which is collected by Bob. This is possible if and only if the dimension of \mathcal{H}_B is larger than the dimension of $\mathcal{H}_E \otimes \mathcal{H}_R$. In other words if

$$\alpha + (1 - p) \leq p.$$

Hence a fraction $(1 - p)$ of the Hawking radiation will be $e^{\alpha S}$ -forgetful so long as $\alpha \leq 2p - 1$.

However, by the subspace decoupling duality, if some fraction $1 - p$ of the Hawking radiation is $2^{\alpha S}$ -forgetful, the remaining fraction p of the Hawking radiation must encode the α -bits of the black hole. Rearranging, we again find that the α -bits of the black hole are encoded in any fraction $p > (1 + \alpha)/2$ of the Hawking radiation. As for the zero-bit case, the assumption of thermality, whenever it is consistent with unitarity, gives an answer that agrees with the random unitary model.

The black hole evaporation therefore compresses the α -bits of the *entire* black hole (consisting of, say, n qubits) into only $(1 + \alpha)n/2$ physical qubits of Hawking radiation. The Hawking radiation forms an α -bit code for the entire black hole that encodes

$$\frac{2}{1 + \alpha} \text{ logical } \alpha\text{-bits per physical qubit.}$$

This is exactly the α -bit capacity of the noiseless qubit channel; black holes give up their α -bits as fast as they possibly can. On the one hand, this is unsurprising since random unitary channels were exactly the strategy used in [12] to originally achieve the α -bit capacity. It is nonetheless in sharp contrast to usual idea of Hawking radiation containing as little information as possible.

However, as we have seen, these two phenomena are not merely reconcilable; they are actually equivalent. The subspace decoupling duality means that if the black hole releases as little information as possible in less than half of its Hawking radiation, then it necessarily also releases as much information as possible in more than half of the Hawking radiation (at least in the specific sense of encoding the α -bits for as large a value of α as possible).

Just like a small diary, if a black hole diary (as before with horizon area $\alpha\mathcal{A}$) is thrown into a black hole that has already partially evaporated, the information within it will be revealed more quickly. (Note that $\mathcal{A} = 4G_N S$ is now the horizon area of the combined black hole after it has both been allowed to partially evaporate and then had the black hole diary thrown in.) We can see this for the random unitary model by making use of results used to prove the achievability of the entanglement-assisted α -bit capacity in [12]. These show that if, when the black hole diary is thrown in, the black hole is already approximately maximally entangled with Hawking radiation of entropy βS , Bob will be able to determine the state of the black hole diary so long as the fraction p of the remaining Hawking radiation that he obtains satisfies

$$\beta S + pS \geq (1 - p)S + \alpha S. \tag{3.1}$$

This inequality can also be derived from the thermality (whenever consistent with unitarity) assumption. The left hand side is the combined entropy of the original Hawking radiation βS together with the newly emitted Hawking radiation pS (i.e. the systems that Bob has access to). The right hand side consists of the entropy of the Hawking radiation that is thrown away $(1 - p)S$ plus the maximum entropy αS of the reference system \mathcal{H}_R that we need to consider according to (2.1) (i.e. the systems that Bob does not have access to). For these systems to look thermal while being purified by those Bob has access to requires (3.1).

Bob will therefore recover the state of the diary once he has access to at least a

$$p = \frac{1 + \alpha - \beta}{2} \tag{3.2}$$

fraction of the remaining Hawking radiation. In the special case $\beta = 1$ and $\alpha = 0$, the fraction required tends to zero. This case was studied in detail in [14]; it is probable that Bob only needs to wait for at least the scrambling time ($O(\log S)$) before he can recover the state of the diary.

It may seem from (3.2) that entangled black holes exceed the entanglement-assisted α -bit capacity of $2/(1 + \alpha)$ for a noiseless channel that was derived in [12]. However, a precise comparison, done in section 8.5, shows that, as with the unentangled case, it merely saturates the capacity.

At the start of this section, we claimed that, unlike for the α -bits of the black holes, we should not expect that even a single qubit of the black hole Hilbert space, no matter what basis we work in, is revealed until almost the end of the Hawking evaporation process. If a qubit was revealed before this point, then there would necessarily exist a pair of states in the black hole Hilbert space that are both maximally entangled with a reference Hilbert space of only one fewer qubit than the black hole Hilbert space, but for which the Hawking radiation produces (almost) orthogonal states after only some fraction $p < 1$ of the Hawking radiation has been produced. Yet for such states, the entropy $(S - \log 2)$ of the reference system, plus the remaining entropy $(1 - p)S$ of the black hole is far larger than the entropy pS of the Hawking radiation. It follows that the Hawking radiation should look thermal for randomly chosen states of this form, and hence (again using the fact that our black hole evaporation is slightly thermodynamically non-reversible) measure concentration will be sufficient to ensure that all such pairs of states should have almost indistinguishable Hawking radiation, contradicting our original assumption.

Finally, it is worth commenting briefly on which of our conclusions are likely to change if a black hole is allowed to evaporate *irreversibly*, and which should continue to be valid. In particular, we would expect that the black hole evaporation will no longer saturate the noiseless α -bit capacity. For example, the Page time, when the zero-bits of the black hole should be revealed, will now occur when the entropy of the Hawking radiation is significantly more than half of the initial black hole entropy (for realistic black holes the figure is approximately 60% [20]).

However, many of the results above should still apply with small adaptations. In particular, it should continue to be the case that the α -bits of the black hole are revealed when

$$\alpha S_0 < S_{\text{rad}} - S_{\text{BH}} \tag{3.3}$$

where S_0 was the initial Bekenstein-Hawking entropy of the black hole, S_{BH} is the Bekenstein-Hawking entropy of the partially evaporated black hole and S_{rad} is the thermodynamic entropy of the Hawking radiation.

To see this, note that the natural generalisation of the random unitary models used above to thermodynamically irreversible evaporation is a sequence of nested random isometries. At each step a few qubits are released as Hawking radiation and then a smaller number of qubits are added back to the black hole using a random isometry. Hence, over time, the number of qubits describing the black hole (the Bekenstein-Hawking entropy) decreases, but the total number of qubits describing the black hole together with the Hawking radiation (the total thermodynamic entropy) increases. This is shown in figure 3 and is an example of a random tensor network; such networks obey a version of the Ryu-Takayanagi formula [21], where the entropy of a subsystem is proportional to the number of legs in the network that need to separate the subsystem from its complement.




Figure 3. A simple two model of an evaporating black hole that incorporates thermodynamic irreversibility. At each time step, two qubits are released from the black hole as Hawking radiation and then a random isometry is applied to the black hole that adds a single qubit. This model is an example of a random tensor network and hence obeys a version of the Ryu-Takayanagi formula.

An initial state that is maximally entangled with a reference system with $\log d_R = \alpha S_{\text{BH}}$ will therefore have maximal entropy (and hence be maximally mixed) on the black hole Hilbert space plus reference system, so long as (3.3) holds. By the subspace decoupling duality, the α -bits of the initial black hole can therefore be reconstructed from the Hawking radiation at the same point in time.

The other results in this section can be similarly extended to irreversible evaporation. For example, all the qubits of the initial black hole state are still revealed simultaneously. However this revelation will now happen somewhat before the black hole has completely evaporated, when

$$S_{\text{rad}} - S_{\text{BH}} > S_0; \tag{3.4}$$

the last part of the Hawking radiation can be completely error corrected.

In general, the inherent inefficiency of an irreversible process prevents the black hole evaporation from saturating noiseless α -bit capacities. However, the random nature of black hole dynamics means that the α -bits of the black hole are still revealed ‘as soon as possible’, subject to these inefficiencies.

4 Alpha-bits in the entanglement wedge

We now turn to developing the main claim of this paper — that there exist bulk regions, containing but not solely consisting of a black hole, for which the α -bits, but only the α -bits, are encoded in a certain boundary region. These bulk regions are bounded on both sides by extremal surfaces with areas \mathcal{A}_1 and \mathcal{A}_2 and they satisfy $\alpha = (\mathcal{A}_2 - \mathcal{A}_1)/\mathcal{A}_0$ where \mathcal{A}_0 is the horizon area of the black hole.

We first introduce the concept of entanglement wedge reconstruction. We then establish a correct version of the entanglement wedge reconstruction conjecture in section 4.1, including an underappreciated subtlety that proves qualitatively important for code spaces whose dimension grows quickly in the limit $G_N \rightarrow 0$. Finally, in section 4.2 we apply our results to geometries containing a single black hole in AdS space, establishing the results mentioned above.

The Ryu-Takayanagi formula [22, 23] states that, to leading order in G_N , the entanglement entropy of a boundary region A is equal to

$$S_{\text{RT}} = \frac{\mathcal{A}}{4G_N} + S_{\text{bulk}}, \tag{4.1}$$

where \mathcal{A} is the area of the bulk minimal surface anchored to the boundary of A and S_{bulk} is the bulk entropy of the bulk region bounded by the boundary region A and the minimal surface. However, as conjectured in [13] and proved in [24], the correct definition of the ‘minimal surface’ that one must consider is not simply the surface of minimal area (the minimal classical extremal surface),¹⁰ but rather the minimal quantum extremal surface, which is the surface anchored on the boundary of A that minimises the total size of (4.1).¹¹ If the size of the code space is held fixed as $G_N \rightarrow 0$, then the second term does not contribute to leading order in G_N and can be safely ignored at this order. However, if the dimension d_{code} of the code space of allowed bulk states exponentially large in $1/G_N$, the two terms can compete even in the semiclassical limit. This is exactly what happens when one considers thermal or two-sided black hole states; it is one way to understand the source of the ‘homology constraint’ in the original Ryu-Takayanagi formula [2].

The conjecture of entanglement wedge reconstruction was developed in [26–28] and then established with increasing rigour in [9, 29, 30]. It states that, if we take a code subspace of states with a fixed geometry, any region A of the boundary acts as a quantum error correcting code for the region of the bulk within its *entanglement wedge*, the bulk region enclosed by the boundary region A and the RT surface (specifically the quantum extremal surface) associated to region A . However, as we shall see, the form of the quantum error correction involved is somewhat more general than the definitions that we have given so far.

Let us split the boundary into two complementary regions A and \bar{A} . Each boundary region has an associated entanglement wedge in the bulk that we label a and \bar{a} , as shown in figure 4. However both the area term (at subleading order in G_N) and the bulk entropy term of (4.1) will in general depend on the state of the system; hence both the entanglement wedges may be state-dependent. If the boundary state is pure, the two entanglement wedges, a and \bar{a} will be complementary bulk regions; their union will contain the entire bulk. However, in the case of mixed states, as well as states that are entangled with a second system, there may a non-empty third region a' contained in neither entanglement wedge; one-sided (thermal) and two-sided black holes are respectively good examples of these two cases.

Every bulk region has an associated von Neumann algebra acting on the bulk (code) Hilbert space; similarly boundary regions are associated with von Neumann algebras acting

¹⁰We assume throughout this paper that we are in the Einstein gravity limit ($\lambda \rightarrow \infty$), where higher curvature corrections can be ignored. Moreover, we only consider static spacetime geometries and hence there is no need to use HRT surfaces [25] which generalises RT surfaces (which require a time-reflection symmetry) to general spacetimes. There will no doubt exist α -bit codes in time-dependent spacetimes as well, but we do not consider them here.

¹¹To avoid ambiguities about whether homology constraints are being applied the term RT surface will always refer in this paper to the quantum extremal surface.




Figure 4. Ryu-Takayanagi surfaces for empty AdS-space and a black hole in AdS. In each case, the boundary is separated into two regions, A and \bar{A} . In empty AdS, the boundary between these regions is also the boundary of a single minimal surface through the bulk, with (divergent) area \mathcal{A} , which is known as the Ryu-Takayanagi or RT surface. This minimal surface separates the bulk into two regions, a and \bar{a} , which we refer to as the entanglement wedges of A and \bar{A} respectively. When a black hole with horizon area \mathcal{A}_0 is introduced to the bulk, it creates infinitely many locally minimal surfaces with the same boundary as A and \bar{A} . In particular we are interested in the minimal surface homologous to \bar{A} with area \mathcal{A}_1 and the minimal surface homologous to A with area \mathcal{A}_2 . They divide the bulk into three regions a , \bar{a} and a' , where a' lies between the minimal surfaces and contains the black hole. For the thermal or canonical ensemble, regions a and \bar{a} form the entanglement wedges of A and \bar{A} respectively.

on the larger boundary Hilbert space. For both bulk and boundary, the algebra of a region forms the commutant of the algebra associated to its complementary region.

Since we are only interested in a single bulk geometry, there are no non-trivial operators in the centres of either the bulk or boundary algebras that are particularly relevant for our purposes. Therefore, for pedagogical reasons, we shall mostly assume that the centres of all the algebras are trivial and hence that we can associate a subsystem Hilbert space to each region. This assumption, although incorrect, is commonly used in the literature for simplicity and clarity. The von Neumann algebra associated to each region is simply the algebra of operators acting on the associated subsystem. The isometric embedding of the code subspace into the larger CFT Hilbert space has the form

$$\mathcal{H}_{\text{code}} \cong \mathcal{H}_a \otimes \mathcal{H}_{\bar{a}} (\otimes \mathcal{H}_{a'}) \subseteq \mathcal{H}_A \otimes \mathcal{H}_{\bar{A}} \cong \mathcal{H}_{\text{CFT}}, \tag{4.2}$$

where $\mathcal{H}_{a'}$ appears if region a' is non-empty. For convenience we will sometimes use $J : \mathcal{H}_{\text{code}} \rightarrow \mathcal{H}_{\text{CFT}}$ to represent the canonical embedding isometry.

In appendix A, we provide a more detailed description of the framework of operator algebra quantum error correction, which is necessary to talk about von Neumann algebras with non-trivial centres. This is particularly important for understanding code spaces with more than one semiclassical geometry, where the area of the RT surface corresponds to a non-trivial operator in the centre of the bulk algebras. All statements made in this section can be translated into statements about operator algebras (thus eliminating the incorrect

assumption about the algebras' centres) simply by replacing tensor product factors with von Neumann algebras and their commutants, and replacing partial traces with restrictions to subalgebras.

If we assume trivial centres and hence a tensor product factorisation, the entanglement wedge reconstruction conjecture can be phrased as follows: the channel $\mathcal{N} = \text{Tr}_{\bar{A}}(J(\cdot)J^\dagger)$ is a subsystem error correcting code for \mathcal{H}_a .¹² This means that there exists some decoding channel $\mathcal{D} : \mathcal{H}_A \rightarrow \mathcal{H}_a$ such that for all states $\rho \in S(\mathcal{H}_a \otimes \mathcal{H}_{\bar{a}} (\otimes \mathcal{H}_{a'}))$,

$$(\mathcal{D} \circ \mathcal{N}) \rho \approx \rho_a. \tag{4.3}$$

In other words, region A contains all the information necessary to approximately reconstruct the reduced density matrix of the state for region a .¹³

More commonly, we tend to think about quantum field theory using the Heisenberg rather than Schrödinger picture. The adjoint decoding channel $\mathcal{D}^\dagger : \mathcal{B}(\mathcal{H}_a) \rightarrow \mathcal{B}(\mathcal{H}_A)$ is a unital completely-positive superoperator that maps bulk operators in \mathcal{H}_a to boundary operators on \mathcal{H}_A . It is defined by

$$\text{Tr}(\mathcal{D}^\dagger(X_a)\rho_A) = \text{Tr}(X_a\mathcal{D}(\rho_A)), \tag{4.4}$$

for any observable $X_a \in \mathcal{B}(\mathcal{H}_a)$. We can therefore use \mathcal{D}^\dagger to reconstruct operators in the bulk using operators in only a subregion of the boundary, so long as the bulk operators are contained in the entanglement wedge of the boundary subregion; the existence of such a map is the most commonly-used definition of entanglement wedge reconstruction.

4.1 Entanglement wedge reconstruction from approximate decoupling

Before considering the specific task of bulk reconstruction in black hole geometries, we first establish some more general facts about entanglement wedge reconstruction. It was argued in [9] that a boundary region A can be used to reconstruct a bulk region a_0 if the bulk region a_0 was not contained in the entanglement wedge of the complementary boundary region \bar{A} . Equivalently, region a_0 must be contained in the entanglement wedge of region A for all *pure* states.¹⁴

For simplicity, the technical arguments in [9] ignored the existence of finite G_N corrections that make any reconstruction at best approximate; however the authors expected that the arguments should generalise to the approximate setting. For the situations considered in [9], where the code space dimension remains fixed in the limit $G_N \rightarrow 0$, this is indeed the case, as was verified in [30].

¹²Without the assumption of trivial centres, it states that the algebra associated to region A forms an operator algebra quantum error correcting code for the algebra associated to region a , see appendix A.

¹³Since the entanglement wedge a is in general state-dependent, we shall see in section 4.1 that we really the intersection of the entanglement wedge a for all states in the code space.

¹⁴We continue to distinguish the entanglement wedge a from the decodable region a_0 because in general the entanglement wedge a may depend on the state we are considering. Note that state-dependent entanglement wedges were not explicitly considered in [9]. However, applying the reconstruction theorem from [9] to code spaces with state-dependent entanglement wedges and ignoring approximation issues, one finds that the decodable region a_0 must be contained in the entanglement wedge a for *all* pure states in $\mathcal{H}_{\text{code}}$.

However, while the technical results of [9] continue to be true for exact error correcting codes, even when the code space dimension is very large (e.g. when the code space contains many black hole microstates), the generalisation to approximate error correction does not.

Recall that we saw in section 2 that being able to universally and exactly decode small subspaces implies being able to exactly decode the entire space, but that the approximate version of this statement was not true, because the quality of the approximation could degrade linearly with the dimension of the decoded space. Furthermore, to obtain dimension-independent bounds on the decoding error, we had to consider the forgetfulness of the environment for input states entangled with a reference system of equal dimension to the subspaces being decoded. No reference system was necessary when the error correction was exact.

The same phenomenon will appear when we try to prove approximate entanglement wedge reconstruction. It is not enough for the entanglement wedge of \bar{A} to never contain the bulk region a_0 ; instead the entanglement wedge of $\bar{A} \cup R$ cannot contain the bulk region a_0 for any (potentially entangled) pure state $|\psi\rangle \in \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$, where R is a reference system with the same dimension as the code space $\mathcal{H}_{\text{code}}$. Equivalently, the entanglement wedge a of A needs to always contain the bulk region a_0 for any pure, or mixed state. Hence region a_0 is really the intersection of the entanglement wedges of A for all pure states $|\psi\rangle \in \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$. The weaker Dong-Harlow-Wall condition that a bulk operator only needs to be contained in the entanglement wedge a for all pure states is only sufficient to show that the *zero-bits* of the operator can be reconstructed in region A .

Our justification for this claim involves a significant amount of technology from quantum information and takes the rest of this section: only the basic conclusions already mentioned will be required to understand the essential arguments in section 4.2. The same results were also reached via very different arguments in [30]. We choose to make a decoupling-based argument here, firstly because it makes explicit exactly when and how the Dong-Harlow-Wall argument from [9] fails and secondly because we can use it to prove that reconstruction can be made exact to all orders in perturbation theory. It appears significantly harder to adapt the argument in [30] to be perturbatively exact.

A version of the argument used by Dong, Harlow and Wall in [9] to justify the entanglement wedge reconstruction conjecture starts as follows. It had been shown in [29] that in the limit $N \rightarrow \infty$ relative entropies in the bulk become equal to relative entropies on the boundary. This means that if any two states, ρ and σ in the code subspace satisfy

$$\rho_{\bar{a}} = \sigma_{\bar{a}}, \tag{4.5}$$

then

$$S(\rho_{\bar{A}} || \sigma_{\bar{A}}) \leq \varepsilon \tag{4.6}$$

for some small ε that tends to zero if $N \rightarrow \infty$. In general the difference between the bulk and boundary relative entropies will be $O(G_N)$. However, $\rho_{\bar{a}} = \sigma_{\bar{a}}$ implies that the quantum extremal surface should be the same for the two states ρ, σ to all orders in perturbation theory. Hence ε will be non-perturbatively small (although still non-zero) for small G_N [24].

If $\varepsilon = 0$, then for all pairs of states $\rho, \sigma \in S(\mathcal{H}_{\text{code}})$ satisfying (4.5), it can be shown that there must exist a channel $\mathcal{D} : S(\mathcal{H}_A) \rightarrow S(\mathcal{H}_a)$ such that

$$\mathcal{D} \circ \text{Tr}_{\bar{A}}(\rho) = \rho_a, \tag{4.7}$$

for all states $\rho \in S(\mathcal{H}_{\text{code}})$. (In fact it would imply that we could additionally reconstruct region a' if such a region existed for some mixed states.) The channel $\text{Tr}_{\bar{A}}(J(\cdot)J^\dagger)$ would therefore form an exact error-correcting subsystem code for \mathcal{H}_a , which is what we wanted to show. This was the main technical result of [9].

While acknowledging that the equality between bulk and boundary relative entropy, and hence entanglement wedge reconstruction, should only be approximate at finite G_N , [9] did not attempt to prove an approximate version of their decoupling theorem, and instead left it as a task for future work. Let us now attempt to do exactly that. To generalise the argument given above to apply even for small but non-zero ε requires a generalisation of (2.1), or (equivalently) a version of Kretschmann et al.'s information-disturbance theorem [19] that works for subsystem error-correcting codes. Fortunately, such a generalisation is relatively straightforward and was done for the even more general structure of operator algebra error correction in [15]. We discuss the general form and its applicability in appendix A, but for now we shall simply apply the result in the special case of subsystem error correction.

As in section 2, we only obtain dimension-independent bounds if we consider states that are entangled with a reference system \mathcal{H}_R that has the same dimension as the space of states we wish to decode.

We first define an approximate subsystem error correcting code as follows. Let the channel $\mathcal{N} : S(\mathcal{H}_{a_0} \otimes \mathcal{H}_{\bar{a}_0}) \rightarrow S(\mathcal{H}_A)$. The channel \mathcal{N} forms an approximate subsystem error correcting code with error δ_1 if

$$\delta_1 = \inf_{\mathcal{D}} \|\mathcal{D} \circ \mathcal{N} - \text{Tr}_{\bar{b}}(\cdot)\|_{\diamond}, \tag{4.8}$$

where the infimum is taken over decoding channels $\mathcal{D} : S(\mathcal{H}_A) \rightarrow S(\mathcal{H}_b)$.

The complementary channel \mathcal{N}^c completely forgets the subsystem \mathcal{H}_{a_0} with uncertainty δ_2 if

$$\delta_2 = \sup_{|\psi\rangle} \|\psi_{\bar{A}R} - \text{Tr}_A(\omega_{a_0} \otimes \psi_{\bar{a}_0 R})\|_1 \tag{4.9}$$

where $\omega_{a_0} \in S(\mathcal{H}_{a_0})$ is maximally mixed and the supremum is again over all states $|\psi\rangle \in \mathcal{H}_{a_0} \otimes \mathcal{H}_{\bar{a}_0} \otimes \mathcal{H}_R$. The dimension d_R of the reference system can be unrestricted; however, it is sufficient to consider a reference system whose dimension is equal to the dimension of $\mathcal{H}_{\text{code}}$.

What does (4.9) mean in the context of entanglement wedge reconstruction? Let a_0 be the intersection of the entanglement wedges of a for all states $|\psi\rangle \in \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$ and let \bar{a}_0 be its bulk complement. Hence we have

$$\mathcal{H}_{\text{code}} \cong \mathcal{H}_{a_0} \otimes \mathcal{H}_{\bar{a}_0}. \tag{4.10}$$

If, as before, we have $\mathcal{N} = \text{Tr}_{\bar{A}}(J(\cdot)J^\dagger)$, then the complementary channel $\mathcal{N}^c = \text{Tr}_A(J(\cdot)J^\dagger)$. We note that because in (4.9) we are considering pure states $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_{\bar{A}} \otimes \mathcal{H}_R$, the entanglement wedge of $\bar{A} \cup R$ is by definition the complement of the entanglement wedge of A . In other words, the entanglement wedge of $\bar{A} \cup R$ is given by $a' \cup \bar{a} \cup R$.¹⁵ By definition this is contained in $\bar{a}_0 \cup R$. Therefore

$$(\omega_{a_0} \otimes \psi_{\bar{a}_0 R})|_{a' \bar{a} R} = \psi_{a' \bar{a} R} \quad (4.11)$$

and, hence, by the approximate equality between bulk and boundary relative entropies [24, 29], we have

$$S(\text{Tr}_A(\omega_{a_0} \otimes \psi_{\bar{a}_0 R}) || \psi_{\bar{A}R}) \leq \varepsilon, \quad (4.12)$$

for some non-perturbatively small ε . Using Pinsker's inequality [31], it follows that $\delta_2 \leq \sqrt{2\varepsilon \ln 2}$ and so this will also be non-perturbatively small.

However the size of the uncorrectable error δ_1 and the uncertainty in the forgetfulness δ_2 are related by [15]

$$\frac{1}{4}\delta_2^2 \leq \delta_1 \leq 2\delta_2^{\frac{1}{2}}, \quad (4.13)$$

and hence both tend to zero simultaneously with dimension-independent bounds. (The equivalent result for general operator algebra error correction is reproduced in appendix A as theorem 2; the subsystem error correction case (4.13) follows as a trivial consequence from this.) We have therefore shown that there exists an approximate subsystem error correcting code for region a_0 with non-perturbatively small error.

As for ordinary subspace error correction (discussed in section 2), such dimension-independent bounds are not possible if we do not include a reference system with the same dimension as the code space in the definition of complete forgetfulness. If the dimension of the code space is fixed in the semiclassical limit $G_N \rightarrow 0$, this is not especially problematic. It may contribute a large constant factor to the size of the decoding error, but it cannot affect with the error tends to zero in the limit $G_N \rightarrow 0$. This will not be true if want our code space S to include a large number of black hole microstates. Since the code space dimension $d_{\text{code}} \rightarrow \infty$ if $G_N \rightarrow 0$, dimension-dependent factors can affect whether the error tends to zero in this limit. As a result, they cannot be safely ignored.

If the code space dimension is fixed, the RT formula is dominated by the classical area term in the semiclassical limit and the RT surface is state-independent. However if the code space dimension diverges sufficiently fast, the bulk entropy term can compete

¹⁵Some readers may be unhappy at the notion of an entanglement wedge for states entangled with a reference system, which is not itself holographic. If so, it may be comforting to imagine the reference system as a second copy of the CFT with the same bulk code subspace so that everything is holographic. Note that, depending on context, R then refers to either the entire boundary or the entire bulk of this second system. We also remind readers that the arguments in [30] give the exact same conclusion we reach here (region a_0 must be contained in the entanglement wedge of A for all pure or mixed states) without ever invoking a reference system. We only need to do so here to make comparisons with the arguments in [9].

with the area term and it matters whether we consider the entanglement wedge of mixed states or only pure states. It is not a coincidence that a divergent code space dimension is required for both; the Dong-Harlow-Wall argument gives qualitatively different conclusions to the conclusions derived here — in exactly the contexts where we have shown that its conclusions should not be trusted.

4.2 Entanglement wedges for code spaces containing black holes

We argue in section 7, that we can construct code spaces

$$\mathcal{H}_{\text{code}} \cong \mathcal{H}_{\text{BH}} \otimes \mathcal{H}_{\text{ext}},$$

where the dimensions d_{code} and d_{BH} satisfy

$$\lim_{G_N \rightarrow 0} 4G_N \log d_{\text{BH}} = \lim_{G_N \rightarrow 0} 4G_N \log d_{\text{code}} = \mathcal{A}_0,$$

for a black hole with horizon area \mathcal{A}_0 . The Hilbert space \mathcal{H}_{ext} describes the degrees of freedom outside the horizon, while \mathcal{H}_{BH} describes the microstate of the black hole itself. Bulk operators outside the black hole horizon act only \mathcal{H}_{ext} , while the degrees of freedom in \mathcal{H}_{BH} are localised to the black hole (in other words the state on a boundary region only depends on \mathcal{H}_{BH} if the entanglement wedge of the boundary region contains the black hole). Moreover, all of the microstates in this code space are thermalised, typical “equilibrium states”.¹⁶

In figure 4b, we show an area A of the boundary together with two extremal surfaces through the bulk whose boundary is ∂A . The first is homologous to \bar{A} (the black hole is between the minimal surface and A) and has area \mathcal{A}_1 . The second is homologous to A (the black hole is between the minimal surface and \bar{A}) and has area \mathcal{A}_2 . We shall assume $\mathcal{A}_1 < \mathcal{A}_2 < \mathcal{A}_1 + \mathcal{A}_0$. We label the bulk region between A and the minimal surface with area \mathcal{A}_2 by a , while the region between \bar{A} and the minimal surface with area \mathcal{A}_1 is called \bar{a} and the region between the two minimal surfaces (which contains the black hole) is labelled a' . Note that, for the thermal state, region a is the entanglement wedge of region A and region \bar{a} is the entanglement wedge of region \bar{A} and so this is consistent with our previous notation. However, for individual microstates, the entanglement wedge of A will be region $a \cup a'$. To avoid confusion, we will keep the definitions of regions a , a' and \bar{a} fixed and state-independent throughout this section, rather than having their definition depend on the state.

The code space has the form

$$\mathcal{H}_{\text{code}} \cong \mathcal{H}_a \otimes \mathcal{H}_{a'} \otimes \mathcal{H}_{\bar{a}}.$$

Since the black hole is contained within region a' , we have

$$\mathcal{H}_{a'} \cong \mathcal{H}_{\text{BH}} \otimes \mathcal{H}_{a'}^{\text{ext}}.$$

¹⁶This last property is only possible because we are not trying to include *all* black hole microstates in our code space, merely sufficiently many microstates to give the Bekenstein-Hawking entropy up to a subleading correction.

Since we chose the code space geometry to include only a single black hole plus perturbative excitations outside it, only $\mathcal{H}_{a'}$ (and not \mathcal{H}_a or $\mathcal{H}_{\bar{a}}$) contains an exponential number of states (w.r.t. $1/G_N$).

Now consider an arbitrary pure state $|\psi\rangle \in \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$. We wish to find the entanglement wedge of $\bar{A} \cup R$ for all such states. If some bulk region a' is not contained in the entanglement wedge of $\bar{A} \cup R$ for any state in $\mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$, then, by our arguments in section 4.1, it should be possible to reduced density matrices for region $a \cup a'$ from reduced density matrices for region \bar{A} . Equivalently it should be possible to simulate operators acting on $\mathcal{H}_a \otimes \mathcal{H}_{a'}$ with operators acting on \mathcal{H}_A .

What can this entanglement wedge be? Since we chose the code space such that the geometry is approximately the same for every state in the subspace, the first term in the Ryu-Takayanagi formula (4.1) is approximately the same for all code states and any fixed surface. As a result the only way state dependence of the entanglement wedge can arise is through the second term in (4.1).¹⁷ Moreover, the only source of bulk entanglement that can be sufficiently large to compete with the area term as $G_N \rightarrow 0$ is entanglement between the black hole and reference system. As a result, the quantum minimal surface will always be one of the two classical extremal surfaces with areas \mathcal{A}_1 and \mathcal{A}_2 . The only question is whether the entanglement wedge of $\bar{A} \cup R$ contains only $\bar{a} \cup R$ or consists of $a' \cup \bar{a} \cup R$, as depicted in figure 5. We can therefore decode $a \cup a'$ (and hence the black hole) from A so long as for all states $|\psi\rangle \in \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$,

$$S(\bar{a}a'R)_\psi + \frac{\mathcal{A}_2}{4G_N} > S(\bar{a}R)_\psi + \frac{\mathcal{A}_1}{4G_N}, \tag{4.14}$$

We have $S(a'\bar{a}R)_\psi = S(a)_\psi = O(1)$ and, by the triangle inequality, $|S(R)_\psi - S(\bar{a}R)_\psi| \leq S(\bar{a}) = O(1)$. Hence, to leading order in G_N , (4.14) is equivalent to

$$4G_N S(R)_\psi < \mathcal{A}_2 - \mathcal{A}_1, \tag{4.15}$$

If

$$\lim_{G_N \rightarrow 0} 4G_N \log d_R = \lim_{G_N \rightarrow 0} 4G_N \log d_S < \mathcal{A}_2 - \mathcal{A}_1, \tag{4.16}$$

then this will be satisfied for any state $|\psi\rangle \in \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$. Conversely if

$$\lim_{G_N \rightarrow 0} 4G_N \log d_R = \lim_{G_N \rightarrow 0} 4G_N \log d_S > \mathcal{A}_2 - \mathcal{A}_1, \tag{4.17}$$

then it will be violated for any maximally-entangled state on $\mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$. We therefore see that in the classical limit we can decode $\mathcal{H}_a \otimes \mathcal{H}'_a$ from \mathcal{H}_A for any subspace $\mathcal{H}_S \subseteq \mathcal{H}_{\text{code}}$ whose dimension is less than $e^{\alpha \frac{\mathcal{A}_0}{4G_N}}$ for

$$\alpha = \frac{\mathcal{A}_2 - \mathcal{A}_1}{\mathcal{A}_0}. \tag{4.18}$$

¹⁷In principle, as with the thermofield double state, there could be a smooth geometry between the horizon of the black hole and R with an associated minimal area; however, we are always free to interpret this as simply bulk entanglement, see [2].




Figure 5. The black hole is now entangled with a reference system R . If the entanglement between the black hole and R has entropy greater than $\frac{A_2 - A_1}{4G_N}$, the entanglement wedge of $\bar{A} \cup R$ is $\bar{a} \cup a' \cup R$. Otherwise the entanglement wedge of $\bar{A} \cup R$ is only $\bar{a} \cup R$.

In other words A contains the α -bits of region a' for the entire code space $\mathcal{H}_{\text{code}}$. This region contains not just the black hole, but also an additional bulk region outside the black hole horizon but between the two minimal surfaces. In contrast the region a can be decoded for the entire code space, since this region will always lie outside the entanglement wedge of region $\bar{A} \cup R$.

In the Heisenberg picture, this means that we can simulate an operator $\mathcal{O}_{a'}$ on a' with an operator \mathcal{O}_A on A so long as we only require that the operator \mathcal{O}_A behave in the same way as $\mathcal{O}_{a'}$ within a subspace $\mathcal{H}_S \subseteq \mathcal{H}_{\text{code}}$ with dimension less than $e^{\alpha \frac{A_0}{4G_N}}$. In other words, for any bulk operator in a' , whether acting on the black hole or outside the horizon, the operator on A must be state-dependent. We discuss the connection with other proposed forms of operator state dependence in quantum gravity in section 8.3.

Finally, we want to show that the value of α given in (4.18) is optimal. In other words, that for any $\alpha' > \alpha$, the α' -bits of region a' are not encoded in region A . We have already shown that (4.15) can be violated when $d_R = e^{\alpha' S}$ for such α' and hence that there exist states in $\mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$ for which the entanglement wedge of $\bar{A} \cup R$ contains region a' . Specifically we can consider a state where \mathcal{H}_R is maximally entangled with \mathcal{H}_{BH} . Acting with bulk operators in region a' but outside the black hole horizon cannot change the entanglement wedge for such a state, and so we can construct a small subspace of states $\mathcal{H}_S \subseteq \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$ with dimension d_S which look identical in region a and for which region a' is never contained in the entanglement wedge of A . Since the dimension of \mathcal{H}_S is small, even if we consider states entangled with a second reference system $\mathcal{H}_{R'}$, the entanglement wedge of $A \cup R'$ will still never contain region a' .

It follows, by the arguments made in section 4.1, that we can recover the state for region a' from $\bar{A} \cup R$ so long as we know that it lies in the subspace \mathcal{H}_S . By the no cloning theorem (or, more formally, Kretschmann et al.'s information-disturbance theorem [19]), we therefore cannot recover the state from A . However the support of any state in \mathcal{H}_S lies

in a subspace $\mathcal{H}_{\text{sup}} \subseteq \mathcal{H}_{\text{code}}$ of dimension at most $d_R d_S$. It therefore cannot be possible to decode region a' using only region A for the subspace \mathcal{H}_{sup} . Since d_S is small and fixed,

$$\lim_{G_N \rightarrow 0} \frac{\log(d_R d_S)}{S_{\text{BH}}} = \alpha'. \tag{4.19}$$

Since we could have chosen α' to be arbitrarily close to α , region A cannot encode the α' -bits of region a' for any $\alpha' > \alpha$.

This argument makes clear that the operator state dependence is unavoidable; it is not simply the product of a particular reconstruction strategy. However, as we make region A larger, not only does the region a where operators can always be decoded become larger, so does the value of α itself. The size of the subspaces \mathcal{H}_S , for which operators in region a' can be reconstructed, grows as region A grows. Eventually, we reach the point where $\alpha = 1$ and hence a single operator will exist that exhibits the correct behaviour for all the black hole microstates in $\mathcal{H}_{\text{code}}$. The entanglement wedge of A will now always contain regions a and a' ; for the purposes of entanglement wedge reconstruction, it no longer matters which of these regions an operator is in.

5 Alpha-bits of BTZ black holes

We now consider the case of an uncharged, non-rotating BTZ black hole in 2+1 dimensions. This provides a sufficiently simple example of the phenomena introduced in section 4 that many of the relevant quantities can be calculated analytically. Most of these calculations have already been done in the literature [32, 33]. In particular the Holevo information χ for an ensemble of black hole microstates in AdS_3 was calculated in [33]; it turns out that the Holevo information has a very simple relation to α which is given by¹⁸

$$\chi = \frac{\mathcal{A}_2 - \mathcal{A}_1}{4G_N} = \alpha S. \tag{5.1}$$

Nonetheless we shall carry out all the calculations here explicitly in the interest of clarity.

In general, the explicit calculations conform with one's intuition; increasing the size of the boundary region A increases α , while increasing the radius of the black hole decreases α . We also explicitly calculate the volume of the region a' and find that it is equal to $2\pi L^2$ — independent of the radius of the black hole. The size of the region is always approximately AdS scale, even for very large black holes. The size is independent of G_N — α -bit codes exist even in the semiclassical limit (in fact that is where they are best defined) — but it cannot be made significantly larger than the AdS scale. The same effect is seen in section 6 in tensor network toy models of holographic α -bit codes.

One possible explanation for this is that α -bit codes in general seem to rely on the properties large random-like unitaries — essentially they rely on and reflect some form of scrambling. However, fast scrambling only happens in large AdS black holes up to the AdS

¹⁸It should be obvious that αS is a lower bound for the Holevo information. The fact that (5.1) is actually an equality, however, is a non-trivial fact about black holes and comes from the universal behaviour of black hole microstates.

scale [34]. More generally, locality above the AdS scale in AdS/CFT comes from locality in the CFT (and locality in energy scale for the radial dimension). However, sub-AdS scale locality is more mysterious and is associated with the large number of local matrix degrees of freedom in $\mathcal{N} = 4$ SYM at large N (or the equivalent degrees of freedom in other examples).

This suggests that α -bit codes, or more specifically α -bit degrees of freedom *outside* the horizon, are a property of CFTs with large N (or large central charge). They are teaching us something about the encoding of degrees of freedom specifically for CFTs with weakly curved gravity duals. It would be interesting to see if the volume of region a' continues to be AdS scale for more complex black holes, or whether, by tuning charges, angular momenta etc., we can make the region a' arbitrarily large.

In Schwarzschild coordinates the BTZ metric is given by

$$ds^2 = -\frac{r^2 - R^2}{L^2} dt^2 + \frac{L^2}{r^2 - R^2} dr^2 + r^2 d\phi^2 \tag{5.2}$$

where $\phi \sim \phi + 2\pi$ and R and L are respectively the horizon radius and the AdS scale.

Since gravity has no local degrees of freedom in $2 + 1$ dimensions, the BTZ black hole can be identified with a quotient of pure AdS₃. We can take advantage of this by calculating the area of minimal surfaces, which are just geodesics in $2 + 1$ dimensions, using formulas for the lengths of geodesics in pure AdS₃.

The geodesic length d between two points in pure AdS satisfies

$$\cosh\left(\frac{d}{l}\right) = T_1 T'_1 + T_2 T'_2 - X_1 X'_1 - X_2 X'_2 \tag{5.3}$$

where the embedding co-ordinates (T_i, X_i) can be identified with Schwarzschild co-ordinates (r, t, ϕ) by


$$\begin{aligned} T_1 &= \frac{1}{R} \sqrt{r^2 - R^2} \sinh \frac{Rt}{L}, \\ T_2 &= \frac{r}{R} \cosh \frac{R\phi}{L}, \\ X_1 &= \frac{1}{R} \sqrt{r^2 - R^2} \cosh \frac{Rt}{L}, \\ X_2 &= \frac{r}{R} \sinh \frac{R\phi}{L}. \end{aligned} \tag{5.4}$$

The relevant geodesics on the BTZ geometry, which travel from $(r, 0, 0)$ to $(r, 0, \phi)$ in the limit $r \rightarrow \infty$, can be identified with the geodesics in pure AdS from $(r, 0, 0)$ to $(r, 0, \phi)$ and $(r, 0, \phi - 2\pi)$, since we identify $\phi \sim \phi + 2\pi$ in the BTZ geometry. We label their lengths by d_1 and d_2 respectively. In the limit of large r , then (5.3) implies,

$$\frac{d_1}{L} = 2 \log\left(\frac{2r}{R}\right) + 2 \log\left(\sinh \frac{R\phi}{2L}\right), \tag{5.5}$$

and

$$\frac{d_2}{L} = 2 \log\left(\frac{2r}{R}\right) + 2 \log\left(\sinh \frac{R(2\pi - \phi)}{2L}\right). \tag{5.6}$$



(a) Plot of α against boundary angle ϕ . (b) Plot of angle ϕ against radius R for various α .

Figure 6. (a) α smoothly increases as we increase the size of the boundary region A (with angle ϕ). The black hole is a BTZ black hole with horizon radius $R = 2L$. The boundary encodes no information about region a' for $\phi < \pi$. For $\phi > \pi$, it contains the α -bits for increasingly large α , until eventually $\alpha = 1$ and the region a' can be decoded without any knowledge of the black hole microstate. (b) The angle ϕ of the boundary region A required to decode the bulk region a' for fixed values of α steadily increases with the radius R of the black hole in AdS units. As the radius becomes large, the required angle ϕ converges to $(1 + \alpha)\pi$.

Since the horizon area is simply $2\pi R$, we find that a boundary region of angle $\phi > \pi$ encodes the α -bits of a BTZ black hole of horizon radius R for

$$\alpha = \frac{d_1 - d_2}{2\pi R} = \frac{\log\left(\sinh\frac{R\phi}{2L}\right) - \log\left(\sinh\frac{R(2\pi-\phi)}{2L}\right)}{\pi\frac{R}{L}}. \quad (5.7)$$

Figure 6a shows how α increases smoothly with the size of the boundary region A for a fixed black hole radius, while figure 6b shows how the size of the region required for any fixed α increases as the black hole radius increases.

If we take the limit where $\frac{R}{L} \gg 1$, while holding ϕ fixed, we find that

$$\alpha = \frac{2\phi}{2\pi} - 1. \quad (5.8)$$

The fraction of the boundary required to encode the α -bits is the inverse of the α -bit capacity of the noiseless qubit channel $(1 + \alpha)/2$. This is identical to the fraction of the Hawking radiation that we learned in section 3 was required to encode the α -bits of a black hole. We discuss this further in section 8.5.

Finally we can calculate the size of the bulk region a' (ignoring the black hole itself) for which the α -bits are encoded in the boundary region A . An explicit calculation in Schwarzschild co-ordinates seems daunting, so we shall instead again take advantage of the fact that the BTZ geometry is a quotient of pure AdS space. If we extend our picture of the BTZ black hole to include both sides of the Einstein-Rosen bridge, we see from figure 7 that two copies of region a' forms the interior of a hexagon in hyperbolic space, bounded by two copies of each of the geodesics A_1 and A_2 together with the edges of the fundamental region of the BTZ quotient.

This hexagon could be broken down in four triangles whose vertices all lie on the boundary of the space. These are known as ideal triangles and by using the symmetries




Figure 7. A spatial slice of a two-sided BTZ black hole can be viewed as a quotient of the Poincaré disk. The fundamental domain is shown unshaded. We split each boundary into regions A and \bar{A} , with an associated geodesic through the bulk (shown in red). The black hole horizon is shown in blue. Between the horizon and the geodesics (and within the fundamental region) there is a copy of region a' on each side. The union of these two regions is bounded by an ideal hexagon on the Poincaré disk.

of hyperbolic space, we can push the three boundary points to any other three boundary points: this shows that the area of the triangles (and hence a') is independent of both the radius R of the black hole and the boundary angle ϕ .

To calculate the volume of the hexagon explicitly we can use the Gauss-Bonnet formula

$$\int_{\mathcal{M}} K \, dA + \int_{\partial\mathcal{M}} k_g \, ds = 2\pi\chi(\mathcal{M}). \tag{5.9}$$

The Gaussian curvature $K = -\frac{1}{L^2}$ and so can be taken out the front of the integral. Since the line segments are geodesics the only contribution to the boundary curvature term comes from the six corners, each of which lies on the boundary and so has angle $\theta \rightarrow \pi$. Finally the hexagon is homeomorphic to the disk and so has Euler characteristic $\chi(\mathcal{M}) = 1$. Evaluating this we obtain,

$$V_{a'} = \frac{1}{2}V_{\text{hex}} = \frac{1}{2}(6\theta - 2\pi)L^2 = 2\pi L^2. \tag{5.10}$$

We see that the region a' is always approximately AdS scale, as discussed above. A weakly curved bulk dual is required for region a' to be sharply defined.

6 Alpha-bits in tensor networks

Tensor networks have been widely studied as a toy model of holography in recent years [8, 21, 35–38]. We find that they also give simple toy models of holographic α -bit codes. Most importantly, they provide very clear intuition for why state dependence appears in the operator reconstruction; to reconstruct operators in the α -bit region to the boundary we have to push them through the black hole itself. In other words, the isometry mapping bulk




Figure 8. A pentagon code with a black hole in the centre. The black hole is modelled as a large random tensor with a single D -dimensional dangling leg creating the space of allowed microstates and n_{tot} d -dimensional legs flowing into the network. The remaining tensors are all perfect tensors with a single dangling d -dimensional leg that forms the bulk degrees of freedom and 5 d -dimensional legs flowing into the network.

operators to boundary operators doesn't just depend on the state of the black hole; part of the isometry consists of the tensor that literally describes the state of the black hole.

We consider a version of the pentagon code developed in [8]. Empty AdS in the bulk is described by perfect tensors, which have the property that they describe a unitary map from any half of their legs (or indices if one is more traditionally inclined) to the other half. Each six-legged tensor has a single dangling bulk leg. The overall network can be thought of as an isometry¹⁹ from a bulk Hilbert space to a large boundary Hilbert space. The image of the bulk Hilbert space can be thought of as the code subspace of the boundary space.

To include a black hole in this network, we simply add a random tensor with a large number of legs at the centre of the network, as shown in figure 8. If we wish to include an entire subspace of black hole microstate (rather than simply a single microstate), this tensor must have also have dangling bulk input leg whose dimension is the number of microstates we wish to consider.

If we divide the boundary into two regions A and \bar{A} as before, there is a natural notion of a bulk surface of ‘minimal area’, which is simply the path through the bulk which cuts through the fewest bulk legs. If we have a black hole in the network, we can find the minimal surface on either side of the black hole. Just as in real AdS/CFT, there may exist tensors outside of the black hole that lie between these two geodesics, giving a region that

¹⁹If we view the tensor network a map from the bulk sites outwards to the boundary, we see that each tensor in the network has at least three legs flowing ‘outwards’ and hence forms an isometry from the dangling and inwards pointing legs to the outwards legs. Additionally so long as $\log D < n_{\text{tot}} \log d$ the black hole random tensor will be approximately an isometry from the dangling microstates leg to the remaining legs. The entire network is therefore an isometry from bulk to boundary.

naturally corresponds to the region a' in figure 4b. Depending on the boundary points in question, the region a' may contain anywhere from zero to two tensors adjacent to the black hole (as well as potentially other tensors further from the black hole). This is in agreement with our calculation in section 5, where we found that volume outside of the black hole horizon for the region a' in the BTZ geometry is of approximately AdS scale, independent of both the radius of the black hole and the angle of the boundary contained in region A .

Since the map from the bulk Hilbert space to the boundary Hilbert space is an isometry, clearly operators in the bulk can be represented (non-uniquely) by operators in the boundary. In fact we can use the properties of perfect tensors to do considerably better than this and represent operators in the bulk with operators on only part of the boundary; the isometry is an error-correcting code. To see why, note that we can view the perfect tensor at our bulk operator site as a unitary map from the bulk leg, as well as two other legs of our choice, to the remaining three legs. By conjugating the operator by this unitary, we can map it to an operator acting on only those three legs of the tensor network. In turn, we can ‘push’ the operator acting on those three legs, through more tensors, so long as we always push the operator onto at least three new legs each time. In this way we can push the operator onto only a subregion of the boundary and achieve a version of entanglement wedge reconstruction.²⁰

What happens if we have a situation equivalent to the one in section 4? In this context that means that the difference between the size of the minimal cuts through the network on each side of the black hole is less than the number of legs coming out of the black hole. They naturally divide the bulk into regions a , a' and \bar{a} , as in figure 4b. An example is shown in figure 9. In this case we see that it is impossible to push an operator in region a' between the two cuts to the boundary without pushing it *through* the black hole. Of course we can only push the operator through the black hole if the black hole tensor is (at least approximately) an isometry from the legs we treat as input to the legs we use as outputs.

A random tensor becomes approximately an isometry with very high precision in the limit of large dimension if the dimension of the input grows less quickly than the output dimension. If n_a legs flow out of the black hole into region a , while $n_{a'}$ flow into region a' and $n_{\bar{a}}$ legs flow into region \bar{a} , then, in the limit of large bond dimension d , we can approximately push the operator through the black hole (with its D allowed microstates) and onto sites in region a so long as

$$\log D + (n_{\bar{a}} + n_{a'}) \log d < n_a \log d. \tag{6.1}$$

In other words we require

$$\log D < (n_a - n_{a'} - n_{\bar{a}}) \log d = \frac{n_a - n_{a'} - n_{\bar{a}}}{n_{\text{tot}}} S \tag{6.2}$$

²⁰In fact the region that can be reconstructed is slightly different to what one might naively guess, since the pentagon code does not always satisfy the Ryu-Takayanagi formula. Instead only the region known as the *greedy entanglement wedge* can be reconstructed. For the same reason, the pentagon code will prove to be a weaker α -bit encoding than true AdS/CFT.




Figure 9. An example of an operator that can only be pushed to region A of the boundary through the black hole. This can only be achieved if the space of microstates is small enough for the black hole to act approximately as an isometry. In this case we have $n_a = 12$, while $n_{\bar{a}} = 6$ and $n_{a'} = 2$. As a result we must have $\log D < 4 \log d$, or equivalently $\alpha = \frac{1}{5}$.

where n_{tot} is the total number of legs flowing out of the black hole and $S = n_{\text{tot}} \log d$ is the entropy of the black hole. We can therefore consider at most $D = e^{\alpha S}$ microstates, where

$$\alpha = \frac{n_a - n_{a'} - n_{\bar{a}}}{n_{\text{tot}}}.$$

How does this compare to our formula for AdS/CFT? We have that

$$\frac{A_2 - A_1}{A_0} = \frac{n_a + n_a^{\text{ext}} - n_{\bar{a}} - n_{\bar{a}}^{\text{ext}}}{n_{\text{tot}}}, \quad (6.3)$$

where n_a^{ext} and $n_{\bar{a}}^{\text{ext}}$ are the number of legs cut by the minimal surfaces which do not connect to the black hole. In the case of figure 9,

$$n_a^{\text{ext}} = n_{\bar{a}}^{\text{ext}},$$

while $n_{a'} = 2$, so we cannot decode as large subspaces in the pentagon code as we expect to be able to decode in AdS/CFT.

Since the pentagon code does not always satisfy the Ryu-Takayanagi formula this should not disturb us; however, in this case the explanation is quite simple. Because the bulk dangling legs in the pentagon code also have dimension d , we cannot ignore their contribution to the bulk entropy. If region a' has k tensors outside the black hole, these can give an additional contribution of $k \log d$ to the bulk entanglement entropy. This means that the quantum extremal surface can be pushed to exclude a' even if we only have

$$\log D > (n_a + n_a^{\text{ext}} - n_{\bar{a}} - n_{\bar{a}}^{\text{ext}} - k) \log d. \quad (6.4)$$




Figure 10. Region A is now sufficiently large that the same operator from figure 9, which is still in region a' , can now be pushed directly to the boundary without passing through the black hole. As a result it can be reconstructed on the boundary without any dependence on the state of the black hole.

In the case of figure 9, we have $k = n_{a'} = 2$ and so this is sufficient to explain the discrepancy; however, it is possible to choose the region A of the boundary such that operators in region a' can only be decoded for smaller α than we expect from AdS/CFT (even taking the bulk entanglement fully into account). Just as with ordinary quantum error correction, the pentagon code gives a weaker α -bit encoding than true AdS/CFT.

As in AdS/CFT, if we make the region A sufficiently large, we are able to simulate the bulk for an arbitrarily large space of microstates, or even for a maximally-entangled black hole. As shown in figure 10, eventually, as A gets larger, it becomes possible to push operators in region a' to the boundary without having to push them through the black hole. This gives a boundary operator that works exactly for any black hole state (or even no black hole at all).

7 The space of black hole microstates

Throughout this paper and particularly (and explicitly) in section 4, we have assumed that there exists a linear subspace

$$\mathcal{H}_{\text{code}} \cong \mathcal{H}_{\text{BH}} \otimes \mathcal{H}_{\text{ext}}$$

with dimensions d_{BH} and d_{code} satisfying

$$\lim_{G_N \rightarrow 0} 4G_N \log d_{\text{BH}} = \lim_{G_N \rightarrow 0} 4G_N d_{\text{code}} = \mathcal{A}_0 \tag{7.1}$$

where we can treat \mathcal{H}_{BH} as describing the microstate of the black hole, while \mathcal{H}_{ext} describes the state outside the horizon. Is this a reasonable assumption?

In this section we try to argue as rigorously as possible that the answer is yes. Anyone who already agrees with this claim may well find the arguments overly detailed and pedantic. We encourage such readers to skip this section. However, since this paper is built around effects that are only possible when the dimension of the code space is exponentially large, it is important to justify that such exponentially large code spaces do in fact exist.

The AdS/CFT dictionary maps pure black hole microstates in the bulk to pure states that are thermalised on the boundary. We would not necessarily expect the subset of a Hilbert space consisting of approximately thermal pure states to form a linear subspace. There may well be ways to take superpositions of thermal states and produce out-of-equilibrium unthermalised states. As a result, it seems likely that there does not exist a linear subspace of CFT states which contains ‘all’ black hole microstates, without also including some states with very different properties to generic microstates.²¹

However attempting to include every black hole microstate in a subspace is much harder than finding a subspace with only black hole microstates and the correct entropy to leading order in G_N . If

$$\log d_{\text{BH}} = (1 - \delta) \frac{\mathcal{A}_0}{4G_N} \tag{7.2}$$

for some fixed very small $\delta > 0$, only a tiny fraction of the microstates, exponentially small in $\frac{1}{G_N}$, need be included. However, restricting to such a subspace only affects α by $O(\delta)$ and can be ignored to leading order.²² It follows that it doesn’t matter whether we are talking about the α -bits of this smaller space or the hypothetical much larger space of *all* black hole microstates (if such a space could even be defined, as discussed above).

Let us start by considering the microcanonical ensemble: the maximally mixed state formed from all energy eigenstates within some small energy window whose width is independent of G_N (we of course need the energy itself to scale as $1/G_N$ in order to keep the size of the black hole constant in AdS units). The dimension of this space will obviously scale as $e^{\mathcal{A}_0/4G_N}$.

By the arguments made in [39], almost all states, in the subspace spanned by eigenstates in this energy window, will look almost indistinguishable from the microcanonical ensemble (which in turn will look like the thermal or canonical ensemble to leading order) so long as we only look at the state restricted to a subregion with less than half the size of the entire boundary. In fact, the measure concentration is so strong that we can use the union bound to show that, with very high probability, this will be true for all states in a

²¹A more precise statement might be that there do not necessarily exist code spaces containing only black hole microstates such that the projection of the thermal ensemble onto the code space approaches the thermal ensemble in the semiclassical limit. It should nonetheless be possible to include all microstates in a sufficiently small energy window without including unthermalised states, as we argue below.

²²Indeed, if we know that code spaces with the correct properties exist in the limit $G_N \rightarrow 0$ for any $\delta > 0$, we can construct a sequence of such code spaces where $\delta \rightarrow 0$ (albeit potentially very slowly) as $G_N \rightarrow 0$. This sequence of code spaces will exactly satisfy (7.1).

randomly-chosen subspace of dimension

$$e^{(1-\delta)\frac{A_0}{4G_N}}$$

in the limit of small G_N for any $\delta > 0$. See, for example, similar arguments in [17].

This is essentially all we need to construct the space \mathcal{H}_{BH} : any such space has the requisite size and all the states are indistinguishable from a thermal state (a black hole) if we look at less than half the boundary, and hence an entanglement wedge that does not include the black hole. It follows that all operators acting on this space approximately commute with any bulk operator in any such entanglement wedge. There is a small entanglement shadow around the black hole horizon which is not contained in any such entanglement wedge, however, it is sufficient for our purposes that the degrees of freedom in \mathcal{H}_{BH} be localised within this shadow, since by definition no entanglement wedge can enter it.

Furthermore, if we assume the eigenstate thermalisation hypothesis [40–42], then the expectation of all simple correlation functions of bulk operators outside the horizon, even ones within the entanglement shadow, will be approximately equal to the thermal expectation for all eigenstates. In fact since the variance of the off-diagonal terms is suppressed by e^{-S} , this will still be true for all superpositions of at most $e^{(1-\delta)S}$ eigenstates. An advantage of this construction is that it allows us to explicitly construct the code space simply by restricting to a very small energy window (of size $e^{-\delta S}$). The argument based on measure concentration only allows us to make statements about ‘most’ randomly chosen subspaces.

We can therefore construct a space of sufficient size with no bulk degrees of freedom outside the black hole horizon. Every state in the space is an “equilibrium state” [43] i.e. they satisfy the KMS condition [44]. Let us assume we have constructed such a subspace, which we shall call $\mathcal{H}_{\text{BH}}^0$. To extend our code space to include \mathcal{H}_{ext} and hence degrees of freedom outside the black hole requires slightly more work, but essentially it is the same procedure used to construct a code space associated to a single microstate, where we simply add states to the code space if they can be produced by acting with bulk operators on the microstate (up to some limit).

Suppose we do this for an entire basis of states $\{|\Omega_i\rangle\}$ in $\mathcal{H}_{\text{BH}}^0$. We get a set of code subspaces $\{\mathcal{H}_{i,\text{ext}}\}$ — one for each basis microstate. However, because we already argued that the expectation of any bulk operator is constant within $\mathcal{H}_{\text{BH}}^0$, these individual code subspaces will all be (approximately) orthogonal within the larger CFT Hilbert space. We also have the canonical isomorphisms

$$a_{\text{bulk}} |\Omega_i\rangle \cong a_{\text{bulk}} |\Omega_j\rangle$$

for all bulk operators a_{bulk} . We can use these to identify

$$\bigoplus_i \mathcal{H}_{i,\text{ext}} \cong \mathcal{H}_{\text{BH}} \otimes \mathcal{H}_{\text{ext}}, \tag{7.3}$$

where

$$\mathcal{H}_{\text{BH}} \cong \mathcal{H}_{\text{BH}}^0 \tag{7.4}$$

and

$$\mathcal{H}_{\text{ext}} \cong \mathcal{H}_{i,\text{ext}}. \tag{7.5}$$

Operators on \mathcal{H}_{ext} act on the code spaces $\mathcal{H}_{i,\text{ext}}$ while operators on \mathcal{H}_{BH} map those spaces into one another in accordance with their action on the basis states $\{|\Omega_i\rangle\}$.

Any low-energy bulk operator outside the horizon preserves the code spaces $\mathcal{H}_{i,\text{ext}}$, and hence corresponds to an operator acting only on \mathcal{H}_{ext} . Meanwhile operators on \mathcal{H}_{BH} change the state of the black hole itself, while leaving the state outside the horizon unchanged. We have constructed $\mathcal{H}_{\text{code}}$ as desired.

It is important to note that, while many of the steps that we have made in this section are only true approximately, in most cases the corrections are suppressed by powers of $\exp(-\delta\mathcal{A}_0/4G_N)$. As a result, almost all the statements in this section should be perturbatively exact to all orders in G_N for any fixed $\delta > 0$.

The only approximation that is not exact at the level of the perturbative expansion is the equivalence between the microcanonical ensemble of all states in a narrow energy range and the canonical or thermal ensemble. We shall briefly explain this observation, mostly to make it clear that it does not prevent our construction from being perturbatively exact. If we accept the eigenstate thermalisation hypothesis, any bulk operator \mathcal{O} can be written in the energy eigenbasis as

$$f_{\mathcal{O}}(E_i)\delta_{ij} + O(e^{-S}).$$

The expectation of \mathcal{O} for the microcanonical ensemble is just $f_{\mathcal{O}}(E)$. For the canonical ensemble it is given by

$$\int dE e^{S(E)-\beta E} f_{\mathcal{O}}(E). \tag{7.6}$$

To leading order in a saddle point expansion this agrees with the microcanonical ensemble. However, higher-order corrections will break this equality and are only polynomially suppressed in $1/S \sim G_N$. Fortunately this disagreement is completely unproblematic for our purposes; all we wanted is for every state in $\mathcal{H}_{\text{BH}}^0$ to be indistinguishable at the perturbative level outside the horizon — and for their leading order geometry to be a black hole. Both the canonical and microcanonical ensembles are described by a black hole geometry to leading order in G_N , and both will have perturbative corrections and fluctuations in the geometry from the excitation of fields (including gravitons) outside the horizon. The fact the bulk correlators have perturbative differences just tells us that those perturbative bulk corrections are different for the microcanonical and canonical ensembles.

For temperatures below the AdS scale, the black hole saddle point is no longer dominant and so the canonical ensemble is dominated by states consisting of thermal excitations on a vacuum-AdS background, as is the microcanonical ensemble at sufficiently low energies. Nevertheless, there still exists a saddle point in the partition function that becomes semiclassical in the limit $G_N \rightarrow 0$. As a result, we expect all the arguments made above to continue to apply, with the caveat that the ensembles being discussed are no longer, strictly speaking, canonical or microcanonical.

8 Discussion

8.1 Upper bounds on error correction accuracy

The results in this paper can be used to put strong upper bounds on the accuracy with which error correction can be achieved in certain circumstances. The essential idea is that, if some combination of code space and boundary region forms an α -bit code, but not an $(\alpha + \eta)$ -bit code, we can place a hard lower bound on the error of the α -bit encoding. Since an α -bit code can be viewed as a family of ordinary error correcting codes, one for every subspace of sufficiently small dimension, this also puts a lower bound on the error of the ordinary quantum error correcting codes associated with each of the subspaces.

The sudden emergence of previously tiny, non-perturbative errors to create dramatic, leading-order effects is reminiscent of the expected breakdown of Hawking’s black hole evaporation calculation at the Page time. Indeed we saw in section 3 that they are two examples of the same phenomenon. No ‘new’ effect causes the change in qualitative behaviour; instead the tiny corrections that have always been present build on top of each other more and more until they suddenly become significant.

We again consider a code space $\mathcal{H}_{\text{code}} \in \mathcal{H}_A \otimes \mathcal{H}_{\bar{A}}$ of states which all have a single black hole of fixed size in the centre of the AdS bulk. We shall use the same notation given in figure 4b. For simplicity we will assume that all states in the code space are identical in region \bar{a} . This avoids the need to deal with the details of subsystem error correction. However, the generalisation of the proof to subsystem or operator algebra error correction is straightforward. Let

$$\delta = \sup_{|\psi\rangle} \|\psi_{R\bar{A}} - \psi_R \otimes \omega_{\bar{A}}\|_1, \tag{8.1}$$

where ω is the maximally mixed state in the code space be the uncertainty in the forgetfulness for a reference system of fixed dimension d_R . We know that δ grows at most linearly with the dimension d_R of the reference system [17].

Let

$$\alpha < \frac{\mathcal{A}_2 - \mathcal{A}_1}{\mathcal{A}_0} < \alpha + \eta,$$

for some small η . Suppose that with $d_R = e^{\alpha S}$, where $S = \mathcal{A}_0/4G_N$ as usual, we have

$$\delta < e^{-\eta' S},$$

for some $\eta' > \eta$. Then, for $d_R = e^{(\alpha+\eta)S}$, we would have to have

$$\delta < e^{(\eta-\eta')S} \ll 1. \tag{8.2}$$

However this would mean that region A forms an $(\alpha + \eta)$ -bit code for region a' , which was shown to be impossible in section 4. It follows that for $d_R = e^{\alpha S}$ we must have

$$\delta \geq e^{-\eta' S}.$$

A lower bound on the trace distance can be converted into a lower bound on relative entropy using Pinsker's inequality. There must therefore exist a state $|\psi\rangle \in \mathcal{H}_{\text{code}} \otimes \mathcal{H}_R$ such that

$$S\left(\psi^{\bar{A}R} \parallel \omega^{\bar{A}} \otimes \psi^R\right) \geq \frac{1}{2} e^{-2\eta' S}. \quad (8.3)$$

Since the entanglement wedge of $\bar{A}R$ does not include the black hole for either state, the bulk relative entropy is zero. As a result this provides a strict upper bound on the accuracy of the equality between bulk and boundary relative entropies (4.6), even when the bulk relative entropy is zero.

Similarly, using (2.1), we see that this provides an upper bound on the accuracy of the error correction for the α -bit code of

$$\|\mathcal{D} \circ \text{Tr}_{\bar{A}} - \text{Id}\|_{\diamond} \geq \frac{1}{64} e^{-2\eta' S}. \quad (8.4)$$

Of course this argument, as we have given it, only puts a lower bound on the largest recovery error for any state in the space. In the Heisenberg picture, it lower bounds the worst-case error in operator recovery. That operator may be some complicated operator involving the black hole Hilbert space \mathcal{H}_{BH} .

However, even if we restrict ourselves to the algebra \mathcal{A} of bulk operators in region a' , but outside the black hole horizon, we can make an almost identical argument to the one above, by using theorem 2 from appendix A. The uncertainty $\delta_{\mathcal{A}}$ (defined formally in (A.13)) with which the algebra \mathcal{A} is forgotten by region \bar{A} still grows at most linearly with the allowed dimension of the reference system. Hence, if the algebra could be error corrected too accurately for all code subspaces of dimension $e^{\alpha S}$ (which would imply that $\delta_{\mathcal{A}}$ is similarly small, even for larger code spaces, so long as we restrict the reference dimension to be at most $e^{\alpha S}$), it could also be corrected for all larger spaces (which we know to be impossible) with the same bounds as above. The same argument can be made even for the algebra generated by any single operator. It follows that we get the same lower bound on the error for every operator in region a' , whether acting on the black hole or only acting in the bulk outside the horizon.

By either tuning the region A or adjusting the value of α , we can make the parameter η (and hence η') arbitrarily small. It follows that there exist cases where error correction is possible if $G_N \rightarrow 0$, but that at large but finite G_N the smallest achievable error is greater than $\exp(-\eta/G_N)$ for any arbitrarily small $\eta > 0$.²³ Similarly there exist states whose bulk relative entropy is zero, but whose boundary relative entropy is at least $\exp(-\eta/G_N)$ for any $\eta > 0$.

The equality between bulk and boundary relative entropies is not true to all orders in G_N ; however, a more precise equality was developed in [24] that is true to all orders in perturbation theory. If the bulk relative entropy is exactly zero, this reduces to the statement that the boundary relative entropy should be perturbatively equal to zero to all

²³Technically, the error should presumably smoothly interpolate between exponentially suppressed and order one as the dimension is increased. However, we do not have semiclassical control over the effects of such fine-grained changes in code space dimension.

orders in G_N . Even this more sophisticated equality was always expected to be broken by non-perturbative effects. However, here we have shown that such effects are absolutely required for the expected error-correction properties of AdS/CFT to hold — at least when a black hole is involved.

If there is no black hole, our arguments do not apply. However, a simple argument [45] based on the Reeh-Schlieder theorem [46] shows that it is still impossible for the error correction to be exact, even for a code space based on perturbations of the vacuum state $|\Omega\rangle$. Divide the boundary into four non-empty disjoint regions A, B, C and such that some bulk operator ϕ lies in the entanglement wedge of AB and the entanglement wedge of BC , but not in the entanglement wedge of B . There exists a boundary operators ϕ_{AB} and ϕ_{BC} acting on AB and BC respectively such that

$$\phi_{AB} |\Omega\rangle \simeq \phi |\Omega\rangle \simeq \phi_{BC} |\Omega\rangle . \tag{8.5}$$

Since ϕ does not lie in the entanglement wedge of B we know that $\phi_{AB} - \phi_{BC} \neq 0$. However, since D is non-empty, then, by the Reeh-Schlieder theorem, we know that

$$(\phi_{AB} - \phi_{BC}) |\Omega\rangle \neq 0 \tag{8.6}$$

and hence the equalities in (8.5) cannot be exact. We cannot ascribe the approximate nature of the error correction in AdS/CFT purely to the presence of particular black hole microstates. However it is only in the presence of a black hole that we are able to place lower bounds on the size of the error.

What causes these non-perturbative corrections? Clearly, it seems to have something to do with the existence of a closely competing candidate for the entanglement wedge. The exponents in the error bounds that we calculate are proportional to the difference between the Ryu-Takayanagi formula, evaluated for each of the two candidate wedges. This has the strong feel of a contribution from a subleading saddle point calculation. Unfortunately, it is not clear how this such contributions should be calculated, at least in the usual approach to deriving the RT formula [23, 47]. Such effects are, nonetheless, widely expected to exist (for example, in order to smooth out phase transitions in the entanglement entropy at finite N).

8.2 State dependence of the entanglement wedge

If we consider a code space $\mathcal{H}_{\text{code}}$ of perturbations about vacuum-AdS, the entanglement wedge of a region A of the boundary is state-independent in the semiclassical limit where gravity decouples and the background geometry is fixed. We can of course consider superpositions of states with different background classical geometries. However, since gravity decouples, these become non-interacting superselection sectors of the theory. The entanglement wedge is fixed for each classical geometry and hence does not depend on the state in any non-trivial sense.

In contrast, at finite coupling, the geometry is dynamical (and the bulk entanglement term in the RT formula may in principle compete with the area term). As a result, the entanglement wedge will in general depend on the state — the size of the fluctuations will be

approximately Planckian scale. We will be unable to reconstruct operators precisely unless they are separated from the entanglement wedge by a large distance in Planck units.

Alpha-bit codes for a black hole in the limit $G_N \rightarrow 0$ are in some sense an intermediate case: the geometry is fixed and smooth at arbitrarily small scales. However, large amounts of bulk entanglement can cause the entanglement wedge to be state-dependent. If we believe in the ER = EPR correspondence [48], this bulk entanglement may itself have a geometrical interpretation.

This provides arguably the most controlled setting in which the entanglement wedge is state-dependent. We do not have to deal with any issues of backreaction or fluctuations in the geometry. By making the coupling G_N small, we can make the size of the α -bit region (which was shown in section 5 to be approximately the AdS scale) arbitrarily large compared to the Planck scale where the classical geometry breaks down. In this way we are able to isolate the issues that arise from state dependence of the entanglement wedge, without having to deal with all the attendant issues of the classical geometry breaking down at the Planck scale.

Now consider the quantum error correction which exists in each of these three cases. In the strict $G_N \rightarrow 0$ limit and the absence of black holes, the quantum error correction approaches zero error. Conversely, at finite coupling, the error correction will always only be approximate. Just like when we looked at the fluctuations of the entanglement wedge, the α -bit encodings discussed in this paper form an interesting intermediate case. The actual error ε tends to zero for any decoding of allowed subspaces in the semiclassical limit. However, the phenomenon is only possible at all because the error is non-zero for any finite dimensional code space and finite coupling — it is still a phenomenon inherent only to approximate quantum error correction.

It seems likely that, if we want to truly understand entanglement wedge reconstruction in the context of a dynamical spacetime geometry, we may well have no choice but to think hard about issues specific to approximate quantum error correction.

8.3 State dependence of operators and the Papadodimas-Raju proposal

Of course, in α -bit codes, it is not only the entanglement wedge that is state-dependent; operators within the ‘ α -bit region’ (region a' in figure 5) are as well. If we wish to construct a decoding map $\mathcal{D} : S(A) \rightarrow S(a \otimes a')$ such that

$$\|\mathcal{D} \circ \text{Tr}_{\bar{A}}(\cdot) - \text{Tr}_{\bar{a}}(\cdot)\|_{\diamond} \tag{8.7}$$

is small and hence we can recover the state on $a \otimes a'$ from the state on A , we must first choose some code subspace of dimension less than d^α that we wish to decode — it is impossible to construct an approximate decoding channel that works for all states at once. In the Heisenberg picture, an operator $\phi_{a'}$ that is local to region a' can be approximately simulated by an operator $\mathcal{D}^\dagger(\phi_{a'})$ on the boundary that is local to A so long as we only require that this operator behave correctly when restricted to some code subspace of dimension less than d^α .

This state dependence of operators is unavoidable: there is no operator that will work for all states at once. This is in contrast to previous work on entanglement wedge recon-

struction where there was apparent state dependence because the method of constructing the operator made use of a choice of fixed state, but where this was merely an artifact of the construction. Boundary operators would still work for the entire code space, even if the boundary operator corresponding to a given bulk operator was not unique and so could have a definition that depended non-trivially on a choice of state.

One open question is to what extent similar state dependence exists for operators reconstructed on boundary subregions, even in the absence of a black hole, but at finite coupling. In this case fluctuations in the entanglement-wedge are only of order the Planck scale and as a result it is not even clear how well-defined it is to talk about bulk operators localised *within* such a region. Nonetheless, brushing aside such issues, one might speculate that it should again prove possible to reconstruct operators near the RT surface with greater accuracy if they only have to act correctly on a single state rather than an entire large code space of states. We will not, however, do so any more here.

More importantly, it has been argued, primarily by Papadodimas and Raju [10, 11], that local operators behind a black hole horizon are necessarily state-dependent. For some fixed choice of KMS-equilibrium black hole state $|\psi_0\rangle$ (essentially a pure state in the space $\mathcal{H}_{\text{BH}}^0$ that we defined in section 7), they use a variant of Tomita-Takesaki theory to construct ‘mirror operators’ that depend on $|\psi_0\rangle$ and that they claim describe the modes behind the black hole horizon. The mirror operators will behave correctly, so long as the state lies in the ‘small Hilbert space’ formed by applying simple operators to the chosen equilibrium black hole state $|\psi_0\rangle$.

These small Hilbert spaces form the equivalent of the decodable code subspaces in α -bit codes. Since their dimension is independent of G_N , in our language it is tempting to say that the Papadodimas-Raju proposal effectively gives a procedure for recovering the zero-bits of the interior of the black hole.²⁴

There has been considerable debate as to whether the Papadodimas-Raju proposal is consistent with a standard quantum mechanical interpretation of the bulk e.g. [49]. In our construction there are no such concerns: the state dependence is simply a phenomenon that arises from this particular form of quantum error correction and because we are restricting ourselves to a subregion of the boundary. If we have access to the entire boundary, the entire bulk (outside the horizon) can be understood without resorting to any state dependence beyond the basic geometry of the bulk itself. In contrast, in the Papadodimas-Raju proposal, the observer always has access to the entire boundary; in this sense there aren’t any errors to correct. If we accept the usual paradigm for error correction in AdS/CFT, where there is an isometry from a bulk code space to a larger boundary Hilbert space, then all operators, including interior operators should be reconstructable on the entire boundary in a state-independent way.

²⁴One apparent difference between the Papadodimas-Raju proposal and zero-bits is that the small Hilbert space is not a completely arbitrary choice of subspace. Instead we are free to choose the equilibrium state $|\psi_0\rangle$, and this choice determines the entire small Hilbert space. However, if we want to reconstruct simple interior operators then we clearly need those simple operators to preserve the code subspace. So this constraint on the small Hilbert space is simply a consequence of the fact that in α -bit codes it is only meaningful to reconstruct operators for a subspace that is preserved by those operators.

However there are important physical reasons, independent of Papadodimas-Raju construction itself, to think that bulk reconstruction is only possible for interior modes in a state-dependent way. Recent work [50] has shown that in a simple toy model of quantum gravity, known as the SYK model, there exists an (over)complete basis of black hole microstates for which simple excitations behind the horizon can be pulled out of the horizon by using a perturbed Hamiltonian which depends on the microstate in question. It follows that we can decode interior operators by evolving exterior operators with this perturbed Hamiltonian. However the reconstructed interior operator that one obtains will depend on the black hole microstate. Even more recently, it was argued in [51] that it should be possible in a general holographic CFT to use the mirror operators \tilde{O} to create a negative energy shockwave that will pull degrees of freedom outside the horizon — provided the state lies in the small Hilbert space defined above.

On the other hand, there clearly cannot be a single Hamiltonian that can pull out interior modes for *any* black hole microstate. If such a Hamiltonian existed, we could presumably use it to evolve one CFT in a thermofield double state and thereby create a traversable wormhole without having any interaction between the two CFTs. From this, it is very natural to conclude that the modes behind the horizon can only be defined on the boundary in a state-dependent way. The Papadodimas-Raju proposal suggests that only code subspaces of fixed dimension can be decoded — that perhaps *only* the zero-bits of the black hole interior are encoded in the boundary CFT.

It is tempting to go further and suggest that these two examples of state dependence — the α -bit codes discussed in this paper and the possible state dependence of operators in the black hole interior — result from the same basic mechanism. In the situations considered in this paper, the region a' where operators are state-dependent lies behind the causal (Rindler) horizon, just as the state dependence in the Papadodimas-Raju proposal appears for operators behind the black hole horizon. Furthermore, if we consider too large a code space of interior bulk states, it seems plausible that the RT surface might jump to near the horizon, for states sufficiently entangled with a reference system. Hence it would follow that interior modes could only be reconstructed for sufficiently small code subspaces, thus ‘deriving’ something like the Papadodimas-Raju proposal from the same sort of argument we have been making throughout this paper.

However, if we really take this argument seriously, we don’t reach quite the same conclusion as Papadodimas and Raju. For a single-sided pure black hole the true Ryu-Takayanagi surface for the entire boundary should trivially be empty since the boundary state is pure. Even if we have a significant amount of bulk entanglement between the region behind the horizon and a reference system, the entropy of the complete boundary CFT will simply be equal to the entropy of the reference system. In order to make the Ryu-Takayanagi surface lie directly on the black hole horizon we would need a reference system of dimension

$$d_R \sim e^S.$$

Hence bulk reconstruction for code spaces whose dimension

$$d_S \sim e^{\alpha S},$$

for any $\alpha < 1$, which would be far larger than the G_N -independent small Hilbert space in the Papadodimas-Raju proposal.

One speculative way to try to reconcile these differences would be to argue that we can include the mirror operators for an increasing number of modes (perhaps because of the reduced backreaction) as we take $G_N \rightarrow 0$, creating a Hilbert space whose dimension scales rapidly with $1/G_N$. An alternative approach, which is perhaps more compelling, would be to conclude that the Papadodimas-Raju proposal is correct but incomplete — that there also exist more complicated reconstructions which work for larger code subspaces. In particular it should be possible to combine the Papadodimas-Raju mirror operators \tilde{O}_S for a large number of orthogonal small Hilbert spaces \mathcal{H}_S with the boundary projectors P_S associated to each small Hilbert space to construct an operator

$$\tilde{O} = \sum_S \tilde{O}_S P_S,$$

that reproduces the bulk interior mode operator for a far larger class of states. This argument will break down when the small Hilbert spaces stop all being approximately orthogonal, but (by similar arguments to those in section 7) we might hope that this should not occur if we only consider $e^{(1-\delta)S}$ small Hilbert spaces for any fixed $\delta > 0$. (Note that even if you can decode operators directly behind the black hole for any code subspace of dimension $e^{(1-\delta)S}$ for arbitrary small $\delta > 0$, it is still not possible to decode operators directly behind the black hole for the thermofield double state, which, as we argued above, should clearly be impossible.²⁵)

Similarly, the method used in [50] to pull operators out from behind the horizon in the SYK model relies on evolving the system with a perturbed Hamiltonian precisely tuned to the black hole microstate in question. However, even if the Hamiltonian were only approximately tuned to the microstate (for example if only some sufficiently large fraction of the terms in the perturbed Hamiltonian were tuned correctly), it should still be possible to successfully pull the mode from behind the horizon. Hence a single Hamiltonian, and thus a single boundary reconstruction, can be used for a larger subspace of microstates. It would be very interesting to try to calculate exactly how large a subspace could successfully be reconstructed in this way.

As our understanding of AdS/CFT has developed, there has been a gradual de-emphasis of the idea of a unitary isomorphism between two distinct bulk and boundary Hilbert spaces, as in the traditional sense of a duality. Since most boundary states correspond to a large black hole with little geometrical bulk interpretation, it is instead more natural to think of an isometry from a code space of semiclassical geometrical bulk states to the larger boundary space.²⁶ By recognising that only a (relatively) small code space of

²⁵The idea that the α -bits are encoded for any $\alpha < 1$, but not for $\alpha = 1$, is strongly reminiscent of the discontinuity in the amortised α -bit capacity at $\alpha = 1$, see [12].

²⁶It is of course still possible that there exists a non-perturbative bulk description of the entire Hilbert space. In this case we would have a true duality between different theories and the code space isometries would simply come from the embedding of semiclassical sectors into the larger bulk Hilbert space. As yet, however, such a complete non-perturbative bulk description remains unknown and it is not clear that one must exist at all.

states have any given bulk geometry, apparent paradoxes where bulk operators appeared to commute with every local boundary operator were resolved [1]. However, if operators behind a black hole horizon are truly necessarily state-dependent — despite the fact that we have access to the entire boundary — then we cannot even have a simple isometry from bulk to boundary, since an isometry can always be error-corrected exactly without any need for state dependence.

One possibility is that we need to consider a general quantum channel, where bulk operators outside the horizon can lie in the image of the space of boundary operators under the adjoint channel, but bulk operators inside the horizon can only be simulated in a state-dependent way. This more general formalism was used in [30], although any physical implications of a noisy bulk-to-boundary channel were not discussed and the possibility was included mostly for the sake of completeness. However a noisy quantum channel does not seem to be quite the right object. It would suggest that pure bulk states should be associated to mixed boundary states. Instead almost the opposite seems to be the case, with multiple bulk descriptions corresponding to the same pure boundary state. The bulk-to-boundary map is more reminiscent of a linear map that is not always exactly isometric. Certainly, such a non-isometric linear map could appear very naturally in a tensor network model of AdS/CFT. For obvious physical reasons, generalisations of quantum error correction to such maps have not really been studied; they seem to be at least worth considering, if we want to understand holography.

There is one final connection between the Papadodimas-Raju proposal and more general zero-bit codes, which is too intriguing not to mention, but whose exact meaning and significance is somewhat unclear to the authors of this paper. The mirror operators in the Papadodimas-Raju proposal give an effective doubling of the CFT degrees of freedom, provided you only consider sufficiently simple operators. This doubling appears to correspond extremely elegantly with the fact that it is possible to encode zero-bits at an asymptotic rate of at most two zero-bits per qubit. In other words, if we only want the boundary state to encode the zero-bits of the bulk state, we are able to encode twice the number of degrees of freedom into the boundary state — potentially the modes both in front of and behind the horizon. Of course, as discussed above, the map from bulk to boundary does not quite seem to be a noisy quantum channel in the usual sense. Hence we cannot really think of mirror operators as giving a capacity-achieving zero-bit code. However the correspondence seems too tantalising to completely ignore.

8.4 Explicit reconstruction of state-dependent bulk operators

Until this point we have contented ourselves with showing the existence of decoding maps with the desired properties, we have not worried about *how* they should be constructed. Fortunately to a large degree this work has already been done for us. For any particular choice of subspace that we wish to decode, finding an explicit decoding map for an α -bit entanglement wedge is no different than the task of finding an explicit decoding map for ordinary entanglement wedge error correction, when the code space is given by the subspace we wish to decode. This task has been achieved recently in [30, 52].

In [30], it was shown that operators in the entanglement wedge can be reconstructed on the boundary using a universal decoding map, built from the twirled Petz map $\mathcal{R}_{\sigma, \mathcal{N}}$, developed in [53]. It is defined for some fixed state σ , which we shall take to be the maximally mixed state on the subspace we wish to decode, and encoding channel \mathcal{N} , which here is simply the partial trace over \bar{A} , by

$$\mathcal{R}_{\sigma, \mathcal{N}} = \int_{\mathbb{R}} dt \beta_0(t) \sigma^{\frac{1}{2}(1-it)} \mathcal{N}^\dagger \left(\mathcal{N}(\sigma)^{-\frac{1}{2}(1-it)} (\cdot) \mathcal{N}(\sigma)^{-\frac{1}{2}(1+it)} \right) \sigma^{\frac{1}{2}(1+it)}, \quad (8.8)$$

where $\beta_0(t) := \frac{\pi}{2}(\cosh(\pi t) + 1)^{-1}$. The argument in [30], like the arguments made here and in [9], relies on the approximate equality between bulk and boundary relative entropies. However, in this case, the relative entropies under consideration are associated to the boundary region A , used to decode the state, rather than its complement \bar{A} . For this argument to go through, we need the bulk region we wish to reconstruct to be contained in the entanglement wedge of A for all states (pure or mixed) in the space we wish to decode. This is exactly equivalent, as we would hope, to the bulk region not being contained in the entanglement wedge of $\bar{A}R$ for any pure state that may be entangled with the reference system R . It follows that all the same arguments used in section 4 still apply and the twirled Petz map can be used to decode correctable subspaces in the α -bit codes. In practice unfortunately the twirled Petz map is hard to evaluate even for simple bulk geometries — trying to do so for a particular subspace of black hole microstates is likely to be essentially impossible. Furthermore, since the argument in [30] is based on the equality of bulk and boundary relative entropies for bulk states that are not identical either within or outside the entanglement wedge of A , it is unclear whether this recovery map is exact to all orders in perturbation theory.

In [52] a different explicit formula for entanglement wedge reconstruction was developed by relating the bulk and boundary modular flows (the evolution of operators using the modular Hamiltonian $K = -\log \rho$). The basic extension of the extrapolate dictionary to bulk and boundary modular evolutions of operators developed in [52] should continue to apply when the bulk contains a black hole. However, because the bulk no longer consists only of free fields to leading order (the reconstruction needs to also somehow depend on the state of the black hole), the bulk modular-evolved fields will no longer be simply a Bogoliubov transformation of the original fields. This was the key property used in [52] to argue that they formed a natural basis to attempt entanglement wedge reconstruction.

It should be essentially unsurprising that explicit construction of the state dependent operators is highly challenging in practice. The whole point of the requirement of state dependence is that we are unable to reconstruct the operators without taking advantage of our detailed knowledge of the subspace of black hole states in question. In section 6 we saw that in tensor network toy models this involves literally pushing the operator through the black hole; all the details of the construction of the boundary operator depend on the black hole states in question. Nonetheless it is comforting that there exist an explicit, even if impractical, reconstruction procedure.

8.5 Black holes are α -bit sup?

We argued in section 3 that the α -bits of a black hole are encoded in any fraction greater than

$$p = \frac{1 + \alpha}{2}$$

of its Hawking radiation so long as the evaporation is approximately thermodynamically reversible. The α -bit capacity of the noiseless qubit channel [12] is equal to

$$\frac{2}{1 + \alpha}.$$

It follows that Hawking radiation is a *capacity-achieving* α -bit encoding of the black hole state.

Although less obvious at first glance, this continues to be true in the ‘entanglement-assisted’ case, when some Hawking radiation (with entropy βS) has been emitted before the unknown black hole is dropped in: in this case we require a fraction greater than

$$p = \frac{1 + \alpha - \beta}{2} \tag{8.9}$$

of the Hawking radiation in order to decode a subspace of dimension $\exp(\alpha S)$. The ‘entanglement-assistance’ has reduced the fraction of the Hawking radiation that is required to decode a subspace of the same size, so this might seem to exceed the α -bit capacity. However, Bob had to know the state of the original black hole, which must have had entropy of at least βS in order to be so entangled with the original Hawking radiation. Even if the remaining ‘unknown’ entropy of the black hole were as large as possible, it would only be equal to $(1 - \beta)S$. As a result, Bob is effectively only recovering the α' -bits of this smaller unknown system for

$$\alpha' = \frac{\alpha}{1 - \beta}. \tag{8.10}$$

He can decode the α' -bits of a system with entropy $(1 - \beta)S$ using Hawking radiation with entropy only

$$pS = \frac{1 + \alpha'}{2}(1 - \beta)S. \tag{8.11}$$

In other words, we find that the α' -bits of the unknown black hole are still being encoded at the noiseless α' -bit capacity rate of $2/(1 + \alpha')$.

If the black hole evaporation is significantly thermodynamically irreversible, it no longer saturates the α -bit capacities. However, this inefficiency is an inevitable consequence of thermodynamic entropy being produced. As we saw in (3.3), the black hole still reveals its α -bits in the Hawking radiation as soon as the Hawking radiation contains sufficient entropy to both purify the black hole and encode the α -bits of the initial state

Similarly, in section 5, we found that the α -bits of a large (in AdS units) BTZ black hole are encoded holographically in any fraction of the boundary greater than

$$\frac{1 + \alpha}{2}.$$

This is in fact also true in higher dimensions, since we can approximate minimal surfaces around a large black hole by the minimal surface connecting the horizon to the boundary, together with part of the horizon itself. Corrections to this approximation will only be of order the AdS scale and so are negligible in the limit of large horizon area. As a result they have no effect on α to leading order in L/R .

If we impose a UV cut-off on the CFT, which cuts off the bulk close to the horizon area, we should expect the effective dimension of the CFT to be only of order e^S . As a result, the boundary subregion forms a capacity-achieving encoding of the black hole state.

Because black holes seem to achieve the α -bit capacity whenever possible, we say that black holes are ‘ α -bit sup’. As discussed in section 3, the subspace decoupling duality shows that anything which looks ‘as thermal as possible’ when looking at less than half of its degrees of freedom is necessarily also α -bit sup when one has access to more than half of the degrees of freedom. It should therefore not surprise us at all that black holes tend to satisfy it. However, to a physicist used to thinking of black holes as hiding information whenever possible, it may seem counterintuitive that we can equivalently think of them as revealing information (specifically their α -bits) as soon as possible.

At this point, there exist almost no known explicit constructions of capacity-achieving α -bit codes;²⁷ it has only been possible to show that they exist by considering randomly chosen unitaries. It is questionable exactly how much more explicit it is to define a space as ‘a subspace of black hole states in AdS/CFT’ than to define it as ‘a subspace that is k -forgetful to the environment (which provably must exist)’, but it is amusing to note from a quantum information perspective that black holes therefore form the first example of an explicit capacity-achieving α -bit code for a noiseless quantum channel (for a forgiving definition of explicit).

Acknowledgments

We would like to thank Ahmed Almheiri, Daniel Harlow, Matthew Headrick, Isaac Kim, Aitor Lewkowycz, Sepehr Nezami, Phil Saad, Grant Salton, Steve Shenker, Lenny Susskind, Aron Wall, Michael Walter and Edward Witten for valuable discussions. In particular, we would like to thank Jonathan Sorce for discussions on the contents of the acknowledgements section. This work was supported by AFOSR (FA9550-16-1-0082), CIFAR and the Simons Foundation.

A Operator algebra quantum error correction and the information disturbance tradeoff

The framework of operator algebra quantum error correction was first introduced in [55]; it is a generalisation of the notion of subsystem quantum error correction that is natural in holography because it allows one to talk about superpositions of different geometries in AdS/CFT [1, 56]. Even though we make essentially no use of the formalism in the rest

²⁷The only known example is for the zero-bit capacity of the noiseless cbit channel [54].

of the paper, we include a brief review here, mainly focussed on providing a statement of (and also something of the context behind) theorem 2, which was first proved in [15].²⁸

Theorem 2, applied to the special case of a subsystem error correcting code as (4.13), lies at the heart of the results of section 4. We include the operator algebraic version here because it has not appeared previously in the literature on AdS/CFT and is the natural generalisation of the Dong-Harlow-Wall condition [9] to approximate reconstruction. Using the argument from [24] that the boundary relative entropy is zero to all orders in perturbation theory if the bulk states agree exactly, theorem 2 is sufficient to prove that entanglement wedge reconstruction can be made exact to all orders in $1/N$. A version of the proof assuming the existence of a tensor product factorization was given in section 4.1. We give a more general argument from an algebraic perspective here, although we still need to assume that all the Hilbert spaces involved are finite-dimensional.

So long as the dimension of the code space grows at most polynomially with N or $1/G_N$, there is no need to consider a reference system and the argument can proceed almost identically to the argument based on assuming exact error correction given in [9]. However, as we have seen in this paper, the argument completely fails for larger code spaces unless we additionally consider states entangled with a reference system. Explicitly considering the approximate case therefore provides important non-trivial insight.

A quantum channel $\mathcal{N} : S(\mathcal{H}_1) \rightarrow S(\mathcal{H}_2)$ is a completely positive trace-preserving linear superoperator and hence maps density matrices to density matrices. The adjoint channel $\mathcal{N}^\dagger : \mathcal{L}(\mathcal{H}_2) \rightarrow \mathcal{L}(\mathcal{H}_1)$ is defined by

$$\text{Tr}(\mathcal{N}^\dagger(X)\rho) = \text{Tr}(X\mathcal{N}(\rho)) \tag{A.1}$$

and hence is a unital, trace-preserving map from operators on \mathcal{H}_2 to operators on \mathcal{H}_1 . When working in the Heisenberg picture, this is sometimes given as the definition of a quantum channel.

The diamond norm between two channels $\mathcal{N}_1, \mathcal{N}_2$ is defined by

$$\|\mathcal{N}_1 - \mathcal{N}_2\|_\diamond = \sup_{\|\rho\|_1 \leq 1} \|[(\mathcal{N}_1 - \mathcal{N}_2) \otimes \text{Id}_R] \rho\|_1, \tag{A.2}$$

where we can take the dimension of the reference Hilbert space \mathcal{H}_R to be equal to the dimension of \mathcal{H}_1 . Equivalently

$$\|\mathcal{N}_1 - \mathcal{N}_2\|_\diamond = \sup_{X \neq 0} \frac{\|[(\mathcal{N}_1^\dagger - \mathcal{N}_2^\dagger) \otimes \text{Id}_R] X\|_\infty}{\|X\|_\infty}, \tag{A.3}$$

where the supremum is taken over non-zero operators X acting on $\mathcal{H}_2 \otimes \mathcal{H}_R$.

A finite-dimensional von Neumann algebra $\mathcal{A} \subseteq \mathcal{L}(\mathcal{H}_1)$ is a subset of linear operators that contains the identity and is closed under multiplication, scalar multiplication, addition and Hermitian conjugation. It can be shown that for such an algebra, there exists an orthogonal decomposition of the Hilbert space

$$\mathcal{H}_1 = \bigoplus_i \mathcal{H}_{A_i} \otimes \mathcal{H}_{\bar{A}_i}$$

²⁸For earlier related results, see [58].

such that the algebra \mathcal{A} consists of all operators of the form

$$X = \sum_i X_i \otimes \mathbb{1}_{\bar{A}_i}$$

for $X_i \in \mathcal{L}(\mathcal{H}_{A_i})$. If P_i is the projector onto $\mathcal{H}_{A_i} \otimes \mathcal{H}_{\bar{A}_i}$, we define the projector $\mathcal{P}_{\mathcal{A}}$ onto the algebra to be the quantum channel

$$\mathcal{P}_{\mathcal{A}}(\rho) = \sum_i \frac{1}{d_{\bar{A}_i}} \text{Tr}_{\bar{A}_i}(P_i \rho P_i) \otimes \mathbb{1}_{\bar{A}_i}, \quad (\text{A.4})$$

Note that $\mathcal{P}_{\mathcal{A}}^2 = \mathcal{P}_{\mathcal{A}} = \mathcal{P}_{\mathcal{A}}^\dagger$. $\mathcal{P}_{\mathcal{A}}$ is therefore both trace-preserving and unital; it is the quantum channel in both the Schrödinger and Heisenberg pictures. For notational convenience we write $\rho_{\mathcal{A}} = \mathcal{P}_{\mathcal{A}}(\rho)$. The image of the channel $\mathcal{P}_{\mathcal{A}}$ (when acting on density matrices) is the intersection of the algebra \mathcal{A} with space of density matrices $S(\mathcal{H})$ and is canonically isomorphic to the space of positive normalised linear functionals on \mathcal{A} , which is the standard abstract definition of states on a von Neumann algebra.

We say that the pair of channels $\mathcal{N} : S(\mathcal{H}_1) \rightarrow S(\mathcal{H}_2)$ and $\mathcal{D} : S(\mathcal{H}_2) \rightarrow S(\mathcal{H}_1)$ form an exact quantum error correcting code for the finite-dimensional von Neumann algebra \mathcal{A} if

$$\mathcal{D} \circ \mathcal{N} = \mathcal{P}_{\mathcal{A}}. \quad (\text{A.5})$$

Note this implies that for all $X \in \mathcal{A}$ then

$$(\mathcal{D} \circ \mathcal{N})^\dagger X = \mathcal{P}_{\mathcal{A}}(X) = X. \quad (\text{A.6})$$

To make this approximate, we simply allow some small separation $\varepsilon_{\mathcal{A}}$ in terms of the diamond norm

$$\|\mathcal{D} \circ \mathcal{N} - \mathcal{P}_{\mathcal{A}}\|_{\diamond} \leq \varepsilon_{\mathcal{A}}. \quad (\text{A.7})$$

Define the von Neumann algebra \mathcal{A}' to be the set of operators that commute with all operators in \mathcal{A} (this is known as the commutant of \mathcal{A}). These are the operators of the form

$$X' = \sum_i \mathbb{1}_{A_i} \otimes X'_i \quad (\text{A.8})$$

for $X'_i \in \mathcal{L}(\mathcal{H}_{\bar{A}_i})$. By considering the Stinespring dilation $V : \mathcal{H}_1 \rightarrow \mathcal{H}_1 \otimes \mathcal{H}_1$ defined by

$$V \left(|\psi_i\rangle^{A_i} \otimes |\phi_i\rangle^{\bar{A}_i} \right) = |\psi_i\rangle_{A_i} \otimes |\omega_i\rangle_{\bar{A}_i A_i} \otimes |\phi_i\rangle_{\bar{A}_i} \quad (\text{A.9})$$

for some maximally entangled state $|\omega_i\rangle \in \mathcal{H}_{\bar{A}_i} \otimes \mathcal{H}_{A_i}$, we see that $\mathcal{P}_{\mathcal{A}'}$ is the complementary channel to $\mathcal{P}_{\mathcal{A}}$.

This relationship allows the application of the following theorem [57], defined in terms of the fidelity between channels

$$F(\mathcal{N}, \mathcal{M}) = \min_{\rho} F((\mathcal{N} \otimes \text{Id}) \rho, (\mathcal{M} \otimes \text{Id}) \rho) \quad (\text{A.10})$$

where $F(\rho, \sigma)$ is the usual fidelity between states.

Theorem 1.

$$\max_{\mathcal{R}} F(\mathcal{R} \circ \mathcal{N}, \mathcal{M}) = \max_{\mathcal{R}'} F(\mathcal{N}^c, \mathcal{R}' \circ \mathcal{M}^c) \quad (\text{A.11})$$

where \mathcal{Q}^c is the complementary channel to \mathcal{Q} .

If we take the channel \mathcal{M} to be the projector $\mathcal{P}_{\mathcal{A}}$ onto the algebra and convert the statements about fidelities into statements about distances with respect to the diamond norm, we obtain the following result, which was first proved in [15].

Theorem 2. *Let*

$$\varepsilon_{\mathcal{A}} = \min_{\mathcal{D}} \|\mathcal{D} \circ \mathcal{N} - \mathcal{P}_{\mathcal{A}}\|_{\diamond} \quad (\text{A.12})$$

and let

$$\delta_{\mathcal{A}} = \|\mathcal{N}^c - \mathcal{N}^c \circ \mathcal{P}_{\mathcal{A}'}\|_{\diamond}. \quad (\text{A.13})$$

Then

$$\frac{1}{4}\delta_{\mathcal{A}}^2 \leq \varepsilon_{\mathcal{A}} \leq 2\delta_{\mathcal{A}}^{\frac{1}{2}}. \quad (\text{A.14})$$

For the purposes of this paper, we are mostly content to consider the case of a single geometry, which could treat as subsystem quantum error correction without causing serious problems.²⁹ This corresponds to the special case of operator algebra error correction where the centre of the algebra is trivial and hence the Hilbert space

$$\mathcal{H}_1 = \mathcal{H}_a \otimes \mathcal{H}_{\bar{a}}.$$

The algebra \mathcal{A} is simply operators on \mathcal{H}_a while the commutant \mathcal{A}' consists of operators on $\mathcal{H}_{\bar{a}}$. The projectors become

$$\mathcal{P}_{\mathcal{A}} = \text{Tr}_{\bar{a}}(\cdot) \otimes \omega^{\bar{a}} \quad \text{and} \quad \mathcal{P}_{\mathcal{A}'} = \omega^a \otimes \text{Tr}_a(\cdot),$$

where ω is maximally mixed in both cases. It follows that

$$\varepsilon_{\mathcal{A}} = \min_{\mathcal{D}} \|\mathcal{D} \circ \mathcal{N} - \text{Tr}_{\bar{a}}(\cdot) \otimes \omega^{\bar{a}}\|_{\diamond} = \min_{\mathcal{D}} \|\mathcal{D} \circ \mathcal{N} - \text{Tr}_{\bar{a}}(\cdot)\|_{\diamond}.$$

If, as in section 4, we define \mathcal{N} to be the restriction of $\text{Tr}_{\bar{A}}(\cdot)$ to $S(\mathcal{H}_{\text{code}})$ (and hence \mathcal{N}^c is the restriction of $\text{Tr}_A(\cdot)$ to $S(\mathcal{H}_{\text{code}})$), then $\varepsilon_{\mathcal{A}}$ and $\delta_{\mathcal{A}}$ become δ_1 and δ_2 , as defined in (4.8) and (4.9), respectively. We see that theorem 2 becomes (4.13).

As we saw throughout this paper, for large code spaces the fact $\delta_{\mathcal{A}}$ is defined using the diamond norm (and hence that we have to consider states entangled with a reference

²⁹In reality, all the statements in section 4 are more precisely interpreted as statements about operator algebras; the use of the subsystem error correction paradigm was purely for pedagogical reasons. Since nothing in section 4 relied on the (false) assumption that the algebras had trivial centres, the entirety of section 4 can, if desired, be converted back into the (less sloppy) language of operator algebras using the dictionary given below.

system) is crucially important. However, if the dimension d of the code space is fixed in the limit $G_N \rightarrow 0$, we can use the bound [17]

$$\|\mathcal{N} - \mathcal{M}\|_{\diamond} \leq d \sup_{|\psi\rangle \in \mathcal{H}_{\text{code}}} \|\mathcal{N}(\psi) - \mathcal{M}(\psi)\|_1 \quad (\text{A.15})$$

to avoid considering any sort of reference system at all.

It was shown in [24] that for any two, up to non-perturbative corrections, the boundary relative entropy,

$$S(\rho_{\bar{A}}|\sigma_{\bar{A}}) = \langle A^{X_{\sigma}} + K_{\text{bulk},\sigma} \rangle_{\rho} - \langle A^{X_{\rho}} + K_{\text{bulk},\rho} \rangle_{\rho}, \quad (\text{A.16})$$

where A^X is the area of the RT surface, K_{bulk} is the bulk relative entropy for region \bar{a} and we have ignored higher curvature corrections (we assume $\lambda \rightarrow \infty$). Since the right hand side of (A.16) depends only on the restriction of the bulk states to algebra \mathcal{A}' ,

$$S(\rho_{\bar{A}}|(\rho_{\mathcal{A}'})_{\bar{A}}) \leq \varepsilon, \quad (\text{A.17})$$

for some ε that is non-perturbatively small in G_N .³⁰ As usual, we convert the relative entropy into a trace distance using Pinsker's inequality, which gives

$$\|(\mathcal{N}^c - \mathcal{N}^c \circ \mathcal{P}_{\mathcal{A}'}) \rho\|_1 \leq \sqrt{2\varepsilon \ln 2}, \quad (\text{A.18})$$

which is also non-perturbatively small. If the dimension d of the code space is held fixed (or even grows polynomially in $1/G_N$ or N), then $\delta_{\mathcal{A}}$ and $\varepsilon_{\mathcal{A}}$ will also be non-perturbatively small and the error correction will be exact to all orders in perturbation theory

As a final technical comment, when working perturbatively to higher orders in G_N , we need to take into account the state dependence of the entanglement wedge \bar{a} . At finite G_N , different states therefore have different algebras \mathcal{A} and \mathcal{A}' associated to the same boundary regions A and \bar{A} . To fix this we should instead work with the bulk algebra \mathcal{A}'_0 associated to the *union* of the entanglement wedge \bar{a} for all states in the code space. Then for any algebra \mathcal{A}' associated to the entanglement wedge \bar{a} of a particular state ρ ,

$$\mathcal{P}_{\mathcal{A}'} \mathcal{P}_{\mathcal{A}'_0} = \mathcal{P}_{\mathcal{A}'}, \quad (\text{A.19})$$

and hence

$$\left\| \mathcal{N}^c \circ (\text{Id} - \mathcal{P}_{\mathcal{A}'_0}) \rho \right\|_1 \leq \left\| \mathcal{N}^c \circ (\text{Id} - \mathcal{P}_{\mathcal{A}'}) \rho \right\|_1 + \left\| \mathcal{N}^c \circ (\text{Id} - \mathcal{P}_{\mathcal{A}'}) (\mathcal{P}_{\mathcal{A}'_0}(\rho)) \right\|_1 \leq \sqrt{8\varepsilon \ln 2} \quad (\text{A.20})$$

where we have used the fact that both ρ and $\mathcal{P}_{\mathcal{A}'_0}(\rho)$ have the same entanglement wedge and hence satisfy (A.18) for the same algebra \mathcal{A}' .

It follows that only the algebra \mathcal{A}_0 associated to the *intersection* of the entanglement wedge a for all states can be decoded from region A . In the limit $G_N \rightarrow 0$, the entanglement wedge for every state converges to the classical entanglement wedge with no backreaction and hence the algebra \mathcal{A}_0 becomes the algebra associated to the classical region a .

³⁰To be as rigorous as possible in our construction we should acknowledge that the boundary theory also does not actually factorise into Hilbert spaces on A and \bar{A} and hence \mathcal{N} and \mathcal{N}^c are really projectors onto the algebras of regions A and \bar{A} respectively, rather than partial traces onto subsystems. However, this makes no difference to the argument.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY 4.0](https://creativecommons.org/licenses/by/4.0/)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] A. Almheiri, X. Dong and D. Harlow, *Bulk Locality and Quantum Error Correction in AdS/CFT*, *JHEP* **04** (2015) 163 [[arXiv:1411.7041](https://arxiv.org/abs/1411.7041)] [[INSPIRE](#)].
- [2] D. Harlow, *The Ryu-Takayanagi Formula from Quantum Error Correction*, *Commun. Math. Phys.* **354** (2017) 865 [[arXiv:1607.03901](https://arxiv.org/abs/1607.03901)] [[INSPIRE](#)].
- [3] T. Hirata and T. Takayanagi, *AdS/CFT and strong subadditivity of entanglement entropy*, *JHEP* **02** (2007) 042 [[hep-th/0608213](https://arxiv.org/abs/hep-th/0608213)] [[INSPIRE](#)].
- [4] M. Van Raamsdonk, *Building up spacetime with quantum entanglement*, *Gen. Rel. Grav.* **42** (2010) 2323 [[arXiv:1005.3035](https://arxiv.org/abs/1005.3035)] [[INSPIRE](#)].
- [5] E. Verlinde and H. Verlinde, *Black Hole Entanglement and Quantum Error Correction*, *JHEP* **10** (2013) 107 [[arXiv:1211.6913](https://arxiv.org/abs/1211.6913)] [[INSPIRE](#)].
- [6] S. Lloyd and J. Preskill, *Unitarity of black hole evaporation in final-state projection models*, *JHEP* **08** (2014) 126 [[arXiv:1308.4209](https://arxiv.org/abs/1308.4209)] [[INSPIRE](#)].
- [7] B. Yoshida and A. Kitaev, *Efficient decoding for the Hayden-Preskill protocol*, [arXiv:1710.03363](https://arxiv.org/abs/1710.03363) [[INSPIRE](#)].
- [8] 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](https://arxiv.org/abs/1503.06237)] [[INSPIRE](#)].
- [9] 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](https://arxiv.org/abs/1601.05416)] [[INSPIRE](#)].
- [10] K. Papadodimas and S. Raju, *An Infalling Observer in AdS/CFT*, *JHEP* **10** (2013) 212 [[arXiv:1211.6767](https://arxiv.org/abs/1211.6767)] [[INSPIRE](#)].
- [11] K. Papadodimas and S. Raju, *State-Dependent Bulk-Boundary Maps and Black Hole Complementarity*, *Phys. Rev. D* **89** (2014) 086010 [[arXiv:1310.6335](https://arxiv.org/abs/1310.6335)] [[INSPIRE](#)].
- [12] P. Hayden and G. Penington, *Approximate quantum error correction revisited: Introducing the alpha-bit*, [arXiv:1706.09434](https://arxiv.org/abs/1706.09434).
- [13] N. Engelhardt and A.C. Wall, *Quantum Extremal Surfaces: Holographic Entanglement Entropy beyond the Classical Regime*, *JHEP* **01** (2015) 073 [[arXiv:1408.3203](https://arxiv.org/abs/1408.3203)] [[INSPIRE](#)].
- [14] P. Hayden and J. Preskill, *Black holes as mirrors: Quantum information in random subsystems*, *JHEP* **09** (2007) 120 [[arXiv:0708.4025](https://arxiv.org/abs/0708.4025)] [[INSPIRE](#)].
- [15] C. Bény, *Conditions for the approximate correction of algebras*, in proceedings of the *Theory of quantum computation, communication and cryptography: 4th workshop, TQC 2009*, Waterloo, Canada, 11–13 May 2009, Springer, *Lect. Notes Comput. Sci.* **5906** (2009) 66.

- [16] A. Winter, *Quantum and classical message identification via quantum channels*, in *Festschrift “A.S. Holevo 60”*, O. Hirota ed., Rinton Press (2004), pp. 171–188, reprinted in *Quant. Inf. Comput.* **4** (2004) 563 [[quant-ph/0401060](#)].
- [17] P. Hayden and A. Winter, *Weak decoupling duality and quantum identification*, *IEEE Trans. Inf. Theory* **58** (2012) 4914.
- [18] D. Kretschmann and R.F. Werner, *Tema con variazioni: quantum channel capacity*, *New J. Phys.* **6** (2004) 26.
- [19] D. Kretschmann, D. Schlingemann and R.F. Werner, *The information-disturbance tradeoff and the continuity of Stinespring’s representation*, [quant-ph/0605009](#).
- [20] D.N. Page, *Time Dependence of Hawking Radiation Entropy*, *JCAP* **09** (2013) 028 [[arXiv:1301.4995](#)] [[INSPIRE](#)].
- [21] P. Hayden, S. Nezami, X.-L. Qi, N. Thomas, M. Walter and Z. Yang, *Holographic duality from random tensor networks*, *JHEP* **11** (2016) 009 [[arXiv:1601.01694](#)] [[INSPIRE](#)].
- [22] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [[hep-th/0603001](#)] [[INSPIRE](#)].
- [23] T. Faulkner, A. Lewkowycz and J. Maldacena, *Quantum corrections to holographic entanglement entropy*, *JHEP* **11** (2013) 074 [[arXiv:1307.2892](#)] [[INSPIRE](#)].
- [24] X. Dong and A. Lewkowycz, *Entropy, Extremality, Euclidean Variations and the Equations of Motion*, *JHEP* **01** (2018) 081 [[arXiv:1705.08453](#)] [[INSPIRE](#)].
- [25] V.E. Hubeny, M. Rangamani and T. Takayanagi, *A Covariant holographic entanglement entropy proposal*, *JHEP* **07** (2007) 062 [[arXiv:0705.0016](#)] [[INSPIRE](#)].
- [26] 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](#)] [[INSPIRE](#)].
- [27] M. Headrick, V.E. Hubeny, A. Lawrence and M. Rangamani, *Causality & holographic entanglement entropy*, *JHEP* **12** (2014) 162 [[arXiv:1408.6300](#)] [[INSPIRE](#)].
- [28] A.C. Wall, *Maximin Surfaces and the Strong Subadditivity of the Covariant Holographic Entanglement Entropy*, *Class. Quant. Grav.* **31** (2014) 225007 [[arXiv:1211.3494](#)] [[INSPIRE](#)].
- [29] D.L. Jafferis, A. Lewkowycz, J. Maldacena and S.J. Suh, *Relative entropy equals bulk relative entropy*, *JHEP* **06** (2016) 004 [[arXiv:1512.06431](#)] [[INSPIRE](#)].
- [30] 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](#)].
- [31] F. Hiai, M. Ohya and M. Tsukada, *Sufficiency, KMS condition and relative entropy in von Neumann algebras*, in *Selected Papers of M. Ohya*, World Scientific (2008), pp. 420–430.
- [32] S.H. Shenker and D. Stanford, *Black holes and the butterfly effect*, *JHEP* **03** (2014) 067 [[arXiv:1306.0622](#)] [[INSPIRE](#)].
- [33] N. Bao and H. Ooguri, *Distinguishability of black hole microstates*, *Phys. Rev. D* **96** (2017) 066017 [[arXiv:1705.07943](#)] [[INSPIRE](#)].
- [34] Y. Sekino and L. Susskind, *Fast Scramblers*, *JHEP* **10** (2008) 065 [[arXiv:0808.2096](#)] [[INSPIRE](#)].

- [35] B. Swingle, *Entanglement Renormalization and Holography*, *Phys. Rev. D* **86** (2012) 065007 [[arXiv:0905.1317](#)] [[INSPIRE](#)].
- [36] B. Czech, L. Lamprou, S. McCandlish and J. Sully, *Tensor Networks from Kinematic Space*, *JHEP* **07** (2016) 100 [[arXiv:1512.01548](#)] [[INSPIRE](#)].
- [37] E.P. Verlinde, *Emergent Gravity and the Dark Universe*, *SciPost Phys.* **2** (2017) 016 [[arXiv:1611.02269](#)] [[INSPIRE](#)].
- [38] T.J. Osborne and D.E. Stiegemann, *Dynamics for holographic codes*, [arXiv:1706.08823](#) [[INSPIRE](#)].
- [39] S. Popescu, A.J. Short and A. Winter, *Entanglement and the foundations of statistical mechanics*, *Nat. Phys.* **2** (2006) 754.
- [40] J.M. Deutsch, *Quantum statistical mechanics in a closed system*, *Phys. Rev. A* **43** (1991) 2046.
- [41] M. Srednicki, *Chaos and quantum thermalization*, *Phys. Rev. E* **50** (1994) 888.
- [42] D. Marolf and J. Polchinski, *Gauge/Gravity Duality and the Black Hole Interior*, *Phys. Rev. Lett.* **111** (2013) 171301 [[arXiv:1307.4706](#)] [[INSPIRE](#)].
- [43] R. Haag, N.M. Hugenholtz and M. Winnink, *On the Equilibrium states in quantum statistical mechanics*, *Commun. Math. Phys.* **5** (1967) 215 [[INSPIRE](#)].
- [44] R. Kubo, *Statistical-mechanical theory of irreversible processes. I. General theory and simple applications to magnetic and conduction problems*, *J. Phys. Soc. Jap.* **12** (1957) 570 [[INSPIRE](#)].
- [45] W.R. Kelly, *Bulk Locality and Entanglement Swapping in AdS/CFT*, *JHEP* **03** (2017) 153 [[arXiv:1610.00669](#)] [[INSPIRE](#)].
- [46] H. Reeh and S. Schlieder, *Bemerkungen zur Unitäräquivalenz von Lorentzinvarianten Feldern*, *Nuovo Cim.* **22** (1961) 1051 [[INSPIRE](#)].
- [47] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, *JHEP* **08** (2013) 090 [[arXiv:1304.4926](#)] [[INSPIRE](#)].
- [48] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, *Fortsch. Phys.* **61** (2013) 781 [[arXiv:1306.0533](#)] [[INSPIRE](#)].
- [49] D. Harlow, *Aspects of the Papadodimas-Raju Proposal for the Black Hole Interior*, *JHEP* **11** (2014) 055 [[arXiv:1405.1995](#)] [[INSPIRE](#)].
- [50] I. Kourkoulou and J. Maldacena, *Pure states in the SYK model and nearly-AdS₂ gravity*, [arXiv:1707.02325](#) [[INSPIRE](#)].
- [51] J. de Boer, R. van Breukelen, S.F. Lokhande, K. Papadodimas and E. Verlinde, *On the interior geometry of a typical black hole microstate*, *JHEP* **05** (2019) 010 [[arXiv:1804.10580](#)] [[INSPIRE](#)].
- [52] T. Faulkner and A. Lewkowycz, *Bulk locality from modular flow*, *JHEP* **07** (2017) 151 [[arXiv:1704.05464](#)] [[INSPIRE](#)].
- [53] 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 [[arXiv:1509.07127](#)] [[INSPIRE](#)].

- [54] O. Fawzi, P. Hayden and P. Sen, *From low-distortion norm embeddings to explicit uncertainty relations and efficient information locking*, *J. ACM* **60** (2013) 44.
- [55] C. Bény, A. Kempf and D.W. Kribs, *Generalization of quantum error correction via the Heisenberg picture*, *Phys. Rev. Lett.* **98** (2007) 100502.
- [56] C. Bény, Z. Zimborás and F. Pastawski, *Approximate recovery with locality and symmetry constraints*, [arXiv:1806.10324](#) [INSPIRE].
- [57] C. Bény and O. Oreshkov, *General conditions for approximate quantum error correction and near-optimal recovery channels*, *Phys. Rev. Lett.* **104** (2010) 120501.
- [58] J. Tyson, *Two-sided bounds on minimum-error quantum measurement, on the reversibility of quantum dynamics, and on maximum overlap using directional iterates*, *J. Math. Phys.* **51** (2010) 092204 [[arXiv:0907.3386](#)].