

The Python's Lunch: geometric obstructions to decoding Hawking radiation

Adam R. Brown,^{a,b} Hrant Gharibyan,^c Geoff Penington^b and Leonard Susskind^{a,b}


^aGoogle, 100 Mayfield Ave, Mountain View, CA 94043, U.S.A.

^bStanford Institute for Theoretical Physics, Stanford University,
450 Serra Mall, Stanford, CA 94305, U.S.A.

^cInstitute for Quantum Information and Matter, Caltech,
1200 E California Blvd, Pasadena, CA 91125, U.S.A.

E-mail: mradambrown@x.team, hrant@caltech.edu, geoffp@stanford.edu,
susskind@stanford.edu

ABSTRACT: According to Harlow and Hayden [[arXiv:1301.4504](https://arxiv.org/abs/1301.4504)] the task of distilling information out of Hawking radiation appears to be computationally hard despite the fact that the quantum state of the black hole and its radiation is relatively un-complex. We trace this computational difficulty to a geometric obstruction in the Einstein-Rosen bridge connecting the black hole and its radiation. Inspired by tensor network models, we conjecture a precise formula relating the computational hardness of distilling information to geometric properties of the wormhole — specifically to the exponential of the difference in generalized entropies between the two non-minimal quantum extremal surfaces that constitute the obstruction. Due to its shape, we call this obstruction the ‘Python’s Lunch’, in analogy to the reptile’s postprandial bulge.



KEYWORDS: AdS-CFT Correspondence, Black Holes, Models of Quantum Gravity

ARXIV EPRINT: [1912.00228](https://arxiv.org/abs/1912.00228)

Contents

1	Introduction	1
2	The shortening of wormholes	2
3	Restricted and unrestricted complexity	5
4	The Python's Lunch	8
4.1	Using ancilla qubits	9
4.2	The complexity of the Python's Lunch	10
4.3	Post-selection is superpolynomial	11
4.4	Covariant Lunches	12
5	Evaporating black holes	14
5.1	Preliminary example	14
5.2	Hawking radiation	15
5.3	The covariant Lunch	21
5.4	The time dependence of the restricted complexity	24
6	Python's Lunches beyond black hole evaporation	24
7	What is holographic complexity?	27
7.1	Measuring radiation and one-sided lunches	29
7.2	Post-selected state complexity	31
7.3	Post-selected complexity can be exponential	32
8	Summary	34
A	Complexity of post-selection	35
A.1	Projecting on a ququit	35
A.2	Projecting on a qudit	36
A.3	Projecting on very unlikely outcomes	37
A.4	Removing the state dependence	38
B	Maximinimax prescription for the bulge surface	40
C	Explicit calculation of the late-time bulge surface	44

1 Introduction

Computational complexity is relevant to the study of quantum gravity in (at least) two ways: in its traditional role as a measure of the difficulty of carrying out tasks [1]; and as a possible holographic dual for properties of the spacetime behind horizons [2–4].

Harlow and Hayden [1] studied the complexity of the task of distilling a single qubit of information from Hawking radiation. They argued that the complexity of distillation grows exponentially with the entropy S of the black hole.

Later, in the context of the AdS/CFT duality, one of us proposed a holographic identification between the computational complexity of an entangled pair of boundary states and the size of the Einstein-Rosen bridge in the dual two-sided black hole [2].

At first sight there seems to be some tension between these two roles of complexity. While the complexity of decoding Hawking radiation is exponential in S , the volume of the wormhole connecting the black hole to its radiation is only polynomial.

The source of the discrepancy is that we are using two different definitions of complexity. The decoding task [1] is only hard because we are restricted to act solely on the radiation outside the black-hole horizon. In ref. [2] there is no such restriction. The distinction between *restricted* and *unrestricted* complexity will be a central theme of this paper. In particular we will be interested in the distinction between the holographic dual of unrestricted complexity, which was the subject of [2–4], and the holographic dual of restricted complexity, which will be a subject we will develop in this paper. The main point of this paper is not to prove the Harlow-Hayden conjecture — like almost everything else in complexity theory this is too hard — but to explain how it may be related to the geometry of wormholes.

Consider one possible decoding ‘strategy’ for distilling information while acting solely on the Hawking radiation.¹ The first step in this strategy is to gather the radiation and collapse it into a second black hole. This new black hole is entangled with the first black hole, and the entanglement can be interpreted, according to ER=EPR [5], as a wormhole connecting them. At the Page time the wormhole would have a volume of order S^2 , far less than the exponential complexity claimed by Harlow and Hayden, but still too large to easily implement the decoding. The next step would be to apply unitary operations to the second black hole in order to shorten the wormhole and bring it to the thermofield-double state. In that state the two horizons have no separation between them, and the structure of the entanglement is especially simple. Once this is accomplished the decoding should be easy. The only potentially hard step in this strategy, therefore, is shortening the wormhole.

Since the only potentially hard step is shortening the wormhole, and since the Harlow-Hayden argument shows that decoding information from the Hawking radiation alone is indeed exponentially hard, we conclude that shortening the wormhole from one side must be exponentially hard. This situation suggests that there must be some kind of obstruction in the wormhole, an obstruction which prevents us from efficiently shortening the wormhole from one side. Moreover this obstruction cannot be large volume, since the volume is not large.

¹We emphasize that this strategy is chosen more for illustrative clarity than engineering practicality.




Figure 1. A spatial slice through a ‘Python’s Lunch’ geometry. On the far left, the wormhole opens up to one asymptotic region with infinite cross-sectional area; on the far right, the wormhole opens up to the other asymptotic region also with infinite cross-sectional area. In AdS-Schwarzschild black holes the cross-sectional area reaches a minimum in the middle of the wormhole, and increases on either side. By contrast, in the Python’s Lunch geometry the cross-sectional area first shrinks, then grows, then shrinks, then grows again, giving rise to a bulge in the middle of the wormhole — the eponymous Lunch. \mathcal{A}_L and \mathcal{A}_R are the areas of the minimal surfaces on each side and \mathcal{A}_{max} is the area of the luncheon bulge.

In this paper we will conjecture that the geometric obstruction is a bulge in the wormhole, which because of its shape we call the “Python’s Lunch”, as depicted in figure 1. We will estimate the complexity of bypassing the Python’s Lunch, and find that, consistent with the Harlow-Hayden claim, it is indeed exponential. In eq. (4.6) we will conjecture that the restricted complexity is dual to the size of the Python’s Lunch via

$$\mathcal{C}_R[U_{PL}] \sim \exp \left[\frac{1}{2} \frac{\mathcal{A}_{max} - \mathcal{A}_R}{4G\hbar} \right], \tag{1.1}$$

where \mathcal{A}_{max} is the maximum cross-section of the wormhole and \mathcal{A}_R is the size of the throat connecting the wormhole to the radiation.² In eq. (4.8) we will make a covariant generalization of this conjecture. This proposal for the geometric dual of the *restricted* complexity is complementary to existing conjectures about the geometric duals to *unrestricted* complexity [2–4].

Despite our focus on restricted complexity, in section 7 we find that one-sided Python’s Lunches can also teach us about unrestricted complexity. We suggest an improvement to the definition of unrestricted holographic complexity conjectured to be dual to volume & action in refs. [2–4]. Specifically, we argue that these conjectures should have defined complexity to permit not only unitary gates but also non-unitary projections.

2 The shortening of wormholes

In much of what follows we will assume that black holes can be modeled as “quantum computers” by which we mean collections of N qubits evolving by means of k -local all-to-

²This conjecture is expected to be valid for states with semi-classical geometry where complexity of the state on bulk degrees of freedom is not very large.

all Hamiltonians or discrete gates. The number of qubits is determined by the entropy of the black hole,³

$$N \sim S. \tag{2.1}$$

We will encounter both (unitary) operator complexity and relative state complexity.

The complexity of a unitary operator U may be defined as the minimal number of 2-qubit gates g needed to prepare it; in other words the smallest n for which

$$U = g_n g_{n-1} \dots g_1. \tag{2.2}$$

There are other definitions but for our purposes this definition will do. The complexity of U will be denoted by $\mathcal{C}(U)$. By construction, it satisfies

$$\mathcal{C}(U) = \mathcal{C}(U^\dagger). \tag{2.3}$$

The relative complexity of two states $|\psi\rangle$ and $|\phi\rangle$ is defined as the complexity of the *least complex* unitary that connects them $|\psi\rangle = U|\phi\rangle$. In other words it is the minimum number of gates required to transform $|\phi\rangle$ to $|\psi\rangle$,

$$|\psi\rangle = U|\phi\rangle = g_n g_{n-1} \dots g_1 |\phi\rangle.$$

Relative complexity is denoted by $\mathcal{C}(\psi, \phi)$. Due to eq. (2.3), it is symmetric in its arguments

$$\mathcal{C}(\psi, \phi) = \mathcal{C}(\phi, \psi). \tag{2.4}$$

Unitary matrices represent operators,

$$U = \sum_{IJ} U_{IJ} |I\rangle\langle J| \tag{2.5}$$

where I, J label a complete basis of N qubit states in the computational basis. The same matrix plays a second role in representing a maximally entangled state of $2N$ qubits,

$$|\Psi\rangle = \frac{1}{2^{N/2}} \sum_{IJ} U_{IJ} |I\rangle|\bar{J}\rangle, \tag{2.6}$$

where $|\bar{J}\rangle$ is the time reversal of $|J\rangle$.

A special case is $U_{IJ} = \delta_{IJ}$ which describes the infinite temperature thermofield-double state,

$$|\text{TFD}\rangle\Big|_{T=\infty} = \frac{1}{2^{N/2}} \sum_I |I\rangle|\bar{I}\rangle. \tag{2.7}$$

The infinite temperature $|\text{TFD}\rangle$ state may also be written as a product of N Bell pairs,

$$|\text{TFD}\rangle\Big|_{T=\infty} = |\text{Bell}\rangle^{\otimes N}. \tag{2.8}$$

The thermofield-double state of a two-sided black hole is a finite temperature state of an infinite number of qubits, but is often modeled as an infinite temperature state of a

³Modeling black holes as a quantum computer has been extremely fruitful in the study of quantum information scrambling [6, 7], onset of random matrix behavior [8] and derivation of the RT formula [9].

finite number of qubits. We too will make that approximation, so by $|\text{TFD}\rangle$ we will mean the state in eq. (2.8). The state $|\text{TFD}\rangle$ is the simplest case of a maximally entangled state.

The two subsystems called A and B are under the control of Alice and Bob respectively. The natural evolution of the system is governed by an overall Hamiltonian which is the sum of two non-interacting terms,

$$H = H_A \otimes \mathbb{1} + \mathbb{1} \otimes H_B. \tag{2.9}$$

For simplicity we will assume that the two Hamiltonians are identical and real (so that we don't have to worry about the details of time reversal).

The natural time evolution operator is a product,

$$U(t) = U_A(t) \otimes U_B(t) = e^{-iH_A t} \otimes e^{-iH_B t}. \tag{2.10}$$

We will define the time-evolved state $|\text{TFD}(t)\rangle$ by

$$|\text{TFD}(t)\rangle \equiv U(t)|\text{TFD}\rangle. \tag{2.11}$$

In the maximally entangled case $|\text{TFD}(t)\rangle$ can be constructed by evolving on only one side for a total time $2t$,

$$\begin{aligned} |\text{TFD}(t)\rangle &\equiv U_A(t) \otimes U_B(t) |\text{TFD}\rangle \\ &= U_A(2t) \otimes \mathbb{1} |\text{TFD}\rangle \\ &= \mathbb{1} \otimes U_B(2t) |\text{TFD}\rangle. \end{aligned} \tag{2.12}$$

As t evolves, the linearly growing complexity of the state is dual to the linearly growing volume of the wormhole. For sub-exponential times the complexity is the sum of the complexity of $U_A(t)$ and $U_B(t)$ which is equal to the complexity of $U_A(2t)$ and of $U_B(2t)$.

Given the evolved state $|\text{TFD}(t)\rangle$ Alice can return it to the initial TFD by applying $U_A^\dagger(2t)$. We will think of doing this in a series of small steps of low complexity,

$$|\text{TFD}\rangle = (U_A^\dagger)^\epsilon (U_A^\dagger)^\epsilon \dots (U_A^\dagger)^\epsilon |\text{TFD}(t)\rangle. \tag{2.13}$$

Pictorially Alice is shortening the long wormhole by a series of incremental small steps, as illustrated in figure 2.

Bob may also accomplish the shortening by acting on $|\text{TFD}(t)\rangle$ with $U_B^\dagger(2t)$, or Alice and Bob may act together with $U_B^\dagger(t_1)U_A^\dagger(2t - t_1)$.

Here are some questions to consider:

- Why might one be interested in shortening a wormhole? There are a number of reasons. One that we have already mentioned is that it would be a step in decoding Hawking radiation.

A second would involve the use of the entanglement as a resource for quantum teleportation. Quantum teleportation requires the use of pre-existing entangled qubits. Not only must these qubits be entangled, they must also be brought to have low complexity in order to successfully teleport. In the language of ER=EPR, if we want to make a wormhole traversable [10, 11], we first need to make it as short as possible.

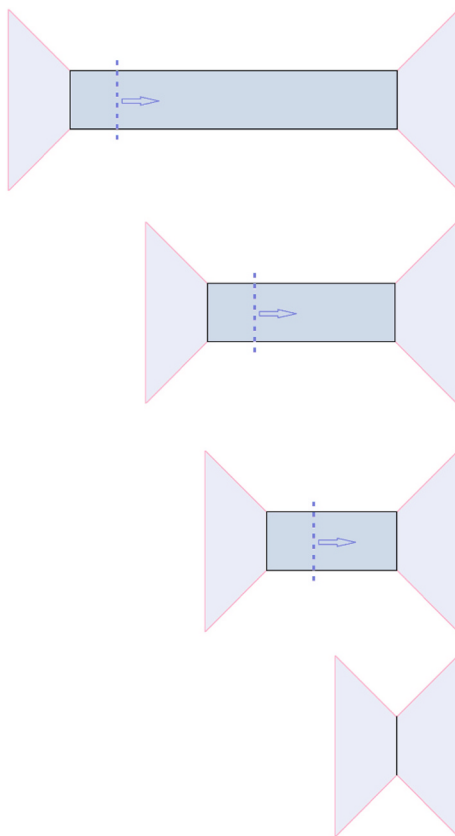


Figure 2. Successive spatial slices through the wormhole. Since the two sides are maximally entangled, Alice is able to shorten the wormhole by unitary operations $U_A \otimes \mathbb{1}$ that act only on her side.

- Why might one want to shorten the wormhole by processes which do not couple the two sides? If the two sides are being used to communicate over a long distance then coupling them quantum-mechanically may be unfeasible. Thus there are practical reasons why one might be interested in the complexity of shortening a wormhole by acting on it from one side.
- Are there situations in which it is easy to shorten a wormhole by interactions which involve both sides, but in which it is extremely difficult to do so from one side?
- Most of all we are interested in whether the answer to the previous question correlates with geometric properties of the wormhole, and if so, what properties?

With regard to this last question, we will argue that there is a particular kind of geometric obstruction which prevents us from efficiently shortening a wormhole by one-sided operations, even though the wormhole has small volume and can be easily shortened by two-sided operations. The shape of the obstruction suggests the name “Python’s Lunch”.

3 Restricted and unrestricted complexity

The restricted complexity \mathcal{C}_R of a maximally entangled state of $2N$ qubits is the number of gates needed to construct it from the TFD state under the restriction that all gates act only

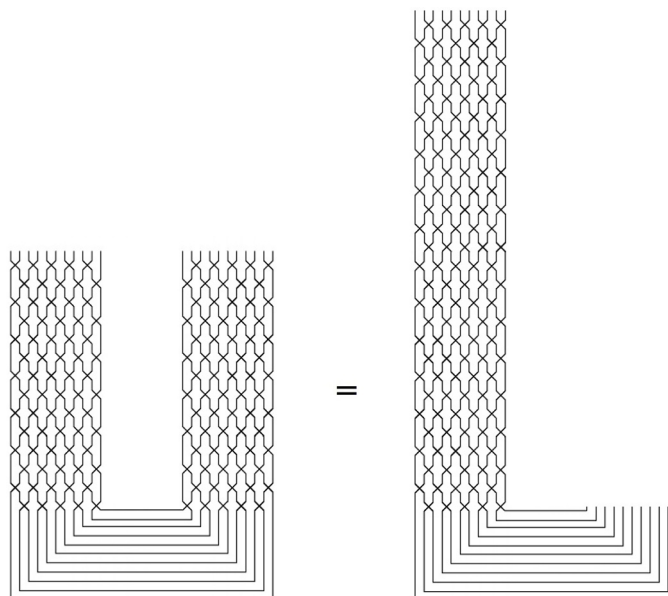


Figure 3. Left: the quantum circuit that prepares $U_A(t) \otimes U_B(t)|\text{TFD}\rangle$. Right: the quantum circuit that prepares $U_A(2t) \otimes \mathbb{1}|\text{TFD}\rangle$. The two states are the same.

on one side.⁴ We will sometime use $C_{R,A}$ to indicate that the restriction is on a specific subsystem (A in this case). Without loss of generality we can assume the gates all act on Alice’s side or we may distribute them symmetrically between the two sides. A useful picture is provided by the tensor network (TN) description. The state $|\Psi\rangle$ is represented as a TN as in figure 3.

The restricted relative complexity of $|\Psi\rangle$ and $|\text{TFD}\rangle$ is also the complexity of the unitary operator U corresponding to $|\psi\rangle$, as defined in (2.5).

By acting with a layer of gates on Alice’s side, a layer can be removed from the TN. By repeating this operation enough times, as in figure 4, the state can be brought to the simple state $|\text{TFD}\rangle$. The minimal number of gates needed to carry out the shortening operation defines⁵ the restricted complexity of $|\Psi\rangle$.

The restricted complexity would be an appropriate measure of the difficulty of the task of shortening the wormhole if the two computers were too far apart to directly couple.

The unrestricted complexity \mathcal{C}_U is the number of 2-qubit gates needed to complete the shortening task, allowing for gates which couple the two computers. Figure 5 shows such an unrestricted circuit.

It is obvious that $\mathcal{C}_R \geq \mathcal{C}_U$. In this paper we will be interested in the conditions under which the restricted complexity may be exponentially larger than the unrestricted

⁴What we call the ‘restricted complexity’, Aaronson [12] calls the ‘separable complexity’.

⁵Note that this definition of complexity is in terms of processes which bring the state back to a simple state rather than processes which prepare the state starting with the simple state. For circuits built from unitary gates the two are the same, but for more general concepts of complexity that may make use of non-unitary elements the two may differ significantly. Later we will consider tensor networks that include non-unitary elements such as projections where this distinction is relevant.




Figure 4. By acting with $U_A(t)^\dagger \otimes \mathbb{1}$ in a series of incremental k -local steps, Alice can undo time evolution, thus mapping $U_A(t) \otimes \mathbb{1}|\text{TFD}\rangle$ back to $|\text{TFD}\rangle$.




Figure 5. A unitary U_{AB} cannot in general be decomposed as $U_A \otimes U_B$. In the example in this figure, the horizontal red links between the left and right sides represent gates which couple qubits on the two sides.




Figure 6. The tensor network that corresponds to the Python’s Lunch geometry. The throats and bulge (where the girth is constant) are composed of unitary gates, whereas the shoulders (where the girth changes) involve projections like those shown in figure 7.

complexity. This subject is not new; it was introduced by Harlow and Hayden [1] in the context of black hole physics, and elaborated on by Aaronson [12]. Our particular interest is to understand this large gap between restricted and unrestricted complexity through the geometry of the wormholes connecting entangled systems. The question is: can we identify the situations in which $\mathcal{C}_R \gg \mathcal{C}_U$ from the shape of the wormhole? To put it another way, can we identify a geometric obstruction to shortening the wormhole from one side?

If the state $|\Psi\rangle$ was prepared by acting with restricted gates on $|\text{TFD}\rangle$, and if the number of such gates is not exponentially large, we expect $\mathcal{C}_U = \mathcal{C}_R$. On the other hand if $|\Psi\rangle$ was prepared from $|\text{TFD}\rangle$ by a circuit that allows interaction between the two computers, then we expect $\mathcal{C}_{R,A} \gg \mathcal{C}_U$ (assuming the circuit is longer than the scrambling length). In particular if the number of unrestricted gates used in preparing $|\Psi\rangle$ is enough to scramble the system then Harlow and Hayden have argued that the restricted complexity $\mathcal{C}_{R,A}$ will be exponential in N , and the same for $\mathcal{C}_{R,B}$. At the same time the unrestricted complexity may be no bigger than polynomial in N .

4 The Python’s Lunch

We now come to the Python’s Lunch geometry: a wormhole with a bulge, as illustrated in figures 1 and 6. For simplicity we will assume that it consists of three regions, all of length polynomial in N , where N denotes the entropy (number of qubits). The two outer regions have area $A_L \approx N \cdot (4G\hbar)$ and $A_R \approx (1 + \gamma)N \cdot (4G\hbar)$, where $\gamma > 1$ is a numerical constant. The bulge between the two outer regions has larger area,

$$A_{\text{max}} = (1 + \alpha)N \cdot (4G\hbar), \tag{4.1}$$

where $\alpha > \gamma$ is a constant independent of N . In order to count as a Python’s Lunch, we will mostly assume that the length must be larger than the scrambling time $t_* \sim \log N$.

An alternative way to look at this geometry is as the tensor network (TN) in figure 6 which prepares a two-sided state.

If all the vertices in the tensor network were unitary gates, the number of qubits would be the same for every vertical cross-section, but tensor networks (unlike standard quantum circuits) allow certain non-unitary vertices called isometries. Inspection of figure 6 shows




Figure 7. An isometry of $|i\rangle \rightarrow |j\rangle|k\rangle$, represented as the projection of a unitary matrix. Such components make up the ‘shoulders’ on the Python’s lunch shown in figure 6, where the $|0\rangle$ ‘leg’ is omitted.




Figure 8. Zooming in on the ‘shoulder’ of the Python’s Lunch in figure 6.

that some of the vertices involve three edges; those are the isometries. They occur in the transition regions where the area varies.

An isometry can be thought of as a unitary gate in which one of the legs has been projected onto the state $|0\rangle$ as shown in figure 7. This allows us to draw the TN as a collection of edges connecting unitary gates, but with a subset of the edges being projected. A portion of the TN with isometries is shown in figure 8. Reading the tensor network from left to right, the tensor network expands when we input an ancilla qubit, and contracts when we post-select on a qubit. In general, the number of input ancilla qubits, $m_L = \alpha N$, can be different from the number of post-selected qubits, $m_R = (\alpha - \gamma)N = \beta N$.

One question is whether the operator defined by the TN in figure 6 is approximately unitary, or equivalently, is the two-sided state that it defines approximately maximally entangled? So long as the right end of the tensor network is larger than the left end (at leading order), the TN will generically be an almost perfect isometry from the left to the right. The state on the left will therefore be almost exactly maximally entangled with the state on the right. We shall assume that this is indeed the case.

With this assumption the TN can be shortened from the right by one-sided unitary operations. But the question is how many one-sided k -local operations are required? If the TN is small — say of polynomial size — one might conclude that the number of gates should also be small, but that is not the case.

4.1 Using ancilla qubits

Let’s consider an initial state $|I\rangle$ in the computational basis and act on it with the TN, inserted from the left side. The output state is

$$|\psi\rangle \propto \langle 0|^{\otimes m_R} U_{\text{TN}} |I\rangle |0\rangle^{\otimes m_L}, \quad (4.2)$$

where U_{TN} is the original map from left to right defined by the tensor network, $|I\rangle$ is the input state on the left, and $|\psi\rangle$ is the normalized output state on the right side after post-selection. We will be interested in the relative complexity of $|I\rangle$ and $|\psi\rangle$, when we allow Alice to prepare m_L ancilla qubits that start at $|0\rangle$ and in the end m_R qubits finish in $|0\rangle$ state.

We begin on the left side of figure 6, and working from left to right, apply the unitary gates. The isometries on the left ‘shoulder’ of the Python’s Lunch are straightforward; we simply couple in the ancilla and treat the isometries as unitary circuit elements.

But when we arrive at the right ‘shoulder’ the isometries correspond to final state projections (post-selections). Final state projections are not physically implementable processes, so we must do something new. That new thing is measurement: as Alice arrives at each isometry she simply measures the dangling qubit in figure 8 in the Z basis.⁶ If she gets 0 she moves on to the next isometry and repeats the process. If at the end of all the isometries all measurements give 0 she moves on to the right of the TN and at the end she has prepared $|\psi\rangle$.

However, if she measures 1 at some point she starts over and repeats the entire process, until she succeeds in obtaining 0 for all post-selected qubits. On average she will have to repeat the procedure 2^{m_R} times, for m_R post-selected qubits.⁷ The total number of gates she will have to apply is of order

$$2^{m_R} \cdot \mathcal{C}_{\text{TN}}$$

where \mathcal{C}_{TN} is the number of nodes in the TN.

If this were the minimal protocol we would say that the complexity $\mathcal{C}(\psi, I)$ is of order $2^{m_R} \mathcal{C}_{\text{TN}}$.⁸ However, in the appendix we show that there is a more efficient quantum procedure using a version of Grover’s algorithm (which was applied to a similar problem by Kitaev and Yoshida in ref. [13]).

4.2 The complexity of the Python’s Lunch

In appendix A.2, we describe a protocol that uses Grover search to prepare the $|\psi\rangle \otimes |0\rangle^{\otimes m_R}$ from an initial state $|I\rangle \otimes |0\rangle^{m_L}$ with a unitary circuit using $2^{\frac{m_R}{2}} \mathcal{C}_{\text{TN}}$ gates. Assuming that the length of the lunch is greater than the scrambling time, there is no reason to think that the task can be performed with fewer gates, thus implying that

$$\mathcal{C}(\psi, I) \approx 2^{\frac{m_R}{2}} \cdot \mathcal{C}_{\text{TN}}. \tag{4.3}$$

Our initial version of this protocol works only for a single fixed input state $|I\rangle$.

In contrast, we are really interested in finding a unitary operation U_{PL} that satisfies

$$|\psi\rangle |0\rangle^{m_R} = U_{\text{PL}} |I\rangle |0\rangle^{m_L} \tag{4.4}$$

⁶Of course, measurements are still not unitary and so are not traditionally allowed when calculating the complexity of a state. In contrast, the faster Grover-search protocol that we discuss below is genuinely unitary.

⁷Here we assume that the bulge of the tensor network is longer than scrambling time. This assumption is made to insure that the distribution over post-selected qubits is uniform and is given by $1/2^{m_R}$.

⁸As before, this assumes that we allow measurements in our definition of complexity.

for any input state $|I\rangle$.⁹ In appendix A.4, we such a ‘state-independent’ protocol by a variant on the usual Grover-search algorithm.

Our state-independent protocol succeeds with high probability for any input state $|I\rangle$, given either the assumption that there exists an exact isometry from left-to-right, or simply the assumption that the right system is parametrically larger than the left system and that the tensor network is scrambling (and so can be modelled using 2-designs). The complexity of this protocol is again given by $2^{\frac{m_R}{2}} \mathcal{C}_{\text{TN}}$.

Since we have no good reason to think that a faster algorithm exists, we conjecture that

$$\mathcal{C}[U_{\text{PL}}] \approx 2^{\frac{m_R}{2}} \cdot \mathcal{C}_{\text{TN}}. \quad (4.5)$$

Since we have assumed that m_R is a finite fraction of N , i.e., $m_R = \beta N$, the complexity of U_{PL} is exponential in N .

That is our main technical result: that the complexity of a TN is expected to be $O\left(2^{\frac{m_R}{2}} \mathcal{C}_{\text{TN}}\right)$ where m_R is the difference between the maximal area of the lunch and the area of the right side (or, more generally, the larger side). In particular when $m_R \sim N$ the complexity is exponential in the number of qubits N at either end. The surprising point about this result is that the TN that prepares $\mathcal{C}(U_{\text{PL}})$ can be as small as $\mathcal{C}_{\text{TN}} \sim N \log N$. We can now use the analogy between tensor networks and bulk geometry to conjecture a relationship between the restricted complexity and the geometry of a Python’s Lunch.

Restricted complexity conjecture. In a Python’s Lunch geometry with min-max-min areas $\mathcal{A}_L, \mathcal{A}_{\text{max}}$ and \mathcal{A}_R , and with the assumption $\mathcal{A}_L < \mathcal{A}_R$, we conjecture that the restricted complexity on the right system is

$$\text{conjecture: } \mathcal{C}_R[U_{\text{PL}}] = (\text{const.}) \times \mathcal{C}_{\text{TN}} \cdot \exp\left[\frac{1}{2} \frac{(\mathcal{A}_{\text{max}} - \mathcal{A}_R)}{4G\hbar}\right], \quad (4.6)$$

where \mathcal{C}_{TN} denotes the size of the tensor network and is related to the volume/action of the wormhole ($\mathcal{C}_{\text{TN}} = V/G\hbar l_{\text{AdS}}$) from the CV/CA conjectures [2–4].

In particular, if $A_L \approx N \cdot (4G\hbar)$, $A_{\text{max}} \approx (1 + \alpha)N \cdot (4G\hbar)$, and $A_R \approx (1 + \gamma)N \cdot (4G\hbar)$ the complexity of decoding from the right system alone is

$$\mathcal{C}_R[U_{\text{PL}}] \approx (\text{const.}) \times \mathcal{C}_{\text{TN}} \cdot e^{(\alpha - \gamma)N/2}, \quad (4.7)$$

where $\mathcal{C}_{\text{TN}} \geq N \log N$.

4.3 Post-selection is superpolynomial

In the previous subsection we argued that we can decode Hawking radiation by projecting out m qubits, and provided a version of a Grover search [14, 15] that allows to do this projection with a unitary that has complexity $\sqrt{2^m}$. Can we rule out the possibility that there is an even faster algorithm that can perform this projection?

On the one hand, we can almost certainly rule out the possibility that there could be an *exponentially* faster algorithm. There cannot be an algorithm that projects onto m

⁹This requirement is not sufficient to specify U_{PL} completely. Instead, our conjecture is that the minimal complexity of *any* unitary operator, satisfying (4.4), is given by (4.5).

qubits in a time that scales polynomially with m . Or — more precisely — if there *were* such an algorithm then it would contradict widely held conjectures about computational complexity theory. The complexity class of decision problems you can solve on a quantum computer if you were allowed post-selections (including post-selections onto states with exponentially small amplitude) is called PostBQP. It has been shown [16] that PostBQP is a fantastically powerful class — it is equal to the class PP. Conversely, if you could use a normal quantum computer to implement exponentially rare projections in polynomial time, this would imply BQP=PostBQP. Taking these results together would imply BQP=PP. But PP is a very large class that contains all sorts of problems not believed to be efficiently soluble on a quantum computer, including NP.¹⁰ It would therefore be in gross violation of widely held complexity assumptions if we could post-select in a time polynomial in m .

On the other hand, it is not obvious how to rule out the existence of a *polynomially* faster algorithm, that would still take a time exponential in m . It is not obvious that there can't be a protocol that would improve (say) the square root to a cube root. (The effect of such a speed up would be to change the coefficient in the exponent of the conjecture eq. (4.6) from 1/2 to 1/3.) It has been proved that Grover search amongst d items cannot be implemented faster than $\frac{\pi}{4}\sqrt{d}$ [17–19],¹¹ but we have an advantage not available in the Grover task, which is that we know in advance which final state we wish to post-select on.

4.4 Covariant Lunches

So far in this paper, to determine whether the spacetime contains a Python's Lunch we have implicitly assumed the existence of some preferred choice of bulk Cauchy slice. This is, in large part, a limitation of the tensor network toy models that we have been using to guide us and which resemble a bulk Cauchy slice rather than a full bulk spacetime. However, for the non-static spacetimes that we will be considering in future sections, it is not obvious how the correct slice should be chosen.

In earlier work, the complexity was conjectured to be dual to the volume of the maximal volume slice. An obvious possibility is to work entirely within this slice.

¹⁰Let us see how the ability to post-select would allow us to solve problems in NP. The classic NP-complete problem is 3SAT. An instance of 3SAT is a map $f(\vec{x})$ from m Boolean variables to one Boolean variable (1=true or 0=false). Supposing there is exactly one assignment \vec{x} that makes $f(\vec{x})$ come up true, our job is to find it. This problem is in NP because given a candidate \vec{x}_{answer} it is easy to check whether $f(\vec{x}_{\text{answer}}) = 1$; on the other hand it may be hard to find \vec{x}_{answer} . On a quantum computer, we could test all the possibilities at once by building a circuit that implements $U_{\text{circuit}}|\vec{x}\rangle|0\rangle = |\vec{x}\rangle|f(\vec{x})\rangle$ and then plugging in an superposition over all possible inputs,

$$U_{\text{circuit}} \left(\frac{1}{\sqrt{d}} \sum_{\vec{x}} |\vec{x}\rangle|0\rangle \right) = \frac{1}{\sqrt{d}} |\vec{x}_{\text{answer}}\rangle|1\rangle + \frac{1}{\sqrt{d}} \sum_{\vec{x} \neq \vec{x}_{\text{answer}}} |\vec{x}\rangle|0\rangle,$$

but this generally doesn't help, because even though the wavefunction 'knows' the answer, linearity means there is no measurement we can do that has more than an $O(1/d)$ chance of success that can induce the wavefunction to tell us what it knows. If we were able to project on the (exponentially in $m = \log d$) unlikely outcome that a measurement of the final qubit is $|1\rangle$, then we could find $|\vec{x}_{\text{answer}}\rangle$ in one step.

¹¹There are also other known lower bounds for more general classes of algorithms that perform amplitude amplification [20–22] and which are closely related to our state-independent protocol from appendix A.4.

However, the covariant surface that is analogous to the minimal cut through a tensor network is the HRT surface [23–25], the minimal area extremal surface homologous to one end of the wormhole. For spacetimes where quantum effects are important, such as evaporating black holes, it is, more precisely, the minimal generalized entropy quantum extremal surface,¹² also known as the Engelhardt-Wall (EW) surface [26, 27].

The existence of a second, locally minimal cut is analogous to the existence of a second extremal surface, satisfying the same homology constraint. In all the cases that we will consider, there will also be a third extremal surface, that lies in between the first two surfaces, and has a larger generalized entropy than either. This third surface has an important qualitative difference compared to the other two surfaces: we cannot choose a Cauchy slice within which any small (but not necessarily local) deformation of this third extremal surface will increase its area (or generalized entropy). In particular, this means that it cannot ever be the HRT (or EW) surface, which is always globally minimal within some Cauchy slice [28, 29]. Instead of corresponding to one of the narrow constrictions at the ends of the python, this third surface is a covariant definition of the maximum size of the bulge in the middle of the lunch.

In general, none of these surfaces will lie in the maximal volume slice. We therefore should not expect the correct covariant definition of a Python’s Lunch to involve the maximal volume slice (although, in many examples, such as evaporating black holes, the maximal volume slice will also look like a Python’s Lunch). Instead, we should think of a Python’s Lunch as being *defined* by this set of three extremal surfaces, the two ‘end surfaces’ and the ‘bulge surface’ in the middle. With this new covariant definition of Python’s Lunch we can modify our conjecture.

Restricted complexity conjecture (covariant version). In a covariant Python’s Lunch geometry with min-max-min generalized entropies $S_L^{(\text{gen})}$, $S_{\text{max}}^{(\text{gen})}$ and $S_R^{(\text{gen})}$, and with the assumption $S_L^{(\text{gen})} < S_R^{(\text{gen})}$, the restricted complexity on the right system is,

$$\text{conjecture: } \mathcal{C}_R[U_{\text{PL}}] = (\text{const.}) \times \mathcal{C}_{\text{TN}} \cdot \exp \left[\frac{1}{2} \left(S_{\text{max}}^{(\text{gen})} - S_R^{(\text{gen})} \right) \right], \quad (4.8)$$

where again \mathcal{C}_{TN} denotes the size of tensor network.

Naively, a covariant Python’s Lunch seems a very specific and unusual feature of a spacetime. It needs to feature three extremal surfaces. Moreover the bulge surface needs to have greater area (or generalized entropy) than either end surface, and, unlike the end surfaces, within any Cauchy slice there should exist small deformations of the bulge surface that decrease its area (or generalized entropy). Nevertheless, every example that we consider of a spacetime with more than one extremal surface will turn out to have a Python’s Lunch.

¹²The generalized entropy of a surface Σ_B is $S^{(\text{gen})}(\Sigma_B) = A(\Sigma_B)/4G\hbar + S_{\text{bulk}}(\Sigma_B)$, where the first term is the area of the surface Σ_B , which should be homologous to boundary region B , and the second is von Neumann entropy of the bulk fields contained in the corresponding entanglement wedge (the region between B and Σ_B). A quantum extremal surface is a surface that is an extremum of the generalized entropy.

In appendix B, we explain this phenomenon. We use ‘maximin’ techniques [28, 29] to sketch an argument that almost all spacetimes with more than one extremal surface will contain a Python’s Lunch. Specifically, we argue that one can generically find a third extremal surface by considering ‘foliations’ of a Cauchy slice from one extremal surface to the other, taking the maximal area (or generalized entropy) surface within that foliation, minimizing the maximum over all foliations, and then maximizing the resulting ‘minimax’ surface over all Cauchy slices. We call this a ‘maximinimax’ prescription for finding the bulge surface.

5 Evaporating black holes

In this section we will see how a Python’s Lunch explains the exponential difficulty of decoding Hawking radiation.

After the Page time an evaporating black hole is maximally entangled with its own Hawking radiation [30]. Harlow and Hayden [1] asked how hard it is to isolate a degree of freedom \mathbf{r} in the radiation that is entangled with a particular quantum \mathbf{b} of Hawking radiation that is about to be emitted by the black hole (the AMPS task [31]). A highly related task is to decode the state of a small unknown diary thrown into the (known) black hole, just from the state of the Hawking radiation.¹³ This is the Hayden-Preskill decoding task [32] and is also expected to be exponentially hard. If we can get the black hole and Hawking radiation into a simple state, both tasks are simple. The difficulty in doing either task comes from the exponentially large restricted complexity of the combined state of the black hole and Hawking radiation.

Building on earlier ideas in [33, 34], it was shown in [35–39] that the information-theoretic *achievability* (or otherwise) of the Hayden-Preskill and Harlow-Hayden tasks could be understood holographically using entanglement wedge reconstruction.¹⁴ After the Page time, a large part of the interior of the black hole is in the entanglement wedge of the early Hawking radiation, and so is encoded in the radiation.¹⁵ This is essentially a holographic derivation of black hole complementarity and ER=EPR [47, 48]. We shall now see that the exponential computational *difficulty* of the Harlow-Hayden and Hayden-Preskill tasks can likewise be understood holographically as coming from the existence of a Python’s Lunch.

5.1 Preliminary example

We will consider a preliminary example. Consider a quantum computer initialized at $t = 0$ in some simple state $|I\rangle$ which then evolves for a time greater than the scrambling time t_* . The qubits are then split into two subsystems, Alice’s and Bob’s shares A and B . The two

¹³If the diary is maximally entangled with a reference system then the Hayden-Preskill task is the same as the Harlow-Hayden task except the degree of freedom \mathbf{r} now needs to purify the reference system instead of a late-time Hawking quantum.

¹⁴The idea of entanglement wedge reconstruction was originally conjectured in [28, 40, 41], and established in [42–44] using the ideas of [45, 46].

¹⁵In this context, the entanglement wedge of the Hawking radiation is defined as the bulk domain of dependence bounded by the EW surface.




Figure 9. A scrambled system of $2N$ qubits can be separated into Alice’s and Bob’s subsystems which due to the scrambling are close to a maximally entangled state. Since they are maximally entangled, it must be possible to shorten the wormhole by restricted unitary operations.

subsystems continue to evolve for a short time but with no coupling between them. The process is illustrated¹⁶ in figure 9. Let us imagine sweeping across the TN by a series of cuts which foliate it as in figure 10. It is obvious that the number of qubits crossed by the cuts first increases and then decreases. At its maximum the number of qubits is at least $N \log N$. Therefore the geometry of the associated wormhole has a Python’s Lunch.

Because the system is scrambled at time t the subsystems A and B are approximately maximally entangled. It follows that Alice can act unitarily on her side in order to bring the system to a state close to the TFD. The arguments of the previous section show that the restricted complexity is exponential in N although the number of vertices in the TN is much smaller.

5.2 Hawking radiation

In this section we will explain how the Python’s lunch geometry appears during the evaporation of a black hole. Note that, for the moment, we are restricting our attention to a single Cauchy slice through this black hole — a generic ‘nice’ Cauchy slice that stays close to the black-hole horizon. For concreteness we can take it to be the maximal volume slice. In section 5.3, we will discuss the full covariant description of this lunch.

A classical one-sided black hole (Bob’s black hole) in a pure state is not connected to any purifying system by a wormhole. But it is connected to a growing “bridge to nowhere” (BTN) whose volume represents the complexity of the state. This is shown schematically in figure 11.

Starting at the horizon and moving into the BTN the area remains constant for most of its length until it quickly shrinks at the far end. The whisker-like lines at the end represent

¹⁶The black-hole dynamics is better characterized by a 2-local circuit *without* geometric locality, but due to the authors’ artistic limitations we will present the illustrations in a geometrically local 1d lattice.




Figure 10. A foliation of the tensor network of figure 9 which interpolates between Alice’s side and Bob’s side. It is clear from the figure that the number of qubits cut by the slices increases and then decreases as the foliation sweeps round.




Figure 11. Successive spatial slices through a one-sided non-evaporating black hole that formed from collapse. The “whiskers” on the left side depict the infalling matter. The ‘bridge to nowhere’ grows and becomes elongated as the complexity increases.




Figure 12. Successive spatial slices through a one-sided evaporating black hole that formed from collapse. The red dots mark where Hawking particles were emitted.

the infalling matter which originally created the black hole. As time increases the BTN grows.

Evaporation modifies this picture and effectively turns the one-sided system into a two-sided system. The black hole becomes entangled with its own Hawking radiation which in effect becomes a second side. The process which was explained in [5] is depicted in figure 12 as a time-sequence at times $t_1 < t_2 < t_3$. As the black hole radiates the area of its horizon decreases. In the figure this is shown as a decrease in the thickness of the BTN as one moves from left to right. The interior modes that purify Hawking radiation are shown as red dots. The partners of Hawking radiation emitted later in the evaporation are at the furthest right of the diagram. These interior modes are entangled with the Hawking radiation, and so they have an entropy that contributes to the generalized entropy of any region containing them, but not containing the Hawking radiation, or vice versa. Alternatively, in the language of ER=EPR, they can be thought as being connected by a Planck-area “micro-wormhole” to the Hawking radiation. The homology constraint forces us to cut these micro-wormholes, which increases the generalized entropy.

Now suppose that Alice collects the radiation in a second system A shown as an elongated ellipse in figure 13. System A will now be connected to the bridge-to-nowhere via either bulk entanglement/micro-wormholes.

Let us assume that, like a tensor network, the entropy of the Hawking radiation is given by the ‘minimal cut’ through this Cauchy slice, where the size of a cut is given by its generalized entropy. Of course, in general, this will only be true if we have chosen our Cauchy slice appropriately. However, we will be able to derive the correct qualitative conclusions by just studying the maximal volume slice.

We construct “cuts” separating A from the boundary at the right end of the “fish tail”. The cuts can be characterized by an area. In figure 14, we see a series of such cuts as we sweep from A to the fishtail.

The generalized entropy of each cut consists of two contribution. One is the portion that cuts through the bulk entanglement between the interior and the Hawking radiation. The second contribution comes from the classical area required to cut the bridge-to-nowhere. Let us track the generalized entropy as the cut proceeds:

- In the first cut in figure 14 the only contribution to the generalized entropy comes from the bulk entanglement. That contribution is proportional to the entropy in the radiation.




Figure 13. The ‘expandable space blimp’. The one-sided wormhole from figure 12, now showing the bulk entanglement connecting the distant Hawking particles to their anchor points on the red dots.




Figure 14. A non-temporal sweep of spatial cuts through the expandable space blimp geometry of figure 13, analogous to the sweep of the spatial slices in figure 10.

- The next cut also cuts the entanglement between the interior and the radiation. However, it also cuts across the largest part of the BTN. We see that there is a quick increase in the generalized entropy of the cut.
- In the third and fourth cuts in figure 14, the cut moves to the right. As it does so the both contributions to the generalized entropy decrease.
- The final cut in figure 14 only cuts across the bridge-to-nowhere near the horizon of the black hole.

The evolution of the generalized entropy of the cut is shown in figure 15.

There are two minima to the generalized entropy. One is the cut through the bulk entanglement represented by the green lines in the top picture. As we sweep across the generalized entropy makes a fairly sudden increase, and then a slow gradual decrease to a second minimum — the horizon — at the fishtail. Up to subtleties involving the choice of Cauchy slice, these two minima correspond to the two quantum extremal surfaces found in [35, 36].

Early on, the horizon area A_{hor} (or, more precisely, the Bekenstein-Hawking entropy $A_{\text{hor}}/4G_N$) is much larger than the bulk entanglement between the interior and the Hawk-




Figure 15. The size of the cut for the slices of figure 14 as a function of the sweep parameter. For $t < t_{\text{page}}$, the minimum generalized entropy cut is at the beginning, when the bulk entanglement between the black hole and the Hawking radiation is being cut. For $t > t_{\text{page}}$, the minimum generalized entropy cut is near the end, when the cut is near the black hole horizon. In both cases, the largest cut comes near the beginning, when the generalized entropy is the sum of the semiclassical entropy of the radiation, plus the initial Bekenstein-Hawking entropy of the black hole.

ing radiation. At a very late time the horizon shrinks to zero while the bulk entanglement becomes very large.

At some point there is a crossover where the two minima are degenerate. This defines the Page time. Because the evaporation is irreversible, this happens when the horizon area A_{hor} is slightly larger than half its initial area A_0 (for Schwarzschild black holes in our universe it happens when $A_{\text{hor}} \sim 0.6A_0$ [49]).

The important point is that the geometry has a Python’s Lunch separating the two minima. The generalized entropy at the maximum of the bulge is

$$\frac{A_0}{4G_N} + S_{\text{rad}}. \tag{5.1}$$

Now, suppose that, just after the Page time, Alice, who controls the Hawking radiation, wants to apply gates or a Hamiltonian to shrink the wormhole to the TFD associated with the black hole of area A_{hor} . Assuming that the analogy between a tensor network and the Cauchy slice holds, the protocol in the appendix she can do so in a time that is $O(S)e^{A_0/8G_N}$. This is consistent with the restricted complexity being exponentially large.

Of course, so far we have only considered one way of sweeping through the Cauchy slice, from one minimal cut to the other. If we could find another way of sweeping through the slice that had a smaller maximal generalized entropy, it would suggest that a more efficient protocol exists, since in a tensor network less post-selection would be required.




Figure 16. At late times, a more efficient way of sweeping between the two minimal cuts is to create two cuts near the horizon, and then to sweep one of these cuts “backwards” along the wormhole.

In fact, as far as we can tell, the way of sweeping through the slice just described should be optimal both before and shortly after the Page time. However, at late times, an alternative way of sweeping through the slice becomes preferable. We now analyze this second way of sweeping through the slice, which is shown in figure 16.

- In the first cut in figure 16 the only contribution to the generalized entropy comes from the bulk entanglement and is proportional to the entropy in the radiation, as before.
- The next cut also cuts the bulk entanglement between the interior and the radiation. However, it also includes an additional “double cut” near the horizon. This gives an additional area contribution equal to $2A_{\text{hor}}/4G_N$.
- In the third and fourth cuts in figure 16, one half of the double cut moves to the left. As it does so, its area increases, but the bulk entanglement decreases. Because black hole evaporation is irreversible, the second effect is slightly larger than the first, and so the generalized entropy slowly decreases.
- Finally, the left-moving cut reaches the end of the bridge-to-nowhere and disappears. The generalized entropy therefore has a sharp decrease by $A_0/4G_N$.

The evolution of the generalized entropy of the cut is shown in figure 17. The generalized entropy quickly increases as the double cut is added, reaching its maximum size of

$$S_{\text{rad}} + 2\frac{A_{\text{hor}}}{4G_N}. \tag{5.2}$$

It then slowly decreases, before a final sudden drop by $A_0/4G_N$.




Figure 17. A comparison of the generalized entropy as a function of sweep parameter for the ‘forwards’ (figure 14) and ‘reverse’ (figure 16) ways of sweeping through the bridge to nowhere. Initially, the generalized entropy of both is given by S_{rad} . The forwards sweep quickly increases by $A_0/4G_N$ and then steadily decreases as the sweep moves along the bridge-to-nowhere. The reverse sweep quickly increases by $2A_{\text{hor}}/4G_N$, slowly decreases as the cut moves backwards along the bridge-to-nowhere, and then finally quickly decreases by $A_0/4G_N$. The reverse sweep has a smaller maximum size, and hence is more efficient, when $A_{\text{hor}} < A_0/2$.

As shown in figure 17, this method of sweeping through the bridge-to-nowhere is therefore more efficient than the forwards sweep when $A_{\text{hor}} < A_0/2$. Note that this transition happens strictly after the Page time, which, although commonly described as happening at the halfway point in the evaporation, happens when $A_{\text{hor}} > A_0/2$.

5.3 The covariant Lunch

Of course, as discussed in section 4.4, the correct covariant definition of a Python’s Lunch is not given by the shape of a single Cauchy slice, but by a set of three quantum extremal surfaces, the two end surfaces and the bulge surface in the middle. However, as discussed in appendix B, there exist Cauchy slices within which the most efficient ‘sweep’ has locally minimal generalized entropy at the end surfaces, and locally maximal generalized entropy at the bulge surface. We can think of these as the Cauchy slices where the tensor network analogy is valid.

For an evaporating black hole, the location of the end surfaces were calculated in [35, 36]. The first end surface is the empty surface, containing no points. Its generalized entropy is simply the semiclassical entropy S_{rad} of the Hawking radiation. The second end surface, which becomes the EW surface after the Page time, lies at a radius that is $O(G_N)$ inside the horizon of the black hole, and at an infalling, or retarded, time that is one scrambling time in the past of the current boundary time. Its generalized entropy is given at leading order by the Bekenstein-Hawking entropy $A_{\text{hor}}/4G_N$.

The bulge surface was obviously less of a focus in [35, 36], because it is never the EW surface. However, the bulge surface that corresponds to the maximum size in the ‘forwards sweep’ was briefly discussed in [36]. For a one-sided black hole formed from collapse, it lies inside the infalling matter that created the black hole. As the black hole forms, a classical apparent horizon appears that moves outwards in a spacelike direction towards the lightlike event horizon (which, being teleological, formed before the infalling matter even arrived). After the infalling matter has passed through and the black hole begins to evaporate, the classical apparent horizon ends up a Planckian radial distance *outside* the event horizon. When no greybody factors are present, we can use eqs. (89) and (90) from [35] to see that there will exist a quantum extremal surface, at sufficiently late times, at a radius

$$r = r_s(v) - \frac{G_N c_{\text{evap}}}{3(d-1)\Omega_{d-1}r_s^{d-2}}, \tag{5.3}$$

where $r_s(v)$ is the Schwarzschild radius (and hence the radius of the apparent horizon) in the ingoing Vaidya metric describing the black hole, c_{evap} is the number of two-dimensional bosonic modes (i.e. number of angular momentum modes in higher dimensions) involved in the evaporation, and Ω_{d-1} is the volume of the unit $(d-1)$ -sphere, and at an infalling time v when

$$r_{\text{hor}}(v) = r_s(v) + \frac{G_N c_{\text{evap}}}{3(d-1)\Omega_{d-1}r_s^{d-2}}. \tag{5.4}$$

What is the infalling time v when is eq. (5.4) satisfied? Since the apparent horizon $r_s(v)$ goes from far inside to a Planckian distance outside the event horizon $r_{\text{hor}}(v)$, a solution must exist (assuming the metric is smooth) somewhere inside the infalling matter. Generally it will be near the end of the infalling matter, when the apparent horizon has approached within a Planckian radial distance of the event horizon.

The early time bulge surface is shown in figure 18. We note that its generalized entropy is indeed approximately equal to the initial Bekenstein-Hawking entropy of the black hole, plus the entropy of the Hawking radiation, as expected from our analysis of the maximal volume Cauchy slice (see [35, 36] for explicit calculations).

What about at late times, when our analysis of the maximal-volume slice suggested that a ‘reverse sweep’ through the bridge-to-nowhere was optimal? In this case, we expect that the dominant bulge surface should consist of the union of two spheres, both close to the late-time horizon. In general, calculating the location of an extremal surface with this topology is considerably harder than finding extremal surfaces consisting of a single sphere. However, it is possible for JT gravity plus free Dirac fermions. We study this theory in appendix C. As hoped, we find an extremal surface that consists of two points (or zero-spheres). Both points lie one scrambling time in the infalling past of the current boundary time, just like the late-time EW surface. However, as shown in figure 19, both points are spacelike separated from, and further inside the horizon than, the EW surface. This is consistent with this surface being the maximum generalized entropy surface in a sweep through a Cauchy slice that goes from the empty surface to the late-time EW surface. Consistent with our expectations based on a single Cauchy slice, the generalized entropy of this surface is approximately $2A_{\text{hor}}/4G_N + S_{\text{rad}}$.




Figure 18. The covariant bulge surface at early times (i.e. for the ‘forwards’ sweep) lies inside the infalling matter that forms the black hole. It lies a Planckian radial distance inside the apparent horizon (dashed line), at an infalling time when the apparent horizon is itself a Planckian radial distance inside the event horizon (solid line). Its generalized entropy is approximately $A_0/4G_N + S_{\text{rad}}$.




Figure 19. The covariant bulge surface (black) at late times (i.e. for the ‘reverse’ sweep) consists of two spheres (or points in two dimensions). Both lie slightly inside the EW surface (red), which determines the Hawking radiation entanglement wedge \mathcal{E}_A (green) and the entanglement wedge \mathcal{E}_B of the system containing the black hole (blue).

5.4 The time dependence of the restricted complexity

Using our covariant restricted complexity conjecture from section 4.4, we can now make a precise conjecture about how the restricted complexity of the evaporating black hole state evolves over the course of the evaporation. We show a plot of $\log \mathcal{C}_R$ against $A_0 - A_{\text{hor}}$ over the course of the entire evaporation in figure 20.

There are three distinct phases to this evolution. Before the Page time, the EW surface is empty, with generalized entropy S_{rad} , while the bulge surface is the forwards-sweep surface, with generalized entropy $S_{\text{rad}} + A_0/4G_N$. Finally, the larger end surface has generalized entropy $A_{\text{hor}}/4G_N$. According to our conjecture in eq. (4.8), the restricted complexity is controlled by the difference in generalized entropy between the bulge surface and the larger end surface. It is given by

$$\mathcal{C}_R = O \left(t \exp \left[\frac{1}{2} \left(S_{\text{rad}} + \frac{A_0 - A_{\text{hor}}}{4G_N} \right) \right] \right). \quad (5.5)$$

The factor of t here comes from the volume of the lunch, which controls \mathcal{C}_{TN} according to the conjectures of refs. [2–4].

The first phase transition happens at the Page time, where the EW surface becomes nonempty. This means that the larger end surface becomes the empty surface, with generalized entropy S_{rad} . The restricted complexity is therefore

$$\mathcal{C}_R = O \left(t \exp \left[\frac{A_0}{8G_N} \right] \right), \quad (5.6)$$

and only changes linearly with time.

Finally, when $A_{\text{hor}} = A_0/2$, we have a second phase transition. This time, it is the bulge surface that changes. It becomes the reverse-sweep surface, with generalized entropy $S_{\text{rad}} + 2A_{\text{hor}}/4G_N$. The restricted complexity begins to decrease, and is given by

$$\mathcal{C}_R = O \left(t \exp \left[\frac{A_{\text{hor}}}{4G_N} \right] \right). \quad (5.7)$$

Importantly, as the black hole completely evaporates the exponent tends to zero, and the restricted complexity becomes $O(t)$. This is exactly what we expect. The black hole has completely evaporated and so we have a one-sided system again. The restricted complexity will therefore be equal to the unrestricted complexity, which is $O(t)$.

6 Python’s Lunches beyond black hole evaporation

Black hole evaporation is an inherently quantum mechanical phenomenon. It violates the null energy condition (NEC), for example, even when all the quantum fields in the theory would classically obey the NEC. One might therefore think that Python’s Lunches themselves are an inherently quantum mechanical phenomenon and cannot exist in classical spacetimes. As we shall see, this is not at all true. Instead, there are numerous important examples, beyond black hole evaporation, of both classical and quantum lunches.



Figure 20. A plot of our conjecture for $\log C_R$ over the course of the black hole evaporation. For the purposes of this plot we have assumed $S_{\text{rad}} = 1.5(A_0 - A_{\text{hor}})$.

It is important to note that, in many of these examples, the size of the lunch is fixed in the semiclassical limit. This means that we would not expect a tensor network toy model to fully scramble the degrees of freedom over the course of the lunch. We should therefore be somewhat circumspect in conjecturing that the restricted complexities are actually exponentially large in these cases.

The first, and simplest, example of a Python’s Lunch is a two-sided black hole with a heavy brane, sitting in the Einstein-Rosen bridge, as shown in figure 21. The backreaction of the brane on the spacetime separates the left and right horizons, creating a Python’s Lunch. The two end surfaces lie on the left and right bifurcation surfaces, while the bulge surface lies at the intersection of the brane world line with the static slice.¹⁷ Unlike an evaporating black hole, this spacetime does not violate the null energy condition. However, it cannot be created from the thermofield-double state using semiclassical unitary Lorentzian evolution, because the entire worldline of the brane lies behind the black hole horizon. It can be created using a simple Euclidean path integral, as shown in figure 21, but Euclidean evolution is not unitary. Both facts are consistent with the restricted complexity of the state being very large.¹⁸

A second classical example of a Python’s lunch is a one-sided black hole, formed from collapse. This is shown in figure 22. The two ends of the python are just two halves of the single boundary. The bulging lunch in the middle is just the bridge-to-nowhere. On each side of the black hole there are classical extremal surfaces, which form the ends of the lunch. What about the bulge surface? At sufficiently early times, it seems likely that

¹⁷Formally, in the classical spacetime, we would need to smear the energy of the brane out slightly in order to have a smooth spacetime and hence an actual extremal surface here. This will inevitably happen once we include quantum effects.

¹⁸In this case, it is not obvious that even the *unrestricted* unitary complexity of the state should be small. In fact, we should probably expect it to be large, because there is a nonempty extremal surface even for the combined left and right boundaries, see section 7.1.




Figure 21. Adding a positive tension brane to the center of the Einstein-Rosen bridge separates the left and right horizons, creating a Python's Lunch. The Lorentzian geometry of such a state is shown in (a). A Cauchy slice is shown in (b). Such a state can be easily constructed using a Euclidean path integral with a heavy operator insertion halfway between the left and right boundaries, as shown in (c). However it cannot be easily constructed using unitary evolution from the thermofield-double state.




Figure 22. A one-sided black hole forms a Python's Lunch between one half of the boundary, region A , and the other half, region \bar{A} . There is an extremal surface either side of the black hole, with the 'bridge-to-nowhere' of the black hole forming the bulging lunch.

the bulge surface will go inside the horizon (for that matter, at sufficiently early times the end surfaces will also go inside the horizon). However, we would expect the area of any extremal surface going inside the horizon to grow with time. At late times, we should instead expect the bulge surface to 'wrap around' the horizon. Indeed, for a BTZ black hole, it will be a self-intersecting geodesic that winds once around the horizon.

It is easy to see that the unrestricted complexity of this one-sided black hole state is small. It will be proportional to the time since the black hole first formed. However, because this time evolution couples the two halves of the boundary, then we should expect the restricted complexity to be very large, at least at late times. Again, the existence of a Python's Lunch corresponds to a large gap between restricted and unrestricted complexity.

A more quantum example, closer in spirit to the single-sided black hole evaporation studied in section 5.2, but with a connected classical geometry everywhere, goes as follows.




Figure 23. Coupled evolution of the thermofield-double state creates a quantum lunch in a classical python. The classical cross-section of the wormhole remains constant in size, but there is long range entanglement between the two sides, which increases the generalized entropy of a cut through the middle of the wormhole.

We start with the thermofield-double state, and then evolve it forwards in time, but with a small coupling between the left and right boundaries. This coupling will mix Hawking radiation between the two exteriors, causing Hawking radiation from the right to end up falling into the black hole from the left and vice versa.

Because the two sides of the black hole are in thermal equilibrium with one another, the size of the black hole will stay approximately constant. There will be no ‘classical’ Python’s Lunch. However the coupling between the two sides will create long-range entanglement between the quantum fields at each end of the wormhole, via Hawking radiation escaping one end and then falling into the other. This creates a quantum lunch (in a classical python). See figure 23.

Again, the unrestricted complexity of the state should be small, because it was created by a simple unitary evolution from the thermofield-double state. However, this evolution coupled the two sides, and so the restricted complexity may well be large.

Our final example of a Python’s Lunch is the AdS_3 vacuum. If we divide the boundary into two connected halves, as with the one-sided black hole discussed earlier in this section, there is no Python’s Lunch. However, if each end of the python itself consists of two disconnected regions, as shown in figure 24, a lunch appears. There are two topologically distinct end extremal surfaces, plus a self-intersecting bulge surface in the middle. This suggests that the restricted complexity of constructing the vacuum state without coupling these two complementary disjoint regions may be very large. However, the small volume of the lunch suggests that it won’t be fully scrambling and so our restricted complexity conjecture (4.6) may not apply. It is not the aim of this paper to revisit the details of complexity = action proposal but rather briefly discuss several examples that will be helpful to distinguish between unrestricted unitary complexity and size of tensor network with non-unitary elements.

7 What is holographic complexity?

There have been various proposed definitions for the bulk quantity that is to be holographically dual to boundary complexity. The two most prominent proposals have been the




Figure 24. Even the AdS_3 vacuum contains a Python’s Lunch when viewed as a python from the boundary region A consisting of two disjoint intervals to its complementary region \bar{A} .

volume of the maximal-volume slice [2], and the action of the Wheeler-de Witt patch [3, 4]. In practice, these two proposals tend to give similar answers.

Dual to this abundance of promising bulk quantities is the abundance of promising boundary quantities. Various definitions of boundary complexity have been considered. The original suggestion was that it should be unitary circuit complexity (the minimal number of simple gates, from a given primitive gate set, required to build the state from a simple starting point). There is also a ‘continuous’ variant on the unitary circuit complexity, which is defined using the Nielsen geometry [50]. Finally, there is the intuition that the volume of a slice is dual to the size of a tensor network required to build the state. As we have seen in this paper, this is different from the unitary circuit complexity, because a tensor network may contain non-unitary elements.

One way to make a highly complex state with a small volume/action is to have the state of the bulk fields be highly complex. This suggest that the holographic complexity should have a ‘quantum correction’, similar to the quantum corrections to the Ryu-Takayanagi formula, given by the complexity of the state of the bulk fields.

As we have seen in this paper, even when the bulk fields are in a simple state, the *restricted* unitary circuit complexity may be much larger than the volume/action, if the geometry contains a Python’s Lunch. This is not a contradiction with the conjectures of refs. [2–4], since the *unrestricted* complexity is still small, and comparable to the volume/action. This has been the case in most of the discussion of this paper. However, as we shall see in this section, there are examples one can construct that the *unrestricted* unitary circuit complexity can be much larger than the volume/action if we consider states that are prepared using non-unitary processes and therefore may contain ‘one-sided Python’s Lunches’. This is true even when the bulk state is very simple. This suggests that the more




Figure 25. A one-sided black hole is allowed to evaporate, and then the Hawking radiation is measured. This produces a pure black hole microstate with interior modes that are in some simple state that depends on the measurement outcome. In the maximal volume slice, there is a local minimum of the generalized entropy near the horizon, where the area is smallest.

general and correct dual quantity to the holographic complexity is the size of the tensor network required to make the quantum state, or equivalently the circuit complexity where non-unitary post-selections onto the outcomes of simple measurements are allowed.

7.1 Measuring radiation and one-sided lunches

Suppose we take a one-sided black hole, and allow it to evaporate as in section 5.2. However, rather than storing the radiation in a second system, we instead measure it in some complete basis. This basis does not have to be complicated: it can be a product basis, for example. The black hole will now be in a pure state; in particular, the interior modes that were previously entangled with the Hawking radiation will now be in a pure state that depends on the measurement outcome.

The resulting bulk geometry can be thought of as a ‘one-sided Python’s Lunch’, with a bridge to nowhere which is largest at the end and then becomes gradually smaller as one approaches a quantum extremal surface near the horizon, as shown in figure 25. The exact location of this quantum extremal surface is hard to calculate, but it is easy to show that it should exist, as argued in figure 26.

The volume of the maximal volume slice, and the action of the Wheeler-de Witt patch will grow linearly with the time the black hole was allowed to evaporate for. We expect that this will also be the size of the minimal tensor network needed to describe the state. However, if this tensor network resembles the bulk geometry, its cross-section will be largest at the end of the wormhole and then become smaller near the horizon. Such a tensor network cannot generically be produced by a unitary circuit of the same size, without allowing post-selection.

Instead, applying our restricted complexity conjecture, in the special case where one system is trivial and so the restricted complexity is actually just the unrestricted complexity, we find that the unitary circuit complexity should be proportional to $\exp[(S_{\max}^{(\text{gen})} - S_R^{(\text{gen})})/2]$, where $S_R^{(\text{gen})} \approx A_{\text{hor}}/4G_N$ is the generalized entropy of the quantum extremal surface that we just discussed and $S_{\max}^{(\text{gen})} \approx A_0/4G_N$ is the generalized entropy of a bulge surface inside the infalling matter that first formed the black hole.




Figure 26. To argue that a nonempty quantum extremal surface should exist for the pure state of a two-dimensional black hole, after measuring Hawking radiation, we consider Cauchy slices that asymptote to the maximal volume slice in the distant past, but which are allowed to vary elsewhere. In particular, we allow the slice to vary at infalling times approximately one scrambling time in the past of the current boundary time (this is where the non-empty extremal surface was found in [35, 36]). As shown in figure 25, such slices should have a local minimum in their generalized entropy near the horizon. If we vary the Cauchy slice too far into the interior of the black hole, this local minimum will become small because the area becomes small. Conversely, if we push the slice too close to the past lightcone, the bulk entropy will become very small, which will also decrease the generalized entropy. By choosing our Cauchy slice to maximize the generalized entropy of this local minima, we would necessarily find a non-empty extremal surface, in a location similar to the surface found in [35, 36].

This is exactly what we should expect. The state was prepared using a measurement, which can only be reproduced deterministically by using post-selection, or by using Grover search, as discussed in section 4 and appendix A. The complexity of this process is indeed exponential in the number of post-selected qubits.

One might worry that this conclusion seems wrong: what about if we just reversed time, ignoring the fact that the measurement had happened, until we got back to a time before the black hole ever formed? Once there was no black hole, it would presumably be easy to get back to a simple state using a simple circuit. The answer (see [33, 35, 36, 51] for similar discussion) is that, by measuring the Hawking radiation, we necessarily create a small positive-energy localized shock that approaches the horizon of the black hole as we go backwards in time. At an infalling time approximately the scrambling time in the past, the backreaction of this shock becomes significant and it turns the black hole into a white hole. If we continue to evolve the boundary backwards in time, the black hole will never disappear and the interior modes will never escape. Instead we will just see a time-reversed version of Hawking radiation coming out of the newly created white hole. This is shown in figure 27.

How does this correspond to a tensor network model? As we evolve time backwards, the simplest tensor network describing the state initially becomes smaller, as the reverse time evolution undoes tensors that were previously added to the network by the forward evolution. However, we cannot ‘undo’ the projections created by the measurement of Hawking radiation. At one scrambling time in the past, these projections will have infected




Figure 27. (a) Penrose diagram for the spacetime of a one-sided black hole, with the Hawking radiation extracted and then measured. (b) If we reverse the time evolution after the measurement is done, the black hole won't disappear as one might naively expect. Instead, at one scrambling time in the past, the backreaction from the measurement (shown in red as a small shockwave) will create a white hole. The interior partners of the Hawking radiation will therefore never escape from behind the horizon.

the entire cross section of the network. Further backwards time evolution cannot remove any more tensors from the network, and so instead will have to add new tensors: the size of the simplest tensor network therefore ‘bounces’ from this minimal size and begins to increase. This corresponds to the appearance of the white hole horizon.

Another example of a one-sided Python’s Lunch is the brane-in-a-wormhole state discussed at the start of section 6. In that section, we used it as an example of a two-sided lunch. However, the union of the left and right extremal surfaces is homologous to the union of the left and right boundaries. We therefore have a non-empty classical extremal surface for the union of the two boundaries, which creates a one-sided Python’s Lunch. This suggests that even the unrestricted unitary complexity of the state may be very large. This is perfectly consistent, since we only know how to prepare such a state using a non-unitary Euclidean evolution.

We have argued that states with a non-empty extremal surface for the entire global boundary (i.e. a one-sided Python’s Lunch) have high unrestricted unitary complexity. A good consistency check on this claim is that such states cannot be created by semiclassical Lorentzian evolution from states with a Cauchy slice that is entirely within the causal wedge of the boundary. This is indeed true: the entanglement wedge, defined using *any* extremal surface (not just the HRT surface) must contain the causal wedge, and hence the entire spacetime, by standard focussing arguments (this requires the generalized second law or quantum focussing conjecture [52] in the case of quantum extremal surfaces). It is therefore entirely consistent that they should always have very high unitary circuit complexity.

7.2 Post-selected state complexity

The black hole with measured Hawking radiation appears to be a counterexample to the idea that unitary circuit complexity equals volume/action. The volume and action only grow linearly with the time the black hole is allowed to evaporate for, but the size of the

simplest unitary circuit required to produce the state appears to grow exponentially. One response to this problem is to note that this state could only be produced by measuring Hawking radiation, and that a measurement really corresponds to entangling the state with an ancilla measurement apparatus. If we include the measurement apparatus, the unrestricted complexity of the state (including the measurement apparatus) will still be small.

This argument is somewhat unsatisfying. It would be nice to have a boundary quantity (such as complexity) that corresponds to volume/action even for states that can only be produced using post-selection, or by Euclidean path integrals as with the first example from section 6.

There is an obvious candidate quantity. We just redefine the notion of state complexity to allow post-selection onto simple states (say $|0\rangle$). Equivalently, we define it as the size of the smallest tensor network required to make the state.

With the usual definitions the complexity of creating a state from a simple state is the same as the complexity of starting with the state in question and returning to the simple state. But when post-selections are allowed, this changes.¹⁹ The relevant complexity is $\mathcal{C}_p(I, \psi)$ i.e. the complexity of creating the state $|\psi\rangle$, from, the simple state $|0\rangle^{\otimes n}$, using both unitary gates and simple projections.

For specific states $|\psi\rangle$, allowing projections may dramatically reduce the number of operations required, as we have seen. However, for typical states, projections don't buy you much. As we shall see below with a counting argument, it still takes $\sim 4^N$ operations to get to the most general state.

As an aside, it is worth noting that it doesn't significantly matter whether we allow post-selection at arbitrary intermediate points in the state preparation process, or only at the end after all the unitary gates have been applied. This is because we can always implement the desired post-selection by using a unitary operator that 'measures' the relevant register into an ancilla quantum register (this is sometimes called a von Neumann measurement of the first kind) and then post-selecting the ancilla register after all the other unitaries have been applied.

7.3 Post-selected complexity can be exponential

The maximal unitary circuit complexity of an N qubit state scales exponentially with N . This can be established with a counting argument. On the one hand Hilbert space is double-exponentially huge

$$\text{number of } \epsilon\text{-balls in } N\text{-qubit Hilbert space} \sim 2^{2^N}. \tag{7.1}$$

On the other hand, the number of states that can be made with \mathcal{C} gates (for definiteness, let's say we have a universal 2-local gate together with a 1-local phase) is merely exponentially large. At each step, we can apply our 2-local gate in one of $\binom{N}{2}$ places, or apply our

¹⁹Note that $\mathcal{C}_p(I, \psi) \neq \mathcal{C}_p(\psi, I)$, where $\mathcal{C}_p(\cdot, \cdot)$ is relative state complexity, when projection into simple states is allowed. In fact, going from $|\psi\rangle$ to $|I\rangle$ is trivial if post-selection is allowed, with $\mathcal{C}_p(\psi, I) \leq N$.

1-local phase in one of $\binom{N}{1}$ places, so

$$\#states \lesssim \left(\binom{N}{2} + \binom{N}{1} \right)^c. \tag{7.2}$$

In order to reach all the states, \mathcal{C} must be exponentially big.

However we have now changed our definition of ‘complexity’. We have given ourselves the power not only to apply 2-local gates, but also to project the first m -qubits to $|0000000\rangle$. Since this increases our power, it decreases the complexity. This gives rise to the state synthesis version of the PostBQP complexity class, and as we saw in section 4.3 this is fantastically powerful — there are many quantum states that would normally be exponentially hard to make that are now easy. Are they, in fact, *all* now easy?

Let’s prove that the answer to this question is ‘no’. We will argue that even granting ourselves the power to post-select, there are still states that are exponentially hard to make. This can again be established with a counting argument. We have more options than before. At each step, as well as applying our 2-local gate or our 1-local phase, we can also project on the first m qubits²⁰ for any $1 \leq m \leq N$. Thus the number of different states we can make is

$$\#states \lesssim \left(\binom{N}{2} + \binom{N}{1} + N \right)^c, \tag{7.3}$$

which — the point is — is still only exponentially big. We still need exponentially large \mathcal{C} to hit all of the double-exponentially numerous ϵ -balls.

It would be nice to have a definition of post-selected state complexity that did not rely on a choice of discretization of the Hilbert space into ϵ -balls. For unitary state complexity, one such definition is the smallest geodesic distance from the identity to a unitary taking one state to the other in the so-called Nielsen geometry [50]. This is a right-invariant (but not left-invariant) metric on the space of unitaries where distances are much smaller in simple directions (generated by k -local Hamiltonians) than in other directions.

However, if we allow post-selection on a single qubit, then any state can be prepared using a unitary that has arbitrarily small complexity as measured by the Nielsen geometry. The reason for this is that the Nielsen metric (like any metric) is continuous. Hence the complexity of a unitary mapping

$$|0\rangle|\psi\rangle \rightarrow |0\rangle|\psi\rangle + \varepsilon|1\rangle|\phi\rangle \tag{7.4}$$

can be made arbitrarily small by making ε sufficiently small. However, we can always post-select onto $|1\rangle$ and produce $|\phi\rangle$, no matter how small ε is.

Instead, it seems like the right continuous measure of post-selected state complexity would be to make the cost of the post-selection be $\log p$, where p is the probability of obtaining the correct measurement outcome. For typical states where the amplitudes of the post-selected outcomes are not exceptionally large or small, this still corresponds to having an $O(1)$ cost for each post-selected qubit.

²⁰Notice that it would not have made a difference had we given ourselves the power to post-select *any* m qubits in the computational basis, since using just m swap gates we can always move the desired qubits to the front. On the other hand, it would have made a huge difference had we given ourselves the power to post-select in *any* basis — then all state synthesis is trivial, since we just post-select onto the desired state!

8 Summary

This paper has addressed an apparent inconsistency between the holographic complexity conjectures [2–4] and the Harlow-Hayden result [1]. The inconsistency is manifest in an evaporating black hole slightly after the Page time: on the one hand, the volume or action of the black hole is only polynomial in the entropy S , and thus the holographic complexity must be moderate; on the other hand, Harlow & Hayden argue that the complexity of decoding the Hawking radiation must be exponentially large. The difference arises from using different definitions of complexity. The holographic complexity conjectures relate the volume/action of the geometry to *unrestricted* complexity, which allows gates that span the entire system; whereas Harlow & Hayden’s result is about *restricted* complexity, which forbids gates that couple the interior of the black hole to the previously emitted Hawking radiation.

This distinction motivated us to ask: if action or volume are the geometric duals of unrestricted complexity [2–4], what is the geometric dual of restricted complexity? We conjectured an answer. Exponentially large restricted complexity corresponds to the existence of a geometrical feature that we call a “Python’s lunch”. In a Python’s lunch, the cross-sectional area of the wormhole grows and then shrinks again, in a min-max-min pattern. The restricted complexity, we conjectured in eq. (4.6), is given by the exponential of the difference between the area of the maximum and the area of the larger of the two flanking minima.

We tested this conjecture in a toy tensor-network model, and found agreement with the Harlow-Hayden estimate. We then made a covariant version of our conjecture, eq. (4.8), by replacing the min and max areas of the Python’s Lunch with generalized entropies of appropriate quantum extremal surfaces. With this generalization, we studied several examples of the Python’s Lunch and estimated the restricted complexities in each case, including evaporating black holes, one-sided pure-state black holes, and empty AdS with two disjoint intervals. In all cases where we were able to test our conjecture, the restricted complexity was consistent with the size of the Python’s Lunch.

Lastly, in section 7 we returned to the subject of unrestricted complexity. We studied the example of black holes that have had all their Hawking radiation measured, and which therefore have been rendered pure. Using this example, we reconsidered exactly which boundary quantity it is that is holographically dual to the volume or action of the wormhole. In refs. [2–4] it was conjectured that this quantity is the unrestricted unitary circuit complexity, which means the allowed primitive gates are all unitary. Instead, we argued that the definition of unrestricted holographic complexity should also permit non-unitary post-selection — holographic complexity should allow projections onto simple states.

Acknowledgments

We thank Andras Gilyen, Patrick Hayden, John Preskill, Stephen Shenker for fruitful discussions. H.G. is supported by the Simons Foundation through the It from Qubit

collaboration. G.P. is supported in part by AFOSR award FA9550-16-1-0082 and DOE award DE-SC0019380. L.S. is supported by NSF Award Number 1316699.

A Complexity of post-selection

A lemma we will use repeatedly in this appendix is that if an N -qubit state $|\psi\rangle$ is simple, then so too is the unitary

$$|\psi\rangle \text{ is simple} \rightarrow U_\psi \equiv \mathbb{1} - 2|\psi\rangle\langle\psi| \text{ is simple} . \tag{A.1}$$

This unitary flips the sign of the $|\psi\rangle$ component of a wavefunction, while leaving all orthogonal components unchanged. Let's explicitly construct a simple circuit that does this. First, note that what it means for $|\psi\rangle$ to be simple is precisely that there is a simple unitary $U_{|\bar{0}\rangle\rightarrow|\psi\rangle}$ that connects it to the reference state, $|\psi\rangle = U_{|\bar{0}\rangle\rightarrow|\psi\rangle}|\bar{0}\rangle$. We can therefore write

$$U_\psi = U_{|\bar{0}\rangle\rightarrow|\psi\rangle} U_{\bar{0}} U_{|\bar{0}\rangle\rightarrow|\psi\rangle}^\dagger; \tag{A.2}$$

this first transforms to a basis in which the $|\psi\rangle$ component of the wavefunction becomes the $|\bar{0}\rangle$ component, then flips the phase of the $|\bar{0}\rangle$ component, then transforms back again. This construction upper bounds the complexity

$$\mathcal{C}[U_\psi] \leq 2\mathcal{C}[U_{|\bar{0}\rangle\rightarrow|\psi\rangle}] + \mathcal{C}[U_{\bar{0}}] = 2\mathcal{C}[|\psi\rangle] + \mathcal{C}[U_{\bar{0}}]. \tag{A.3}$$

We can understand the factor of 2 in this equation as arising from the fact that to make $\mathbb{1} - 2|\psi\rangle\langle\psi|$ the protocol sweeps twice over the circuit that manufactures the state, first to unmake it (the $\langle\psi|$ part) then to make it again (the $|\psi\rangle$ part).

A.1 Projecting on a ququit

Suppose we have a simple unitary U that maps a simple state to a superposition

$$U|s\rangle|1\rangle = \frac{|\alpha\rangle|1\rangle + |\beta\rangle|2\rangle + |\gamma\rangle|3\rangle + |\delta\rangle|4\rangle}{2}. \tag{A.4}$$

Since U is simple, and $|s\rangle|1\rangle$ is easy to make, the right-hand side must also be easy to make.

- Question: how complex can it be to make $|\alpha\rangle|1\rangle$?

As we will see, the answer is “not complex”.

One simple strategy to make $|\alpha\rangle|1\rangle$ is just to make $U|s\rangle|1\rangle$ and then measure the last ququit. Sometimes we'll find the last ququit to be $|2\rangle$, $|3\rangle$, or $|4\rangle$; if that happens we throw the state away and start over. Other times we'll find the last ququit to be $|1\rangle$ and can declare victory. In this way we can make an $|\alpha\rangle$ -factory that has efficiency $\frac{1}{4}$.

For some situations, this simple strategy suffices. But if the projection is a step buried deep within a larger circuit, starting again might not be so easy. Or if our initial state is entangled with another state that we do not control, then starting again might be impossible. And if we want to project not 1 ququit but N ququits, the probability of success falls like 2^{-2N} .

The solution to this problem is to use a close cousin of Grover’s algorithm, as we will now explore. First, let’s define the ‘pre-image’ of the four possible answers we could have gotten

$$|\tilde{\alpha}\rangle \equiv U^\dagger|\alpha\rangle|1\rangle; \quad |\tilde{\beta}\rangle \equiv U^\dagger|\beta\rangle|2\rangle; \quad |\tilde{\gamma}\rangle \equiv U^\dagger|\gamma\rangle|3\rangle; \quad |\tilde{\delta}\rangle \equiv U^\dagger|\delta\rangle|4\rangle, \quad (\text{A.5})$$

so that $|s\rangle|1\rangle = \frac{1}{2}(|\tilde{\alpha}\rangle + |\tilde{\beta}\rangle + |\tilde{\gamma}\rangle + |\tilde{\delta}\rangle)$. Next observe that it is simple to make the operator $U_{\tilde{\alpha}}$ that flips the sign of the $|\tilde{\alpha}\rangle$ term in a wavefunction while leaving the other terms invariant. We can do this by evolving the state with U , doing a sign-flip controlled on the last ququit being $|1\rangle$, and then evolving back with U^\dagger

$$V \equiv \mathbf{1} - 2|\tilde{\alpha}\rangle\langle\tilde{\alpha}| = U^\dagger\left(\mathbf{1} \otimes (\mathbf{1} - 2|1\rangle\langle 1|)\right)U. \quad (\text{A.6})$$

Finally, since $|s\rangle|1\rangle$ is by assumption easy to build, it must also be easy to build the operator that flips everything except $|s\rangle|1\rangle$,

$$U_s \equiv 2|s\rangle\langle s| \otimes |1\rangle\langle 1| - \mathbf{1}. \quad (\text{A.7})$$

Now we concatenate these easy operations to give the desired projection

$$V|s\rangle|1\rangle = \frac{1}{2}(-|\tilde{\alpha}\rangle + |\tilde{\beta}\rangle + |\tilde{\gamma}\rangle + |\tilde{\delta}\rangle) \quad (\text{A.8})$$

$$U_s V|s\rangle|1\rangle = |\tilde{\alpha}\rangle \quad (\text{A.9})$$

$$U\left[U_s V\right]|s\rangle|1\rangle = |\alpha\rangle|1\rangle. \quad (\text{A.10})$$

Note that the only reason this worked was because there was an easily addressed ququit to diagnose the final branch. If we just had $U|s\rangle = |\alpha\rangle + |\beta\rangle + |\gamma\rangle + |\delta\rangle$ then this method wouldn’t have worked, and indeed couldn’t have worked since there is such a decomposition for *any* $|\alpha\rangle$, even if U is the identity.

A.2 Projecting on a qudit

The case of the last subsection was misleadingly easy: we needed only a single implementation of $U_s V$ to hit the target state exactly. More generally, we may wish to project onto the value of a qudit, where we can think of $d = 2^m$ for m qubits or $m/2$ ququits. And it may be that the amplitudes of the states are not evenly distributed. Let’s start with an initial state $|s\rangle|0\rangle^{\otimes n}$ on total of k qubits with n -ancillary qubits at state $|0\rangle$. Then, we wish to project (post-select) onto the outcome of m qubits. The output state after the unitary U is applied can be expressed as

$$U|s\rangle|0\rangle^{\otimes n} = \sin\theta|\alpha\rangle|0\rangle^{\otimes m} + \cos\theta|\beta\rangle, \quad (\text{A.11})$$

where $|\beta\rangle$ is any normalized state with $\langle 0|^{\otimes m}|\beta\rangle = 0$. Repeating the procedure of the last subsection gives

$$U\left[U_s V\right]|s\rangle|0\rangle^{\otimes n} = \sin 3\theta|\alpha\rangle|0\rangle^{\otimes m} + \cos 3\theta|\beta\rangle. \quad (\text{A.12})$$

where we denote by

$$U_s = 2|s\rangle\langle s| \otimes |0\rangle\langle 0|^{\otimes n} - \mathbf{1} \quad (\text{A.13})$$

$$V = U^\dagger\left(\mathbf{1} \otimes (\mathbf{1} - 2|0\rangle\langle 0|^{\otimes m})\right)U \quad (\text{A.14})$$

for n ancillary and m postselected qubits. We see that in the last subsection we got lucky, since $\sin \theta = \frac{1}{2} \rightarrow \sin 3\theta = \sin 3\frac{\pi}{6} = 1$. For more general θ , a single iteration will not yield the desired projection. Iterating l times gives

$$U[U_s V]^l |s\rangle|0\rangle^{\otimes n} = \sin[(2l+1)\theta]|\alpha\rangle|0\rangle^{\otimes m} + \cos[(2l+1)\theta]|\beta\rangle. \quad (\text{A.15})$$

Thus the number of iterations to implement the projection is given by $(2l+1)\theta = \frac{\pi}{2}$ so

$$\text{number of implementations of } [U_s V] = \frac{\pi}{4\theta} - \frac{1}{2}. \quad (\text{A.16})$$

When $U|s\rangle|0\rangle^{\otimes n}$ is an equal superposition of $d(=2^m)$ states, we have $\theta = \arcsin[1/\sqrt{d}]$ and this gives the celebrated large- d Grover scaling $l \sim \frac{\pi}{4}\sqrt{d}$.

Finally, we must confront the possibility that eq. (A.16) does not give a whole number. This is problematic since in general even if a U is easy to implement, \sqrt{U} may be hard. We could lower our ambitions by implementing the integer part of n and settling for being approximate. But we can do better. We first introduce a fresh qubit, and then use it to bleed some of amplitude out of $\sin \theta|\alpha\rangle|0\rangle^{\otimes m}$

$$|\bar{s}\rangle = U_\phi U|s\rangle|0\rangle^{\otimes n}|0\rangle = \cos \phi \sin \theta|\alpha\rangle|0\rangle^{\otimes m}|0\rangle + \sin \phi \sin \theta|\alpha\rangle|0\rangle^{\otimes m}|1\rangle + |\text{other}\rangle. \quad (\text{A.17})$$

Then we repeat the iterative procedure using $|\bar{s}\rangle$ instead of $|s\rangle|0\rangle^{\otimes n}$ and I instead of U , and carefully choose ϕ to land on the next integer greater than $\frac{\pi}{4\theta} - \frac{1}{2}$,

$$\text{number of implementations of } U_{\bar{s}} V = \left\lceil \frac{\pi}{4\theta} - \frac{1}{2} \right\rceil. \quad (\text{A.18})$$

Even though it is in general not easy to implement a fractional power of an easy unitary, for the specific unitaries we are considering it is.

Let us see what this analysis means for the simplest possible case, that of an evenly split qubit, with $\theta = \frac{\pi}{4}$. Equation (A.12) made the situation look hopeless — we would just cycle in a loop forever, $\theta = \frac{\pi}{4} \rightarrow \frac{3\pi}{4} \rightarrow \frac{\pi}{4} \rightarrow \frac{3\pi}{4} \rightarrow \dots$, never getting any closer to the target state. But now we see that the correct procedure is to first add an extra qubit, and then use $\phi = \frac{\pi}{4}$ to transform the pair of qubits to an even superposition of a ququit, returning us to the exactly implementable example of section A.1.

A.3 Projecting on very unlikely outcomes

Suppose we wish to project onto a final state that has tiny amplitude, $\theta \ll 1$. For example, we may have an equal superposition over m qubits with large m , giving $\theta = 2^{-m/2}$. How complex is this projection? Let us examine three possible methods:

- Measure-and-pray. If we measure the qubits and hope for the right answer, the probability that we get lucky is θ^2 , giving

$$\text{measure-and-hope method: } \text{complexity} \sim \theta^{-2} = d = 2^m. \quad (\text{A.19})$$

- All-at-once Grover-style projection. Using the method of section A.2, we saw in eq. (A.18) that we can effect a square-root speed up,

$$\text{all-at-once Grover: complexity} \sim \theta^{-1} = \sqrt{d} = 2^{m/2}. \quad (\text{A.20})$$

- Step-by-step Grover-style projection. In eq. (A.20), we simultaneously projected onto the values of all the target qubits. An alternative strategy would be to project on each target ququit in turn, so that each individual projection is then onto a state that is not particularly unlikely. As discussed in appendix A.1, the complexity of projecting onto the first ququit is $3\mathcal{C}(U)$, because we have to implement the unitary U three times to do the post-selection. What about the projection onto the second ququit? Now the unitary U has been replaced by the unitary U' that also does the projection onto the first ququit. So the complexity of this second projection is $3\mathcal{C}(U') = 9\mathcal{C}(U)$. It should be clear that the complexity of projecting onto each ququit grows exponentially. The complexity of the full projection is therefore

$$\text{step-by-step Grover: complexity} \sim 3^{m/2}, \quad (\text{A.21})$$

since there are $m/2$ ququits.

Summary: to project on m qubits, guess-and-check costs 2^m , all-at-once-Grover costs $O(2^{m/2})$, and one-by-one-Grover costs $O(3^{m/2})$. The winning strategy is to project on all the qubits at once, as in section A.2. The complexity of doing so is $O(2^{m/2})$.

A.4 Removing the state dependence

So far we have only tried to construct a unitary that produces a single output state $|\alpha\rangle$, given input $|s\rangle|0\rangle^{\otimes n}$ and unitary U . We used a unitary sequence

$$U \left[\left(2|s\rangle\langle s| \otimes |0\rangle\langle 0|^{\otimes n} - \mathbb{1} \right) U^\dagger \left(\mathbb{1} \otimes (\mathbb{1} - 2|0\rangle\langle 0|^{\otimes m}) \right) U \right]^l \quad (\text{A.22})$$

that works only for the particular input state $|s\rangle$.

Our actual task is somewhat more complicated. We need to construct a unitary circuit that produces the same output as our post-selected circuit for *any* input state $|s\rangle$. In other words, we want to find a unitary \tilde{U} , such that for any input state $|s\rangle$ we have

$$\tilde{U} |s\rangle |0\rangle^{\otimes n} = C |0\rangle^{\otimes m} \langle 0|^{\otimes m} U |s\rangle |0\rangle^{\otimes n}, \quad (\text{A.23})$$

where U is a simple unitary, C is a numerical constant. It turns out that we can easily adapt our construction from previous section to produce such a unitary if such a unitary \tilde{U} exists at all. Our construction is very closely related to *robust oblivious amplitude amplification*, which was independently introduced in [20–22]. We became aware of this work after this section of the manuscript had been completed.

Importantly, for the moment, we shall assume that an *exact* unitary \tilde{U} exists that exactly satisfies (A.23). We shall discuss what happens when \tilde{U} can only approximately satisfy (A.23) at the end of this subsection.

Let $|i\rangle$ be a computational basis for the input Hilbert space. By our assumption that there exists some unitary \tilde{U} satisfying (A.23), it follows that

$$\langle i | \langle 0 |^{\otimes n} U^\dagger | 0 \rangle^{\otimes m} \langle 0 |^{\otimes m} U | j \rangle | 0 \rangle^{\otimes n} = \frac{1}{C^2} \delta_{ij}. \quad (\text{A.24})$$

This implies that

$$U | i \rangle | 0 \rangle^{\otimes n} = \sin \theta |\alpha_i\rangle | 0 \rangle^{\otimes m} + \cos \theta |\beta_i\rangle, \quad (\text{A.25})$$

where θ is independent of i and $\sin \theta = 1/\sqrt{C}$, $\langle \alpha_i | \alpha_j \rangle = \langle \beta_i | \beta_j \rangle = \delta_{ij}$ and

$$\langle 0 |^{\otimes m} |\beta_i\rangle = 0 \quad (\text{A.26})$$

for all i . We can get rid of state dependence of the original protocol by replacing U_s in (A.13) with

$$W = \mathbb{1} \otimes \left(2|0\rangle\langle 0|^{\otimes n} - \mathbb{1} \right). \quad (\text{A.27})$$

This is independent of the input state $|s\rangle$. For an arbitrary state $|s\rangle$, we again have

$$U | s \rangle | 0 \rangle^{\otimes n} = \sin \theta |\alpha\rangle | 0 \rangle^{\otimes m} + \cos \theta |\beta\rangle, \quad (\text{A.28})$$

and one can check that

$$U \left[WV \right] | s \rangle | 0 \rangle^{\otimes n} = \sin 3\theta |\alpha\rangle | 0 \rangle^{\otimes m} + \cos 3\theta |\beta\rangle. \quad (\text{A.29})$$

The key step is that (A.25) implies

$$\langle 0 |^{\otimes n} U^\dagger \left[\cos \theta |\alpha\rangle | 0 \rangle^{\otimes m} - \sin \theta |\beta\rangle \right] = 0. \quad (\text{A.30})$$

Repeating the process N times gives

$$U \left[WV \right]^l | s \rangle | 0 \rangle^{\otimes n} = \sin[(2l+1)\theta] |\alpha\rangle | 0 \rangle^{\otimes m} + \cos[(2l+1)\theta] |\beta\rangle. \quad (\text{A.31})$$

The rest of the argument is identical to that in appendix A.2, proving that

$$U \left[\left(\mathbb{1} \otimes (2|0\rangle\langle 0|^{\otimes n} - \mathbb{1}) \right) U^\dagger \left(\mathbb{1} \otimes (\mathbb{1} - 2|0\rangle\langle 0|^{\otimes m}) \right) U \right]^l \quad (\text{A.32})$$

will output state $|\alpha\rangle | 0 \rangle^{\otimes m}$ with very high probability for any $|s\rangle$ and $l = \pi/4\theta$ iterations of the Grover step. For typical values of $\theta \sim 2^{-m/2}$, the complexity is $O[\mathcal{C}(U)2^{-m/2}]$ as before.

It is important to note that the protocol we just constructed relied crucially on the assumption that there exists an exact unitary \tilde{U} satisfying eq. (A.23). If it is instead only approximate (a more realistic assumption), we can run into problems because we are applying an exponentially long circuit: small errors at each stage can add up to become very large.

We now argue that this will not be the case so long as $n \gg m \gg 1$ and U is scrambling. (If U is not scrambling we expect that much more efficient circuits may well exist

anyway.) Scrambling unitaries are well modeled by typical elements of unitary 2-designs. We want to show that a Grover search using (A.32) and a fixed number of repetitions (independent of $|\psi\rangle$) will work for an arbitrary input state $|\psi\rangle$. This implies that the first and second moment calculations will be exactly given by Haar averages using Weingarten coefficients [53]. Specifically, we first compute the mean of the $\sin^2 \theta$,

$$\overline{\sin^2 \theta} = \mathbb{E}_{U \in \text{Haar}} \left\| \langle 0|^{\otimes m} U |s\rangle |0\rangle^{\otimes n} \right\|_2 = \frac{1}{2^m}. \tag{A.33}$$

In addition, we would like to estimate the variance of the $\sin^2 \theta$. We can also compute this quantity since it only involves two copies of U and two copies of U^\dagger ,

$$\Delta_\theta = \text{var}[\sin^2 \theta] = \mathbb{E}_{U \in \text{Haar}} \left\| \langle 0|^{\otimes m} U |s\rangle |0\rangle^{\otimes n} \right\|_2^2 - \frac{1}{4^m} = \frac{1}{2^{k+m}} \tag{A.34}$$

where k is the total number of qubits. This implies that,

$$\sin^2 \theta \approx 2^{-m} + O(2^{-(k+m)/2}), \tag{A.35}$$

where k is the total number of qubits. It follows that we can assume $\theta^2 = 2^{-m}$ for the state with exponentially small error.

In the state-independent protocol, we use W instead of the state-dependent projector U_s . This results in additional errors if \tilde{U} is not an exact unitary. However, this error in each Grover step can be bounded. The error for a fixed unitary U is

$$\epsilon_U = \left\| \left((|s\rangle \langle s| - \mathbf{1}) \otimes |0\rangle \langle 0|^{\otimes n} \right) U^\dagger \left(\mathbf{1} \otimes (\mathbf{1} - 2|0\rangle \langle 0|^{\otimes m}) \right) U |s\rangle |0\rangle^{\otimes n} \right\|_2.$$

Again we can use that U is an element of a 2-design and calculate the mean of the ϵ_U . Again, one can do those integrals using Weingarten coefficients [53] and arrive at,

$$\bar{\epsilon} = \mathbb{E}_{U \in \text{Haar}} (\epsilon_U) = \frac{1}{2^n}. \tag{A.36}$$

The amplitude of being mapped into a wrong state in each Grover step is $\sqrt{\bar{\epsilon}} = 2^{n/2}$. Our state independent protocol requires application of (A.32) $2^{m/2}$ times, so the total accumulated error will be given by $\sqrt{2^{m-n}}$, which is exponentially small in the limit of interest, when $n \gg m \gg 1$. This completes our argument.

B Maximinimax prescription for the bulge surface

In this appendix, we argue that spacetimes with more than one extremal surface generically contain Python's Lunches. To do so, we use a variant on the maximin arguments introduced by Wall in ref. [28]. Such arguments have numerous subtleties and require considerable effort to rule out as many edge cases as possible. We shan't worry too much about such details here; instead we will just give physics-level arguments that justify our construction. We shall restrict our attention to classical spacetimes obeying the null energy condition (NEC), although it is possible to generalize maximin arguments to include quantum effects [29].

Our starting assumption is the existence of two spacelike-separated extremal surfaces, the HRT surface χ_1 and an additional surface χ_2 , homologous to the same boundary region. The HRT surface χ_1 can be found by a maximin prescription, where one finds the globally minimal area surface within some Cauchy slice, and then maximises that area over all Cauchy slices. The second extremal surface χ_2 must have equal or larger area. Generically it will have larger area; we will assume that this is indeed the case.

Let us temporarily assume that there exists a Cauchy slice, within which any sufficiently small deformation of χ_2 , which preserves the homology constraint but which is not necessarily local, will increase its area. (We shall consider the alternative possibility below.) If we deform C_2 in a sufficiently small neighborhood of χ_2 , we should then be able to find a new Cauchy slice, still containing $\chi_1 \cup \chi_2$, on which χ_2 is minimal within a small neighborhood and χ_1 is still globally minimal. This will be important later.

We define the Wheeler-de Witt patch $W_{1,2}$ as the bulk domain of dependence of any spacelike slice bounded by $\chi_1 \cup \chi_2$. We can then construct a new surface χ_3 by the following maximinimax procedure.

First we choose some Cauchy slice C_3 for the Wheeler-de Witt patch $W_{1,2}$.²¹

Next, we choose a smooth non-degenerate function $\phi_3 : C_3 \rightarrow [0, 1]$, where $\phi_3(\chi_1) = 0$ and $\phi_3(\chi_2) = 1$. Morally, the level sets of the function ϕ_3 define a foliation of C_3 . However, it is somewhat more general than this because the topology of the level set can change if ϕ_3 has critical points. Formally, this is known as a ‘sweepout’ of C_3 . It is necessary both for physical reasons, since in general an extremal surface may have arbitrary topology, so long as it satisfies the homology constraint, and for mathematical reasons, to prevent the appearance of singularities in the surface if we insist that it have the ‘wrong’ topology.

Finally, we choose the level set $\phi_3^{-1}(x_3)$ for $x_3 \in [0, 1]$ of *maximal* area. Note that the level set $\phi_3^{-1}(x)$ will be singular if x is a critical value of ϕ_3 , but so long as ϕ_3 is non-degenerate, the singularities will be at isolated points and the area of the surface should still be well defined.

Having found the surface $\chi_3 = \phi_3^{-1}(x_3)$ of maximal area, we minimise that maximal area over all allowed functions ϕ_3 . Finally, we maximize that minimax area over all Cauchy slices for W . We call the resulting surface the maximinimax surface χ_3 . In other words, we have

$$\chi_3 = \phi_3^{-1}(x_3), \tag{B.1}$$

where the level set $\phi_3^{-1}(x_3)$ is defined by the maximinimaximization

$$\max_{C_3} \min_{\phi_3} \max_{x_3} A(\phi_3^{-1}(x)). \tag{B.2}$$

Provided a unique maximinimax surface χ_3 exists²² and does not lie at the boundary of any of the spaces we are optimizing over, it will be extremal, since, at linear order, an arbitrary variation of the surface χ_3 can be achieved by a linear combination of variations in C_3 , ϕ_3 and x_3 .

²¹Note that this is a Cauchy slice for the patch $W_{1,2}$ but not a Cauchy slice for the entire spacetime.

²²If the surface was non-unique, one would have to worry about stability issues.

We shall not try to rigorously prove that this will be true (even generically). After all, we have not even tried to rigorously define our construction — the geometric measure theory [54] required to do so would be well beyond the scope of this paper. However we shall make a few comments about why we expect that a maximinimax surface should exist and not lie at the various boundaries of the optimization space we are searching over.

So long as ϕ_3 is smooth and non-degenerate, the area of the level set should be a continuous function on a compact interval, and so a maximal area surface should exist.

One might worry that the minimization over functions ϕ_3 could approach a function that is not smooth and non-degenerate, for which a maximal area level set is not well defined. Our understanding of the results of Almgren-Pitts min-max theory [55–57] is that this will not end up being the case. Instead, the minimax surface $\chi_3 \in C_3$ should be a well defined varifold and will be a smooth (possibly self-intersecting) submanifold if the spacetime dimension is less than seven. Intuitively, it is reasonable to expect that, as the function ϕ_3 becomes more badly behaved, the maximal area surface should only increase, rather than decrease in area.

Similarly, we expect any bad behavior in the Cauchy slice C_3 will tend to decrease the area of the minimax surface. Hence a maximinimax surface χ_3 should exist and be well behaved. For more detailed arguments in this direction, see [28]. For known examples, involving timelike de Sitter boundaries, where maximin surfaces do not exist, see [58].

How could the maximinimax surface χ_3 end up on the boundary of the space of surfaces we are searching over? Firstly, the maximinimax surface χ_3 could have nonzero intersection with $\chi_1 \cup \chi_2$. Suppose this intersection were not a connected component of χ_3 (and $\chi_1 \cup \chi_2$). Since the surface χ_3 cannot ever go outside the Wheeler-de Witt patch $W_{1,2}$ and χ_1 and χ_2 are extremal, there must be some point where χ_3 has nonzero mean curvature, within the Cauchy slice, where it bends ‘inwards’ into W . We could then decrease the max area of a level set in ϕ_3 by deforming ϕ_3 slightly to make the surface χ_3 moves slightly inwards at this point, which gives the desired contradiction.

What about if the intersection is a connected component? In that case, the entire connected component will already be extremal and we don’t need to worry about it. One might, of course, worry that χ_3 could end up being the same as either χ_1 or χ_2 , in which case we wouldn’t have really found a new extremal surface. However this is impossible, since, by assumption, there exists a Cauchy slice where any small deformation of χ_1 or χ_2 will increase their area, and hence neither can have maximal area within any ϕ_3 . The min-max surface in this Cauchy slice will have a larger area than either χ_1 or χ_2 , which rules out either χ_1 or χ_2 being maximinimax.

Finally, one might worry that points in χ_3 might end up lightlike separated from either other points in χ_3 , or points in χ_1 or χ_2 . The first cannot happen, because of arguments similar to those in [28]. If a) the minimax surface contained a null segment, the area of the minimax could always be increased by a sufficiently small deformation of the Cauchy slice near this null segment. However, if b) the minimax surface did not contain a null segment, it could not be extremal within the Cauchy slice (using focussing arguments for generic spacetimes), which the minimax should be since its variation is unconstrained so long as it doesn’t intersect χ_1 or χ_2 . The second cannot happen because (using focussing in a generic

spacetime) we could then increase the area of the minimax surface by deforming ϕ_3 so that the level set is locally deformed along the lightcone towards χ_1 or χ_2 .

We also note that the maximin construction automatically guarantees that the bulge surface χ_3 has larger area than either χ_1 or χ_2 .

Having shown that an intermediate bulge surface exists between χ_1 and χ_2 whenever χ_2 is minimal with respect to any small deformation within some fixed Cauchy slice, we now consider the opposite case, where there exist small deformations of χ_2 which decrease the area of χ_2 within any Cauchy slice. In this case, χ_2 cannot be an end surface and so must instead itself be the bulge surface. Without loss of generality, we assume that χ_1 is contained in the interior $\text{Int}[\chi_2]$ of χ_2 .²³

We shall also assume that, in any Cauchy slice, there exist small deformations of χ_2 that a) decrease the area and b) lie entirely in the exterior $\text{Ext}[\chi_2]$ (defined as the complement of the interior $\text{Int}[\chi_2]$). The alternative possibility, where there exist Cauchy slices where only deformations that enter the interior can decrease the area, but none where no deformations can decrease the area, should be non-generic and can be interpreted as the bulge surface and one end surface degenerating into one another.

We can now use the usual maximin construction, to find a second end surface χ_3 . We simply constrain our search to Cauchy slices containing χ_2 , and to surfaces within that Cauchy slice that are entirely in the exterior $\text{Ext}[\chi_2]$. As before, to show extremality, we just need a) for the maximin surface to exist and b) for variations of the Cauchy slice be sufficient to freely vary the maximin surface.

The arguments for both are essentially identical to those for the original maximin construction. The only new potential obstructions that need to be ruled out are the maximin surface either a) intersecting, or b) being lightlike separated from, the surface χ_2 . In the first case, if the intersection was not a connected component the maximin surface would have to bend inwards somewhere, which contradicts its minimality within the Cauchy slice. Any intersection on a connected component will be automatically extremal. Finally, the surface χ_3 cannot simply be equal to χ_2 , since, by assumption, χ_2 does not have globally minimal area within any Cauchy slice.

What about the possibility of lightlike separation from χ_2 ? By focussing arguments in a generic spacetime, the change in area from a lightlike deformation of the surface χ_3 towards χ_2 (in direction k^a) would have to be positive at linear order. Since χ_3 is minimal within the Cauchy slice, there must be some spacelike direction r^a pointing away from the lightcone for which the change in area is nonnegative at linear order. However, this implies a deformation in a timelike direction t^a (that makes χ_3 spacelike separated from χ_2) must increase the area at linear order, in contradiction with the maximality of the Cauchy slice.²⁴

Finally, we note that, since χ_3 is minimal within a slice containing χ_2 , it must have smaller area than χ_2 . We therefore conclude that the generic situation when more than

²³Recall that the two surfaces cannot ever be timelike separated. They also can't intersect, because then generic perturbations would presumably make them be timelike separated.

²⁴We are ignoring issues about stability here.

one extremal surface exists if to have a Python’s Lunch: three extremal surfaces, with the middle bulge surface having larger area than either end surface.

C Explicit calculation of the late-time bulge surface

In this appendix, we explicitly calculate the location of the extremal surface that forms the bulge surface at late times in a particular theory. The theory is JT gravity with c Dirac fermions, and we consider a black hole that is evaporating using transparent boundary conditions, as in [36, 37]. Dirac fermions are the only conformal field theory for which the calculation is possible, since they are the only conformal field theory for which the two-interval von Neumann entropy is known analytically.

We note that this surface only becomes the bulge surface when the horizon area, which, in the case of JT gravity is the horizon dilaton value $\phi + \phi_0$, is less than half of its initial value.²⁵ This means that the initial black hole cannot have been in the regime $\phi \ll \phi_0$ where JT gravity is justified as the dimensional reduction of a near extremal black hole. Nonetheless, a) there is no obvious problem (other than UV issues which are unimportant for this calculation) with defining JT gravity as a theory in its own right when $\phi \gtrsim \phi_0$ and b) it provides a calculable example of an extremal surface that should also exist in more general examples of evaporating black holes.

The JT gravity action is given by

$$S = \frac{\phi_0}{16\pi G_N} \left[\int_M d^2x \sqrt{-g} R + 2 \int_{\partial M} K \right] + \frac{1}{16\pi G_N} \left[\int_M d^2x \sqrt{-g} \phi (R + 2) + 2 \int_{\partial M} \phi_b K \right] + S_{\text{CFT}}[g], \tag{C.1}$$

where the scalar field ϕ is called the dilaton and $S_{\text{CFT}}[g]$ is the action for the CFT (in this case c Dirac fermions) in the gravitational background. We also impose boundary conditions

$$g_{tt}|_{\text{bdy}} = \frac{1}{\varepsilon^2}, \quad \phi = \phi_b = \frac{\bar{\phi}_r}{\varepsilon}, \tag{C.2}$$

where t is the physical boundary time, $\bar{\phi}_r$ is the fixed renormalized boundary dilaton value and ε is small.

The metric of a static black hole in JT gravity is given by

$$ds^2 = -\frac{4\pi^2 T^2 du dv}{\sinh^2[\pi T(u - v)]}, \tag{C.3}$$

with the dilaton profile given by

$$\phi = 2\bar{\phi}_r \pi T \coth[\pi T(u - v_0)]. \tag{C.4}$$

²⁵In fact, if we started with a two-sided black hole, which is generally the case in JT gravity, then the horizon area needs to be less than half of the *increase* in horizon area from the initial energy thrown into the two-sided black hole, before the evaporation began.

Here u and v are the advanced and retarded times respectively. The Bekenstein-Hawking entropy S_{BH} is given by

$$S_{\text{BH}} = \frac{\phi_{\text{hor}} + \phi_0}{4G_N} = \frac{2\bar{\phi}_r\pi T}{4G_N} + S_0 \tag{C.5}$$

where ϕ_{hor} is the horizon dilaton value and $S_0 = \phi_0/4G_N$ is the extremal entropy.

To extend our coordinate system behind the horizon we simply define the Kruskal-like coordinate $U = -\exp(-2\pi T u)$. We find

$$ds^2 = -e^{2\pi T v} \frac{8\pi T dU dv}{(1 + U e^{2\pi T v})^2}, \tag{C.6}$$

and

$$\phi = 2\bar{\phi}_r\pi T \frac{1 - U e^{2\pi T v}}{1 + U e^{2\pi T v}}. \tag{C.7}$$

In the near-horizon region $U e^{2\pi T v} \ll 1$, these simplify to

$$ds^2 = -e^{2\pi T v} 8\pi T dU dv, \tag{C.8}$$

and

$$\phi = 2\bar{\phi}_r\pi T [1 - 2U e^{2\pi T v}]. \tag{C.9}$$

In the semiclassical limit, an evaporating black hole is well approximated by an ingoing Vaidya metric, where we simply promote the temperature T to be a slowly varying function of the infalling time v . The change in temperature is determined by the rate of energy loss from the black hole, where we have

$$\frac{2\pi\bar{\phi}_r}{4G_N} \frac{dT}{dv} = \frac{dS_{\text{BH}}}{dv} = \frac{1}{T} \frac{dM}{dv} = -\frac{\pi c T}{12}. \tag{C.10}$$

Here the first equality uses (C.5), the second equality is the first law of black-hole thermodynamics and the last equality is the $(1+1)$ -dimensional Stefan-Boltzmann law. It follows that, in the near horizon region, we have

$$\frac{1}{4G_N} \frac{\partial\phi}{\partial v} = -\frac{\bar{\phi}_r(2\pi T)^2 2U e^{2\pi T v}}{4G_N} - \frac{\pi T c}{12}, \tag{C.11}$$

and

$$\frac{1}{4G_N} \frac{\partial\phi}{\partial U} = -\frac{4\pi T \bar{\phi}_r e^{2\pi T v}}{4G_N}. \tag{C.12}$$

We now briefly review the calculation of the location of the nonempty ‘end surface’, which is the EW surface after the Page time. This surface consists of a single point (U_0, v_0) . This calculation was previously done for JT gravity in [36], although our strategy is more similar to [35].

Let the ‘current’ boundary time be $t = 0$. The outgoing modes are in the vacuum state with respect to U . Outgoing modes in the interval $U_0 \geq U \geq -1$ are included in the entanglement wedge of the boundary; modes with $U \geq U_0$ are outside the wedge, while modes with $U \leq -1$ have already escaped the spacetime. The entropy of the outgoing modes is therefore given by

$$S_{\text{bulk}}^{(\text{out})} = \frac{c}{6} \log \frac{U_0 + 1}{\sqrt{\varepsilon_U^{(\text{ext})} \varepsilon_U^{(\text{bdy})}}} \tag{C.13}$$

where $\varepsilon_U^{(\text{ext})}$ and $\varepsilon_U^{(\text{bdy})}$ are the cut-offs in units of U at the extremal surface (U_0, v_0) and the boundary respectively. Since the near horizon metric (C.8) is independent of U , we can choose $\varepsilon_U^{(\text{ext})}$ and the entropy $S_{\text{bulk}}^{(\text{in})}$ of the ingoing modes to both be independent of U . We therefore find

$$\frac{\partial S_{\text{bulk}}}{\partial U_0} = \frac{c}{6(U_0 + 1)}. \tag{C.14}$$

What about the ingoing modes? These are in the vacuum state with respect to the physical infalling time v . It follows that their entropy is given by

$$S_{\text{bulk}}^{(\text{in})} = \frac{c}{6} \log \frac{|v_0|}{\sqrt{\varepsilon_v^{(\text{ext})} \varepsilon_v^{(\text{bdy})}}}, \tag{C.15}$$

where $\varepsilon_v^{(\text{ext})}$ and $\varepsilon_v^{(\text{bdy})}$ are the cut-offs in units of v on the ingoing modes at the extremal surface (U_0, v_0) and boundary respectively. Since in the semiclassical limit we will have $v_0 \rightarrow -\infty$, one might think that the gradient $\partial S_{\text{bulk}}/\partial v$ vanishes.

However, this argument would be naive. To correctly renormalize the entropy we need to keep the proper cut-off

$$\varepsilon_{\text{prop}} = \sqrt{g(\varepsilon_U^{(\text{ext})} \partial/\partial U, \varepsilon_v^{(\text{ext})} \partial/\partial v)} \tag{C.16}$$

fixed as we vary the surface. From (C.8), we therefore need

$$\varepsilon_U^{(\text{ext})} \varepsilon_v^{(\text{ext})} \propto e^{-2\pi T v}. \tag{C.17}$$

Hence

$$\frac{\partial S_{\text{bulk}}}{\partial v} = \frac{\pi T c}{6}. \tag{C.18}$$

The location of the extremal surface now follows from extremizing the generalized entropy

$$S^{(\text{gen})} = (\phi + \phi_0)/4G_N + S_{\text{bulk}}. \tag{C.19}$$

We find

$$\frac{\partial S^{(\text{gen})}}{\partial U} = -\frac{4\pi T \bar{\phi}_r e^{2\pi T v}}{4G_N} + \frac{c}{6(U_0 + 1)} = 0 \tag{C.20}$$

and

$$\frac{\partial S^{(\text{gen})}}{\partial v} = -\frac{\bar{\phi}_r(2\pi T)^2 2U_0 e^{2\pi T v}}{4G_N} + \frac{\pi T c}{12} = 0 \quad (\text{C.21})$$

which implies

$$U_0 = \frac{1}{3} \quad v_0 = -\frac{1}{2\pi T} \log \frac{16(S - S_0)}{c}, \quad (\text{C.22})$$

in agreement with the results from [36].

We are now ready to attempt our actual task: calculating the location of the late-time bulge surface. This is the union of two points (U_1, v_1) and (U_2, v_2) , where we assume $U_2 > U_1$ and $v_2 < v_1$ (since the points need to be spacelike separated). In fact, if we started with a two-sided black hole, the bulge surface really consists of three points, where the third point (U_3, v_3) lies close to the other ‘end surface’, near the initial bifurcation surface of the two-sided black hole. In the semiclassical limit we have $U_3 = \exp(O(1/G_N))$ and $v_3 = -O(1/G_N)$. The corrections to the generalized entropy gradient for (U_3, v_3) from the existence of the additional points (U_1, v_1) and (U_2, v_2) are therefore highly suppressed and we can treat (U_3, v_3) as lying exactly on the quantum extremal end surface.

The outgoing entropy is now the entropy of the union of the two intervals $[-1, U_1]$ and $[U_2, U_3]$, which for c Dirac fermions is given by (see [59])

$$S_{\text{bulk}}^{(\text{out})} = \frac{c}{6} \log \left[\frac{(U_3 + 1)(U_1 + 1)(U_3 - U_2)(U_2 - U_1)}{(U_3 - U_1)(U_2 + 1) \sqrt{\varepsilon_{U_3} \varepsilon_{U_2} \varepsilon_{U_1} \varepsilon_U^{(\text{bdy})}}} \right], \quad (\text{C.23})$$

where ε_{U_3} , ε_{U_2} , ε_{U_1} and $\varepsilon_U^{(\text{bdy})}$ are the outgoing mode cut-offs in units of U at U_3 , U_2 , U_1 and the boundary respectively. Since $U_3 = \exp(O(1/G_N))$, terms involving U_3 do not contribute to the gradient of the entropy.

Formally, the ingoing entropy should be calculated by a similar formula for the two intervals $[v_3, v_2]$ and $[v_1, 0]$. However, since $v_1 = -O(\log G_N)$ and $v_2 - v_3 = O(1/G_N)$, while we shall find that $v_1 - v_2 = O(1/T)$, for our purposes we can approximate

$$S_{\text{bulk}}^{(\text{in})} = \frac{c}{6} \log \frac{v_1 - v_2}{\sqrt{\varepsilon_{v_2} \varepsilon_{v_1}}} + \dots, \quad (\text{C.24})$$

where ε_{v_1} and ε_{v_2} are the ingoing cut-offs in units of v at v_1 and v_2 respectively, and we have elided constant terms. As before, to correctly renormalize the entropy, we need

$$\varepsilon_{v_i} \varepsilon_{U_i} \propto e^{2\pi T v_i}. \quad (\text{C.25})$$

Extremizing the generalized entropy, using (C.11), (C.12), (C.23) and (C.24), we find

$$\frac{\partial S^{(\text{gen})}}{\partial U_1} = -\frac{c}{6(U_2 - U_1)} + \frac{c}{6(U_1 + 1)} - 2(S - S_0)e^{2\pi T v_1} = 0, \quad (\text{C.26})$$

$$\frac{\partial S^{(\text{gen})}}{\partial U_2} = \frac{c}{6(U_2 - U_1)} - \frac{c}{6(U_2 + 1)} - 2(S - S_0)e^{2\pi T v_2} = 0, \quad (\text{C.27})$$

$$\frac{\partial S^{(\text{gen})}}{\partial v_1} = -2(2\pi T)(S - S_0)U_1 e^{2\pi T v_1} + \frac{\pi T c}{12} + \frac{c}{6(v_1 - v_2)} = 0, \quad (\text{C.28})$$

$$\frac{\partial S^{(\text{gen})}}{\partial v_2} = -2(2\pi T)(S - S_0)U_2 e^{2\pi T v_2} + \frac{\pi T c}{12} - \frac{c}{6(v_1 - v_2)} = 0. \quad (\text{C.29})$$

This set of equations can be solved numerically. We find

$$(U_1, v_1) = \left(0.874, v_0 - \frac{0.410}{2\pi T}\right), \quad (U_2, v_2) = \left(28.8, v_0 - \frac{5.81}{2\pi T}\right). \quad (\text{C.30})$$

As expected, we have $U_2 > U_1 > U_0$ and $v_2 < v_1 < v_0$.

Finally, we note that the classical contribution to the generalized entropy for this surface is $2(\phi_0 + 2\pi T \bar{\phi}_r)$.²⁶ Meanwhile, up to subleading corrections, (C.23) is equal to $c/6 \log U_3$, which is the entropy of the interior partners of the Hawking radiation. Hence, at leading order, the bulk von Neumann entropy is simply the entropy S_{rad} of the Hawking radiation. We therefore find that the total generalized entropy at leading order is $2S_{\text{BH}} + S_{\text{rad}}$, where S_{BH} is the final Bekenstein-Hawking entropy of the black hole.

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] D. Harlow and P. Hayden, *Quantum Computation vs. Firewalls*, *JHEP* **06** (2013) 085 [[arXiv:1301.4504](https://arxiv.org/abs/1301.4504)] [[INSPIRE](#)].
- [2] L. Susskind, *Computational Complexity and Black Hole Horizons*, *Fortsch. Phys.* **64** (2016) 24 [*Addendum ibid.* **64** (2016) 44] [[arXiv:1403.5695](https://arxiv.org/abs/1403.5695)] [[INSPIRE](#)].
- [3] A.R. Brown, D.A. Roberts, L. Susskind, B. Swingle and Y. Zhao, *Complexity, action, and black holes*, *Phys. Rev. D* **93** (2016) 086006 [[arXiv:1512.04993](https://arxiv.org/abs/1512.04993)] [[INSPIRE](#)].
- [4] A.R. Brown, D.A. Roberts, L. Susskind, B. Swingle and Y. Zhao, *Holographic Complexity Equals Bulk Action?*, *Phys. Rev. Lett.* **116** (2016) 191301 [[arXiv:1509.07876](https://arxiv.org/abs/1509.07876)] [[INSPIRE](#)].
- [5] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, *Fortsch. Phys.* **61** (2013) 781 [[arXiv:1306.0533](https://arxiv.org/abs/1306.0533)] [[INSPIRE](#)].
- [6] A. Nahum, S. Vijay and J. Haah, *Operator Spreading in Random Unitary Circuits*, *Phys. Rev. X* **8** (2018) 021014 [[arXiv:1705.08975](https://arxiv.org/abs/1705.08975)] [[INSPIRE](#)].
- [7] V. Khemani, A. Vishwanath and D.A. Huse, *Operator spreading and the emergence of dissipation in unitary dynamics with conservation laws*, *Phys. Rev. X* **8** (2018) 031057 [[arXiv:1710.09835](https://arxiv.org/abs/1710.09835)] [[INSPIRE](#)].
- [8] H. Gharibyan, M. Hanada, S.H. Shenker and M. Tezuka, *Onset of Random Matrix Behavior in Scrambling Systems*, *JHEP* **07** (2018) 124 [*Erratum JHEP* **02** (2019) 197] [[arXiv:1803.08050](https://arxiv.org/abs/1803.08050)] [[INSPIRE](#)].

²⁶We will not worry about the time-dependence of the temperature here, since we are only interested in the answer at leading order.

- [9] 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](#)].
- [10] P. Gao, D.L. Jafferis and A.C. Wall, *Traversable Wormholes via a Double Trace Deformation*, *JHEP* **12** (2017) 151 [[arXiv:1608.05687](#)] [[INSPIRE](#)].
- [11] J. Maldacena, D. Stanford and Z. Yang, *Diving into traversable wormholes*, *Fortsch. Phys.* **65** (2017) 1700034 [[arXiv:1704.05333](#)] [[INSPIRE](#)].
- [12] S. Aaronson, *The Complexity of Quantum States and Transformations: From Quantum Money to Black Holes*, [arXiv:1607.05256](#) [[INSPIRE](#)].
- [13] B. Yoshida and A. Kitaev, *Efficient decoding for the Hayden-Preskill protocol*, [arXiv:1710.03363](#) [[INSPIRE](#)].
- [14] L.K. Grover, *A fast quantum mechanical algorithm for database search*, in *STOC '96*, proceedings of the *Twenty-Eighth Annual ACM Symposium on Theory of Computing*, Philadelphia, Pennsylvania, U.S.A., May 1996, Association for Computing Machinery, New York NY U.S.A. (1996), pp. 212–219.
- [15] L.K. Grover, *Quantum mechanics helps in searching for a needle in a haystack*, *Phys. Rev. Lett.* **79** (1997) 325 [[quant-ph/9706033](#)] [[INSPIRE](#)].
- [16] S. Aaronson, *Quantum computing, postselection, and probabilistic polynomial-time*, [quant-ph/0412187](#).
- [17] C.H. Bennett, E. Bernstein, G. Brassard and U. Vazirani, *Strengths and weaknesses of quantum computing*, *SIAM J. Comput.* **26** (1997) 1510 [[INSPIRE](#)].
- [18] M. Boyer, G. Brassard, P. Høyer and A. Tapp, *Tight bounds on quantum searching*, *Fortschr. Phys.* **46** (1998) 493.
- [19] C. Zalka, *Grover's quantum searching algorithm is optimal*, *Phys. Rev. A* **60** (1999) 2746 [[quant-ph/9711070](#)] [[INSPIRE](#)].
- [20] D.W. Berry, A.M. Childs, R. Cleve, R. Kothari and R.D. Somma, *Simulating Hamiltonian dynamics with a truncated Taylor series*, *Phys. Rev. Lett.* **114** (2015) 090502 [[arXiv:1412.4687](#)].
- [21] A. Gilyén, S. Arunachalam and N. Wiebe, *Optimizing quantum optimization algorithms via faster quantum gradient computation*, in proceedings of the *30th ACM-SIAM Symposium on Discrete Algorithms (SODA 2019)*, San Diego, California, U.S.A., 6–9 January 2019, SIAM (2019), pp. 1425–1444 [[arXiv:1711.00465](#)].
- [22] A. Gilyén, Y. Su, G.H. Low and N. Wiebe, *Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics*, in *STOC 2019*, proceedings of the *51st Annual ACM SIGACT Symposium on Theory of Computing*, Phoenix, AZ, U.S.A., 23–26 June 2019, Association for Computing Machinery, New York NY U.S.A. (2019), pp. 193–204.
- [23] S. Ryu and T. Takayanagi, *Holographic derivation of entanglement entropy from AdS/CFT*, *Phys. Rev. Lett.* **96** (2006) 181602 [[hep-th/0603001](#)] [[INSPIRE](#)].
- [24] V.E. Hubeny, M. Rangamani and T. Takayanagi, *A Covariant holographic entanglement entropy proposal*, *JHEP* **07** (2007) 062 [[arXiv:0705.0016](#)] [[INSPIRE](#)].
- [25] A. Lewkowycz and J. Maldacena, *Generalized gravitational entropy*, *JHEP* **08** (2013) 090 [[arXiv:1304.4926](#)] [[INSPIRE](#)].

- [26] N. Engelhardt and A.C. Wall, *Quantum Extremal Surfaces: Holographic Entanglement Entropy beyond the Classical Regime*, *JHEP* **01** (2015) 073 [[arXiv:1408.3203](#)] [[INSPIRE](#)].
- [27] X. Dong and A. Lewkowycz, *Entropy, Extremality, Euclidean Variations, and the Equations of Motion*, *JHEP* **01** (2018) 081 [[arXiv:1705.08453](#)] [[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] C. Akers, N. Engelhardt, G. Penington and M. Usatyuk, *Quantum Maximin Surfaces*, to appear.
- [30] D.N. Page, *Average entropy of a subsystem*, *Phys. Rev. Lett.* **71** (1993) 1291 [[gr-qc/9305007](#)] [[INSPIRE](#)].
- [31] A. Almheiri, D. Marolf, J. Polchinski and J. Sully, *Black Holes: Complementarity or Firewalls?*, *JHEP* **02** (2013) 062 [[arXiv:1207.3123](#)] [[INSPIRE](#)].
- [32] P. Hayden and J. Preskill, *Black holes as mirrors: Quantum information in random subsystems*, *JHEP* **09** (2007) 120 [[arXiv:0708.4025](#)] [[INSPIRE](#)].
- [33] A. Almheiri, *Holographic Quantum Error Correction and the Projected Black Hole Interior*, [arXiv:1810.02055](#) [[INSPIRE](#)].
- [34] P. Hayden and G. Penington, *Learning the Alpha-bits of Black Holes*, *JHEP* **12** (2019) 007 [[arXiv:1807.06041](#)] [[INSPIRE](#)].
- [35] G. Penington, *Entanglement Wedge Reconstruction and the Information Paradox*, [arXiv:1905.08255](#) [[INSPIRE](#)].
- [36] A. Almheiri, N. Engelhardt, D. Marolf and H. Maxfield, *The entropy of bulk quantum fields and the entanglement wedge of an evaporating black hole*, *JHEP* **12** (2019) 063 [[arXiv:1905.08762](#)] [[INSPIRE](#)].
- [37] A. Almheiri, R. Mahajan, J. Maldacena and Y. Zhao, *The Page curve of Hawking radiation from semiclassical geometry*, *JHEP* **03** (2020) 149 [[arXiv:1908.10996](#)] [[INSPIRE](#)].
- [38] G. Penington, S.H. Shenker, D. Stanford and Z. Yang, *Replica wormholes and the black hole interior*, [arXiv:1911.11977](#) [[INSPIRE](#)].
- [39] A. Almheiri, T. Hartman, J. Maldacena, E. Shaghoulian and A. Tajdini, *Replica Wormholes and the Entropy of Hawking Radiation*, *JHEP* **05** (2020) 013 [[arXiv:1911.12333](#)] [[INSPIRE](#)].
- [40] M. Headrick, V.E. Hubeny, A. Lawrence and M. Rangamani, *Causality & holographic entanglement entropy*, *JHEP* **12** (2014) 162 [[arXiv:1408.6300](#)] [[INSPIRE](#)].
- [41] 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](#)].
- [42] 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](#)].
- [43] X. Dong, D. Harlow and A.C. Wall, *Reconstruction of Bulk Operators within the Entanglement Wedge in Gauge-Gravity Duality*, *Phys. Rev. Lett.* **117** (2016) 021601 [[arXiv:1601.05416](#)] [[INSPIRE](#)].
- [44] 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](#)].

- [45] T. Faulkner, A. Lewkowycz and J. Maldacena, *Quantum corrections to holographic entanglement entropy*, *JHEP* **11** (2013) 074 [[arXiv:1307.2892](#)] [[INSPIRE](#)].
- [46] A. Almheiri, X. Dong and D. Harlow, *Bulk Locality and Quantum Error Correction in AdS/CFT*, *JHEP* **04** (2015) 163 [[arXiv:1411.7041](#)] [[INSPIRE](#)].
- [47] L. Susskind, L. Thorlacius and J. Uglum, *The Stretched horizon and black hole complementarity*, *Phys. Rev. D* **48** (1993) 3743 [[hep-th/9306069](#)] [[INSPIRE](#)].
- [48] J. Maldacena and L. Susskind, *Cool horizons for entangled black holes*, *Fortsch. Phys.* **61** (2013) 781 [[arXiv:1306.0533](#)] [[INSPIRE](#)].
- [49] D.N. Page, *Time Dependence of Hawking Radiation Entropy*, *JCAP* **09** (2013) 028 [[arXiv:1301.4995](#)] [[INSPIRE](#)].
- [50] M.A. Nielsen, M.R. Dowling, M. Gu and A.C. Doherty, *Quantum computation as geometry*, *Science* **311** (2006) 1133.
- [51] I. Kourkoulou and J. Maldacena, *Pure states in the SYK model and nearly-AdS₂ gravity*, [arXiv:1707.02325](#) [[INSPIRE](#)].
- [52] R. Bousso, Z. Fisher, S. Leichenauer and A.C. Wall, *Quantum focusing conjecture*, *Phys. Rev. D* **93** (2016) 064044 [[arXiv:1506.02669](#)] [[INSPIRE](#)].
- [53] B. Collins, *Moments and Cumulants of Polynomial random variables on unitary groups, the Itzykson-Zuber integral and free probability*, *Int. Math. Res. Not.* **17** (2003) 953 [[math-ph/0205010](#)].
- [54] H. Federer, *Geometric measure theory*, Springer (2014).
- [55] W.K. Allard, *Notes on the theory of varifolds*, in *Théorie des variétés minimales et applications*, *Astérisque* **154–155**, Société Mathématique de France (1987), pp. 73–93 and online at http://www.numdam.org/item/AST_1987_154-155_73_0.
- [56] J.T. Pitts, *Existence and Regularity of Minimal Surfaces on Riemannian Manifolds. (MN-27)*, Princeton University Press (1981).
- [57] T.H. Colding and C.D. Lellis, *The min-max construction of minimal surfaces*, [math.AP/0303305](#).
- [58] S. Fischetti, D. Marolf and A.C. Wall, *A paucity of bulk entangling surfaces: AdS wormholes with de Sitter interiors*, *Class. Quant. Grav.* **32** (2015) 065011 [[arXiv:1409.6754](#)] [[INSPIRE](#)].
- [59] H. Casini, C.D. Fosco and M. Huerta, *Entanglement and alpha entropies for a massive Dirac field in two dimensions*, *J. Stat. Mech.* **0507** (2005) P07007 [[cond-mat/0505563](#)] [[INSPIRE](#)].