Non-Isometric Quantum Error Correction in Gravity

Arjun Kar

University of British Columbia

12th Joburg Workshop on String Theory, Gravity, and Cosmology

December 8, 2022

Based on 2210.13476

Motivation

The holographic approach to quantum gravity implies that there are two descriptions of gravitational physics.

The "semiclassical" description involves a curved spacetime with dynamical gravity and propagating quantum fields, and is only approximate.

The "microscopic" description consists of a lower-dimensional quantum theory without gravity. This description is exact and constitutes a non-perturbative definition of quantum gravity.

Motivation

The existence of two descriptions of the same physics, one of which is missing some microscopic details, is a familiar situation. We have tools like the renormalization group to make such ideas precise in quantum field theory.

What is unfamiliar in gravity is the existence of black holes. Black holes in the semiclassical description lead to apparent inconsistencies which do not arise in standard situations without gravity. Two related examples of such inconsistencies are the information paradox and firewall paradox.

Motivation

Ideas from quantum information theory have been helpful in sharpening the apparent disconnect between the semiclassical and microscopic descriptions.

In particular, quantum error correction serves as a link between the two descriptions. The basic idea is that there is a linear map V which embeds a subspace $\mathcal{H}_b$ of the semiclassical Hilbert space into $\mathcal{H}_B$, the microscopic Hilbert space.

$$V : \mathcal{H}_b \to \mathcal{H}_B$$

Motivation

The error correction properties of the map V come from the fact that $\mathcal{H}_b$ is encoded redundantly in $\mathcal{H}_B$. Erasing part of $\mathcal{H}_B$ leaves behind a remainder which can still be enough to "reconstruct" some part of $\mathcal{H}_b$.

To ensure the consistency and functionality of this encoding structure, it is necessary to impose some restrictions on $\mathcal{H}_b$ as a subspace.

A common choice of subspace is a small energy band of local semiclassical excitations around a particular background geometry. This usually ensures that there are enough states in $\mathcal{H}_B$ to encode all of $\mathcal{H}_b$ in a useful manner.

Motivation

In some situations with black holes, choosing a small energy band will not even ensure the most basic requirement of standard error correction: $|b| < |B|$.

In fact, in the information paradox, one explicitly encounters situations where semiclassical physics in the black hole interior is expected to be valid for a very large number of states $|b| \gg |B|$ compared to the black hole entropy. But in general it is not possible to faithfully encode a large Hilbert space into a small one.

The goal of this talk is to explain, in a toy model of black hole evaporation, how gravity gets around this issue and manages to encode semiclassical states (from a larger space) into a smaller microscopic space. The semiclassical states are effective field theory excitations in the interior, while the microscopic states are the black hole microstates.

Outline

1. Isometric and non-isometric error correction
2. Black hole interior code from Euclidean path integrals
3. Properties of the non-isometric interior code

Isometric codes

Standard quantum error correction begins with an isometry L between a code Hilbert space $\mathcal{H}_{\text{code}}$ and a physical Hilbert space $\mathcal{H}_{\text{phys}}$. We distinguish these from V, $\mathcal{H}_b$, and $\mathcal{H}_B$ which we will only use in the gravity context.

$$L : \mathcal{H}_{\text{code}} \to \mathcal{H}_{\text{phys}} \ , \quad L^\dagger L = 1_{\text{code}}$$

The isometry condition implies $|\text{code}| \leq |\text{phys}|$, essentially by counting. If $|\text{code}| > |\text{phys}|$, there must exist a nonzero $|\alpha\rangle \in \mathcal{H}_{\text{code}}$ with $L|\alpha\rangle = 0$, which makes the isometry condition impossible to satisfy.

Isometric codes

The sense in which L corrects errors is then that L allows "simulation" of unitary operators W on $\mathcal{H}_{\text{code}}$ on the encoded state in $\mathcal{H}_{\text{phys}}$.

$$\|\widetilde{W}L|\psi\rangle - LW|\psi\rangle\| \leq \epsilon \ , \quad |\psi\rangle \in \mathcal{H}_{\text{code}}$$

This is trivially satisfied, with zero error ϵ, by a "global reconstruction".

$$\widetilde{W} = LWL^{\dagger}$$

The nontrivial error correction properties of L are related to the case when $\widetilde{W}$ can be defined using only a subspace of $\mathcal{H}_{\text{phys}}$. If we are able to satisfy the simulation condition even when $\widetilde{W}$ acts only on a subspace of $\mathcal{H}_{\text{phys}}$, we can erase the rest of $\mathcal{H}_{\text{phys}}$ and still simulate W.

Non-isometric codes

When L is far from an isometry, it is not clear what to expect from a theory of error correction. Even the global reconstruction is guaranteed to fail on certain states.

Trying the global reconstruction $\widetilde{W} = LWL^\dagger$ in the error correction criterion yields an error proportional to:

$$\|L^\dagger L|\psi\rangle - |\psi\rangle\|$$

As we saw before, when $|\text{code}| > |\text{phys}|$, there must be a state $|\alpha\rangle \in \mathcal{H}_{\text{code}}$ with $L|\alpha\rangle = 0$. Considering $|\alpha\rangle$ in the above criterion gives a large error (1).

Non-isometric codes

A proposal for non-isometric error correction has been developed recently. [Akers, Penington 21] [Akers, Engelhardt, Harlow, Penington, Vardhan 22]

The main ingredients are a discrete set S of states in $\mathcal{H}_{\text{code}}$ and a notion of approximate state-specific reconstruction which should be possible for states in S and operators W which preserve S.

$$\|\widetilde{W}(\psi)L|\psi\rangle - LW|\psi\rangle\| \leq \epsilon\ , \quad |\psi\rangle, W|\psi\rangle \in S$$

One asks only for error correction on this discrete set S in the code space and for operators which preserve it. In this way, non-isometric codes do not faithfully encode the entire space $\mathcal{H}_{\text{code}}$, but only a discrete subset thereof.

Non-isometric codes

A useful method to verify the state-specific error correction criterion is the decoupling principle.

$$|\langle\psi|W^\dagger L^\dagger LW|\psi\rangle - \langle\psi|L^\dagger L|\psi\rangle| \leq \epsilon^2$$

This implies the existence of $\widetilde{W}(\psi)$.

$$\|\widetilde{W}(\psi)L|\psi\rangle - LW|\psi\rangle\| \leq \epsilon$$

The decoupling principle is fairly intuitive: if $L|\psi\rangle$ and $LW|\psi\rangle$ have approximately the same norm, they are approximately related by a unitary rotation $\widetilde{W}(\psi)$. The state-dependence is present because we cannot guarantee the same rotation will work for any choice of $|\psi\rangle \in S$.

Non-isometric codes

This gives a natural type of set S for any non-isometric code L: a set of states in $\mathcal{H}_{\text{code}}$ which have approximately preserved overlaps or norms after the action of L.

$$|\langle\psi_2|L^\dagger L|\psi_1\rangle - \langle\psi_2|\psi_1\rangle| \leq \epsilon \ , \quad |\psi_1\rangle, |\psi_2\rangle \in S$$

So, the discrete sets S we will consider are those upon which L acts almost like an isometry.

The most basic question about non-isometric codes L then is to understand the properties of such sets S, including their typical size and the allowed subspace reconstructions in $\mathcal{H}_{\text{phys}}$.

Black hole interior code

A useful toy model for black hole evaporation is AdS_2 Jackiw-Teitelboim gravity with non-dynamical end-of-the-world branes. [Penington, Shenker, Stanford, Yang 19]

$$I = S_0 \chi(\mathcal{M}) + \frac{1}{2} \int_{\mathcal{M}} \sqrt{g} \phi (R+2) + \int_{\partial \mathcal{M}} \sqrt{h} \phi (K-1) + \int_{\text{brane}} \sqrt{h} (\phi K - \mu)$$

$$\partial \mathcal{M}: \quad \mathrm{d}s|_{\partial \mathcal{M}} = \frac{\mathrm{d}\tau}{\epsilon}\,, \quad \phi|_{\partial \mathcal{M}} = \frac{1}{\epsilon}\,, \quad \epsilon \to 0$$
$$\text{brane:} \quad K = 0\,, \quad n^a \partial_a \phi = \mu$$

We will construct a non-isometric dictionary V sending semiclassical to microscopic states by using the Euclidean path integral of this gravity theory and its holographic dual.

Black hole interior code

A brief review of Euclidean path integral technology will be helpful.

The simplest Euclidean path integral prepares the thermal partition function.

$$Z(\beta) = \mathrm{Tr}\, e^{-\beta H}$$

In the microscopic theory, the Euclidean manifold is a circle with circumference β. The line segment represents the Euclidean evolution $e^{-\beta H}$, and joining the endpoints generates the trace.

In the semiclassical theory, the Euclidean manifold is a disk with the hyperbolic metric which fills in a circle with renormalized circumference β.

Black hole interior code

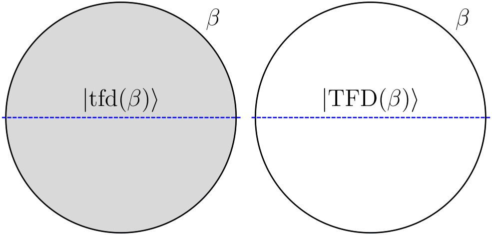

To create states, we cut open the path integrals along a Cauchy slice. This leads to the semiclassical $|\mathrm{tfd}(\beta)\rangle$ and microscopic $|\mathrm{TFD}(\beta)\rangle$ thermofield double states.

$$|\mathrm{tfd}(\beta)\rangle = \int_0^\infty \mathrm{d}E \; \rho(E) e^{-\beta E/2} |E\rangle \; , \quad \rho(E) = \frac{1}{2\pi^2} \sinh(2\pi\sqrt{2E})$$

$$|\mathrm{TFD}(\beta)\rangle = \sum_{n=0}^{\infty} e^{-\beta E_n/2} |E_n\rangle \otimes |E_n\rangle$$

Black hole interior code

The dictionary V acts by "hollowing out" the interior of the semiclassical path integral and leaves behind the microscopic path integral.

$$V|\mathrm{tfd}(\beta)\rangle = |\mathrm{TFD}(\beta)\rangle$$

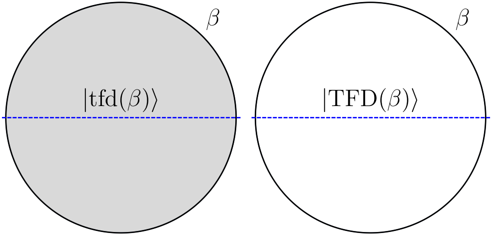

Black hole interior code

To create an excitation in the interior, we insert a brane boundary.

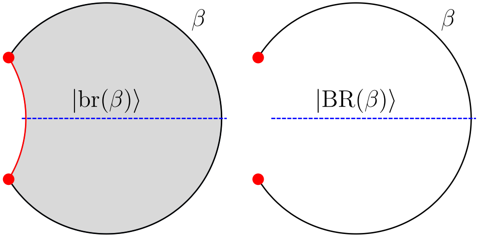

The corresponding states are the semiclassical and microscopic brane states.

$$|\mathrm{br}(\beta)\rangle = \int_0^\infty \mathrm{d}E \; f_\mu(E)\rho(E)e^{-\beta E/2}|E\rangle \; , \quad f_\mu(E) = \Gamma(\mu + \frac{1}{2} + \mathrm{i}\sqrt{2E})$$

$$|\mathrm{BR}(\beta)\rangle = \sum_{n=0}^{\infty} f_\mu(E_n)e^{-\beta E_n/2}C_n|E_n\rangle \; , \quad V|\mathrm{br}(\beta)\rangle = |\mathrm{BR}(\beta)\rangle$$

Black hole interior code

The coefficients C_n are complex Gaussian random variables with zero mean and unit variance.

By adding $|b|$ flavors of branes, we can create different types of interior excitations and we define $\mathcal{H}_b$ as the span of these states. Since the different flavors are non-interacting, these states are orthogonal in the semiclassical description.

In the microscopic theory, this extends the set C_n of complex Gaussian random coefficients to a set $C_{n\xi}$ where ξ labels the brane flavor.

Passing to the microcanonical ensemble (fixing the coarse-grained energy instead of the inverse temperature β) eliminates the tension μ and smooth energy E dependence in the state coefficients, and makes $|B|$ finite.

Black hole interior code

The upshot of all of this is that the dictionary V in this toy model of the black hole interior is proportional to a complex Gaussian random matrix with independent entries.

$$V = \frac{1}{\sqrt{|B|}} C \ , \quad C \sim \mathbb{C}\mathrm{Normal}(0,1)^{|B||b|}$$

$$\dim C = |B| \times |b|$$

The fact that an ensemble appears is related to recent discussions of ensemble averaging in low-dimensional gravity. [Saad, Shenker, Stanford 19]

[Many others 19-22]

Black hole interior code

To make the black hole evaporate, we add a radiation reservoir $\mathcal{H}_R$ in both the semiclassical and microscopic descriptions. This allows us to consider entangled states between the interior excitations (brane flavors) and the radiation. We extend the dictionary by the identity on $\mathcal{H}_R$.

$$\mathcal{H}_{\text{code}} = \mathcal{H}_b \otimes \mathcal{H}_R \ , \quad \mathcal{H}_{\text{phys}} = \mathcal{H}_B \otimes \mathcal{H}_R \ , \quad L = V \otimes 1_R$$

We now turn to studying the properties of this ensemble of dictionaries.

$$L = \frac{1}{\sqrt{|B|}} C \otimes 1_R \ , \quad C \sim \mathbb{C}\text{Normal}(0,1)^{|B||b|}$$

$$\dim C = |B| \times |b|$$

Non-isometric interior code properties

We will use the ensemble to study averaged error correction properties. Specifically, our goals are to study the typical size of discrete state sets S which have preserved overlaps and the allowed subsystem reconstructions (on $\mathcal{H}_B$ or $\mathcal{H}_R$) using the decoupling principle.

The latter question more explicitly is the following: when a unitary operator acts only as W_b or W_R, when can we reconstruct it using $\widetilde{W}_B(\psi)$ or $\widetilde{W}_R(\psi)$?

The subsystem reconstruction results must be consistent with entanglement wedge reconstruction.

Non-isometric interior code properties

The ensemble we have found is a high-dimensional Gaussian distribution on a complex Euclidean manifold. To study its average properties as a family of non-isometric codes, we use measure concentration. [Akers, Engelhardt, Harlow, Penington, Vardhan 22]

Roughly speaking, measure concentration is a property of certain probability distributions on Riemannian manifolds which allows one to use low moments of sufficiently well-behaved functions to put strong bounds on the probability of deviations from the mean.

Non-isometric interior code properties

The main theorem we need states that a κ-Lipschitz function $G(C)$ has bounded deviations with probability that is exponentially suppressed in $1/\kappa$.

$$\Pr\left[G(C) \geq \int \mathrm{D}C\ G(C) + \epsilon\right] \leq \exp\left(\frac{-\epsilon^2}{\kappa^2}\right)$$

The notion of a κ-Lipschitz function formalizes our "well-behaved" criterion. A κ-Lipschitz function $G(C)$ is one which has changes bounded by the geodesic distance of the underlying Riemannian manifold.

$$|G(C_1) - G(C_2)| \leq \kappa \|C_1 - C_2\|_2$$

Non-isometric interior code properties

Let $F(C)$ measure the norm of a state after applying the encoding map L.

$$F(C) = \| V \otimes 1_R |\psi\rangle \| \ , \quad |\psi\rangle \in \mathcal{H}_{\text{code}} = \mathcal{H}_b \otimes \mathcal{H}_R$$

Computing the Lipschitz constant for F and applying measure concentration and the union bound, we find a bound on pairwise overlap preservation in a state set S. [AK 22]

$$\Pr\left[\max_{|\psi_1\rangle,|\psi_2\rangle\in S}\left|\langle\psi_2|L^\dagger L|\psi_1\rangle - \langle\psi_2|\psi_1\rangle\right| \geq \sqrt{18}|B|^{-\gamma}\right] \leq 12\binom{|S|}{2}\exp\left(-\frac{|B|^{1-2\gamma}}{2}\right)$$

This holds for any $0 < \gamma < 1/2$ and $|B| \geq 4$.

Non-isometric interior code properties

$$\Pr\left[\max_{|\psi_1\rangle,|\psi_2\rangle\in S}\left|\langle\psi_2|L^\dagger L|\psi_1\rangle - \langle\psi_2|\psi_1\rangle\right| \geq \sqrt{18}|B|^{-\gamma}\right]$$
$$\leq 12\binom{|S|}{2}\exp\left(-\frac{|B|^{1-2\gamma}}{2}\right)$$

The meaning of this bound is the following. As long as $|S|$ is is parametrically smaller than $\exp(|B|^{1-2\gamma})$, the right hand side is small when $|B| \gg 1$.

$$|S| \leq \exp(|B|^{\alpha})\ , \quad \alpha < 1 - 2\gamma$$

When the right hand side is small, there is a very high probability that a randomly chosen encoding L will approximately preserve all pairwise overlaps in S to within an error $|B|^{-\gamma}$.

Non-isometric interior code properties

So the typical black hole interior code allows for the encoding of a discrete set of states which is in fact quite large. It is subexponentially large in the black hole Hilbert space dimension $|B|$.

$$|S| \leq \exp(|B|^{\alpha})$$

We can compare this to the total number of states in $\mathcal{H}_B$ itself. In the Haar metric, a projective Hilbert space (thought of as a unit sphere of normalized vectors) has finite volume. Discretizing this volume with ϵ-balls, a Hilbert space with dimension $|B|$ has roughly $\exp(|B| \log(1/\epsilon))$ states.

The interior code manages to make fairly efficient use of the microscopic Hilbert space while preserving semiclassical overlaps parametrically well.

Non-isometric interior code properties

Having estimated the number of states $|S|$ for which error correction could possibly be supported by the typical interior code, we turn to the problem of subsystem reconstruction.

The most interesting case is when W_b should be reconstructed as $\widetilde{W}_B(\psi)$ or $\widetilde{W}_R(\psi)$. We begin by defining functions $K_R(C)$ and $K_B(C)$ which are essentially the functions appearing in the associated decoupling criteria.

$$K_R(C) = \frac{\| \operatorname{Tr}_B(LW_b|\psi\rangle\langle\psi|W_b^\dagger L^\dagger) - \operatorname{Tr}_B(L|\psi\rangle\langle\psi|L^\dagger)\|_1}{\|LW_b|\psi\rangle\| + \|L|\psi\rangle\|}$$

The key point in the subsystem decoupling principle is to take the partial trace over the system upon which we want to study reconstruction.

Non-isometric interior code properties

Applying measure concentration gives deviation bounds for $K_B(C)$ and $K_R(C)$.

$$\Pr[K_R(C) \geq (2|R|/|B|)^{1/4} + |B|^{-\gamma}] \leq \exp(-|B|^{1-2\gamma}/36)$$

$$\Pr[K_B(C) \geq (2|B| \operatorname{Tr} \psi_b^2)^{1/4} + |B|^{-\gamma}] \leq \exp(-|B|^{1-2\gamma}/36)$$

When $K_R(C)$ is small, we have a reconstruction $\widetilde{W}_B(\psi)$, and when $K_B(C)$ is small, we have a reconstruction $\widetilde{W}_R(\psi)$.

This is consistent with entanglement wedge reconstruction, which states that reconstruction on B should be possible when $S_\psi(R) \ll \log|B|$ and reconstruction on R should be possible when $S_\psi(R) \gg \log|B|$.

Technical summary

After reviewing expectations for isometric and non-isometric codes, we constructed a non-isometric code in dilaton gravity and studied its properties.

Using measure concentration for the Gaussian distribution, we found global operator reconstruction can be supported for discrete state sets S which can be as large as subexponential $\exp(|B|^{\alpha})$ in the black hole Hilbert space dimension. Furthermore, we used the decoupling principle to show that subsystem reconstructions are possible in regimes consistent with entanglement wedge reconstruction.

Discussion

There are several interesting aspects of the non-isometric interior code.

Most mysteriously, the semiclassical Hilbert space of interior excitations is “re-quantized” or discretized in the fundamental description, and certain states in the semiclassical description have no fundamental meaning. The number of states that can be described semiclassically is still very large, but the structure of truly typical interior states remains unclear, and may involve firewalls.

The fundamental averaging we used to study typical properties is not expected to disrupt standard gravitational features like diffeomorphism invariance, but removing the averaging will require a better understanding of interior dynamics. [Blommaert, Usatyuk 21]

Discussion

The non-isometric error correcting structure represents the current limit of semiclassical physics in the black hole interior.

It would be very interesting to prove that the dilaton gravity code discussed here is asymptotically optimal, at least when applying measure concentration techniques.

If true, it would mean that there is a parametric difference between the total number of states (not just in a basis!) in a microcanonical window and states which have reasonable semiclassical descriptions (operator expectation values match up to e^{-1/G_N} corrections).

Whether the non-isometric structure can be extended to capture more of the microscopic Hilbert space is unclear.

Discussion

In a more practical direction, the improvement in state size $|S|$ of the gravitational code over Haar random tensor network methods is parametrically large by a polynomial power of 12. [Akers, Engelhardt, Harlow, Penington, Vardhan 22]

The theory of non-isometric codes deserves further study from a quantum information standpoint, and the fact that this improvement was possible may mean that they have the potential to be useful tools in the NISQ era of small quantum devices.